




ALEXANDER MANN

CALORIMETER-BASED TRIGGERS AT THE ATLAS DETECTOR FOR SEARCHES FOR SUPERSYMMETRY IN ZERO-LEPTON FINAL STATES

DISSERTATIONSSCHRIFT



The three graphical elements on the title page symbolize the three major projects documented in this thesis: The left one is derived from a comparison of predicted and observed rates for missing transverse energy triggers with different thresholds in 2010. The complete plot is shown in Figure 6.12. The one in the middle is derived from a two-dimensional turn-on curve of the combined jet + \cancel{E}_T trigger EF_j75_a4tc_EFFS_xe45_loose_noMu obtained from the Muons stream. The complete plot is shown in Figure 5.32. The one to the right is derived from exclusion limits as result of a search for Supersymmetry in zero-lepton final states, interpreted in a minimal supergravity model. A similar plot can be found in Figure 7.29.

Calorimeter-Based Triggers at the ATLAS Detector for Searches for Supersymmetry in Zero-Lepton Final States

Dissertation
zur Erlangung des mathematisch-naturwissenschaftlichen Doktorgrades
„Doctor rerum naturalium“
der Georg-August-Universität Göttingen

vorgelegt von
Alexander Mann
aus Göttingen

Göttingen, Januar 2012

Referent: Prof. Dr. Arnulf Quadt, Universität Göttingen
Korreferent: Dr. Carsten Hensel, Universität Göttingen
Tag der mündlichen Prüfung: 16. Februar 2012

Referenznummer: II.Physik-UniGö-Diss-2012/01

The existence of Supersymmetry is almost proven:
Nearly half of the particles of the Minimal Supersymmetric
Standard Model have been observed already.

–anonymous

Contents

1	Introduction	1
2	Theory	3
2.1	The Standard Model of Particle Physics	3
2.1.1	From Classical Mechanics to the Standard Model	3
2.1.2	Particle Content	9
2.1.3	Fundamental Interactions	11
2.1.4	Open Questions and Problems	15
2.2	Possibilities for Beyond Standard Model Physics	17
2.2.1	Technicolor	17
2.2.2	Extra Dimensions	18
2.3	Supersymmetry	19
2.3.1	General Idea	19
2.3.2	Motivations	20
2.3.3	Theoretical Background	20
2.3.4	The Minimal Supersymmetric Standard Model	22
2.3.5	Breaking of Supersymmetry	25
2.3.6	Higgs Mass Bounds and the MSSM	26
2.3.7	Important Supersymmetric Models	27
3	Experimental Setup	29
3.1	Conventions and Units	29
3.2	Collision Detectors in High-Energy Physics	30
3.2.1	Tracking Devices	30
3.2.2	Calorimeters	32
3.2.3	Muon Spectrometers	36
3.2.4	Luminosity	36
3.2.5	Pile-up	37
3.3	CERN and the LHC	39
3.3.1	L'Organisation Européenne pour la Recherche Nucléaire	39
3.3.2	The Large Hadron Collider	40
3.3.3	The 7 Detectors at the LHC	46
3.4	The ATLAS Detector	47
3.4.1	Detector Components	48
3.4.2	Inner Detector and Solenoid Magnet	48
3.4.3	Calorimeters	50
3.4.4	Muon Spectrometers and Toroidal Magnet System	53
3.4.5	Performance Goals	54
3.4.6	Minimum-Bias Trigger Scintillators	54

3.4.7	Forward Detectors	55
4	Data Acquisition and Data Processing	57
4.1	Purpose of a Trigger	57
4.2	Terminology of Triggers	58
4.3	The ATLAS Trigger System	59
4.3.1	General Architecture	59
4.3.2	\cancel{E}_T Trigger	62
4.3.3	Jet Trigger	63
4.3.4	Other Types of Triggers	64
4.3.5	Trigger Nomenclature	66
4.4	Event Reconstruction	66
4.4.1	Leptons	67
4.4.2	Jets	70
4.4.3	Missing Transverse Energy and Sum of Transverse Energy	73
4.4.4	Primary Vertices	75
4.4.5	Effective and Stransverse Mass	76
4.4.6	Runs and Good Runs Lists	77
4.4.7	Luminosity Measurements	78
4.5	Monte Carlo Generation	79
4.5.1	Hadronization and Jet Formation	80
4.5.2	The k -factor	81
4.5.3	Minimum Bias	82
5	Trigger Studies	83
5.1	Trigger Rates	83
5.1.1	Measurements	83
5.1.2	Benefits of Combined Triggers	85
5.2	Trigger Efficiencies: Methodology	91
5.2.1	Introduction	91
5.2.2	Terminology of Trigger Efficiencies	93
5.2.3	Estimation of Efficiencies	94
5.2.4	Use Cases and Trigger Strategies	95
5.2.5	Trigger Efficiencies and the β -Distribution	96
5.2.6	Overview of Methods	97
5.2.7	Bootstrapping	99
5.2.8	Computation of Trigger Efficiencies on Weighted Samples	107
5.3	Measurements of Trigger Efficiencies	109
5.3.1	Measurements of Jet + \cancel{E}_T Trigger Efficiencies in Monte Carlo	109
5.3.2	Measurements of Trigger Efficiencies in Data	114
5.4	Measurements of Trigger Efficiencies in Data Taken in 2010	115
5.4.1	Combined Jet + \cancel{E}_T Triggers in the Trigger Menu	115
5.4.2	Sample Triggers for Measuring Jet + \cancel{E}_T Trigger Efficiencies	116
5.4.3	Data Selection	117
5.4.4	Jet Triggers	118
5.4.5	Missing Transverse Energy Triggers	119
5.4.6	Combined Jet + \cancel{E}_T Triggers	121

5.5	Conclusions for 2010	129
5.6	Measurements of Trigger Efficiencies in Data Taken in 2011	130
5.6.1	Overview of Changes and New Jet + \cancel{E}_T Triggers	130
5.6.2	Jet Triggers	131
5.6.3	Combined Jet + \cancel{E}_T Triggers	137
5.7	Conclusions for 2011	141
6	A Model for \cancel{E}_T Trigger Rates	143
6.1	Motivation	143
6.2	Overview	146
6.3	Model Building	147
6.3.1	\cancel{E}_T Shape	147
6.3.2	$\sum E_T$ Shape	149
6.3.3	Trigger Rates	152
6.3.4	Summary	152
6.3.5	Implementation	153
6.4	Determination of Input Parameters	154
6.4.1	Number of Concurrent Collisions	154
6.4.2	$\sum E_T$ Slope	157
6.4.3	Missing Energy Resolution	158
6.5	Model Predictions	160
6.5.1	Performance in 2010	160
6.5.2	Performance in 2011	163
6.6	Discussion and Outlook	166
6.6.1	Extended \cancel{E}_T and $\sum E_T$ Spectra	166
6.6.2	Other Sources of \cancel{E}_T	169
6.6.3	Pile-up Noise Suppression	170
6.6.4	\cancel{E}_T Significance Triggers	173
7	Search for Supersymmetry	177
7.1	Introduction	177
7.1.1	The Zero-Lepton Channel	177
7.1.2	Outline of the Analysis	178
7.2	Integrated Luminosity, Data Samples and Trigger Selection	179
7.3	Event Selection	182
7.3.1	Jet Cleaning	184
7.3.2	Object Identification	185
7.3.3	Jet Multiplicity and Sub-Channel Selection	188
7.3.4	Signal Regions	189
7.4	Standard Model Backgrounds	190
7.5	Monte Carlo for Standard Model Backgrounds	192
7.5.1	Generators	192
7.5.2	Cross Sections and Background Normalization	192
7.5.3	Normalization of the QCD Background	195
7.5.4	Pile-up Effects in Monte Carlo	196
7.6	Signal Grids	198
7.6.1	mSUGRA Grid	198

7.6.2	MSSM Grid	200
7.6.3	Final-State Dependent Cross Sections	201
7.7	Statistical and Systematic Uncertainties	205
7.7.1	Statistical Uncertainty	206
7.7.2	Uncertainty on the Jet-Energy Scale	206
7.7.3	Uncertainty on the Jet-Energy Resolution	207
7.7.4	Uncertainty on the Integrated Luminosity	207
7.7.5	Uncertainty from Pile-Up	208
7.7.6	Theoretical Uncertainties on the Background Normalization	208
7.7.7	Theoretical Uncertainties on the Signal Normalization	209
7.8	Analysis Results in Numbers	212
7.9	Kinematic Distributions	214
7.10	Overview of Statistical Methods	217
7.10.1	Nomenclature	223
7.10.2	Frequentist and Bayesian Statistics	224
7.10.3	Methods for Setting Exclusion Limits	225
7.10.4	Comparison of the CL_s and PLR Methods	232
7.11	Interpretation of Analysis Results	234
7.11.1	Model-Agnostic Limits	234
7.11.2	Exclusion of Specific SUSY Scenarios	236
7.12	Summary and Conclusion	246
A	Appendix	A-1
A.1	Trigger Emulation	A-1
A.1.1	Implementation	A-2
A.1.2	Validation	A-3
A.2	Production of Trigger Rate Plots from COOL	A-6
A.3	Implementation of the CL_s and PLR Methods	A-8
A.3.1	CL_s Method	A-8
A.3.2	PLR Method	A-9
A.4	Computation of LHC Parameters	A-10
A.4.1	Required Magnetic Field Strength	A-10
A.4.2	Impact of the Center-of-Mass Energy on the Revolution Frequency	A-10
A.5	Analytic Forms of $\sum E_T$ and \cancel{E}_T	A-11
A.5.1	Distribution of $\sum E_T$ for Several Concurrent Interactions	A-11
A.5.2	Distribution of \cancel{E}_T from Resolution Effects	A-13
A.6	Supplementary Plots: \cancel{E}_T Model	A-14
A.7	Supplementary Plots: Trigger Efficiencies in Monte Carlo	A-17
A.8	Overview of Monte Carlo Datasets	A-21
A.8.1	Overview of Final States in the Monte Carlo Signal Grids	A-27
A.8.2	Impact of Unidentified Events	A-27
A.9	Jet Spectrum in Pythia QCD Samples	A-28
A.10	Computing	A-29
A.10.1	ROOT	A-29
A.10.2	The Athena Framework	A-31
A.10.3	Data formats	A-32
A.10.4	Nomenclature of Dataset Names in ATLAS	A-33

A.10.5 Databases in ATLAS	A-34
Image Credits	A-37
Bibliography	A-39
Web Citations	A-53
Acronyms	A-55
Index	A-57

Abstract

This thesis consists of three closely related parts. An analysis of data recorded by the ATLAS detector in 2010 in proton-proton collisions at a center-of-mass energy of $\sqrt{s} = 7$ TeV with an integrated luminosity of 33.4 pb^{-1} is performed, searching for supersymmetric final states containing jets and missing transverse energy and no electrons or muons (zero-lepton channel). No excess over the Standard Model background expectation is observed. Using the CL_s and PLR methods, exclusion limits are set in a minimal supergravity model with $\tan \beta = 3$, $A_0 = 0$ and $\mu > 0$ and in a simplified supergravity model. These considerably extend the excluded parameter ranges from earlier experiments. The analysis validates the official analysis carried out within the ATLAS Supersymmetry group and additionally takes into account the uncertainty from pile-up effects.

The rates and efficiencies of triggers based on combined signatures with jets plus missing transverse energy in ATLAS are studied, which are the primary triggers for the search for Supersymmetry in the zero-lepton channel. Methods to measure the efficiencies on data are tested and optimized to obtain sufficient statistics. For the measurement of the efficiencies in data collected in 2010 and 2011, the bootstrap method is applied to correct for the sample trigger bias. Different sample triggers based on jets and missing transverse energy are compared and their efficiencies are measured. A reweighting approach is studied and used to correct for the bias from the propagation of the uncertainties in the bootstrap method. The resulting efficiency estimates for the primary triggers allow to determine the onset of the plateau of the two-dimensional turn-on curves and are input to the official analyses in the ATLAS Supersymmetry group in 2010 and 2011.

A universal model is developed to describe the contribution of fake missing transverse energy from resolution effects to the rates of missing transverse energy triggers as function of the level of in-time pile-up, i. e. the number of concurrent proton-proton interactions. The input parameters to tune the model to the properties of the ATLAS triggers are determined, and the model predictions are compared to measurements of trigger rates in ATLAS. Good agreement is found for missing transverse energy triggers with low thresholds for which the rates are dominated by resolution effects, whereas the rates for higher thresholds are underestimated due to additional sources of fake and real missing transverse energy which are not incorporated in the model.

Kurzfassung

Diese Arbeit besteht aus drei eng zusammenhängenden Teilen. Es wird eine Analyse von Daten des ATLAS-Detektors aus dem Jahr 2010 durchgeführt, die in Proton-Proton-Kollisionen bei einer Schwerpunktsenergie von $\sqrt{s} = 7$ TeV aufgezeichnet wurden und eine integrierte Luminosität von 33.4 pb^{-1} umfassen, mit dem Ziel, supersymmetrische Endzustände mit Jets, fehlender transversaler Energie und ohne Elektronen und Myonen nachzuweisen. Dabei wird kein über den erwarteten Standardmodelluntergrund hinausgehender Überschuss

an Ereignissen beobachtet. Auf die Massenparameter von zwei supersymmetrischen Szenarien mit gravitationsvermittelter Symmetriebrechung, von mSUGRA mit $\tan\beta = 3$, $A_0 = 0$ und $\mu > 0$ und von einem vereinfachten Modell, werden Ausschlussgrenzen mithilfe der CL_s - und der PLR-Methode gesetzt, welche die Ausschlussgrenzen früherer Experimente deutlich übertreffen. Die Analyse validiert die offiziellen Ergebnisse der ATLAS-Supersymmetrie-Gruppe und berücksichtigt zusätzlich Unsicherheiten infolge der Überlagerung mehrerer gleichzeitiger Proton-Proton-Wechselwirkungen.

Die Raten und Effizienzen der im ATLAS-Detektor verwendeten Trigger, die auf kombinierten Signaturen mit Jets und fehlender transversaler Energie beruhen, werden untersucht. Diese Trigger sind die Primärtrigger für die vorgestellte Suche nach Supersymmetrie im Null-Lepton-Kanal. Methoden zur Bestimmung der Triggereffizienzen auf Daten werden getestet und auf größtmögliche Statistik hin optimiert. Für die Messung der Effizienzen auf Daten aus 2010 und 2011 wird die Bootstrap-Methode verwendet, um den Bias der Sample-Trigger zu korrigieren. Verschiedene, auf Jets und fehlender transversaler Energie basierende Sample-Trigger werden verglichen und ihre Effizienzen werden bestimmt. Der Bias, der der Fehlerfortpflanzung in der Bootstrap-Methode inhärent ist, wird untersucht und korrigiert. Die hierdurch mögliche Bestimmung der Lage des Plateaus in den erhaltenen zweidimensionalen Triggereffizienzkurven ist Bestandteil der offiziellen Analysen der ATLAS-Supersymmetrie-Gruppe 2010 und 2011.

Um den Anteil der durch Auflösungseffekte hervorgerufenen scheinbaren fehlenden transversalen Energie an den Triggerraten als Funktion der Anzahl der gleichzeitigen Proton-Proton-Wechselwirkungen beschreiben zu können, wird ein allgemeingültiges Modell entwickelt. Die Modellparameter werden an die Eigenschaften des ATLAS-Triggers angepasst und die erhaltenen Vorhersagen mit Messungen der entsprechenden Raten des ATLAS-Triggers verglichen. Eine gute Übereinstimmung zeigt sich für Trigger mit niedrigen Schwellwerten, deren Raten von Auflösungseffekten dominiert werden, während die Triggerraten für höhere Schwellwerte unterschätzt werden. Dies wird begründet mit dem Auftreten zusätzlicher Quellen von echter und scheinbarer fehlender transversaler Energie, die im Modell nicht abgebildet werden.

Curriculum Vitae

Personal Details

Name	Alexander Mann
Date of Birth	26 December 1981
Place of Birth	Göttingen, Deutschland
Nationality	German

Education

2008 – 2012	PhD (in Physics), Georg-August-Universität Göttingen
2004 – 2008	Diplom (in Physics), Georg-August-Universität Göttingen
2002 – 2004	Vordiplom (in Physics), Georg-August-Universität Göttingen
2001 – 2002	Civilian Service, Krankenhaus Göttingen-Weende e.V.
1995 – 2001	Abitur, Felix-Klein-Gymnasium Göttingen

Acknowledgements

It is my pleasure to thank the many people who made this thesis possible.

In particular, I wish to express my gratitude to my supervisor, Dr. Carsten Hensel, head of our Emmy Noether young researcher group at Göttingen, which focuses on the search for supersymmetric new physics models. He gave me the opportunity to work on the very interesting, experimentally challenging and theoretically appealing topic of Supersymmetry. I'd like to thank him for his continued encouragement and invaluable suggestions during my work.

I would like to thank Prof. Dr. Arnulf Quadt, chair of the II. Institute of Physics at Göttingen, as co-referee of my thesis and member of my thesis committee. He managed, with great ambitions, to bring particle physics to this city renowned for its excellent physicists. Moreover, I thank Prof. Dr. Detlev Buchholz as the third member of my thesis committee.

I am very grateful to Dr. Julien Morel. It has been a great pleasure to work with him during my time at CERN and afterwards.

The results presented in this thesis are based on the efforts of a huge number of people, who designed, built and operate the LHC machine and the ATLAS detector. I wish to extend my gratitude to everybody in our collaboration who contributed to making the LHC and ATLAS projects a great success, and in particular to my colleagues for the friendly and warm atmosphere in our group at Göttingen and the working groups I was allowed to be part of.

For many useful comments and remarks I would like to thank everybody who was so kind to read one of the various draft versions of parts of my thesis.

I would like to express my deep and sincere gratitude to all my friends who accompanied me on my way, for the fun we had together, their help and inspiration.

Lastly, and most importantly, I wish to thank my parents for their support.

This thesis has been supported by the German Research Foundation (DFG) and the Federal Ministry of Education and Research (BMBF). The financial support is gratefully acknowledged.

1 Introduction

Elementary particle physics deals with the building blocks of matter on the smallest scale of size. It describes the constituents of matter and how they interact with one another. The theoretical framework of elementary particle physics is the Standard Model. The Standard Model is a very satisfactory theory of the electromagnetic, weak and strong interactions of particles, which allows to make predictions and to compute quantities like cross sections and decay rates at high precision and in agreement with experimental findings. However, there is a number of questions it cannot answer, and it is believed to be the low-energy limit of a deeper theory which would at some point also incorporate a quantum theory of the fourth fundamental interaction, gravity. Moreover, the Standard Model has theoretical imperfections. One of these is the fine-tuning or hierarchy problem, which requires a remarkably precise cancellation of physical parameters in order to protect the Higgs mass from quadratic divergences.

To address these shortcomings, an extension of the Standard Model is needed. A very promising candidate is Supersymmetry. Supersymmetric extensions of the Standard Model introduce a new symmetry which unifies fermions, the constituents of matter, and bosons, the particles which mediate the forces between them. This not only solves the fine-tuning problem in a very elegant way, but there are also other quantitative indications in favor of Supersymmetry, including gauge coupling unification and the existence of natural candidates for cold dark matter. The introduction of Supersymmetry doubles the particle content of the Standard Model: every fermion is accompanied by a bosonic partner and vice versa. Inevitably, this leads to the question how all of these new particles could have evaded observation so far. A possible explanation is that Supersymmetry is a broken symmetry so that the supersymmetric partners are much heavier than the Standard Model particles. They can then only be produced at collider experiments which reach sufficiently high energies.

All previous searches for evidence of Supersymmetry have yet been unsuccessful, but with the recent start-up of the Large Hadron Collider at CERN, with its unprecedented center-of-mass energy of 7 TeV, it is now possible to search for new physics in an energy regime which has never been accessible before at a particle collider. The confrontation of theory with the upcoming new experimental results will either confirm or disprove the expectations of theorists and may reshape the foundations of particle physics. In this thesis, data will be evaluated which is collected with the ATLAS detector, one of the two large multi-purpose detectors at the Large Hadron Collider. The ATLAS detector is a device to detect the particles which are produced in high-energy collisions of protons. Comparing its data output to simulations will allow to draw conclusions on what has happened in the collision events and to search for signatures of the production of supersymmetric particles and their decay products.

In many supersymmetric models, the lightest supersymmetric particle has a large mass and evades detection due to its weak interactions, leading to an imbalance in the sum of all measured particle momenta. Therefore, a typical supersymmetric signature involves a considerable amount of missing transverse energy from the lightest supersymmetric par-

ticle, plus high-energetic jets from decays of heavy supersymmetric particles. One of the challenges in a collider experiment is to spot these signatures among the collision events, the overwhelming majority of which contains only well-known Standard Model processes. In order to collect a sufficient number of the rare new physics events, very high collision rates need to be employed. This yields an enormous data output, many orders of magnitudes larger than what can be stored. Hence, a dedicated detector component, the trigger, is needed, to do fast online measurements and to decide automatically in real-time whether or not a given collision event is interesting, and thus kept, or discarded. As in the resulting dataset only events can be examined which have been accepted by the trigger, whereas everything else is irretrievably lost, it must be ensured that the trigger actually does what it is supposed to do. A vital part of data analysis are therefore trigger studies. These are concerned with the efficiency of the trigger selection, but also with the trigger rates and their dependencies.

The final goal and unifying element of the work documented in this thesis is the search for Supersymmetry, which is carried out on data taken with the ATLAS detector. In accordance with the explanations from above, the analysis of this data is complemented by detailed studies of the efficiency of the trigger selection which is used to collect the dataset being analyzed. A further aspect under study in this thesis is the output rate of the trigger for a certain type of collision events with missing transverse energy.

The outline of this thesis is the following. In Chapter 2, the theoretical background is given. Both the Standard Model of Particle Physics and the fundamentals of Supersymmetry are covered, and the resulting particle content and the phenomenological implications are introduced. Afterwards, in Chapter 3, the experimental setup is presented, in particular the collider and the detector which are used to produce the particle collisions and to record the data evaluated in this thesis. As studies concerned with the trigger system constitute an important part, Chapter 4 will be dedicated to the trigger and data-acquisition system of the ATLAS detector. In this context, also the processing of data and event reconstruction are described, thereby complementing the previous description of the hardware of the ATLAS detector. After these introductory texts on the setting, Chapter 5 comprises the studies of trigger rates and trigger efficiencies, which together quantify the selection quality of the trigger system. The focus of these studies lies on the trigger selection which is based on a combined signature of jets and missing transverse energy. Both parts, the jets and the missing transverse energy, are reconstructed and computed from calorimeter measurements. In addition to the rates, the efficiencies and the methods to measure those, the benefits of using this combined signature for the selection of collision events by the trigger are discussed. The trigger rates based on measurements of the missing transverse energy exhibit a non-trivial behavior as function of the activity in the collision events. A simple model to explain this behavior is developed in Chapter 6, and its predictions are compared to the measured rates of the ATLAS trigger. In the final Chapter 7, the search for Supersymmetry is presented. After applying a suitable selection of events, a comparison of event counts in data and simulations is employed to set limits on the parameters of two supersymmetric models. Two different methods are used to compute these limits and the results are compared to results from earlier experiments.

2 Theory

This chapter describes the theoretical foundations and motivations of the studies which are carried out in this thesis. First of all, a review of the Standard Model of Particle Physics, the theoretical framework that comprises the current understanding of experimental results in elementary particle physics, will be given in three steps. In the first step, the theoretical background of the Standard Model is provided. Afterwards, the particle content of the Standard Model is summarized, and the fundamental interactions among the particles and their phenomenology are described in the third step. The ensuing discussion of open questions and problems will make clear the need for an extension of the Standard Model. Two important alternative extensions, technicolor and extra dimensions, will be briefly discussed, before Supersymmetry is introduced as a theoretically appealing model, which can provide a number of answers. Supersymmetry is the physics model which is searched for in the analysis of collision data in this thesis. The theory of Supersymmetry, the particle content of the Minimal Supersymmetric extension of the Standard Model and its implications for the phenomenology form the last part of this chapter.

2.1 The Standard Model of Particle Physics

2.1.1 From Classical Mechanics to the Standard Model

The following summary of the theory of the Standard Model is given in terms of gauge theories. It starts with principles of classical mechanics and leads up to the Standard Model, including the Higgs mechanism to generate the masses of the particles consistent with observation. No derivations will be given for the sake of brevity. It is a very fast-paced introduction^a, needed as a foundation for the following motivation and explanation of the supersymmetric extensions of the Standard Model.

Lagrangians in Classical Mechanics

In classical mechanics, the Lagrangian formalism is based on the Lagrangian function L (often called Lagrangian for short), which is comprised of the kinetic energy T and the potential energy U of the system,

$$L = T - U. \quad (2.1)$$

Using Hamilton's principle, also known as the principle of stationary action, the variational principle

$$\delta S[\vec{q}, \delta\vec{q}] \stackrel{!}{=} 0 \quad (2.2)$$

^a A very clear presentation can be found in [8], on which this summary is based in large parts, together with [9]. Quantum Field Theory is covered e. g. in [13, 15].

is applied to the action S ,

$$S[\vec{q}] = \int_{t_i}^{t_f} L(\vec{q}, \dot{\vec{q}}, t) dt. \quad (2.3)$$

$S[\vec{q}]$ is the integral over the Lagrangian and thus a functional of the generalized coordinates q_i spanning the configuration space of the mechanical system. Conventionally, their derivatives with respect to time t are written as \dot{q}_i . Assuming the variation $\delta\vec{q}$ to vanish at the integration boundaries t_i and t_f , the variational principle gives the Euler-Lagrange equations

$$\frac{d}{dt} \left(\frac{\partial L}{\partial \dot{q}_i} \right) = \frac{\partial L}{\partial q_i} \quad \forall i. \quad (2.4)$$

They relate the generalized coordinates q_i and their derivatives \dot{q}_i and yield the equations of motion.

In field theory, the point-like particles of classical mechanics with trajectories $\vec{x}_i(t)$ in three-dimensional space are replaced by fields which are functions $\phi_i(\vec{x}, t)$ of position and time. In this generalization, the Lagrangian becomes a Lagrangian density \mathcal{L} , which, integrated over space, gives the Lagrangian function. Therefore, integration over all four spacetime coordinates x^μ gives the action

$$S[\phi_i] = \int \mathcal{L}(\phi_i, \partial_\mu \phi_i, x_\mu) d^4x. \quad (2.5)$$

The Lagrangian density \mathcal{L} fully describes the physics of the theory. This is why theoretical models are often specified by stating their Lagrangian density, which is also called Lagrangian for short. The Euler-Lagrange equations become

$$\partial_\mu \left(\frac{\partial \mathcal{L}}{\partial (\partial_\mu \phi_i)} \right) = \frac{\partial \mathcal{L}}{\partial \phi_i}, \quad (2.6)$$

where space and time are now treated on equal footing, as befits for a relativistic theory. The Lagrangian (density) in field theory is constructed axiomatically to give the desired field equations when evaluating the Euler-Lagrange equations. If the constructed \mathcal{L} describes a Lorentz scalar field, the resulting equations of motion are automatically covariant.

Some important Lagrangians for free (non-interacting) fields commonly met in elementary particle physics describe particles with spin 0, $1/2$ and 1. For a single scalar (spin-0) field ϕ describing a free particle of mass m , the Lagrangian is given by

$$\mathcal{L}_{\text{Klein-Gordon}} := \frac{1}{2} (\partial_\mu \phi) (\partial^\mu \phi) - \frac{1}{2} m^2 \phi^2, \quad (2.7)$$

yielding the Klein-Gordon equation

$$\partial_\mu \partial^\mu \phi + m^2 \phi = 0 \quad (2.8)$$

by Equation (2.6). The Dirac Lagrangian for a spinor (spin- $1/2$) field ψ describing a particle of mass m is

$$\mathcal{L}_{\text{Dirac}} := i \bar{\psi} \gamma^\mu \partial_\mu \psi - m \bar{\psi} \psi, \quad (2.9)$$

which leads to the Dirac equation

$$i\gamma^\mu \partial_\mu \psi - m\psi = 0. \quad (2.10)$$

Finally, the Lagrangian of a vector (spin-1) field A^μ ,

$$\mathcal{L}_{\text{Proca}} := -\frac{1}{16\pi} F^{\mu\nu} F_{\mu\nu} + \frac{m^2}{8\pi} A^\nu A_\nu \quad (2.11)$$

with

$$F^{\mu\nu} := \partial^\mu A^\nu - \partial^\nu A^\mu, \quad (2.12)$$

leads to the corresponding Proca equation for a particle of spin 1 and again mass m ,

$$\partial_\mu F^{\mu\nu} + m^2 A^\nu = 0. \quad (2.13)$$

In the limit $m \rightarrow 0$, this yields Maxwell's equations for the electromagnetic field in empty space.

Symmetries and Noether's Theorem

Starting from the above Lagrangians for free fields, the principle of local gauge invariance, i. e. the symmetry of the free Lagrangian under local phase transformations, will lead to the appearance of additional fields mediating fundamental interactions. Symmetries in general play an important role in physics. In this context, a symmetry is understood to be an operation that leaves a system invariant, i. e. carries it into a configuration which cannot be distinguished from the original one. The properties of symmetry operations (closure, associativity and existence of the identity and inverse) are the defining properties of a mathematical group. This motivates the importance of group theory for physics, and in particular Lie groups in case of continuous symmetries. Every group can be represented by a group of matrices, which can then be used to carry out concrete computations. The Poincaré group [7] includes all transformations of type

$$x'^\mu = \Lambda^\mu_\nu x^\nu + a^\mu, \quad (2.14)$$

i. e. translations by a vector a^μ and Lorentz transformations Λ^μ_ν . Poincaré transformations have ten independent parameters, four for the translation, three for rotations in space and three describing Lorentz boosts. They leave the distance between two points in Minkowski space invariant.

Noether's (first) Theorem, published 1918 by Emmy Noether [85], states that every continuous symmetry of nature yields a conservation law, and every conservation law reflects an underlying symmetry. In a more precise formulation, if a system (or its Lagrangian density) exhibits a continuous symmetry of the action, then there are corresponding quantities whose values are conserved in time. Noether's Theorem connects invariance under translations in space and time with the conservation of momentum and energy, respectively, and rotational invariance with the conservation of angular momentum. In addition to these, there are internal symmetries in elementary particle physics, which are independent of the spacetime coordinates, i. e. the transformations commute with the spacetime components of the wave

functions. The simplest example for such an internal symmetry is the invariance of a wave function ψ under a global phase shift $\theta \in \mathbb{R}$,

$$\psi \rightarrow \exp(i\theta) \psi, \quad (2.15)$$

which also holds for the Dirac Lagrangian in Equation (2.9). The family of transformations described by Equation (2.15) form a unitary Abelian (commutative) group $U(1)$, and the corresponding invariance, a global gauge invariance, through Noether's theorem gives rise to the conservation of the electromagnetic charge.

Quantum Electrodynamics

In contrast to the global phase shift, the Dirac Lagrangian is no longer invariant under a local phase transformation, where the phase depends on the spacetime coordinates x^μ ,

$$\psi \rightarrow \exp(i\theta(x)) \psi. \quad (2.16)$$

The local gauge invariance of the Lagrangian can be restored by replacing the normal derivative ∂_μ by the covariant derivative

$$\mathcal{D}_\mu := \partial_\mu + iqA_\mu, \quad (2.17)$$

which entails the introduction of an additional gauge field A^μ . A constant factor q has been pulled out. This vector field requires that an additional free field term given by Equation (2.11) be included in the Lagrangian, describing the kinetic energy of the field. As $A^\nu A_\nu$ is not invariant under the local phase transformation (2.16), from Equation (2.11) it can be seen that the particle described by A^μ must be massless. A^μ can be identified with the electromagnetic potential, i. e. the photon field, and transforms as

$$A_\mu \rightarrow A_\mu - \frac{1}{q} \partial_\mu \theta. \quad (2.18)$$

Imposing the local phase invariance on the free Dirac Lagrangian thus leads to the Lagrangian of Quantum Electrodynamics (QED),

$$\mathcal{L}_{\text{QED}} = \bar{\psi} \left[i\gamma^\mu \partial_\mu - m \right] \psi - q \bar{\psi} \gamma^\mu A_\mu \psi - \frac{1}{16\pi} F^{\mu\nu} F_{\mu\nu}, \quad (2.19)$$

where q is the charge of the Dirac particle.

Quantum Chromodynamics

The same strategy can be applied to derive the Lagrangian of Quantum Chromodynamics (QCD), starting from three color fields, which are combined into one ψ vector to simplify the notation, $\psi = (\psi_r, \psi_g, \psi_b)^T$. Invariance under local $SU(3)$ gauge transformations,

$$\psi \rightarrow S\psi \quad \text{with} \quad S := \exp\left(i\vec{a}(x) \cdot \vec{\lambda}\right), \quad (2.20)$$

is imposed on the Dirac Lagrangian, with eight real numbers a_1, \dots, a_8 as parameters and the Gell-Mann 3×3 matrices $\lambda_1, \dots, \lambda_8$ as generators of the transformation. The invariance can again be restored by introducing the covariant derivative

$$\mathcal{D}_\mu = \partial_\mu + ig\vec{\lambda} \cdot \vec{A}_\mu, \quad (2.21)$$

which now involves eight gauge fields \vec{A}_μ corresponding to the gluons and one new coupling constant g . Due to the non-commutativity of $SU(3)$, the transformation behavior of \vec{A}_μ is different. For infinitesimal transformations it is given by

$$\vec{A}_\mu \rightarrow \vec{A}_\mu - \frac{1}{g} \partial_\mu \vec{a} - 2f_{abc} \vec{e}_a a_b A_\mu^c, \quad (2.22)$$

with the (real) structure constants f_{abc} of $SU(3)$, \vec{e}_i being the i th unit vector and summing over repeated indices implied. This leads to the QCD Lagrangian, which describes quarks and gluons and their interactions:

$$\mathcal{L}_{\text{QCD}} = \bar{\psi} \left[i\gamma^\mu \partial_\mu - m \right] \psi - g \left(\bar{\psi} \gamma^\mu \vec{\lambda} \psi \right) \cdot \vec{A}_\mu - \frac{1}{16\pi} \vec{F}^{\mu\nu} \cdot \vec{F}_{\mu\nu}. \quad (2.23)$$

Like the photon, the eight gluon fields have to be massless because the Proca mass term proportional to $\vec{A}^\nu \cdot \vec{A}_\nu$ is excluded by local gauge invariance. Due to the different form of Equation (2.22), also $\vec{F}^{\mu\nu}$ takes a slightly more complicated form than in Equation (2.12),

$$\vec{F}^{\mu\nu} := \partial^\mu \vec{A}^\nu - \partial^\nu \vec{A}^\mu - 2g f_{abc} \vec{e}_a A_b^\mu A_c^\nu. \quad (2.24)$$

This has the important consequence that the kinetic energy term in Equation (2.23) is not purely kinetic, but includes (three- and four-point) self-interaction terms for the gluons. Gluons can thus not only interact with quarks, but also among themselves.

Higgs Mechanism

The derivation of the QED and QCD Lagrangian is straightforward, using the Dirac Lagrangian for the participating particles and imposing the principle of local gauge invariance. In this procedure, the gauge fields turn out to be massless, which is adequate for the gluons and the photon, but the mediators of the weak interaction, the vector bosons W^\pm and Z^0 , are massive. A possible mechanism to allow for massive gauge fields is spontaneous symmetry-breaking, or, specifically, the Higgs mechanism, which is the spontaneous symmetry-breaking of a local gauge invariance^b.

The main difference is that, unlike for the Lagrangians formulated above, the potential energy now is chosen such that the ground state (or vacuum state) no longer is the trivial one of vanishing fields $\phi = 0$. A common choice for such a potential is the ‘‘Mexican hat’’ potential. For a complex scalar field ϕ , it is given by

$$V(\phi) = -\frac{\rho^2}{2} \phi^* \phi + \frac{\lambda^2}{4} (\phi^* \phi)^2, \quad (2.25)$$

with $\rho^2 > 0$ and $\lambda^2 > 0$. $V(\phi)$ has a circle of minima in the complex ϕ plane at

$$\phi^* \phi = \frac{\rho^2}{\lambda^2}, \quad (2.26)$$

which is easier to recognize writing $\phi := \phi_1 + i\phi_2$, thus $\phi^* \phi = \phi_1^2 + \phi_2^2$. The corresponding Lagrangian

$$\mathcal{L}_{\text{Mexican}} = \frac{1}{2} \left(\partial_\mu \phi \right)^* \left(\partial_\mu \phi \right) + \frac{\rho^2}{2} \phi^* \phi - \frac{\lambda^2}{4} (\phi^* \phi)^2 \quad (2.27)$$

^b It is not possible to put in the mass terms by hand, thereby breaking local gauge invariance, because the resulting theory would be unrenormalizable.

possesses a global $U(1)$ symmetry $\phi \rightarrow \exp(i\theta)\phi$. In order to find the masses of the participating fields, the Lagrangian needs to be rewritten in different field variables and to be expanded about a fixed vacuum ground state. Choosing a particular vacuum state breaks the symmetry and allows to identify the mass terms of the fields. This is called spontaneous symmetry-breaking. Nevertheless, both Lagrangians, before and after rewriting in different field variables, describe the same physical system, and the manifest symmetry only gets hidden by rewriting the Lagrangian. Moreover, according to a theorem by 't Hooft [108], the theory remains renormalizable.

Breaking the global symmetry of the Lagrangian (2.27) is not sufficient, because the Goldstone theorem states that massless scalars (called Goldstone bosons) must occur whenever a continuous global symmetry is spontaneously broken. No massless scalars have been found in nature, and thus breaking the global symmetry does not describe reality. Instead, the Lagrangian (2.27) must be made invariant under a local $U(1)$ gauge symmetry, following the recipe from above for the derivation of the QED and QCD Lagrangian. This implies the introduction of an additional gauge field. Again, a vacuum ground state is fixed and the Lagrangian rewritten in different field variables in an expansion around the chosen ground state. The local $U(1)$ symmetry can be exploited to pick a particular gauge, in which it becomes apparent that the Lagrangian now describes two interacting massive fields, one of which is a massive vector gauge boson and the other a massive scalar, which is named the Higgs particle. The massless Goldstone boson has been turned into the longitudinal polarization of the massive gauge particle [61]. This mechanism is called the Higgs mechanism [9]. Note that the precise form of the Higgs potential in Equation (2.25) is unknown, but it has to be quartic in the fields for the theory to be renormalizable. The Higgs particle has not yet been found, but the Higgs mechanism, which allows to avoid the appearance of massless particles, is believed to be the true mechanism for generating masses in the Standard Model.

Weinberg-Salam Model

To obtain three massive gauge fields for the charged W^\pm bosons and the neutral Z^0 boson in the Standard Model, while retaining a massless photon field, the electroweak Lagrangian is obtained from Equation (2.19) by replacing the $U(1)$ electromagnetic interaction with two interaction terms. They contain the generator T of the $SU(2)_L$ gauge group of the weak isospin and the generator Y of the $U(1)_Y$ gauge group of the weak hypercharge. The underlying symmetry group of the electroweak interactions is thus $SU(2) \otimes U(1)$. The interaction terms involve an isospin triplet of weak vector fields \vec{W}^μ and a fourth single vector field B^μ coupling to the weak isospin and hypercharge current, respectively [9]. These fields mix to give the physical fields W^\pm , Z^0 and the photon γ .

The Higgs mechanism is applied by introducing four scalar fields and choosing a suitable vacuum expectation value which leaves the $U(1)_{\text{em}}$ symmetry with generator $Q = T_3 + Y$ unbroken, so that the photon remains massless. This gives the Weinberg-Salam or minimal model of electroweak interactions. It includes one neutral scalar Higgs field as a remainder of the four scalar fields originally introduced as the Higgs isospin doublet. The other three degrees of freedom have been absorbed by the now massive weak gauge bosons. The Higgs doublet can also be used to give masses to all fermions appearing in the Standard Model, quarks and leptons. Mass terms like $-m_e(\bar{e}_R e_L + \bar{e}_L e_R)$ for the left- and right-chiral electron fields $e_{L,R}$, which would otherwise break gauge-invariance, can be included in the Lagrangian in a gauge-invariant way through the Higgs mechanism [9]. This leads to addi-

tional terms representing couplings of the fermions to the Higgs field H itself, e. g.

$$-\frac{m_e}{v}(\bar{e}_R e_L + \bar{e}_L e_R)H = -\frac{m_e}{v}\bar{e}eH, \quad (2.28)$$

where the vacuum expectation value of the Higgs field $v = 246 \text{ GeV}$ can be calculated from the mass of the W [8]. These couplings are chosen to be proportional to the mass of the fermion^c. The mass of the Higgs m_H itself is given by

$$m_H = v\sqrt{\frac{\lambda^2}{2}} \quad (2.29)$$

at tree-level, where λ^2 is the strength of the Higgs self-interaction in Equation (2.25). Like all other particle masses, the mass of the Higgs boson is therefore not predicted by the theory. The resulting Lagrangians after spontaneous symmetry breaking are quite lengthy already for the simple model with the Lagrangian given in Equation (2.27), and therefore are not spelled out here. The complete Lagrangian of the Standard Model, before and after spontaneous symmetry-breaking, is given in [68].

The above has lead naturally to the Higgs mechanism as a procedure for generating the masses of the particles which appear in the Standard Model. In the next section, these particles and the fundamental interactions will be described. Afterwards, it will be necessary to discuss some of the problems with the Higgs mechanism. This leads to the need for an extension of the Standard Model and to the introduction of Supersymmetry.

2.1.2 Particle Content

Depending on their spin quantum number, particles belong to one of two distinct groups: Fermions with spin $1/2$ or in general half-numbered spin, which are the building blocks of matter, and bosons with integer spin, which are the mediators of the interactions in the Standard Model of particle physics. There are two groups of fermions within the Standard Model, quarks and leptons, which all have spin $1/2$. Six quarks q are known to exist, which bear the names up u , down d , charm c , strange s , top t and bottom (or beauty) b . These six flavors are organized in three generations (also called families) in the order of increasing masses. Each generation has a quark with charge $+2/3$ (u, c, t) and one with charge $-1/3$ (d, s, b) in units of the electron charge magnitude e . Quarks carry color charge and can therefore participate in strong interactions, in contrast to leptons. There are six leptons l in the Standard Model, which are also grouped into three generations, each with one charged and one neutral particle. The charged leptons are, with increasing mass, the electron e , muon μ and tau τ , which have negative charge of $-e$. Every charged lepton is associated with a neutral partner, the electron neutrino ν_e , muon neutrino ν_μ or tau neutrino ν_τ , respectively. The twelve fermions in the Standard Model are complemented by twelve antiparticle partners with opposite charge^d, which reflects the fact that the Dirac equation has positive and negative energy solutions. Two additive conserved quantities are associated with quarks and leptons. The baryon number B is $1/3$ for quarks, $-1/3$ for antiquarks and 0 for all other particles. The lepton number L is 1 for leptons, -1 for antileptons and 0 for all other particles. All observations are consistent with B and L conservation in Standard Model processes. Except for neutrino mixing, L is conserved for the three lepton families separately.

^c This will play an important role later on when discussing problems of the Standard Model.

^d Neutrino and antineutrino are both neutral, of course.

Field	Spin	$SU(3)_C$	$SU(2)_L$	$U(1)_Y$	T_3
$q_L = \begin{pmatrix} u_L \\ d_L \end{pmatrix}, \begin{pmatrix} c_L \\ s_L \end{pmatrix}, \begin{pmatrix} t_L \\ b_L \end{pmatrix}$	$\frac{1}{2}$	3	2	$+\frac{1}{6}$	$+\frac{1}{2}$ $-\frac{1}{2}$
$(u_R)^c = \bar{u}_L, (c_R)^c, (t_R)^c$	$\frac{1}{2}$	$\bar{\mathbf{3}}$	1	$-\frac{2}{3}$	0
$(d_R)^c = \bar{d}_L, (s_R)^c, (b_R)^c$	$\frac{1}{2}$	$\bar{\mathbf{3}}$	1	$+\frac{1}{3}$	0
$l_L = \begin{pmatrix} \nu_{eL} \\ e_L^- \end{pmatrix}, \begin{pmatrix} \nu_{\mu L} \\ \mu_L^- \end{pmatrix}, \begin{pmatrix} \nu_{\tau L} \\ \tau_L^- \end{pmatrix}$	$\frac{1}{2}$	1	2	$-\frac{1}{2}$	$+\frac{1}{2}$ $-\frac{1}{2}$
$(e_R)^c = \bar{e}_L, (\mu_R)^c, (\tau_R)^c$	$\frac{1}{2}$	1	1	+1	0
g	1	8	1	0	0
$W = \begin{pmatrix} W^+ \\ W^0 \\ W^- \end{pmatrix}$	1	1	3	0	+1 0 -1
B^0	1	1	1	0	0

Table 2.1: Field content of the Standard Model of particle physics [1, 16]. T_3 is the third component of the weak isospin corresponding to the $SU(2)_L$ symmetry. For $U(1)_Y$, the weak hypercharge Y quantum number is given, for $SU(3)_C$ and $SU(2)_L$, the dimension of the representation. The Higgs field is not included.

The bosons in the Standard Model are the massless neutral photon γ , the heavy charged W^\pm bosons and the heavy neutral Z^0 boson, and the eight massless neutral gluons g . All of them have spin 1 and are called gauge bosons because they arise from gauge symmetries, as discussed in the previous section. In addition, there is also the hypothesized Higgs boson H , which has spin 0.

Table 2.1 summarizes the observed particle content of the Standard Model in terms of fields. The indices L and R give the chirality of the fields with respect to the weak interaction, which only acts on left-chiral particles. The left-chiral quarks and leptons which belong to the same generation are written as isospin doublets. Right-chiral particles are singlets under $SU(2)_L$, i. e. they do not transform under the weak interaction. $(u_R)^c$ denotes the charge conjugate of the right-handed up quark field, which is identical to the left-handed anti-up quark field \bar{u}_L and so on. Following the usual convention, antiparticles are indicated by a bar over the symbol for the respective particle. The table also lists the W and B bosons, which yield the photon and weak bosons as linear combinations. The electric charge can be calculated as $Q = T_3 + Y$, where T_3 is the third component of the weak isospin, and Y is the weak hypercharge quantum number. (Note that other conventions use twice the value for Y .) The mass spectrum of the Standard Model particles is summarized in Table 2.2. For the neutrinos, the upper limit on the mass from tritium decay is given. More details are given in the discussion of neutrinos in the Standard Model at the end of Section 2.1.3.

	Name	Symbol	Charge	Mass	
Fermions: Leptons	electron	e	-1	0.5110	MeV
	electron neutrino	ν_e	0	< 2	eV
	muon	μ	-1	105.7	MeV
	muon neutrino	ν_μ	0	< 2	eV
	tau (lepton) / tauon	τ	-1	1776.82(16)	MeV
	tau neutrino	ν_τ	0	< 2	eV
Fermions: Quarks	up quark	u	+2/3	1.7 – 3.3	MeV
	down quark	d	-1/3	4.1 – 5.8	MeV
	charm quark	c	+2/3	1270^{+70}_{-90}	MeV
	strange quark	s	-1/3	101^{+29}_{-21}	MeV
	top quark	t	+2/3	$172 \pm 0.9 \pm 1.3$	GeV
	bottom quark	b	-1/3	4190^{+180}_{-60}	MeV
Bosons	Photon	γ	0	< 10^{-18}	eV
	Gluon	g	0	(assumed to be 0 eV)	
	W boson	W^\pm	± 1	80.399(23)	GeV
	Z boson	Z^0	0	91.1876(21)	GeV

Table 2.2: Masses of the Standard Model particles [87]. Charges are given in terms of the electron charge magnitude. Masses of quarks except for the top quark are given in the $\overline{\text{MS}}$ scheme (see text). The masses of the electron and muon are known better than to 1 ppm, thus no uncertainty is given.

2.1.3 Fundamental Interactions

Only three of the four fundamental interactions of physics are so far included in the Standard Model of Particle Physics: the weak, the electromagnetic and the strong interaction. The fourth, gravitation, has therefore not been mentioned in the previous section, but for interactions of elementary particles it can be safely ignored due to its relative weakness at the energy scales which are within reach at current collider experiments.

Electromagnetic Interaction

Quantum Electrodynamics (QED) describes the interactions among charged particles like leptons and quarks, via the exchange of photons. The photon is electrically neutral and therefore cannot interact with other photons at tree level. The coupling constant that describes the strength of electromagnetic interactions is [8]

$$g_e = -q\sqrt{4\pi} = \sqrt{4\pi\alpha} \quad (\text{for } q = -e), \quad (2.30)$$

in Gaussian units, where q is the electric charge of the particle and α is the fine-structure constant. In fact, q and α are not constant, but depend on the energy scale $|k^2|$ of the interaction that is considered, where $|k^2| = -k^2$ and k is the virtual 4-momentum transferred by the photon. α is roughly $1/137$ at $|k^2| = 0$ and increases slowly to $1/128$ at energy scales that correspond to the mass of the W boson, $|k^2| = m_W^2$ [87]. This running of the coupling constant can be explained in terms of a vacuum polarization effect, leading to a screening of the charge.

Weak and Electroweak Interaction

The weak interaction is mediated by three massive vector bosons, two of which, the W^+ and W^- bosons, are electrically charged and one, the Z^0 boson, is neutral. Weak interactions respect the lepton generations, i. e. there is no cross-generational coupling of the type $e^- \rightarrow \nu_\mu + W^-$. However, it is possible to change a quark into a quark from another generation in a vertex involving a W boson^ε. This is called quark mixing. The Cabibbo-Kobayashi-Maskawa (CKM) matrix, a complex unitary 3×3 matrix V that can be parametrized using three angles and one phase factor, relates the physical quark states to the weak eigenstates:

$$\begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} = V \begin{pmatrix} d \\ s \\ b \end{pmatrix} \quad \text{with} \quad |V_{ij}| = \begin{pmatrix} 0.974 & 0.225 & 0.003 \\ 0.225 & 0.973 & 0.041 \\ 0.009 & 0.040 & 0.999 \end{pmatrix} \quad (2.31)$$

The uncertainties on the magnitudes of the values in the CKM matrix are about one per mille or smaller [87]. The non-zero off-diagonal elements of V cause transitions between generations of quarks in weak processes. The values in the CKM matrix have to be determined experimentally and are not predicted by the Standard Model. Weak processes violate parity conservation (in case of W bosons maximally), i. e. the symmetry of physical processes under inversion of spatial coordinates, because the weak interaction only couples to left-chiral fermions and right-chiral antifermions. The observed small violation of invariance under the combined charge conjugation and parity (CP) operation can be accommodated by a complex phase factor in the CKM matrix. Violation of invariance under time reversal and CP (TCP) is theoretically impossible within the current framework of Quantum Field Theory (QFT).

In the Glashow-Weinberg-Salam (GWS) model of unified weak and electromagnetic (electroweak) interactions, the weak and electromagnetic coupling constants and the mass ratio of the W and Z bosons are completely determined by the weak mixing angle $\theta_W = 28.75^\circ$ [87]. This angle is again a parameter that is not predicted by the Standard Model. It defines the mixing of the neutral W and B fields, which give the photon field and the Z boson field. For the coupling constants of the weak and the electromagnetic interaction one obtains

$$g_W \sin \theta_W = g' \cos \theta_W = g_e \quad \text{and} \quad g_Z = \frac{g_e}{\sin \theta_W \cos \theta_W}, \quad (2.32)$$

where g_W , g_Z and g_e are the coupling constants for the weak and electromagnetic couplings, respectively, and g' is the coupling to the weak hypercharge. The Standard Model with its Higgs doublet also predicts the weak mixing angle θ_W to appear in the ratio of the masses of the weak bosons,

$$\cos \theta_W = \frac{m_W}{m_Z}. \quad (2.33)$$

Strong Interaction

Quantum Chromodynamics (QCD) describes the interactions of particles which carry color charge by the exchange of gluons. Only quarks and gluons carry color charge and can participate in strong interactions. Unlike photons, the eight massless gluons can also interact

^ε In the Standard Model, the couplings of the Z boson to fermions are flavor diagonal at tree level and suppressed at higher orders due to the so-called GIM mechanism [54]. Flavor changing neutral currents, i. e. processes which change the flavor of a fermion without altering its electric charge, are however predicted in many new physics models including Supersymmetry [11].

among themselves. This self-interaction leads to two intimately related properties of the strong interaction, which are very different from QED: asymptotic freedom and confinement. Asymptotic freedom describes the fact that the strong interaction gets weaker at short distances. This is the opposite of the running of the coupling constant in QED. While the strong interaction is weak at short distances, at distances above approximately 1 fm quarks experience a confining potential $V(r) \propto \alpha_s r$ [11]. The energy stored in the color field (flux tube) when separating quarks against this non-diminishing force at some point becomes large enough to create new quark-antiquark pairs, which then form hadronic states with the original quarks.

Color is an internal three-valued quantum number. The possible values of the color degree of freedom are labelled red, green and blue. Gluons carry one unit of color plus one unit of anticolor, quarks carry only one unit of color and antiquarks carry one unit of anticolor. Due to confinement, no isolated free quarks have been observed so far, and all observed particles are colorless. This rule limits the possibilities to combine quarks and antiquarks into hadrons, composite particles made of quarks. Colorless combinations of quarks are mesons, $q_r \bar{q}_r, q_g \bar{q}_g, q_b \bar{q}_b$, and baryons, where a red, a green and a blue quark, $q_r q_g q_b$, or three antiquarks with different anticolors, $\bar{q}_r \bar{q}_g \bar{q}_b$, come together.

As quarks so far have not been observed isolated, their masses, with the exception of the top quark as explained below, can only be inferred indirectly from hadronic properties. To define quark masses, typically one uses the modified minimal subtraction ($\overline{\text{MS}}$) scheme [11, 15], a renormalization scheme for QCD perturbation theory which yields the running mass $\bar{m}(\mu)$, where μ is a dimensionful scale parameter [87]. Quarks that are assigned a mass of less than 1 GeV conventionally are called light quarks. This includes the up, down and strange quark. Quarks above this threshold are called heavy quarks, including thus the charm, bottom and top quark. The top quark is the heaviest known elementary particle with a mass of 172 GeV [87]. It is special among the six quarks because unlike all of the others it decays by electroweak decay before it can hadronize and form a bound state. From its decay products, one can therefore deduce properties of the top quark, and in particular measure its mass relatively precisely, with an uncertainty of less than 1% (see Table 2.2).

The Parton Model, Factorization, Hadronization

According to the simple quark parton model [11], a proton consists of three quarks, uud , which yields a proton charge of $2 \cdot \frac{2}{3} - \frac{1}{3} = +1$. An antiproton consists of three antiquarks, $\bar{u}\bar{u}\bar{d}$, yielding a charge $-(2 \cdot \frac{2}{3} - \frac{1}{3}) = -1$. Going one step beyond this simple description, in addition to the three valence quarks in the proton, a sea of quark-antiquark pairs and gluons exists, which are produced by soft QCD processes. Due to quark-mass effects, the probability for heavy quark pairs to occur is smaller than for lighter pairs. The momentum of the proton is shared amongst its constituents,

$$\sum_i \int_0^1 x[q_i(x) + \bar{q}_i(x)] dx + \int_0^1 xG(x) dx = 1, \quad (2.34)$$

where the sum runs over all quark flavors, and q_i and \bar{q}_i are the parton distribution functions (PDFs). $q_i(x)$ is the probability density for finding a quark with flavor i carrying a longitudinal momentum fraction x of the proton (or more generally nucleon) momentum. $G(x)$ is the probability density for a gluon with a longitudinal momentum fraction x of the

proton. The parton distribution functions can be measured by deep-inelastic scattering of leptons with nucleons [87]. The above is written in terms of parton distribution functions that are independent of the transferred momentum $|k^2|$. In fact, the parton distribution functions depend on both x and $|k^2|$, and the evolution of the PDFs with $|k^2|$ is described by the Dokshitzer-Gribov-Lipatov-Altarelli-Parisi (DGLAP) equations [6]. The evolution kernel in the DGLAP equations is determined by splitting functions, which give the probability for a parton to emit another parton carrying a fraction of the original parton's momentum.

The parton distribution functions have been derived by various groups. The most prominent are CTEQ [45] and MRS/MRST/MSTW [84]. They are an important input to simulations of particle reactions and computations of cross sections. As they are universal, the parton distribution functions for proton-proton collisions are the same as for deep inelastic scattering. The cross section can therefore be calculated through the factorization theorem, according to which for a process in which two hadrons $H_{1,2}$ produce a specific observed final state F plus proton rests X , the cross section reads [6]

$$\sigma(H_1 H_2 \rightarrow F + X) = \sum_{i,j \in \{q, \bar{q}, g\}} \int dx_1 dx_2 f_i^1(x_1, \mu^2) f_j^2(x_2, \mu^2) \hat{\sigma}_{ij \rightarrow F}(x_1 P_1, x_2 P_2, \alpha_S(\mu^2), Q^2 / \mu^2), \quad (2.35)$$

where the sum runs over all quark and antiquark flavors and gluons. The cross section thus factors into a partonic cross section $\hat{\sigma}$ averaged over spin and color for the relevant subprocess, and the distribution functions $f^{1,2}$ for the respective parton defined at a factorization scale μ . The factorization scale is a scale parameter that separates the long- and short-distance physics, i. e. the low transverse momentum physics below μ that is considered as part of the hadron structure, and the high transverse momentum physics that goes into the partonic cross section. μ is often chosen to be equal to the renormalization scale, and often both are set to the characteristic scale of the hard scattering process Q [6].

Quarks and gluons, the production of which can be described by Equation (2.35), will subsequently hadronize into mesons and baryons, which in the detector appear as clustered energy depositions or jets (cf. Section 4.5.1). The hadronization is a long-distance (low-energy) process, i. e. it cannot be described by perturbative QCD due to the large values of the coupling constant α_S at these energies. Instead, there are phenomenological models tuned to fit experimental data and non-perturbative lattice QCD models [11]. The computation of the cross sections from the factorization theorem and the phenomenological treatment of the hadronization gives a separation of perturbative and non-perturbative effects, which is often used in simulations in high-energy physics (cf. Section 4.5).

Neutrinos in the Standard Model

In many descriptions of the Standard Model, neutrinos are taken to be massless, although by now it is established experimentally that neutrinos can oscillate between different flavors, ν_e , ν_μ and ν_τ , which is only possible if they have different masses. The values of the neutrino masses are unknown. As the neutrino oscillations are only sensitive to differences in neutrino masses via the conversion probability

$$P(\nu_e \rightarrow \nu_\mu) = \sin^2(2\theta) \sin^2 \left(\frac{(m_2^2 - m_1^2) d}{4E} \right), \quad (2.36)$$

they can only be used to set experimental constraints on the mass splittings. Here, E is the neutrino energy, $m_{1,2}$ are the masses, d is the distance traveled by the neutrino, and θ is the mixing angle relating the mass and flavor eigenstates. Equation (2.36) is written for two-neutrino mixing, but in general all three neutrino flavors mix. Direct measurements from the high energy cut-off of the beta-decay spectrum of tritium can only set upper limits on neutrino masses. The upper limit for the electron neutrino mass is 2 eV [16]. In addition, lower bounds from atmospheric neutrino oscillations exist.

It is not yet clear whether neutrinos are their own antiparticles (Majorana particles). This could be proven by establishing the existence of the neutrinoless double beta decay [8]. Majorana neutrinos would be required by the “see-saw” mechanism, which predicts very heavy neutrino partners at high mass scales [3].

2.1.4 Open Questions and Problems

The Standard Model of particle physics has been in place and constantly improved over many decades now, and it describes and explains many observations at high precision, in particular in the domain of QED, for example from the agreement of different measurements of the fine-structure constant which are able to test QED [11]. It is thus a very successful theoretical framework. However, there are several indications which suggest that the Standard Model needs to receive modifications or extensions, which can roughly be grouped into two classes:

First, there are experimental observations which the Standard Model fails to describe or which directly contradict the respective theoretical prediction from the Standard Model. One example is the anomalous magnetic moment of the muon, which can be measured very precisely [29] and differs from the Standard Model expectation by 3.4σ [19]. Using a different determination (from hadronic τ decays) for one of the correction terms reduces the discrepancy to 0.9σ . A possible explanation are additional contributions as they are found in supersymmetric extensions of the Standard Model [11].

Second, despite its success, the Standard Model cannot provide answers to a large number of open questions^f:

- Why is charge quantized in thirds of the charge of the electron? Why is (only) the weak interaction chiral? Why is the underlying symmetry group a complicated product of three groups, $SU(3) \times SU(2)_L \times U(1)$? Is there a unification of these interactions?
- Why are there three generations of quarks and leptons? Why are the masses of the elementary particles so different, going from a fraction of an electronvolt for neutrinos to almost $2 \cdot 10^{11}$ electronvolts for the top quark? Why can quark generations mix, whereas the lepton number, apart from neutrino oscillations, is conserved for each generation individually?
- What has caused the prevalence of matter over antimatter as observed today?

In addition to these open questions, the Standard Model has a large number of free parameters that are not predicted by theory. Another two open questions will be discussed in more detail in the following.

^f A comprehensive overview of interpretations and suggestions of new physics models that provide possible explanations is given in [11].

The Hierarchy Problem of the Higgs Mass

One of the most unsatisfactory aspects of the Standard Model is known as the fine-tuning or hierarchy problem [1, 106] and concerns the mass of the Higgs particle, as predicted in the Standard Model. The Higgs mass, as given at tree level in Equation (2.29), receives contributions from loop corrections. For a Dirac fermion f which couples to the Higgs through a term $-\lambda_f H \bar{f} f$ in the Lagrangian (cf. Equation (2.28)), this contribution turns out to be [106]

$$\Delta m_H^2 = -\kappa |\lambda_f|^2 \Lambda_{UV}^2, \quad (2.37)$$

with a constant prefactor $\kappa = 1/8\pi^2$. Λ_{UV} is a cut-off representing the scale at which new physics appears and is believed to be indicated by the Planck mass

$$m_p \simeq 1.22 \cdot 10^{19} \text{ GeV} \quad (2.38)$$

at most, which is the scale at which quantum gravity is expected to become important. The loop corrections to the Higgs mass are thus many orders of magnitude larger than the phenomenologically required value of the Higgs mass, which is of the order of a few hundreds of GeV[§]. The required low value for the Higgs mass can still be retained by adjusting the parameter $-\rho^2$ of the (unknown) Higgs potential in Equation (2.25) with a one-loop corrected physical value, but this would rely on a very precise cancellation of two very large numbers. Although not excluded in principle, this fine-tuning is considered very unlikely and therefore unnatural. Note that these quadratic divergences of the loop integrals are a problem that only affects scalar fields. The fermion masses in the Standard Model are protected by the chiral symmetry, reducing the quadratic divergences to logarithmic divergences [1]. The fact that the hierarchy problem can be solved by introducing Supersymmetry, as shown below, is one of the main indications in favor of Supersymmetry, although Supersymmetry historically has not been introduced to solve this problem.

Dark Matter and Dark Energy

Another unsatisfactory aspect is that the Standard Model only covers a few percent of the energy and matter content of the universe. Measurements of velocities of galaxies first done in 1933 using the Doppler effect [118, 119], and later more precise measurements of the rotation curves of galaxies [102, 103], i. e. the rotational velocity as function of the distance from the center of the galaxy, revealed that there must be a lot more matter in galaxies than the matter in visible stars. This additional non-luminous and non-absorbing matter, which is evident only through its gravitational effects, is called dark matter. Today, from a number of cosmological observations like the cosmic microwave background radiation, galactic lensing or the acceleration and large scale distribution of matter in the universe, it is found that the universe is flat and about 21 % of its energy content are dark matter. Only about 5 % are ordinary matter as it is described by the Standard Model, and the remaining 74 % are attributed to dark energy [11]. This hypothesized dark energy, the nature of which is unknown [8], may explain, via its negative pressure, the observed acceleration of the expansion of the universe.

The term dark matter implies that this substance is thought to be made up from elementary particles. Most of the dark matter content must be non-relativistic, so-called Cold Dark

[§] Electroweak precision data favor rather low values of the Higgs mass, below 167 GeV at 95 % confidence level [11]. The even lower current limits on the Higgs mass will be discussed below.

Matter (CDM), so that it can clump together and become gravitationally bound on large scales. This rules out neutrinos, which are also too light to contribute more than a small fraction to dark matter. In general, no Standard Model particle has the right properties to explain cold dark matter, which is supposed to be comprised of Weakly Interacting Massive Particles (WIMPs) [22]. In physics models which go beyond the Standard Model, there are many candidates for WIMPs, for example the lightest supersymmetric particles, similar stable particles in little Higgs or universal extra dimension models, or axions associated to a potential solution of the strong CP problem [11].

2.2 Possibilities for Beyond Standard Model Physics

Before discussing Supersymmetry (SUSY), this section gives an overview of two important theoretical models, which, like Supersymmetry, are capable of addressing shortcomings of the Standard Model, in particular the hierarchy problem of the Higgs mass. The two models are technicolor and (universal) extra dimensions. In some cases, they may give phenomenological signatures which are very similar to Supersymmetry.

2.2.1 Technicolor

Technicolor [107] provides a natural solution to the hierarchy problem by not introducing a scalar Higgs boson in the first place. Instead, a dynamical breaking of electroweak and flavor symmetry is employed, in which a bilinear of fermions acquires a vacuum expectation value. To achieve this, in technicolor models a new non-Abelian gauge group is introduced, which is modelled on QCD: asymptotically free at very high energies and strong and confining at low energies. In addition, massless fermions called technifermions are introduced. The technifermions can form a condensate, which spontaneously breaks the global chiral symmetry of the fermions. Following the Goldstone theorem, the breaking of the symmetry procures Goldstone bosons, which are referred to as technipions. Three of them become longitudinal components of W and Z , which thereby acquire mass.

Additional interactions are needed to give mass to the Standard Model quarks and leptons, which cannot have bare mass terms. This is achieved in extended technicolor by extending the new gauge group to include color, technicolor and flavor symmetries, thereby allowing technifermions to couple to quarks and leptons. After the breaking of the larger gauge symmetry, the resulting massive gauge bosons can mediate transitions between Standard Model and technicolor fermions, giving rise to couplings that make the Standard Model fermions massive. The interactions of the technicolor massive gauge bosons can also raise the masses of technipions, so that they can have escaped detection so far, mediate their decays to Standard Model fermions, and induce mixing of quarks. A successful technicolor model would not only predict the masses and mixings of quarks and leptons, but also why there are three families of each. Technicolor can also provide dark matter candidates in terms of the lowest-lying bound state of technifermions [23].

Similar in spirit to technicolor models, in the sense that they avoid introducing Higgs particles, are a number of other Higgsless models which use extended gauge groups. There are also little Higgs models with composite Higgs states, which use extended gauge groups combined with novel mechanisms for breaking the gauge symmetry and naturally give rise to light Higgs bosons without relying on Supersymmetry. An overview and discussion of

these models, detector signatures and searches for technicolor is given in [87] and [180].

2.2.2 Extra Dimensions

The concept of extra dimensions was first formulated by Kaluza and Klein in the 1920s in an attempt to unify gravitation and electromagnetism, known as the Kaluza-Klein theory. It was revived in the 1980s in the context of string theory. Extra dimensions are needed in consistent quantum theories of strings, which predict ten or eleven spacetime dimensions [98], as new degrees of freedom for the fundamental objects: the one-dimensional strings. Strings can be thought of as a generalization of point-like particles. The extra dimensions must be compactified, i. e. instead of being infinite, they are curled up in a circle of small radius, thus having periodic boundary conditions and a finite length of the order of inverse TeV, so that they can evade direct observation. The scale of the extra dimensions, which is given by the compactification radius, is constrained by experimental data. In the simplest possible compactification scheme, the extra-dimensional part of spacetime is a toroidal structure with one characteristic radius R [16]. The periodic boundary conditions in compactified dimensions imply that fields can be expanded as a Fourier series along this dimensional direction, which results in an infinite tower of massive states, the Kaluza-Klein modes, collectively called the Kaluza-Klein tower.

The theory of Universal Extra Dimensions (UED) is a phenomenological model based on the concept of extra dimensions similar to the first string theories, with the assumption that the extra dimensions have a flat (rather than a warped) metric, and that all fields can propagate in the compact dimensions. Momentum conservation in the extra dimensions translates into the conservation of the number of Kaluza-Klein modes at tree level (KK-parity). Each interaction vertex involves at least two Kaluza-Klein excitations, so that the direct production of Kaluza-Klein states is only possible in pairs and the lightest KK-state is absolutely stable, making it a good candidate for cold dark matter. The implications of KK-parity for the phenomenology of UED at particle colliders are therefore very similar to supersymmetric models with conserved R -parity, as will become clear below. The mass bounds for the lowest excited modes are relatively low. The limits on R^{-1} for one extra dimension are 300 to 500 GeV, depending on the Higgs mass. They lie within the region up to $R^{-1} \sim 1.5$ TeV which is experimentally accessible at the LHC [16]. By introducing large compact dimensions, where the term large refers to any dimension which is larger than the Planck scale, the hierarchy problem can be traded for the problem of a large volume of the extra dimensions, which relates the Standard Model scale and the Planck scale.

Mixed extra dimensions are interpolations between the two extreme cases of universal extra dimensions, where the Standard Model fields can propagate in all compact dimensions, meaning that the Standard Model particles also acquire a tower of KK-excitations, and scenarios, where the Standard Model fields are confined to a D3-brane, i. e. a region of spacetime on which open strings can end with three spatial dimensions. In mixed extra dimensions, only some fields are constrained to a brane, e. g. the gauge bosons can propagate in the extra dimensions, whereas the fermions are confined to the Standard Model D3-brane, which gives different collider bounds. Minimal UED is a five-dimensional model with only one additional flat dimension, where all particles can propagate in all dimensions.

The experimental signatures of universal extra dimensions at hadron colliders are very similar to those of Supersymmetry (cf. Section 2.3) [37, 79]:

- Both UED and SUSY have additional massive particles with the same couplings and quantum numbers as their Standard Model partners, only the spin being different by half a unit for the supersymmetric partners of the Standard Model particles. In UED, there is a tower of additional states, but phenomenologically only the lowest is of interest at this stage.
- In both UED and SUSY, a multiplicatively conserved quantum number can be defined, KK-parity in UED and R -parity in SUSY. In both cases, this implies that the partner particles can only be produced in pairs, and that the lightest partner particle at the end of the decay chain is stable so that it is a candidate for cold dark matter.

This has two important consequences: First, analyses optimized for Supersymmetry may also be sensitive to UED. Second, this shows that detector signatures are not unique and can be interpreted in many different models, so that when indications for beyond Standard Model physics are found, more detailed studies of the produced new particle states are imperative.

An important difference between UED and SUSY lies in the degeneracy of the masses, which for Standard Model particles and their supersymmetric partners are very different after breaking of Supersymmetry, whereas in universal extra dimensions, the masses of the first levels of KK-excitations at tree level are almost degenerate. This is because their splitting is provided by Standard Model mass terms, and most Standard Model particles are relatively light. Loop corrections can lift the degeneracy in UED, but the mass differences will still be small, so that the decay products will be difficult to identify at hadron colliders. The situation can be improved by allowing the violation of KK-parity through gravity, which can yield final states with one or more heavy jets and missing transverse energy, similar to the signature searched for in the analysis in this thesis in Section 7 [79].

2.3 Supersymmetry

This section summarizes the ideas and concepts needed as theoretical background for the evaluation of data from particle colliders in analyses concerned with Supersymmetry. Supersymmetry is theoretically appealing because it is based on a simple and natural extension of the existing theoretical framework of particle physics. Moreover, it follows in many string models, and it can solve a number of problems of the Standard Model outlined above, without entailing predictions which are inconsistent with existing experimental observations.

2.3.1 General Idea

The basic idea of Supersymmetry is to introduce a new discrete symmetry beyond those already included in the Standard Model. This symmetry relates fermions and bosons, and thus every fermion has a bosonic supersymmetric partner particle and vice versa. Supersymmetry therefore doubles the Standard Model particle spectrum, with the exception of the Higgs sector, where the (still elusive) scalar and neutral Higgs boson from the Standard Model is replaced by five Higgs particles in minimal supersymmetric models. The supersymmetric particles have, apart from spin, the same quantum numbers and masses as their Standard Model partners. Experimental constraints therefore require that Supersymmetry is broken so that the masses of the supersymmetric partners of the Standard Model particles are shifted

to much higher values. Otherwise the superpartners would have been observed already. Different models for breaking of Supersymmetry exist and are explained in Section 2.3.5.

2.3.2 Motivations

Supersymmetry provides a natural solution to the hierarchy problem of the Higgs mass by introducing supersymmetric partner particles, because the loop corrections to the Higgs mass arising from bosons and fermions are of opposite sign: The correction for a scalar particle S coupling to the Higgs field H via $-\lambda_S |H|^2 |S|^2$ is given by [106]

$$\Delta m_H^2 = \frac{\kappa}{2} \lambda_S \Lambda_{UV}^2. \quad (2.39)$$

If the coupling constants in Equations (2.37) and (2.39) were equal, $\lambda_S = |\lambda_f|^2$, the quadratic divergences would exactly cancel to all orders in perturbation theory due to the opposite sign^h. This would, of course, be an extraordinary coincidence if not enforced by a symmetry [113]. In fact, Supersymmetry does not only cancel the quadratic divergences, but also divergences which are logarithmic in Λ_{UV} and contribute to the mass term in the Higgs potential [1]. After the necessary (“soft”, see below) breaking of Supersymmetry, the quadratic divergences still cancel, but there are finite contributions to the Higgs mass of the order of the Supersymmetry breaking scale [11].

Another motivation for Supersymmetry is that it provides a natural candidate for dark matter in terms of a weakly interacting massive particle, which does not exist in the Standard Model. Most supersymmetric models postulate the conservation of a new quantum number which is called R -parity. This conservation makes the lightest supersymmetric particle (LSP) absolutely stable. If, as is the case in most supersymmetric scenarios, the LSP is colorless, neutral and heavy, which are all requirements met by the supersymmetric neutralinos, the LSP is an attractive candidate for dark matter.

In addition, there is another motivation in favor of a supersymmetric extension of the Standard Model, which comes from the extrapolation of the three gauge coupling constants to the energy scale of Grand Unified Theories (GUTs) around 10^{16} GeV. At this scale, the inverse gauge couplings are predicted to meet in GUTs, but assuming the Standard Model they only approximately do so. By introducing the superpartners of Standard Model particles, the energy dependence of the running coupling constants is modified such that, in the Minimal Supersymmetric Standard Model (MSSM), it is possible to make them meet exactly in one point at the GUT scale [11]. This behavior is demonstrated in Figure 2.1 in a comparison between the Standard Model and the MSSM evolution of the running of the coupling constants of the three gauge groups $U(1)_{Y/2}$, $SU(2)_L$ and $SU(3)_C$ [69].

2.3.3 Theoretical Background

As was explained in Section 2.1.1, symmetries play a very important role in physics, because they allow to deduce physical laws from very simple assumptions about the invariance of a physical system under a symmetry operation (cf. Noether’s theorem). The introduction of a new symmetry beyond those already implemented in the Standard Model seems to fail in the first place, because the Coleman-Mandula theorem states that the only conserved

^h The additional factor 2 accounts for the fact that for each Dirac fermion two complex scalar partners are introduced.




Figure 2.1: Comparison of the extrapolation of the gauge couplings to the GUT scale, assuming the Standard Model (left) and the MSSM (right). The vertical axis is α_i^{-1} , where α_i is the coupling of the group $U(1)_{Y/2}$, $SU(2)_L$, $SU(3)_C$ for $i = 1, 2, 3$, respectively. μ is the energy scale.

charges which transform as tensors under the Lorentz group are P_μ and $M_{\mu\nu}$, the generators of the Poincaré group. This means that charges associated with symmetries besides those of the Poincaré group are Lorentz scalars. However, the Coleman-Mandula theorem can be circumvented by using fermionic generators [55], which under Lorentz transformations transform as spinors. This is required for generators of supersymmetric transformations because they transform an integer spin field into a spinor field, and therefore must carry a spinorial index. It follows that they do not commute with Lorentz transformations, and follow anticommutation relations due to the (anti-)commutation rules of bosons and fermions. Restrictions from the Coleman-Mandula theorem fix the algebra of the fermionic generators, making Supersymmetry unique in this sense [3].

The fermionic Supersymmetry generators of the Supersymmetry transformations (“supertranslations”) are denoted Q_a . Their algebra for the case of one generatorⁱ is given by [1]

$$\{Q_a, Q_b^\dagger\} = 2(\sigma^\mu)_{ab} P_\mu, \quad (2.40)$$

where the brackets denote the anticommutator and σ^μ , $\mu \in \{0, \dots, 3\}$, are the Pauli matrices [14], including the 2×2 identity matrix as σ^0 . P_μ is again the generator of translations. This is written in the Weyl formalism, where the Supersymmetry charges Q_a are anticommuting quantum field operators with two components^j, $a \in \{1, 2\}$. They transform under Lorentz transformations as spinors, i. e.

$$[Q_a, M_{\mu\nu}] = (\sigma_{\mu\nu})_a{}^b Q_b, \quad (2.41)$$

with $M_{\mu\nu}$ being the generators of the Lorentz transformations and $\sigma^{\mu\nu} = \frac{i}{2} [\gamma^\mu, \gamma^\nu]$. Furthermore, being a symmetry operator, they commute with the Hamiltonian of the system.

ⁱ One can also consider representations with more than one SUSY generator, but making a supersymmetric extension of the Standard Model then becomes difficult because the different behavior of the left- and right-chiral components of the lepton fields cannot be easily accommodated [1].

^j The Hermitian conjugation [†] converts left-chiral spinors into right-chiral spinors, which are conventionally also denoted with dotted indices and a bar over the symbol, i. e. $Q_b^\dagger = \bar{Q}_b$.

The anticommutation relation in Equation (2.40) is complemented by the relation

$$[Q_a, P_\mu] = [Q_a^\dagger, P_\mu] = 0, \quad (2.42)$$

stating that the SUSY generators and the generators of translations commute, and so does P^2 . This means that the states in a supermultiplet which are connected by the actions of the Supersymmetry generators must all have the same 4-momentum, and thus in particular the same mass. When the generators Q_a and Q_a^\dagger act on a single particle state, they either create a state with spin differing by half a unit, or they annihilate it if it is at the “end” of a supersymmetric multiplet.

Three different massless supersymmetric multiplets (supermultiplets) are of interest:

- The chiral supermultiplet contains a massless complex scalar field and a massless Weyl fermion field. For example, the left-chiral electron and neutrino form an $SU(2)_L$ doublet, which is partnered by a corresponding doublet of scalars to form a left-chiral supermultiplet. Similarly, the right-chiral electron field, which is an $SU(2)_L$ singlet, will be partnered by a corresponding supersymmetric state to form a right-chiral supermultiplet.
- The vector or gauge supermultiplet consists of a massless spin-1 (vector) state and a massless Weyl spin-1/2 fermion partner. The gauge bosons of the Standard Model are assigned to gauge supermultiplets.
- The gravity supermultiplet combines a spin-2 graviton and its spin-3/2 gravitino partner. This supermultiplet is needed in theories of supergravity.

Only the first two, the massless chiral and vector supermultiplets, are needed to construct the Minimal Supersymmetric (extension of the) Standard Model, which is the supersymmetric theory including all Standard Model particles with the minimal additional particle content. To account for the fact that the supersymmetric transformations mix states with different spin, it is useful to define a notational extension of the four-dimensional spacetime, which includes two additional anticommuting coordinates in terms of Grassmann variables. The operators which are functions of the coordinates in this superspace are then called superfields [11].

2.3.4 The Minimal Supersymmetric Standard Model

Table 2.3 shows the field content of the Minimal Supersymmetric Standard Model (MSSM). Every Standard Model particle is accompanied by a superpartner, conventionally denoted with a tilde. The fermions in the Standard Model are the lepton fields, e.g. the $SU(2)_L$ doublet consisting of the neutrino ν_{eL} and electron field e_L , and the quark fields, which are triplets under the $SU(3)_C$ color gauge group. The left-chiral lepton fields are accompanied by a doublet of scalar fields, e.g. the sneutrino $\tilde{\nu}_{eL}$ and selectron \tilde{e}_L , where the prefix “s” stands for scalar rather than for supersymmetric. The quarks have corresponding squark fields, where the $SU(2)_L$ doublet (u_L, d_L) is partnered by $(\tilde{u}_L, \tilde{d}_L)$. The same holds for the right-chiral fields, which are $SU(2)_L$ singlets, \tilde{e}_R partners e_R , \tilde{u}_R partners u_R and so on. The indices of the scalar fields, e.g. of \tilde{q}_L , refer to the chirality of their superpartners and show what their $SU(2)_L \times U(1)$ quantum numbers are. For the spin 1 particles of the Standard Model, the gluons, W bosons and the B boson, there is an $SU(3)_C$ octet of Weyl fermions

Name	Symbol	Spin 0	Spin 1/2	SU(3) _C	SU(2) _L	U(1) _Y
Squarks, Quarks	Q	$\begin{pmatrix} \tilde{u}_L \\ \tilde{d}_L \end{pmatrix}$	$\begin{pmatrix} u_L \\ d_L \end{pmatrix}$	3	2	$+\frac{1}{6}$
	\bar{u}	$\tilde{u}_L = \tilde{u}_R^\dagger$	$\bar{u}_L = (u_R)^c$	$\bar{\mathbf{3}}$	1	$-\frac{2}{3}$
	\bar{d}	$\tilde{d}_L = \tilde{d}_R^\dagger$	$\bar{d}_L = (d_R)^c$	$\bar{\mathbf{3}}$	1	$+\frac{1}{3}$
Sleptons, Leptons	L	$\begin{pmatrix} \tilde{\nu}_{eL} \\ \tilde{e}_L \end{pmatrix}$	$\begin{pmatrix} \nu_{eL} \\ e_L \end{pmatrix}$	1	2	$-\frac{1}{2}$
	\bar{e}	$\tilde{e}_L = \tilde{e}_R^\dagger$	$\bar{e}_L = (e_R)^c$	1	1	$+1$
Higgs, Higgsinos	H_u	$\begin{pmatrix} h_u^+ \\ h_u^0 \end{pmatrix}$	$\begin{pmatrix} \tilde{h}_u^+ \\ \tilde{h}_u^0 \end{pmatrix}$	1	2	$+\frac{1}{2}$
	H_d	$\begin{pmatrix} h_d^0 \\ h_d^- \end{pmatrix}$	$\begin{pmatrix} \tilde{h}_d^0 \\ \tilde{h}_d^- \end{pmatrix}$	1	2	$-\frac{1}{2}$
		Spin 1/2	Spin 1			
Gluinos, Gluons		\tilde{g}	g	8	1	0
Winos, W bosons		$\tilde{W} = \begin{pmatrix} \tilde{W}^+ \\ \tilde{W}^0 \\ \tilde{W}^- \end{pmatrix}$	$W = \begin{pmatrix} W^+ \\ W^0 \\ W^- \end{pmatrix}$	1	3	0
Bino, B boson		\tilde{B}	B	1	1	0

Table 2.3: Field content of the Minimal Supersymmetric extension of the Standard Model [1, 106]. The fields with tildes are the supersymmetric partners of the Standard Model fields. Note that only the first generation is shown for the quarks and leptons. The Higgs of the Standard Model is replaced by two Higgs doublets in the MSSM. \tilde{u}_L and \tilde{u}_R are independent scalar fields. The index refers to the chirality of their superpartners.

called gluinos (\tilde{g})[‡], an $SU(2)_L$ triplet of Weyl fermions called winos ($\tilde{W}^{\pm,0}$), and finally a $U(1)_Y$ bino (\tilde{B}).

The Higgs sector in supersymmetric models needs to be extended with respect to the Standard Model to include two Higgs doublets instead of one because Supersymmetry does not allow the needed Yukawa couplings to generate masses for both u and d quarks with only one Higgs doublet. The two independent Higgs chiral supermultiplets are $H_u = (H_u^+ H_u^0)^T$ and $H_d = (H_d^0 H_d^-)^T$ with scalar and fermionic components such as h_u^+ and \tilde{h}_{uL}^+ for H_u^+ . After symmetry breaking, this leaves three neutral and one conjugate pair of charged Higgs fields. In analogy to the mixing of the three Standard Model W bosons and the B boson, in the MSSM their supersymmetric partners mix to give the winos, a zino \tilde{Z}^0 and a massless photino $\tilde{\gamma}$. But also the Higgsinos can mix with the winos and the bino to produce two mass eigenstate Dirac charginos ($\tilde{\chi}_{1,2}^\pm$, so four particles in total) and four mass eigenstate Majorana neutralinos ($\tilde{\chi}_{1,\dots,4}^0$ with $m_{\tilde{\chi}_1^0} < \dots < m_{\tilde{\chi}_4^0}$).

The supersymmetric interactions are fixed by specifying the superpotential, which is a holomorphic function of the left-chiral superfields only and thus itself a left-chiral superfield. In the MSSM, the superpotential W is [1]

$$W = y_u^{ij} \bar{u}_i Q_j \cdot H_u - y_d^{ij} \bar{d}_i Q_j \cdot H_d - y_e^{ij} \bar{e}_i L_j \cdot H_d + \mu H_u \cdot H_d, \quad (2.43)$$

where the symbols for the chiral supermultiplet fields from Table 2.3 have been used[‡]. The color indices for the terms involving quarks are not written out. The dimensionless couplings are given in terms of 3×3 matrices y in generation space, which are exactly the same as the couplings in the Standard Model. μ is thus the only additional parameter needed in this supersymmetric extension of the Standard Model. From the superpotential, the interaction terms of the Lagrangian can be obtained from the so-called F component, which automatically makes the Lagrangian invariant under supersymmetric transformations. Note that the full Lagrangian of the MSSM will contain auxiliary complex fields F , which do not correspond to physical particles, but which are needed to ensure a supersymmetric action. When the Higgs fields acquire vacuum expectation values, the first three terms with y couplings give masses to the quarks and leptons and their supersymmetric partners. The last term in the superpotential yields supersymmetric masses for the Higgs and Higgsino fields. It is connected with the μ -problem [67], which is a problem of naturalness within supersymmetric models.

R-Parity

The superpotential of the MSSM in Equation (2.43) is a specific choice which does not include all gauge-invariant and renormalizable terms. The terms which have been left out are [1]

$$W_{\Delta L=1} = \lambda_e^{ijk} L_i \cdot L_j \bar{e}_k + \lambda_L^{ijk} L_i \cdot Q_j \bar{d}_k + \mu_L^i L_i \cdot H_u, \quad (2.44)$$

[‡] In models with supergravity, the symbols for gluons and gluinos are often chosen to be G and \tilde{G} to reserve g and \tilde{g} for the graviton and gravitino with spin $3/2$.

[‡] Note that $Q = \begin{pmatrix} U \\ D \end{pmatrix}$, with e. g. \bar{u}_L being the scalar and u_L the fermion spinor component of the left-chiral supermultiplet U .

which would violate lepton number conservation, and

$$W_{\Delta B=1} = \lambda_B^{ijk} \bar{u}_i \bar{d}_j \bar{d}_k, \quad (2.45)$$

which would violate baryon number conservation. λ and μ represent all possible coupling constants. In the Standard Model, no renormalizable terms in the Lagrangian are possible, that would violate lepton or baryon number conservation. In the MSSM, however, such terms are possible. Including both non-zero couplings λ_L^{ijk} of the leptoquark term and λ_B^{ijk} of the diquark term would allow the proton to decay. This is in conflict with observation, because the lower limit on the lifetime of the proton is of the order of 10^{31} years [87]. It can thus be ruled out that both terms appear with non-vanishing couplings, but there is no fundamental reason not to include any of these terms. To forbid lepton number and baryon number violating terms, an additional symmetry is postulated [106]. It is called R -parity,

$$R = (-1)^{2s+L+3B}, \quad (2.46)$$

where s is the spin of the particle, L the lepton number and B the baryon number, yielding $R = +1$ for Standard Model particles and $R = -1$ for supersymmetric particles. R -parity is multiplicatively conserved in supersymmetric processes and forbids all terms in $W_{\Delta L=1}$ and $W_{\Delta B=1}$. It has furthermore two important phenomenological consequences: First, supersymmetric particles are always produced in pairs at collider experiments, which means that in supersymmetric events there will always be at least two overlaid decay chains. Second, the Lightest Supersymmetric Particle (LSP) is absolutely stable, because there are no lighter supersymmetric particles, and a decay to only Standard Model particles would violate R -parity. In many supersymmetric models, the LSP is a neutralino, and therefore, if it is stabilized by R -parity, it gives an attractive candidate for non-baryonic cold dark matter. Due to their weakly interacting nature, these LSPs would be detected as missing momentum in collider experiments.

2.3.5 Breaking of Supersymmetry

Because no degenerate supermultiplets are observed, some form of breaking of Supersymmetry is inevitable in order to bring the masses of the sparticles into the not yet excluded range of several hundred gigaelectronvolts. It can also be shown that the Lagrangian of the MSSM, being invariant under Supersymmetry, cannot accommodate electroweak symmetry breaking. Like other symmetries, Supersymmetry can be broken explicitly or spontaneously. In any case, the Supersymmetry breaking needs to be “soft” in order to avoid the reintroduction of the hierarchy problem. “Soft” means that the breaking only appears via terms in the Lagrangian with positive mass dimension, and in general that it maintains the cancellation of quadratically divergent radiative corrections to masses of scalar fields. As a consequence, Supersymmetry breaking terms can only be mass terms for scalars or gauginos or cubic scalar couplings. Spontaneous symmetry breaking is favored for a number of reasons [11], and can occur for example as so-called D -type or F -type breaking: either one of the auxiliary fields D , which need to be introduced for the gauge supermultiplets, or one of the auxiliary fields F mentioned above acquires a symmetry-breaking non-zero vacuum expectation value.

A generic feature of spontaneous Supersymmetry breaking is the appearance of a Weyl fermion, the Goldstino^m, which in supergravity becomes a component of the gravitino in the super-Higgs mechanism (see below). Furthermore, spontaneous symmetry breaking in general implies a sum rule for tree-level mass squares of fermionic and bosonic degrees of freedom. This sum rule requires some superpartners to be light, which means that tree-level spontaneous symmetry breaking cannot be phenomenologically viable for the MSSM. To evade these constraints, the breaking of Supersymmetry must occur in a sector which is only weakly coupled to the chiral supermultiplets of the MSSM, and therefore called hidden sector. The mechanism responsible for the mediation from the hidden sector to the MSSM particles can be of various types which are outlined in Section 2.3.7. Independently of the exact mechanism, the parametrization of the possible gauge invariant terms in the now effective Lagrangian at low energy scales introduces a large number of additional free parameters in the MSSM. In the MSSM with conserved R -parity, there are 124 free parameters [49], without including neutrino masses and right-handed scalar neutrinos, i. e. 105 new parameters compared to the Standard Model [1]. This makes the model flexible, but also makes it difficult to transcribe results from searches for Supersymmetry into conclusions for the MSSM. Thus, simplified models are constructed within the MSSM, which enforce phenomenological constraints, and in the same step reduce the number of free parameters of the breaking mechanism significantly. Some examples are presented in Section 2.3.7.

2.3.6 Higgs Mass Bounds and the MSSM

Assuming an appropriate form of Supersymmetry breaking, the known Standard Model particle content seems to be fully compatible with a possible extension to a supersymmetric model. The question is whether this does also hold for the Higgs boson, which is believed by many to exist, but at the time of writing, this is not yet experimentally proven. In the Standard Model, there is one Higgs doublet with four real scalar degrees of freedom. After electroweak symmetry breaking, three of these become longitudinal modes of the massive vector bosons of the weak interaction, while the remaining fourth becomes the neutral scalar Higgs boson. In the Minimal Supersymmetric Standard Model, there are two complex Higgs doublets with thus eight scalar degrees of freedom. Again, three of them are needed as longitudinal modes of the massive vector bosons, which leaves five Higgs bosons in the MSSM, three of which are neutral (A^0 , h^0 , H^0) and two are charged (H^\pm). If the existence of a Higgs boson is established experimentally, this does not rule out that there are more Higgs bosons and therefore does not exclude the MSSM. However, the opposite case of a light Higgs boson being excluded might be fatal for the MSSM, which sets quite stringent upper bounds on the mass of the lightest Higgs boson h^0 ,

$$m_{h_0} \leq m_Z |\cos(2\beta)| \leq m_Z \quad (\text{tree-level}), \quad (2.47)$$

where m_Z is the mass of the Z boson and $\tan \beta := v_u/v_d$ is the ratio of the vacuum expectation values of the H_u^0 and H_d^0 Higgs fields [106]. The experimental lower bound on the mass of the (Standard Model) Higgs boson from the Large Electron-Positron Collider (LEP) is 114.4 GeV at 95 % confidence level (C. L.) [70]. However, the lower bound in Equation (2.47) receives loop-corrections which may be of sufficient size to shift the theoretical upper bound

^m The Goldstino is not the fermionic superpartner of a scalar Goldstone boson, but itself a fermionic Goldstone mode [1].

on m_{h_0} close to 140 GeV, safely above the lower bound from LEP [47]. Recent results of Higgs searches in ATLAS and CMS exclude (Standard Model) Higgs masses between 130 and 235 GeV [152] and 127 and 600 GeV [123] at 95 % C. L., respectively, which still leaves a gap that would allow for an MSSM Higgs boson.

2.3.7 Important Supersymmetric Models

This section concludes the introduction to Supersymmetry by presenting Minimal Supergravity (mSUGRA) and several other prominent supersymmetric models. The attractiveness of these models lies in the small number of parameters compared to the 124 parameters of the full MSSM. They rely on breaking of Supersymmetry in a hidden sector and mediation of this breaking to the visible sector of MSSM particles via different mechanisms.

Minimal Supergravity

The relevant supersymmetric model for the studies in this thesis is the mSUGRA model, which is also one of the most frequently encountered models in Supersymmetry analyses, and often used for investigations of the phenomenological consequences of weak scale Supersymmetry [36, 56]. As pointed out in [3], making the theory of special relativity invariant under local coordinate transformations leads to the theory of general relativity, and along the same lines, local Supersymmetry is expected to lead to a theory of gravity called supergravity. Indeed, when making the Supersymmetry transformations local, i. e. when the parameter in the transformations generated by the supersymmetric charges become space-time dependent, this necessitates the introduction of a graviton field and its superpartner, the gravitino, and thus gravity, in analogy to other localized gauge symmetries. Supergravity is not a full quantum theory of gravity, but an effective theory which is nonrenormalizable. Again, in analogy to spontaneous breaking of a continuous gauge symmetry (cf. Section 2.1.1), a Goldstone mode in terms of the Goldstino provides the missing degrees of freedom of the gravitino. This is known as the super-Higgs mechanism. The graviton remains massless. In supergravity, the breaking of Supersymmetry in the hidden sector is mediated by gravitational couplings, i. e. the effective Lagrangian contains nonrenormalizable terms for the mediation, which are suppressed by powers of the Planck mass [106].

Models of minimal supergravity have five free parameters, of which four are continuous and one is discrete, and which are conventionally denoted with the following symbols [3]:

- $m_{1/2}$: a common mass for all gauginos at the GUT scale,
- m_0 : a common mass for all scalar partners of leptons and quarks and the Higgs doublets at the GUT scale,
- $\tan \beta$: the ratio of Higgs vacuum expectation values,
- $\text{sign } \mu$: the sign of the Higgs mixing parameter and
- A_0 : a common value for the soft SUSY-breaking trilinear couplings at the GUT scale.

Subsets of these parameters can be traded for others, but this set is commonly used for the parametrization of mSUGRA. Thus, all squarks and sleptons are degenerate in mass at the GUT scale, which allows to eliminate mixings of states with the same electroweak

quantum numbers by unitary transformations. The couplings and masses evolve according to Renormalization Group Equations (RGEs) from the GUT scale down to the electroweak scale. One prediction of the RGE evolution common to most supersymmetric models is that the gluino is expected to be heavier than the states from the electroweak sector [1]. The mass of the gravitino $m_{3/2}$ in minimal supergravity models is of the order of the masses of the other supersymmetric particles [106] and may be set equal to m_0 [86]. As its couplings are of gravitational strength, it is not relevant for collider physics.

Other Possibilities

In [17], an overview of a number of models is given, which all have fewer parameters than the full MSSM and can be tested at current collider experiments. What is referred to in this thesis as the mSUGRA model is called the Constrained MSSM (CMSSM) there, and mSUGRA is defined instead as a subset of the CMSSM models with additional assumptions.

Another related model which is introduced in [17] is the Non-Universal Higgs Mass Model (NUHM). It generalizes the CMSSM by relaxing the boundary conditions at the unification scale, thus introducing additional parameters to allow for non-universal Higgs masses m_H or m_{H_u} and m_{H_d} at the GUT scale. Such extensions of mSUGRA are needed, for example, to make mSUGRA predictions compatible with the observed neutralino relic density [22].

Furthermore, there are a number of models, where instead of gravity other mediators communicate between the visible MSSM and the hidden sector. In Gauge-Mediated Supersymmetry Breaking (GMSB) models, new chiral supermultiplets are introduced as messenger fields that interact with the hidden sector. They are also charged under the Standard Model gauge group and thus couple to MSSM particles through gauge boson and gaugino interactions. In GMSB, the gravitino often is the LSP. This makes the phenomenology quite different from gravity-mediated models, and typically the Next-Lightest Supersymmetric Particle (NLSP) decays to a photon and the gravitino.

The simplest supersymmetric extension of the Standard Model with a scale-invariant superpotential beyond the MSSM is the Next-to-Minimal Supersymmetric Standard Model (NMSSM): It solves the μ -problem by introducing an additional singlet superfield. In addition to an extended Higgs sector with five neutral Higgs states, the NMSSM also has an extended neutralino sector with five neutralinos including the singlino. This leads to significant modifications of all sparticle decay cascades compared to the MSSM if the LSP is singlino-like. Again, constraints can be made to reduce the now even larger number of free parameters to a smaller set, as is done in the models derived from the MSSM.

As said above, the superpotential of the MSSM does not include terms which violate the conservation of the multiplicative quantum number R , which was introduced to forbid terms which would lead to a decay of the proton. There are possibilities, though, to replace the R -parity by other symmetries like baryon-triality or lepton-parity [17], which then allow for violation of the conservation of either the lepton number L or baryon number B , while still guaranteeing stability of the proton. In these R -Parity Violating (RPV) models, additional couplings are introduced in the superpotential, which open up new phenomenological signatures. In general, R -parity violation also means that the LSP no longer is stable and may decay, potentially giving displaced vertices in the detector.

3 Experimental Setup

The data evaluated in this thesis has been collected with the ATLAS detector. The ATLAS detector is one of the four large-scale detectors used to probe proton-proton collisions that are produced at the Large Hadron Collider (LHC) at CERN. The first part of this section gives a general overview of particle detectors, with a focus on selected topics which are particularly relevant for this thesis, like calorimeters, missing transverse energy and pile-up. It is intended to provide definitions of the fundamental notions which are used in the following. In the second part, after a short introduction with respect to the rôle of CERN in particle physics, a description of the Large Hadron Collider, the largest collider ring currently in operation, is given. Afterwards, the hardware of the ATLAS detector is described.

3.1 Conventions and Units

In particle physics, it is common to set $c = \hbar = k_B = 1$, where c is the speed of light in vacuum, $\hbar := h/2\pi$ is the reduced Planck constant and k_B is the Boltzmann constant. This is called “natural units” and will be adopted in the following, in particular with respect to masses and momenta of particles, which will be given in electron volts as a unit of energy ($1 \text{ eV} = 1.602 \cdot 10^{-19} \text{ J}$). The additional factors of c , \hbar and k_B will not be written explicitly, but can be found from dimensional analysis, e. g. the unit of momentum is eV/c , for masses it is eV/c^2 and for lengths it is $\hbar c/\text{eV}$. The typical energy scale for particles produced at current colliders is GeV. In nuclear physics, lower energy ranges are studied (MeV range), astroparticle physics covers a broad spectrum up to very high energies in cosmic rays with 10^{20} eV .

An important unit in particle physics and nuclear physics is barn, which is a unit of area and used to express cross sections and, by its inverse, luminosities (cf. Section 3.2.4). It is defined to be 10^{-28} m^2 . In spite of the large negative value of the exponent, this is still a very large cross section and therefore typically used together with the SI prefixes for pico ($1 \text{ pb} = 10^{-40} \text{ m}^2$) or femto ($1 \text{ fb} = 10^{-43} \text{ m}^2$).

The following symbols are conventionally used: E is the energy of a particle and \vec{p} its vector three-momentum, which together constitute the contravariant four-momentum $P = (P^\mu) = (E, \vec{p})$, $\mu = 0, 1, 2, 3$. The transverse component of the three-momentum, given by its projection onto the x - y plane, is denoted by \vec{p}_T . Its absolute value is p_T . In analogy, the transverse energy E_T is defined as $E_T := E \cdot p_T / |\vec{p}|$. The invariant mass of a system of N particles is given by the sum of their energies and momenta as

$$m^2 = \left(\sum_i^N E_i \right)^2 - \left(\sum_i^N \vec{p}_i \right)^2. \quad (3.1)$$

The invariant mass of a decaying particle is equal to the invariant mass of its decay products, and in the center-of-mass frame, the invariant mass is the total available energy and therefore gives the kinematic boundary for particle production. The center-of-mass energy in

high-energy physics is usually denoted by the square root of s , one of the Lorentz-invariant Mandelstam variables, defined as

$$\sqrt{s} = \sqrt{P_\mu P^\mu}, \quad (3.2)$$

where $P = P_1 + P_2$ is the sum of the four-vector momenta of the two initial particles in a two-body reaction and summation over repeated indices is understood (Einstein convention). In the center-of-mass system, the total momentum is zero and \sqrt{s} is again the energy available for particle production. In a fixed target experiment, the center-of-mass energy scales as $\sqrt{s} \sim \sqrt{2Em_t}$, where m_t is the mass of the resting target particle and E is the energy of the beam particle in the laboratory frame [12]. In a collider experiment, where both of the colliding particles are accelerated to have an energy E and collided head-on, much higher energies can be reached, because here $\sqrt{s} = 2E$.

To account for the rotational symmetry around the beam axis in collisions of elementary particles at a zero crossing angle, usually a cylindrical or spherical coordinate system is used, which will be introduced in Section 3.4.

3.2 Collision Detectors in High-Energy Physics

The purpose of the detector is to determine the kinematic properties of the particles produced in high-energy particle collisions. The final goal of the measurement is to correctly identify as many of the particles as possible. Typical quantities that are used to describe the particles produced in collision experiments are their positions, momenta and energies, and the impact times of the particles. The impact time is normally used to group particles according to the collisions from which they emerged.

Basically, there are two types of instruments in a detector. Tracking devices measure the momentum \vec{p} of charged particles from their deflection in a magnetic field. Calorimeters measure the energy E of particles, which is only possible by stopping them. They must thus be optimized to maximize the energy loss and surround the tracking devices. Unlike all other charged Standard Model particles, the energy loss of muons is usually too small for them to be stopped in the calorimeters. Muons can then be identified in muon spectrometers surrounding the calorimeters, where a supplementary measurement of their momentum can be done. The resulting sequence of instrumentation, tracking devices, calorimeters, muon spectrometers, which is typical of general multi-purpose particle detectors, is also reflected in the design of the ATLAS detector, which is described in detail in Section 3.4.

3.2.1 Tracking Devices

Charged particles can be detected by the ionization they produce when passing through a medium, or by the production of photons as scintillation or Čerenkov light or as transition radiation. Tracking devices aim at determining the space and time coordinates of charged particles and are mostly used in magnetic spectrometers, which additionally allow a measurement of the momentum of the particle via

$$p = |\vec{p}| = 0.2998 \frac{\text{GeV}}{\text{T m}} \cdot q B r, \quad (3.3)$$

if the strength of the magnetic field B and the bending radius r of the particle trajectory are known. Equation (3.3) is written down for a particle with charge q in units of the electron

charge magnitude e . From the sign of the right-hand side in Equation (3.3), the sign of the charge of the particle can be determined. The relative error of the momentum measurement asymptotically scales as $\sigma(p)/p \sim p^{-1}$ [10]. For very fast or heavy particles, the deflection of the track in the magnetic field is small and thus the curvature of the track difficult to measure. In addition to the momentum measurement, tracking devices are needed for the reconstruction of primary and secondary interaction vertices, where the measurement of the distance of the secondary vertices allows the determination of lifetime and thus identity of sufficiently slowly decaying particles. Tracking devices are usually the innermost subdetector close to the beam pipe and interaction point, which requires a high radiation hardness. They also need to be optimized for an energy loss of the particles as small as possible so that they do not interfere with the energy measurement in the calorimeters.

The two main techniques for tracking devices are gaseous detectors and solid state detectors [16]. When an ionizing particle passes through a gas-filled detector, it creates electron-ion pairs. If the electric field applied to the gas volume is sufficiently strong, the ejected electrons can further ionize the medium and create an avalanche and thus a high signal gain. For small electric fields, the electrons only drift along field lines to be captured by anodes. By choosing the voltage appropriately, gaseous detectors can be run in proportional mode, in which the signal amplitude is proportional to the primary charge produced within the sensitive volume. Multiwire proportional chambers consist of a planar cathode and tautened anode wires in a plane parallel to the cathode. A segmentation of the cathode gives additional information along the wire direction. Alternatively, the charge division over resistive sense wires can be employed. In a drift chamber, the constant drift velocity in the region with a homogeneous electric field is exploited to find the position orthogonal to the cathode plane. Straw-tube chambers are long cylindrical chambers with a central anode wire and are often combined in modules or layers. Resistive plate chambers are low-cost, non-proportional detectors with flat, high-resistivity electrodes, between which violent discharges are induced by ionizing particles. They have an excellent timing-resolution and provide short pulses that can be used for triggering.

Solid state detectors use a semiconductor material, usually silicon or germanium, which is doped to have a pn -junction. Through application of a bias voltage, a carrier-free depletion zone is created, which acts as a kind of ionization chamber, in which charged particles create electron-hole pairs. For this to happen, much less energy is needed than in gaseous detectors, the energy being given by the band gap of the semiconductor material, which is of the order of a few electron volts. The charge is collected in an external electric field. Spatial information can be obtained from a segmentation of the electrodes, which may be reflected in the doped regions as well. If the electrodes are approximately quadratic, the device is called a pixel detector; if it is a very elongated rectangular shape, it is called a strip detector. Using orthogonal strips yields two-dimensional track information. The irradiation of the doped semiconductor material will at some point lead to an n to p type inversion, and thus a reversed and with time increasing bias voltage is needed. Recent research on solid state detectors includes using diamond as semiconductor material because of its much higher radiation tolerance and the development of three-dimensional detectors. These have electrodes going through the bulk of material, which allows to improve the charge collection efficiency and to lower the depletion voltage without changing the sensor thickness [16].

3.2.2 Calorimeters

A calorimeter is basically an instrumented block of dense absorbers, in which the produced particles are stopped and (ideally) all of their energy is converted in a very short time into a measurable quantity (light or charge), which is then detected by embedded sensors. This is thus a destructive measurement. It proceeds in a cascade of interactions and the depth needed for a complete absorption grows logarithmically with the energy of the particle.

Calorimeters are important for two reasons: They can measure charged and neutral particles by detecting the charged secondaries, and the asymptotic behavior of the uncertainty of the energy measurement for high energies is in principle $\sigma(E)/E \sim E^{-1/2}$ due to the stochastic nature of the absorption process. This means the more energy the particle carries, the smaller the relative uncertainty because of the higher number of particles produced in the cascade^a. This is complementary to the momentum measurement. Not only the energy deposited in calorimeters is of interest, but also the timing and the pulse shape as a measure of the quality of the energy determination. Depending on the segmentation and granularity of the calorimeter, also the direction and impact point of the incident particle can be measured.

In general, highly energetic particles will have enough energy to initiate the production of secondary particles and thus to start the creation of an avalanche or shower of particles. The shower continues until the energy of the particles produced in the showering process falls below the threshold of energy needed to create further particles, or until other processes of energy dissipation become dominant. The energy of the particles that have been produced up to this point is then absorbed via low-energy processes.

It is common to equip detectors with two types of calorimeters, one specialized to electrons and photons in front of one specialized to hadrons. Photons, electrons and positrons require less material to be stopped compared to hadrons, which interact mainly through the strong interaction. This is also expressed by the radiation length X_0 , which is the characteristic length for the development of an electromagnetic shower by pair production and bremsstrahlung, and is smaller than the nuclear absorption or interaction length λ , which gives the mean free path or, equivalently, the probability of an inelastic collision over a certain distance for a hadron passing through the material.

Electromagnetic showers are initiated by electrons, positrons or photons and are alternating cascades of photon emission via bremsstrahlung and creation of electron-positron pairs. Electrons and positrons above a critical energy mainly lose energy via bremsstrahlung, and at low energies through collisions with material, leading to ionization and excitation. Photons at high energies mainly lose energy via e^+e^- -pair production, and at low energies through Compton scattering and the photoelectric effect. The expected energy loss for an electron above the critical energy is given by the radiation length X_0 , $\langle E(x) \rangle = E_0 \exp(-x/X_0)$. For photons, the intensity goes as $\langle I(x) \rangle = I_0 \exp(-7x/9X_0)$ due to pair production. The scale of electromagnetic showers is therefore determined by X_0 . Hadronic showers are initiated by hadrons like pions, kaons or protons and consist of secondary hadrons produced via strong interactions in inelastic scattering with the nuclei of the calorimeter material. The scale of the hadronic cascade is governed by the nuclear absorption length λ . The charged component of the cascade can be measured through its ionization and excitation effects, and photons from the de-excitation of excited nuclei can be detected directly.

^a See below for a full discussion of the parametrization.

The signal response of a calorimeter to hadrons will on average be lower than that to an electron with the same energy due to several effects [10]: In a hadronic shower, some of the produced particles like neutrinos or muons simply escape from the calorimeter, other particles like neutrons, fragments from spallation or recoiling nuclei are only partially detectable. A significant fraction of the energy in hadronic showers is removed by the production of neutral pions, which will decay predominantly electromagnetically, $\pi^0 \rightarrow \gamma\gamma$, thus leading to a one-way conversion to an electromagnetic shower component, which will give the same signal yield as an electron. The ratio of the signal responses to pions (as representative of hadronic particles) and electrons is usually denoted π/e and of the order of $\pi/e \approx 0.8$. If $\pi/e \neq 1$, this also makes the calorimeter response to hadrons a non-linear function of the energy, because the electromagnetic component increases with larger shower lengths [16].

Calorimeters fall into two groups: In a sampling calorimeter, an absorber material as passive component is sandwiched with active layers containing signal sensors. Sampling calorimeters in general suffer from increased signal fluctuations due to their structure (sampling fluctuations). On the other hand, in a sampling calorimeter it is possible to tune several variables, and by increasing the hadronic and decreasing the electromagnetic response to bring the response ratio π/e back to one. Calorimeters designed to this aim are called “compensating” calorimeters. In the simpler but often larger or more expensive homogeneous calorimeter, the absorber material also serves as detecting medium, for example in the calorimeter of the CMS detector at the LHC. For the energy measurement itself, there are two primary detection techniques. In inorganic or organic scintillators, molecules are excited through interactions with the incident particle and subsequently emit a small fraction of the absorbed energy in form of photons, which can then be detected using photodetectors, usually photomultipliers. In liquefied noble gases or semiconductors, the ionization in terms of electron-ion or electron-hole pair creation can be measured directly by applying an external electric field to collect the charge created by a particle passing through the material.

The conventional parametrization for the relative energy resolution of a calorimeter consists of three terms with different scaling behavior with E [87],

$$\frac{\sigma(E)}{E} = \frac{a}{\sqrt{E}} \oplus \frac{b}{E} \oplus c, \quad (3.4)$$

where \oplus stands for the quadratic sum, $x \oplus y := \sqrt{x^2 + y^2}$. The parameter a gives the stochastic term, accounting for intrinsic stochastic fluctuations in the number of particles produced in the shower process, which are proportional to $\sqrt{N} \sim \sqrt{E}$, but also e. g. for sampling fluctuations. The constant term b , dominating the resolution for low energies, is the noise term. The systematic term c , dominating for high energies, comes from instrumental contributions to the energy resolution like detector non-uniformities, incomplete shower containment (leakage) or calibration uncertainties which scale as $\sigma(E) \sim E$. For measurements of hadronic shower energies, non-compensation also contributes to this term.

Jets

The hadronic jets, which are measured as clustered energy depositions in the calorimeter, are initiated by partons, gluons and quarks, but the particles which interact with the detector material are hadrons and not the partons themselves, which due to confinement cannot leave the interaction point as isolated entities (cf. Section 2.1.3). A jet can thus be defined at




Figure 3.1: Schematic view of the hadronization and jet formation process in a style similar to a Feynman diagram. The curly lines represent gluons, the straight lines quarks and anti-quarks. The exact process of hadronization is unknown and therefore drawn as grey blobs, from which hadrons emerge as dotted lines. Some examples for hadrons have been put.

different levels: At parton level, where every high-energy (hard) parton corresponds to a jet, or at particle level as the collection of hadrons produced from a parton in the hadronization process. Finally, from the point of view of the calorimeter measurement, a jet is defined as a clustered energy deposition in the cells of the calorimeter, which is due to one original particle, be it a prompt electron or photon in case of electromagnetic jets or a parton from the parton level jet in case of hadronic jets. The algorithms that are used to reconstruct the original jet structure from the energy depositions in the calorimeters are described in Section 4.4.2.

Missing Transverse Energy

Many models of physics beyond the Standard Model like Supersymmetry predict the existence of stable and heavy particles which only interact weakly (and gravitationally) with matter. Besides the precise measurement of the particles like electrons, muons or hadrons in terms of jets, the detection of particles which only very rarely interact with the detector material is therefore of great importance. The direct detection of these particles usually is not possible, unless the detector is specifically designed for this purpose and uses very large volumes and sophisticated methods for background suppression. But although these weakly interacting particles are invisible to general purpose particle detectors like ATLAS, in the sense that they almost always evade direct detection, they show up through the momentum and energy they carry away. Due to the principle of momentum conservation, in symmetric particle collisions the sum of all momenta must be zero because the total momentum of the initial particles is zero. If a particle leaves the detector unnoticed, its momentum will be missing in a measurement of the total momentum computed as a vectorial sum of the momentum of all detected particles, and lead to a non-zero result. Of course, a lot of detector effects like noise and mismeasurements will also contribute to the apparent deviation in the sum from zero as discussed below. At hadron colliders, there is an additional complication. As the partons in the initial state of the interaction carry only a fraction of the proton momentum (cf. Section 2.1.3), their momentum along the beam axis direction is unknown, and only the conservation of the transverse component of the momenta can be used, by assuming that the transverse momentum of the initial state partons is zero. The missing transverse

momentum is then defined as

$$\vec{p}_T := - \sum_i \vec{p}_T(i), \quad (3.5)$$

where the sum runs over all visible particles in the final state [87]. The missing transverse momentum is identified with the sum of the transverse momenta of all invisible particles.

The momentum of a particle can be measured in tracking devices only for charged particles, therefore instead of the momentum usually the energy is employed. In the high energy approximation, according to which particle masses are negligible, the magnitude of the momentum of a particle is equal to its energy, which can be inferred from a measurement in the calorimeter. The missing transverse energy (\cancel{E}_T , also abbreviated MET rather than MTE) is then defined as the absolute value of the vectorial sum of all contributions to the transverse energy,

$$\cancel{E}_T := \left| - \sum_i \vec{E}_T(i) \right|. \quad (3.6)$$

In practice, this is computed from a weighted projection of all calorimeter cells onto the transverse plane (cf. Section 4.4.3). Muons as minimum ionizing particles^b typically only deposit a small fraction of their energy in the calorimeters. To account for this, the measurement of the calorimeter \cancel{E}_T can be corrected by using the momentum information from the muon spectrometers. A quantity related to \cancel{E}_T is the sum of the absolute values of the transverse energies,

$$\sum E_T := \sum_i \left| \vec{E}_T(i) \right|. \quad (3.7)$$

The value of $\sum E_T$ gives an indication of the total activity in an event.

In collisions in which no invisible particles are produced, an ideal detector would measure $\cancel{E}_T = 0$. In reality however, there are many detector effects which contribute to the \cancel{E}_T measurements. In particular due to the definition of \cancel{E}_T as the sum of squares of the x and y components of the missing energy,

$$\cancel{E}_T = \sqrt{\cancel{E}_x^2 + \cancel{E}_y^2}, \quad (3.8)$$

although the mean of the noise of the energy measurement in the calorimeter is zero for both components individually, it will give a non-zero contribution to \cancel{E}_T . This is discussed in detail in Section 6. Additional contributions to \cancel{E}_T come from non-instrumented parts of the detector, “dead material” such as cooling, cables for read-out and powering and support structures and crack regions in the calorimeter, where deposited energy cannot be measured and therefore is lost, and mismeasurements of the \cancel{E}_T from the limited resolution and calibration effects. All these detector effects will be referred to as fake \cancel{E}_T in the following as opposed to real \cancel{E}_T , which is induced by non-interacting particles like the Standard Model neutrinos or neutralinos from Supersymmetry and is the quantity which is of interest with respect to physics analyses. In addition, there is also the notion of true \cancel{E}_T , which refers to the Monte Carlo information about the actual amount of real \cancel{E}_T in simulated events. Note that the presumed conservation of the transverse momentum neglects a number of small

^b A minimum ionizing particle can be defined as a charged particle with a such a velocity that its energy loss is near the minimum of the Bethe equation describing the stopping power $\langle -dE/dx \rangle$ as function of $\beta\gamma = p/mc$ [87].

effects, giving rise to a non-zero transverse momentum even if all particle momenta were known exactly. Such effects are, for example, that the partons in the initial state of the collisions have a small, but non-zero p_T , or the non-zero crossing angle between the colliding beams (cf. Section 3.3.2), leading to a residual momentum in the y direction of the order of a GeV [141].

3.2.3 Muon Spectrometers

Muon spectrometers are tracking devices, too, but specifically designed for muons. As said above, they typically lie outside the calorimeter, and they usually cover a large volume with low material density to avoid scattering of the muons. Different types of muon chambers are described in the context of the muon spectrometer of the ATLAS detector in Section 3.4.4.

3.2.4 Luminosity

The luminosity governs the rate of interactions in a collider experiment. It has unit of a flux, Hz/cm², often also written as cm⁻²s⁻¹. The instantaneous luminosity in a collider experiment is given by

$$\mathcal{L}_{\text{inst}} = fn \frac{N_1 N_2}{A}, \quad (3.9)$$

where N_i is the number of particles in the bunches of the two beams which are brought to collision, f the revolution frequency and n the number of particle bunches in each of the beams. While the number of particles in each bunch can differ, especially if the two beams consist of two different types of particles, n usually is the same for both beams. A is the cross-sectional area of the beams at the interaction point, which can be expressed as $A = 4\pi\sigma_x\sigma_y$. The transverse beam profile can in many cases be assumed to be Gaussian [87] and is characterized by the widths $\sigma_{x,y}$ in the horizontal and vertical direction. In order to make the dependence on the collider machine parameters explicit, this can be written as [132]

$$\mathcal{L}_{\text{inst}} = fn \frac{N_1 N_2 \gamma_r}{4\pi\epsilon_T \beta^*}, \quad (3.10)$$

where ϵ_T is the transverse emittance of the beams, β^* is the wavelength of the betatron oscillations of the beams [6], and γ_r is the relativistic factor of the colliding particles.

The integral of the instantaneous luminosity over time gives the integrated luminosity

$$\mathcal{L}_{\text{int}} = \int_{t_1}^{t_2} \mathcal{L}_{\text{inst}} dt. \quad (3.11)$$

The amount of data collected over a certain interval of time is specified in terms of the integrated luminosity. The number of events of a given type N , which is expected to be contained in a data sample, is related to the corresponding integrated luminosity via the total inclusive cross section σ ,

$$\langle N \rangle = \sigma \cdot \mathcal{L}_{\text{int}}. \quad (3.12)$$

The product of the instantaneous luminosity and the cross section for the relevant scattering process gives the reaction rate $\langle R \rangle = \sigma \cdot \mathcal{L}_{\text{inst}}$. Equation (3.12) is important for the normalization of samples with Monte Carlo simulated events to the number of events expected in




Figure 3.2: Plot of the instantaneous luminosity as recorded by ATLAS, for fill number 1225 of the LHC.

the sample of real detector data. The number of events N_0 in the Monte Carlo sample needs to be scaled to the integrated luminosity \mathcal{L}_{int} of the data sample, using the respective cross section σ computed by the Monte Carlo generator. The event weight serving as scaling factor is given by $f = \langle N \rangle / N_0$.

The instantaneous luminosity decreases exponentially with time during a fill of the collider due to several effects with different time constants. Normally, the decrease of the instantaneous luminosity should be dominated by the reduction of the number of particles in the bunches due to collisions rather than due to any other losses like beam-gas interactions (cf. Section 3.3.2). Figure 3.2 shows exemplary the exponential decay of the instantaneous luminosity for one of the first long runs of the Large Hadron Collider (LHC) and the ATLAS detector in 2010. The plot shows in three different colors the luminosity delivered by the LHC, the luminosity delivered during the time when the beams were declared to be stable by the LHC operators, and the luminosity that was actually recorded by the ATLAS detector. After about 15 hours, the beams were dumped and the data taking ended.

3.2.5 Pile-up

In the context of trigger rates, and in particular in the discussion of the \cancel{E}_T model in Section 6, the notion of pile-up will be important. Pile-up means that the particles produced in more than one particle-particle collision reach the detector at the same time, or more generally that their signals are overlaid in a way that they cannot be disentangled. It does, however, not refer to the effect that due to the large extent of the detector, the small temporal bunch spacing and the finite speed of light, the products emerging from several subsequent bunch crossings coexist in the detector volume, as will become clear below from the explanation of the two different types of pile-up. Beforehand, it is useful to introduce the expected average number of interactions per bunch-crossing.

Average Number of Interactions per Bunch-crossing

When bunches of particles are collided, i. e. they run through each other, the probability for an interaction is proportional to the particle densities, or better, including the particle velocity, the flux, which is expressed in terms of the instantaneous luminosity. The number of actual particle collisions taking place when two particle bunches cross (dubbed a *bunch crossing*) is a random variable that follows a Poisson distribution. For low instantaneous luminosities, in most bunch crossings nothing will happen; for high instantaneous luminosities, in most bunch crossings several particle collisions will take place at the same time. In the following, $\langle\mu\rangle$ is defined as the (expected) average number of (inelastic) proton-proton interactions per bunch-crossing. It is an important quantitative measure for the activity within an event and can be calculated from the instantaneous luminosity as shown in Section A.2.

Like the instantaneous luminosity, $\langle\mu\rangle$ depends on the number of protons in the colliding bunches, which is not necessarily the same for all bunches in the collider ring because it may vary depending on the injection quality conditions. Nevertheless, in the following it is a good approximation to assume for simplicity that it is the same for all bunches.

In-time Pile-up

Depending on the value of $\langle\mu\rangle$, in a significant fraction of bunch crossings more than one hard proton-proton interaction will take place concurrently, so that the particles produced in these collisions reach the surrounding detector at the same time. Due to its large inclusive cross section, most of the interactions will be minimum-bias events (cf. Section 4.5.3), and only very rarely will two processes with small cross sections (e. g. W production, cf. Figure 3.6) take place in the same bunch crossing. In Monte Carlo simulations, in-time pile-up is thus incorporated by overlaying the signals from the physics process of interest with signals corresponding to a given number of additional independent minimum-bias collisions.

Depending on the subdetector and type of measurement, it may or may not be possible to distinguish between particles coming from different concurrent proton-proton interactions. For the measurement of \cancel{E}_T and in the trigger, it is in general not possible. For charged particles, it may be possible to trace back their origin to spatially separated production vertices. The impact of in-time pile-up on the rates of the \cancel{E}_T trigger is studied in detail in Section 6.

Out-of-time Pile-up

Out-of-time pile-up subsumes effects which originate from the time the detector needs to return to its stand-by state being longer than the temporal bunch spacing. This means that signals from particles produced in a subsequent collision reach the detector before it is ready to start the next measurement. While in-time pile-up will always lead to an increase of the activity in the event, to more particles appearing in the detector and to more energy being deposited in the calorimeter, out-of-time pile-up may also lead to the opposite effect in calorimeter measurements. Due to the pulse shape of the calorimeter signals, which exhibit an undershoot before coming back to the zero baseline, energy depositions which occur at the same position in the calorimeter, but which are separated by a certain time interval, will lead to a destructive interference in the measurement if the second pulse starts during the undershoot of the first pulse. This may lead to a considerable negative bias in the energy measurements that needs to be taken into account when running at high luminosities with




Figure 3.3: Calibration plot from September 2011 showing the pulse shape of one read-out channel of the electromagnetic barrel calorimeter of ATLAS (side A). The horizontal axis is the time in nanoseconds, the vertical axis the amplitude of the pulse in Analog-to-Digital Converter (ADC) counts.

short bunch distances that allow for out-of-time pile-up. Figure 3.3 gives an example of a typical pulse shape in the electromagnetic liquid-argon calorimeter of ATLAS. It can be seen that the peak of the signal is followed by a long undershoot, which spans tens of bunch crossings at the nominal bunch spacing of 25 ns (see below).

Pile-up and Luminosity

Equations (3.10) and (3.11) show what possibilities there are to increase the luminosity. These have different implications for the amount of pile-up. Increasing the integration time (“physics time” with stable beam conditions) or the number of colliding bunches n in the collider does not increase the level of in-time pile-up because the average number of interactions per bunch crossing stays the same. However, changes in any of the other parameters which govern the instantaneous luminosity per bunch, i. e. an increase in the number of particles in the bunches $N_{1,2}$ or a reduction of the beam parameters ϵ_T and β^* , all do increase in-time pile-up. (The revolution frequency f and relativistic factor γ are fixed.)

3.3 CERN and the LHC

3.3.1 L’Organisation Européenne pour la Recherche Nucléaire

CERN is an international organization with 20 member states and is one of the largest centers for scientific research. It is situated near Geneva on the border of Switzerland and France and operates the world’s largest particle physics laboratory. Its official name today is “European Organization for Nuclear Research” or « L’Organisation Européenne pour la Recherche Nucléaire » in French. The acronym CERN stood for the French “Conseil Européen pour la Recherche Nucléaire” (European Council for Nuclear Research), which was the name of the

council founded in February 1952 with the mandate to establish a fundamental physics research organization in Europe. At the end of September 1954, after the ratification by France and Germany, the European Organization for Nuclear Research officially came into being and the council was dissolved, but the name CERN was retained.

A variety of particle accelerators and colliders have been built on CERN grounds since its foundation, some of which, like the Proton Synchrotron (PS) started up in 1959, are still in use despite their old age in terms of a machine used for research. The PS is an important part of the accelerator chain which feeds the LHC (see below). Besides accelerating particles, also slowing down particles is a scientific occupation, for example in the ATRAP experiment, which aims at comparing hydrogen atoms with their antimatter equivalents, for which the ingredients, antiprotons and positrons need to be cooled down so that they can form anti-hydrogen. So far, five Nobel Prizes were awarded to CERN physicists, in connection to the discovery of the weak bosons, the J/ψ meson, the muon neutrino and detector technology and methods. CERN also claims to be the place where the World Wide Web has been invented, originally for sharing information between scientists working in different places all over the world.

3.3.2 The Large Hadron Collider

The Large Hadron Collider (LHC) is a circular (synchrotron) proton-proton accelerator and collider, designed for a peak instantaneous luminosity of 10^{34} Hz/cm² and an energy of 7 TeV per beam, giving a center-of-mass energy of $\sqrt{s} = 14$ TeV available in collisions [72]. It is currently the world's largest particle accelerator and located about 100 m underground near Geneva, where it spans the border between Switzerland and France. Its primary experimental goals are to study the mechanism of electroweak symmetry breaking, find the Higgs boson and search for indications of physics beyond the Standard Model. A summary of relevant parameters of the LHC is given in Table 3.1.

The LHC occupies a tunnel with a circumference of 27 km which was originally excavated for the Large Electron-Positron Collider (LEP) that ran successfully until the end of 2000. The LEP required such a large radius due to the synchrotron radiation of the electrons and positrons. The major differences with respect to the older Tevatron, the hadron collider with the second highest energy in the world operated in the U.S. at Fermilab, are the higher center-of-mass energy and the fact that at the LHC proton pairs are collided, whereas at the Tevatron protons and antiprotons are collided. This leads to significant differences in the composition of the initial states. In addition to colliding protons with protons, there are also (shorter) phases in the operation of the LHC with dedicated heavy ion runs, in which lead ions (Pb^{82+}) are collided with lead ions. This allows to give an even higher total center-of-mass energy, although the energy per nucleon of 2.76 TeV is lower than in proton-proton collision mode.

The center-of-mass energy at a hadron collider like the LHC is currently limited by the available magnet technology. A short calculation (cf. Equation (A.14)) shows that per TeV of beam energy, the magnetic field needs to be increased by about 1 Tesla to bend the protons onto a track with the radius of the LHC tunnel. In fact, the magnetic field actually needs to be higher, as the LHC is not a perfect circle, but sections equipped with dipole magnets are interspersed with straight segments. The effective radius is thus smaller and the bending force needs to be larger. In total, the LHC has 9300 magnets with superconducting Niobium-titanium (NbTi) coils, which are cooled with liquid nitrogen and superfluid helium, in which




Figure 3.4: Layout of the LHC infrastructures. The ATLAS detector is housed at Point 1, one of the four interaction points. (Note that this image is from 1993 and the “future constructions” by now have been completed.)

they are immersed, to an operating temperature of 1.9 K. The magnetic field of the 1232 main bending magnets is designed to be 8.33 T at the maximum current through the magnets of 11.8 kA. They make up about two thirds of the length. In the remaining length, there are quadrupole magnets for focusing, the beam injection points, the beam dump, beam cleaning stations and long straight sections, at which also the four interaction regions with the main detectors ALICE, ATLAS, CMS and LHC-B are located. The interaction regions, in which the two proton beams cross, are distributed over the eight octants^c into which the LHC is divided (cf. Figure 3.4). In ATLAS and ALICE, the beams cross in the vertical plane, in CMS and LHC-B in the horizontal plane. At Points 2 and 8, the proton beams from the SPS, the last step of the accelerator chain, are injected into the LHC.

The accelerator chain for the injection of protons into the LHC is the following (for lead ions the chain it is slightly different). Protons for each beam are accelerated in the Linac2, a 30 m long linear accelerator, to 50 MeV, then transferred into the Proton-Synchrotron-Booster (PSB) to be accelerated to 1.4 GeV, extracted to be accelerated to 25 GeV in the Proton-Synchrotron (PS), and then to 450 GeV in the Super-Proton-Synchrotron (SPS), from which they are injected into the LHC to be ramped up to the final center-of-mass energy. The acceleration of the protons in the LHC is done using a sinusoidal electric field at a frequency of 400.8 MHz in eight superconducting radio frequency cavities per proton beam at Point 4 with a maximum power of 4.8 MW.

The beam crossing frequency of the LHC is $f_{\text{LHC}} = 40.079$ MHz, which results in an approximate 25 ns width of the 3564 time buckets and a revolution frequency of 11.245 kHz. The number of time buckets is directly related to the relations of the circumferences of the preceding accelerators, the PS and SPS [16]. Due to hardware restrictions, in particular from the rise times of the injection kicker magnets, after subtracting abort and injection gaps, only

^c Correspondingly, there are eight special locations, one in the middle of each octant, which are numbered clockwise around the ring and called Point 1 through Point 8. Points 1, 2, 5 and 8 house the ATLAS, ALICE, CMS and LHC-B detector, respectively.




Figure 3.5: Integral of the delivered and recorded luminosity for ATLAS in 2011 as function of time. Over 5 fb^{-1} have been collected in 2011 at $\sqrt{s} = 7 \text{ TeV}$.

a maximum of 2808 of the 3564 time buckets can be filled with proton bunches. Buckets which contain a proton bunch are often called filled bunches in the following, buckets without protons are called empty bunches. Finally, proton bunches in buckets that are brought to collision at the interaction point in ATLAS are called colliding bunches. The distribution of the proton bunches over buckets is usually given in terms of bunch trains, which describe a certain repeated pattern of filled and empty buckets.

The operation of the LHC requires a very good vacuum to reduce beam-gas interactions. Otherwise, proton bunches would be rapidly depleted and detector measurements would suffer from high background noise levels. The vacuum also determines the single-beam lifetime, which is the time in which the number of particles in stable orbit is reduced by a factor of $1 - \exp(-1)$ due to beam-gas interactions, and is designed to be significantly larger than the luminosity lifetime of 45 hours due to the proton-proton collisions [131]. The required equivalent hydrogen density to achieve this is 10^{15} m^{-3} , close to the detectors it will be even a 100 times better.

After an incident in September 2008 in the commissioning phase of the LHC [191], which revealed problems with the bus bars, electrical interconnectors between the magnets, it was decided that it is only safe to run the LHC at a reduced center-of-mass energy of 7 TeV, 3.5 TeV per beam, until all the superconducting interconnect splices will have been overhauled, which will require a long consolidation shutdown [176]. In March 2010, for the first time proton-proton beams were collided at a center-of-mass energy of 7 TeV in the LHC. Since then, the instantaneous luminosity has been constantly increasing. At the time of writing (November 2011), the peak instantaneous luminosity is already above $3.5 \cdot 10^{33} \text{ Hz/cm}^2$ [200], so there is only little needed to reach the design luminosity of the LHC. Part of this can be achieved by reducing the bunch spacing from 50 ns to 25 ns, which will double the instantaneous luminosity. In 2010, the LHC has delivered 49 pb^{-1} of integrated luminosity, in 2011 over 5.6 fb^{-1} . A plot of the evolution of the integrated luminosity over time is shown in Figure 3.5 for 2011. The increase in the instantaneous luminosity is reflected by the increase

Parameter	Value
Proton energy	7 TeV
Relativistic γ factor of protons	7461
Number of proton bunches	2808
Number of protons per bunch	$1.15 \cdot 10^{11}$
Transverse emittance	3.75 μm
β function at IP1 and IP5	0.55 m
Half crossing-angle	$\pm 142.5 \mu\text{rad}$
Peak luminosity at IP1 and IP5	$1 \cdot 10^{34} \text{ Hz/cm}^2$
per bunch crossing	$3.56 \cdot 10^{30} \text{ Hz/cm}^2$
pp collisions per bunch crossing	19
Revolution frequency	11.245 kHz
Bunch spacing	24.97 ns
Beam crossing frequency	40.079 MHz
RF frequency	400.790 MHz
Stored energy per beam	362 MJ
Circulating beam current	0.582 A
Beam current lifetime	14.9 hours
Ring circumference	26 658.883 m
Number of main bends	1232
Field of main bends	8.33 T
Bending radius	2803.95 m

Table 3.1: Summary of relevant LHC parameters [131]. The parameters given in this table are the design parameters. At the moment (2011), the energy per beam is only 3.5 TeV and the design luminosity not yet reached (see text). IP1 and IP5 are the interaction points at the center of the ATLAS and CMS detector (cf. Figure 3.4).

of the slope in the integrated luminosity.

Cross Sections at the LHC

Figure 3.6 compares the cross sections of several important processes as function of the center-of-mass energy. The center-of-mass energies of the LHC and the Tevatron are indicated by the dashed vertical lines. Some of the curves in the plot exhibit a discontinuity, which has nothing to do with the center-of-mass energy being 4 TeV. Instead, these are processes which have different cross sections at the Tevatron, being a proton-antiproton collider, and the LHC, being a proton-proton collider, due to the nature of the dominant production mode. For processes where gluon fusion is the dominant production mode, like for Higgs bosons, the cross section curves have a smooth transition between the Tevatron and LHC. The vector bosons W and Z get significant contributions to their production from quark-antiquark annihilation, and therefore the cross section at a fixed center-of-mass energy is slightly larger at the Tevatron because the antiquark here can be a valence quark from the antiproton, whereas at the LHC antiquarks only appear as sea quarks, which carry smaller fractions of the proton momentum.




Figure 3.6: Cross sections for several important processes at the LHC compared to the Tevatron, the center-of-mass energies of which are indicated by the dashed vertical lines. (For the LHC the design center-of-mass energy of 14 TeV is assumed.) The number of events per second on the right vertical axis is computed for the design luminosity of the LHC.

Plans for LHC Upgrades

As the LHC is not yet running at its design parameters, in particular only at half the center-of-mass energy \sqrt{s} it is designed for, the next step naturally is to ramp up \sqrt{s} to 14 TeV, which is planned for a two-year shutdown in 2013 and 2014 [192]. It may be that only $\sqrt{s} = 13$ TeV will be reached. After three years of running at design luminosity, the peak instantaneous luminosity shall be increased to 2 or $3 \cdot 10^{34}$ Hz/cm², to collect about 300 fb⁻¹ of integrated luminosity until 2021.

Afterwards, the LHC may be replaced by the Super Large Hadron Collider (SLHC) or High Luminosity LHC, raising the peak instantaneous luminosity to 10^{35} Hz/cm². Among the possibilities to achieve this (cf. Section 3.2.4), current plans foresee to increase the beam intensities by a factor 5 and by a stronger squeezing of the beams at the interaction points 1 and 5 to decrease β by a factor 2 [16]. With the increased luminosity and higher event rates, also the radiation doses will be much higher, which will be a challenge especially for the detector systems close to the beam pipe. The SLHC shall deliver 200 to 300 fb⁻¹ per year to collect a total of 3000 fb⁻¹.

After the LHC Era

Considering the long time-scales involved in the research and development for a machine as large and complex as the LHC — from the first discussions in 1984 [53] to the official start in 2008 more than 20 years passed — it is already time to think about what will come afterwards. Some of the projects being discussed in the scientific community are outlined here.

There are two competing projects for building an electron-positron collider, the ILC and the CLIC project. As the energy loss per revolution due to synchrotron radiation for cyclic collider rings goes as $E^4 r^{-1} m_0^{-4}$ with the energy E , orbital radius r and rest mass m_0 of the accelerated particle and is thus much more severe for electron colliders than for proton colliders, both projects are based on linear machine layouts. The International Linear Collider (ILC) is a proposed 31 km long linear electron-positron collider [186], which is supposed to have a total energy of 500 GeV, with an option to upgrade to 1 TeV. It would use 16,000 superconducting RF cavities with an accelerating gradient of 31.5 MeV/m, allowing to scan a center-of-mass energy range of 200 up to 500 GeV. The Compact Linear Collider (CLIC) is a proposed electron-positron collider with a design center-of-mass energy of 3 TeV. To reach this energy with a machine of reasonable physical dimensions, CLIC is supposed to use very high accelerating gradients of 150 MeV/m at room-temperature using a two-beam technology with a high-current, low-energy drive beam parallel to the main beam.

Although the above electron-positron colliders enter the TeV regime, the highest center-of-mass energies are still reserved to hadron colliders. The Very Large Hadron Collider (VLHC) is a (very) theoretical proton-proton collider, with a foreseen center-of-mass energy of 100 TeV [16]. Another LHC-related project is lepton-nucleon scattering in the LHeC, colliding the existing LHC proton or heavy ion beam with an electron beam, possibly from an additional beamline in the LHC tunnel and synchronous to the proton-proton collisions, or from a linear accelerator which could give a lower instantaneous luminosity but a higher center-of-mass energy [193].

3.3.3 The 7 Detectors at the LHC

In addition to ATLAS, which will be the subject of a section of its own, there are three other large-scale detectors hosted at the LHC: ALICE, CMS and LHCb; and three smaller ones, LHCf, MOEDAL and TOTEM.

ALICE (A Large Ion Collider Experiment) is specifically designed to probe lead-ion collisions, in which a quark-gluon plasma is expected to be created, like it is supposed to have existed a very short time after the Big Bang. In this plasma, the energy density is high enough so that quarks and gluons are deconfined. The aim of the ALICE detector is to study the properties and cooling of the quark-gluon plasma and the formation of particles which constitute the matter content of the universe today. ALICE is 26 m long, 16 m high and wide and weighs 10,000 tonnes. It has an asymmetric geometry with a central barrel and a dimuon spectrometer in the forward direction. ALICE will record data at a much lower rate than ATLAS, but the size of the recorded data per collision is much larger, giving a challenging data flux of 1.25 GB/s.

The LHCb (Large Hadron Collider beauty) detector is another specialized detector. Its main purpose is to answer the question, why there is such an asymmetry between the matter and antimatter content in the universe. To do so, it studies the properties of b quarks in B mesons and their decay products. LHCb is 21 m long, 10 m high and 13 m wide and weighs 5600 tonnes. Its design is very different from general-purpose detectors in that it does not enclose the interaction region, but is a forward spectrometer with a number of planar detectors stacked one after the other on one side of the interaction point.

Like ATLAS, the Compact Muon Solenoid (CMS) detector is a general-purpose detector with the typical onion-like layout of its subdetectors. CMS shares many of the scientific goals of the ATLAS detector, and having these two detectors will allow to do cross-checks of the results of the respective other detector. The structure of CMS is quite different from the ATLAS detector though. A huge solenoid magnet with a cylindrical coil of superconducting cable generates a magnetic field of 4 T and surrounds both the silicon tracker as well as the electromagnetic and hadronic calorimeters. The electromagnetic calorimeter uses crystals of lead tungstate, PbWO_4 , which are very dense but optically clear, whereas the hadronic calorimeter is a sampling calorimeter made of brass and steel interleaved with plastic scintillators. Outside the solenoid, a huge iron return yoke interspersed with muon chambers guides and confines the magnetic field. CMS is smaller, but much heavier than ATLAS due to the weight of the iron yoke. At a size of 21 m in length and 15 m diameter, it weighs 12,500 tonnes. The trigger system of CMS is basically a two-level system with respect to its hardware, consisting of a Level 1 and a High Level Trigger, which roughly correspond to the hardware-based Level 1 and software implementation of the Event Filter of the ATLAS detector. Like in ATLAS, only muon and calorimeter information are available at Level 1. The second level at CMS can take a much higher rate than the Event Filter in ATLAS. However, the processing in the High Level Trigger of CMS is also subdivided into three steps, of which the first one uses only calorimeter and muon information, and the second only partial tracker information [38].

LHCf (Large Hadron Collider forward) are two very small detectors ($0.3 \times 0.8 \times 0.1 \text{ m}^3$), sharing the interaction region at Point 1 with the ATLAS detector, but at a distance of 140 m to either side of the interaction point exactly at an angle of zero degrees in a Y-shaped transition of the beam tube [73]. They measure the number and energy of neutral pions produced in the forward directions in the proton-proton collisions, with the aim of understand-

ing ultrahigh-energy cosmic ray events which happen in the upper atmosphere. TOTEM (TOTal Elastic and diffractive cross section Measurement) also studies forward particles. It consists of three sub-detectors extending over 440 m in length at Point 5 and is 5 m high and wide and weighs 20 tonnes. Using Roman pots (specially designed detectors housed in cylindrical vessels), gas electron multipliers and cathode strip chambers, its physics goals are a precise measurement of the proton-proton interaction cross section, as well as the study of the proton structure. Also, all of the LHC detectors will use TOTEM's measurement to calibrate their luminosity monitors. MOEDAL (Monopole and Exotics Detector At the LHC) sits in the cavern of LHCb and is comprised of an array of plastic nuclear track detectors. It searches for highly ionizing stable massive particles such as magnetic monopoles or the lightest stable states in R -parity conserving Supersymmetry or models of universal extra dimensions with Kaluza-Klein parity.

3.4 The ATLAS Detector

The ATLAS detector^d is a general-purpose detector for particle physics, with a forward-backward symmetric cylindrical geometry and near 4π coverage in solid angle, built for probing proton-proton and heavy-ion collisions. It is a unique machine of 25 m in height and 44 m in length, weighing approximately 7000 tonnes in total. The ATLAS detector is housed at Point 1, one of the interaction points of the LHC ring adjacent to the CERN main entrance, in a cavern about 100 m below the surface.

In order to be able to describe the location of the different subsystems in the following, first the detector coordinate system will be introduced, which will be used throughout this thesis. The right-handed ATLAS coordinate system has its origin at the nominal interaction point at the center of the detector. The direction of the beam pipe defines the z -axis and the x - y plane is transverse to the beam direction. The positive x -axis points from the interaction point to the center of the LHC ring, and the positive y -axis points upwards. The side of the ATLAS detector with positive z is side "A", the other side is side "C". Polar coordinates with a radial distance r and an azimuthal angle ϕ are used in the transverse plane. In spherical coordinates, the third coordinate can either be expressed through the polar angle θ with respect to the z -axis or the pseudorapidity

$$\eta := -\ln \tan(\theta/2). \quad (3.13)$$

The x - y plane at $z = 0$ is defined by $\eta = 0$. In the directions along the beam pipe the pseudorapidity diverges. For particles at very high energies so that their mass is negligible, the pseudorapidity is equal to the rapidity y , which is defined as

$$y := \frac{1}{2} \ln \left(\frac{E + p_z}{E - p_z} \right). \quad (3.14)$$

Geometrical distances of reconstructed objects are often expressed in terms of

$$\Delta R := \sqrt{\Delta\eta^2 + \Delta\phi^2}. \quad (3.15)$$

The following description of the design of the ATLAS detector is based on [21, 43, 129].

^d ATLAS is an acronym standing for "A Toroidal LHC Apparatus".




Figure 3.7: The ATLAS detector (computer-generated cut-away view). The different sub-systems are explained in the text.

3.4.1 Detector Components

ATLAS has a cylindrical geometry with the beam pipe as axis of its rotational symmetry and the nominal interaction point at its very center, where the proton beams delivered by the LHC are brought to collision. The beam pipe has a diameter of 58 mm and is constructed from 0.8 mm thick beryllium. The detector components are described as being part of the barrel if they are located in the central pseudorapidity region or as part of one of the two end-caps if they are in the forward or backward region. In the following, the detector components will be described, going from the center to the outermost components. Three main components can be distinguished, which give complementary information and consist of several subsystems each. Closest to the interaction point at the center of the detector is the Inner Detector, which is a tracking device used to measure precisely the tracks of charged particles in the pseudorapidity region up to $|\eta| < 2.5$. It is immersed in a 2 T solenoidal magnetic field and surrounded by the calorimeter system, which measures the energy of easily stopped particles. The overall dimensions of the detector are defined by the muon spectrometer, which consists of a toroidal magnet system and several different types of muon chambers. They measure the tracks of particles such as the highly penetrating muons, which cannot be stopped in the calorimeter. A computer-generated view of the detector and its components is shown in Figure 3.7.

3.4.2 Inner Detector and Solenoid Magnet

The Inner Detector (ID) consists of three independent and complementary tracking systems. The innermost is the silicon pixel detector (Pixel), surrounded by another precision silicon tracker, the Silicon Microstrip Tracker (SCT), and the Transition Radiation Tracker (TRT).




Figure 3.8: The ATLAS Inner Detector (computer-generated cut-away view).

The Pixel and SCT cover the region $|\eta| < 2.5$, while the TRT, being only the outer part of the cylinder, covers $|\eta| < 2$, as can be seen in Figure 3.8, which shows a detailed overview of the Inner Detector only. In the barrel region, the components are arranged as concentric cylinders, in the endcap region as disks perpendicular to the beam axis. The design goals of the Inner Detector are to provide a high momentum resolution, precisely determine the primary and secondary vertex position and detect charged tracks with p_T above 0.5 GeV.

The silicon pixel detector needed for the vertex detection has the highest granularity. It consists of three layers in the barrel and three disks in each of the endcaps, with a minimum pixel size of $50 \mu\text{m} \times 400 \mu\text{m}$. The intrinsic accuracies of the layers in the barrel are $10 \mu\text{m}$ in $R-\phi$ and $115 \mu\text{m}$ in the z -direction and vice versa in the disks. The pixel detector is operated at a temperature of -5 to -10°C to reduce electronic noise and has 80.4 million readout channels in total. The silicon microstrip tracker has four coaxial cylindrical layers in the barrel and nine disks in each endcap, which are arranged such that a straight track with $|\eta| < 2.5$ crosses at least four modules, giving four space-point measurements with two hits per module. The two times two sensors on each module are rotated against each other and have a mean pitch of the strips of $80 \mu\text{m}$. The rotation by a small angle allows to do stereo measurements of both coordinates in the plane. The intrinsic accuracies per module in the barrel are $17 \mu\text{m}$ in $R-\phi$ and $580 \mu\text{m}$ in the z -direction and vice versa in the disks. In total the SCT has 6.3 million readout channels. The Transition Radiation Tracker (TRT) is a combination of a straw tracker and a transition radiation detector. It is made of gaseous straw tube elements with 4 mm diameter, which are thin-walled proportional drift tubes, interleaved with foils and fibres of transition radiation material. The straw tubes are filled with a $\text{Xe}/\text{CO}_2/\text{O}_2$ gas mixture. In contrast to the semiconductor trackers, the TRT is designed to operate at room temperature. Its intrinsic accuracy is $130 \mu\text{m}$ per straw. In the barrel, the straws are mounted parallel to the beamline with a length of 144 cm in up to 73 layers. In

the end-cap, the straws run radially and are only 37 cm long, organized in wheels of 160 planes. The TRT has 351,000 readout channels in total. Straw hits in outer layers contribute significantly to the momentum measurement, in spite of their lower resolution due to the long lever arm of the measurement. In addition to the tracking abilities, the TRT provides for electron identification by the detection of transition radiation photons: If a charged particle travels through a medium at a velocity higher than the local speed of light c/n , where n is the index of refraction of the medium, the medium radiates photons at a characteristic angle. This Čerenkov transition radiation is used to identify electrons with energies between 0.5 and 150 GeV. The distinction between transition radiation and tracking signals is made using separate low and high thresholds in the front-end electronics.

The Inner Detector is fully contained in the solenoid magnet consisting of one thin superconducting coil with a length of 5.8 m and a diameter of 2.56 m. It weighs 5.7 tonnes and provides a 2 T axial field, in which the Inner Detector components are immersed, to bend the tracks of charged particles and thereby make possible the momentum measurements. At nominal current, the energy stored in the magnet field is 40 MJ. The flux is returned by the steel of the ATLAS hadronic calorimeter (see below) and its girder structure. It can be cooled down to the operating temperature of 4.5 K in one day. In order not to deteriorate the calorimeter performance, the construction of the solenoid magnet had to be as light-weight and thin as possible. It contributes approximately 0.66 radiation lengths.

A major problem for the Inner Detector components is that, being so close to the beam-pipe, they are subject to a high radiation dose which is absorbed by the material over the running time of the experiment. It therefore needs to have a good radiation hardness. The pixel inner vertexing layer must therefore be replaced after approximately three years of running at the design luminosity. Also, the bias voltage will need to be increased from initially 150 V to up to 600 V to achieve a good charge collection efficiency after several years of running.

3.4.3 Calorimeters

The ATLAS calorimeter system covers the region up to $|\eta| < 4.9$ and consists of electromagnetic calorimeters and hadronic calorimeters which are located directly behind (“down-stream”) the electromagnetic calorimeters. The calorimeter system is based on different materials, but all components are sampling calorimeters. In general, the inner cylinder is made up of calorimeters which use liquid argon (LAr) as active medium, surrounded by an outer concentric cylinder of tile calorimeters, as shown in Figure 3.9. The calorimeter system has a full ϕ -symmetry and coverage around the beam axis. The thickness of the electromagnetic calorimeters is > 22 radiation lengths in the barrel and > 24 radiation lengths in the end-cap, and for the hadronic calorimeters the thickness amounts to 9.7 interaction lengths in the barrel and about 10 in the end-caps. Over the η region that is covered by the Inner Detector, the electromagnetic calorimeter has the highest granularity to provide precision measurements of the energy deposited by electrons and photons.

The LAr electromagnetic calorimeter is a sampling calorimeter with accordion-shaped kapton electrodes and lead absorber plates and LAr as the active medium. It consists of a barrel part which extends up to $|\eta| < 1.475$, with a small gap of 4 mm at $z = 0$ between two symmetric halves (EMA and EMC), and two end-cap components in the region $1.375 < |\eta| < 3.2$. The end-cap components (EMEC) are again mechanically divided into two wheels with $1.375 < |\eta| < 2.5$ and $2.5 < |\eta| < 3.2$. The barrel and the inner wheels have




Figure 3.9: The ATLAS Calorimeters (computer-generated cut-away view).

together at least three segmentations in depth (sampling layers), while the outer wheels have only two. The granularity of the electromagnetic calorimeter in η is very fine, down to $\Delta\eta = 0.025/8$ in the first layer in the range up to $|\eta| < 1.8$. The granularity in ϕ is mostly 0.1 in the first and 0.025 in the second and third layers, in which most of the energy is collected for high energy electrons (see [21] for details). A presampler, consisting of an instrumented active LAr layer, is installed in front of the electromagnetic calorimeter and covers the region up to $|\eta| < 1.8$. It corrects for the energy lost by electrons and photons before reaching the calorimeter.

The hadronic calorimeter uses liquid argon as well as scintillating tiles as active components. It is comprised of three parts: the tile calorimeter, the LAr hadronic calorimeters in the end-caps and the hadronic part of the LAr forward calorimeter. (Note that there is also an electromagnetic calorimeter in the forward region, see below.) The tile calorimeter (Tile) consists of three parts, one barrel and two extended barrel parts, which cover the region up to $|\eta| < 1.7$. It uses steel as absorber and scintillating tiles as active component, and has three segments in depth. The LAr hadronic calorimeters in the end-caps (HEC) cover $1.5 < |\eta| < 3.2$ to have a slight overlap with the tile and forward calorimeters. They use copper plates interleaved with LAr as active medium and consist of two independent wheels per end-cap, each divided into two segments in depth.

The LAr forward calorimeter (FCal) covers the forward and backward regions closest to the beam with $3.1 < |\eta| < 4.9$. It consists of three modules in each end-cap, of which the first uses copper and is optimized for electromagnetic measurements, whereas the second and third use tungsten and measure predominantly energy from hadrons. Again, liquid argon is the active medium flowing within the metal matrix. The longitudinal segmentation helps to do electromagnetic shower identification.

All calorimeter systems together have roughly $1.9 \cdot 10^5$ readout channels. There are two




Figure 3.10: The muon system of the ATLAS detector (computer-generated cut-away view).

separate readout paths: One with coarse granularity, for which calorimeter cells occupying the same area in η and ϕ are combined to trigger towers. This is used by the Level 1 trigger (cf. Section 4). The other has a fine granularity and is used for the final readout by the High-Level Trigger and in the offline reconstruction. The performance in the transition region between the barrel and end-cap electromagnetic calorimeters, $1.37 < |\eta| < 1.52$, is expected to suffer from the larger amount of material in front of the first active calorimeter layers. The shapes of the calorimeter signals are compared to nominal shapes, allowing to determine a χ^2 -like quality factor. Also the timing of signal is measured and can be used in analyses.

LAr Hole

The “LAr hole” is the commonly used term within the ATLAS collaboration for a detector problem which has occurred in the first half of 2011 data taking^e and which affects part of the electromagnetic liquid-argon calorimeter in the barrel. Due to the failure of a crate controller, the information from six front-end boards is lost, four of which could be recovered starting with period 2011 I. The region affected by the LAr hole is $[0.0, 1.45] \times [-0.78847, -0.59213]$ in $\eta \times \phi^f$. For physics analyses, a way to deal with this problem is to reject events which include jets that fall into the LAr hole region. Due to the high number of jets in QCD events, this may make up a significant fraction of events and therefore refined veto strategies are recommended. The analysis presented in Section 7 uses only 2010 data and is thus not affected by the problem. The impact on the efficiencies of jet + \cancel{E}_T triggers in 2011 is studied in Section 5.6.3.




Figure 3.11: Cross-sectional view in the y - z plane showing the positions of the four muon systems and the coils of the toroid magnet in one quadrant of the ATLAS detector.

3.4.4 Muon Spectrometers and Toroidal Magnet System

The muon spectrometer of the ATLAS detector is a large tracking system and measures the charge and momentum of charged particles leaving the calorimeters. It also serves for the identification of particles as muons. All electromagnetically or strongly interacting Standard Model particles except for the muon are expected to be stopped within the calorimeters, unless they have exceedingly high energies such that they can punch through the calorimeters.

The four subsystems of the muon spectrometer are mounted in and around air-core toroids, which generate the necessary magnetic bending field as can be seen in Figure 3.10. They use four different technologies and can be classified into two subsystems which do precision tracking and another two which are specialized to provide information to the Level 1 trigger at a high time-resolution. The systems are organized in three concentric cylindrical layers at approximately 5, 7.5 and 10 m in the barrel. In the transition and end-cap region, they form large wheels installed orthogonal to the beam direction, as can be seen in Figure 3.11.

The high-precision tracking chambers measure the track coordinates in the bending plane. The Monitored Drift Tubes (MDTs), covering $|\eta| < 2.7$ (2.0 for the innermost of the three concentric layers), consist of up to eight layers of pressurized drift tubes per chamber, organized in two multilayers, and achieve an average resolution of $80 \mu\text{m}$ per tube and about $35 \mu\text{m}$ per chamber. They are called monitored, because their deformations and relative positions are monitored by 12,000 alignment sensors based on optical monitoring. The MDTs are complemented by Cathode Strip Chambers (CSCs), which cover $2.0 < |\eta| < 2.7$ and have a higher rate capability up to counting rates of 1000 Hz/cm^2 , compared to 150 Hz/cm^2 for the MDTs. The CSCs are multiwire proportional chambers with cathodes segmented into strips in orthogonal directions, thus allowing to measure both coordinates. Their resolution is $40 \mu\text{m}$ in the bending plane and about 5 mm in the transverse plane.

^e 30.04.2011, starting with the beginning of data-taking period 2011 E.

^f No official documentation on this issue appears to be publicly available at the time of writing.

The muon trigger chambers have a time resolution of 1.5 to 4 ns. Two different technologies are used as trigger chambers, Resistive Plate Chambers (RPCs) for $|\eta| < 1.05$ and Thin Gap Chambers (TGCs) for $1.05 < |\eta| < 2.7$ (2.4 for triggering). They are also used for measuring the muon coordinate in the direction orthogonal to that determined by the precision chambers with a resolution of several millimeters. The TGCs are also multiwire proportional chambers, with the wire-to-wire distance being smaller than the wire-to-wall distance. This allows for running in a saturated mode with a very high gas gain. The ϕ -symmetry of the muon chamber systems follows the structure of the toroids, consisting of eight octants. There are gaps in the coverage of the muon system to allow for services to the solenoid magnet, the calorimeters and the Inner Detector and the detector support structure. The total number of readout channels of the muon systems is $1.1 \cdot 10^6$. For details see [21].

Toroid Magnets

The muon spectrometer relies on a set of large superconducting air-core toroidal magnets. The light and open structure of the magnets is needed to minimize multiple scattering effects, which would degrade the muon identification performance and momentum resolution. The toroidal magnet system consists of a long barrel toroid in the region up to $|\eta| < 1.4$ and two smaller end-cap magnets inserted into both ends of the barrel toroid which cover $1.6 < |\eta| < 2.7$. The strength of the magnetic field is approximately 0.5 and 1 T in the central and end-cap regions, respectively. Each of the three toroids consists of eight coils, the end-cap toroid coil system being rotated by $\pi/8$ rad with respect to the barrel toroid coil system to improve the overlap of the magnetic fields. The toroid magnets have an axial length of 25.3 and 5.0 m and an outer diameter of 20.1 and 10.7 m for the barrel and end-cap parts, respectively. The total weight is 830 plus 2×239 tonnes, and the total energy stored in the magnetic field is 1.58 GJ, the region where the magnetic field exceeds 50 mT being 12,000 m³. The strength of the magnet field is monitored by 1800 Hall sensors, which are distributed throughout the spectrometer volume. Note that the direction of deflection for charged particles is different from the one effected by the field of the solenoidal magnet, in which the Inner Detector is immersed.

3.4.5 Performance Goals

In Table 3.2, an overview is given of the performance goals with respect to the design of the ATLAS detector components that have been described above. The numbers are taken from [21] and mainly for illustrational purposes. The resolutions actually achieved may differ, cf. Section 4.4.

3.4.6 Minimum-Bias Trigger Scintillators

For the initial phase of data taking at low instantaneous luminosities, the Minimum-Bias Trigger Scintillators (MBTS) can be used to trigger on a minimum collision activity during proton-proton collisions [161]. The MBTS detector consists of 32 scintillator paddles, which are organized in two disks perpendicular to the beam pipe and installed at either side of the detector at a distance of $z = \pm 356$ cm from the interaction point. Each disk covers an η range of $[2.09, 3.84]$ and is split into two rings with eight independent sectors each. The light from the scintillators is collected and converted into an electrical signal in photomultiplier tubes,

Detector component	Required resolution
Tracking	$\sigma(p_T)/p_T = 0.05\% p_T \oplus 1\%$
Muon spectrometer	$\sigma(p_T)/p_T = 10\%$ at $p_T = 1$ TeV
EM calorimetry	$\sigma(E)/E = 10\%/\sqrt{E} \oplus 0.7\%$
Hadronic calorimetry (jets)	
barrel and end-cap	$\sigma(E)/E = 50\%/\sqrt{E} \oplus 3\%$
forward	$\sigma(E)/E = 100\%/\sqrt{E} \oplus 10\%$

Table 3.2: Design goals of the ATLAS detector with respect to energy and momentum resolution of the different subdetectors. E and p_T are given in GeV. The numbers are taken from [21] and may differ from the resolutions actually achieved.

which is passed on to the Central Trigger Processor (CTP) (see Section 4), where the number of hits per scintillator is evaluated by comparing the signal to a given discriminator threshold. Triggering on minimum-bias events can be done by requiring either single, multiple or coincident hits in both disks.

3.4.7 Forward Detectors

The forward detectors are a group of special detectors [21] located relatively far away from the interaction point at very high pseudorapidities. The LUCID detector, where LUCID stands for luminosity measurement using a Čerenkov integrating detector, is a detector which is primarily dedicated to relative luminosity measurements. It is located at ± 17 m from the interaction point and can record the luminosity for each bunch crossing separately. The Zero-Degree Calorimeter (ZDC) is ± 140 m away from the interaction point and detects forward neutrons in heavy-ion collisions. The ALFA (Absolute Luminosity For ATLAS) detector at a distance of ± 240 m consists of scintillating-fibre trackers which are located inside Roman pots.

4 Data Acquisition and Data Processing

Following the description of the ATLAS detector hardware, this chapter describes the acquisition and the processing of data in ATLAS, and is mainly concerned with software issues. The central part of the data-acquisition system is the trigger. Its purpose in general and the necessary terminology are discussed first, before the ATLAS trigger system itself is described. When the data has been recorded, it needs to be suitably processed. This is called event reconstruction and is explained in the following part. In the last part of this chapter, a brief introduction to Monte Carlo event generation is given and important related notions are introduced. Additional technical background can be found in Section A.10.

4.1 Purpose of a Trigger

The trigger system of a detector can be thought of as a filter that, out of the large number of collision events taking place in a collider experiment, selects those that are stored for later processing and analysis. The need for such a filter becomes evident when looking exemplarily at the following numbers: The LHC is designed to deliver proton-proton collisions at a rate of up to 40 MHz, so that every 25 ns a collision of proton bunches takes place at all interaction points. The output rate, at which the measurements of the collision events can be written out, is limited mainly by the allocated storage capacity in long terms and the output bandwidth of data-taking system in short terms. For ATLAS, the latter amounts to a few hundred megabytes per second. Assuming an average event size of 1.5 MB, this yields an allowed output rate of 200 to 300 Hz, which may be raised up to 600 Hz for limited periods of time [111]. (Of course, there are also limitations from the speed at which the detector systems can be read out at full granularity, but those are less severe.) This means that a rate reduction of the order of 10^5 is needed, implying that most of the collision data will be discarded. In fact, it is not a problem to achieve this reduction from the point of view of analysis, as most of the collisions will only procure events which contain processes that are well known and studied. Of interest are those events which contain rare processes, like the production and decay of some so far unknown particle. The task is therefore to identify those events among the huge number of collisions taking place every second, which are interesting in some sense. This is done by the trigger system: It uses a suitably reduced set of information to identify, in real time during data taking (“online”), the most interesting events to retain for detailed analysis (“offline”), ideally by rejecting all uninteresting events while keeping the interesting ones at a high efficiency. The trigger system will usually comprise several more or less independent triggers running in parallel. If any of these triggers fires, i. e. decides to accept the collision event, the event is stored.

A well-designed trigger system ought to meet a number of requirements:

- **Robustness:** The operation of the trigger should be reliable and stable. Technical problems leading to missing, incomplete or invalid data should not lead to a crash of the whole trigger system or parts of it.

- **Redundancy:** Redundancy is needed both with respect to online data taking and offline aspects. During data-taking, it may become necessary to disable a particular trigger because its output rate is too high, or because it is misconfigured and disturbs the rest of the data-taking process. With respect to offline, it may be found afterwards that a particular trigger has selected other types of events than it was intended to. In both cases, loss of data can only be prevented if the task of the disabled or malfunctioning triggers can be taken over by another. This redundancy normally should not increase the output rate a lot.
- **Inclusiveness:** Keeping the selection criteria a trigger applies to the data rather general allows to use the data taken with this trigger for a lot of different analyses. It also reduces the risk of retrieving only a biased subsample of all events or to miss unexpected but interesting collision events.
- **Simplicity:** Defining the triggers in such a way that their decisions are easy to comprehend simplifies the commissioning and the maintenance of the trigger system. It also helps understanding the trigger (in-)efficiencies, an aspect that is discussed in detail in Section 5.2.

In order to be able to run within the constraints of the trigger system, in addition to these requirements, the triggers need to have a precision that is high enough to make an efficient decision and at the same time be fast enough to process all events within the given time. The performance of a trigger can be quantified in terms of its rate and efficiency, i. e. the quality of its selection. The optimization of the trigger is then done by finding the best working point as a compromise between smaller rates and higher efficiency. This is discussed in Section 5.1.2 using the example of jet + \cancel{E}_T triggers. The main challenge is that all events which are rejected by the trigger are lost forever. Thus, it is possible to confirm that the triggered events are interesting, but not that there are no interesting events among those which have been discarded. The solution to this problem is to keep a few additional events, regardless of the trigger decision (pass-through events, see below), which allows to search a small sample of discarded events for anything interesting that may have been missed.

4.2 Terminology of Triggers

Before the ATLAS trigger system will be described in the next section, some basic notions and concepts will be introduced, which are needed in the discussion of triggers in general.

The rate of a particular trigger can be adjusted, for example to follow the constant decay of the instantaneous luminosity within a fill of the collider, by applying different *prescales*. Applying a prescale factor N to a trigger means that out of N events in which this trigger has given a positive decision only one randomly selected event is actually stored. The rate at which the trigger fires is therefore directly proportional to the inverse of the prescale. A trigger is called *prescaled* if it has a prescale larger than one. In some sense the opposite of the prescale factor is the *pass-through* factor. A pass-through factor of N means that one in N events is treated as if it had been accepted by the trigger regardless of the actual trigger decision. A use case for pass-through factors is to obtain a control sample of discarded events as mentioned above. Triggers can also be run in *transparent mode*, which means that the event is accepted without executing the trigger. Finally, an error within the trigger system

may lead to a *forced accept* of the event without further execution of triggers. The accept signal of a trigger may not only be inhibited by prescales, but also due to a dead time veto or busy veto of the trigger system. *Dead time* occurs always after an event has been accepted (see below). A busy signal can be issued when some part of the trigger or data-acquisition system cannot cope with the current trigger rate.

Triggers can be categorized with respect to their type and their function. With respect to their type, a distinction is made between single object triggers, multiple object triggers, combined triggers and topological triggers. *Single* and *multiple object triggers* are used for final states with at least one or two or more characteristic objects of the same type. *Combined triggers* select events with two or more objects of different types, whereas the decision of *topological triggers* is based on the combined properties of objects of different types. With respect to their purpose, there are primary, supporting and monitoring or calibration triggers. *Primary triggers* are used to select events for physics analysis, which potentially contain the searched for signal. *Supporting triggers* are also used in physics analysis, but have a supplementary function. They select events, for example, for studies of the background. The data collected with *monitoring* and *calibration triggers* is used to ensure the correct operation of the trigger and detector.

The configuration of the trigger system is stored in the *trigger menu*, which includes information about which triggers are executed and their parameters. Different triggers are rarely exclusive, so that some events will be triggered, i. e. accepted, by more than one trigger. The consequence of this overlap between triggers is that deactivating a particular trigger will not reduce the total output rate of the trigger system by the full rate of this trigger, but only by the *unique rate* of this trigger. This is the rate of events in which exclusively this trigger fires. The overlap of triggers makes it difficult to predict the impact of changes in the trigger menu on the total rate of the system.

4.3 The ATLAS Trigger System

4.3.1 General Architecture

The ATLAS detector is equipped with a three-level trigger system [43], filtering events in three subsequent steps, where each level confirms and refines the decision of the previous level. The three trigger levels are called Level 1 (L1), Level 2 (L2) and Event Filter (EF). A schematic diagram of the ATLAS trigger and data-acquisition system is shown in Figure 4.1. The dataflow between the different components and the propagation of the trigger decisions are indicated by arrows. Details can be found in [140].

Level 1 is implemented as a hardware-based system [130], using fast custom electronics and algorithms which are programmed into the firmware. This part of the trigger is located underground close to the detector and only uses coarse information from the calorimeters and the dedicated muon trigger chambers, to be able to make a decision within an average time of 25 ns. The main element of Level 1 is the Central Trigger Processor (CTP), which computes the trigger decision, using as input up to 256 different trigger signal bits from the calorimeter and muon trigger processors. In addition, bunch group requirements are applied in order to restrict some of the triggers at Level 1 to only filled or empty bunch crossings. The maximum latency of Level 1 is 2.5 μ s. As this time spans up to 100 bunch crossings, the detector data needs to be kept in pipeline memories until the decision for the respective

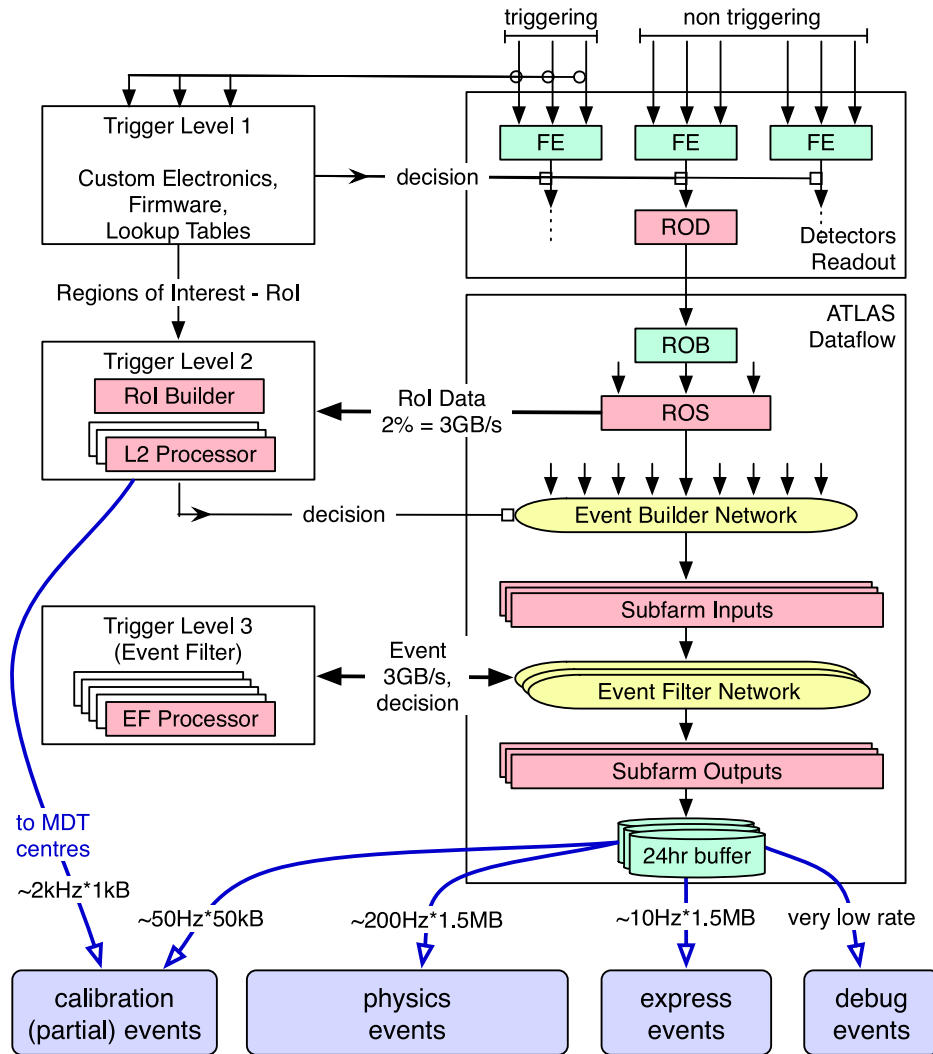


Figure 4.1: Schematic diagram of the ATLAS trigger and data-acquisition system. On the left, the data-processing components are shown, on the right, the data-acquisition components.

collision event has been computed. Upon a Level 1 Accept signal, the data is transferred to the detector-specific Read-Out Buffers (ROBs), where it is stored until requested by Level 2 algorithms or by the Event Builder upon a positive Level 2 decision. The maximum designed output rate of Level 1 is 75 kHz, depending on the trigger menu. In 2010, the maximum rate was 30 kHz. In addition to the rate reduction, the main purpose of Level 1 is to define Regions of Interest (RoIs), which describe geometrical regions of the detector in η - ϕ space and are used in the next step by Level 2.

Level 2 and the Event Filter constitute the High-Level Trigger (HLT) which is implemented in software. It runs on a cluster of commercially available computers and networking hardware near the detector, but at the surface. In 2010, the computer farm comprises 1100 nodes in total, each node having 8 Central Processing Units (CPUs), most of which are running at 2.4 GHz [111]. For the design luminosity, the number of nodes will be increased to 500 for L2 and 1800 dedicated to the EF.

Level 2 is seeded by the RoIs defined at Level 1. It can read out the detector data with full granularity, but only in a limited region around the defined RoIs. Level 2 can also access the Inner Detector data in addition to the calorimeter and muon spectrometer information, which is not used before, because no track finding can be done within the time available at Level 1. In 2010, the output rate of Level 2 was 3 kHz and the average processing time 50 ms per event. If an event is accepted by Level 2, the full detector data is collected into an event record by the Event Builder and passed to the Event Filter. At the Event Filter level, it is possible to evaluate the full data in the events with more sophisticated software algorithms. These algorithms basically are the same algorithms which are used in the offline processing of data. The maximum design output rate of the Event Filter is 200 Hz. In 2010, the average processing time was about 0.4 s per event. The maximum time available to the Event Filter to process an event depends on the number of nodes in the computer farm and the Level 2 output rate.

The trigger configuration of the ATLAS trigger system is organized in terms of chains. Each chain spans all three trigger levels and specifies the sequence of reconstruction and selection steps at each level. (Often the terms trigger and trigger chain are used interchangeably.) The process that controls the execution of the trigger chains is referred to as the trigger steering [62]. To minimize the average time spent on the computation of the trigger decision for every event, the trigger steering is designed to reject events as early as possible (early-rejection principle). In order to achieve this, each chain is processed in a step-wise and seeded manner and all chains are processed in parallel. Higher levels use the output of lower levels and from previous steps within the chains at the different levels. To be accepted and stored, an event needs to pass at least one of the trigger chains.

The HLT algorithms which the trigger chains are comprised of can be separated into two classes with different tasks. The Feature Extraction (FEx) algorithms perform the time-consuming unpacking of data and reconstruction of trigger objects like calorimeter clusters. These objects are called features in the trigger, hence the name. The hypothesis algorithms perform selection cuts based on the reconstructed objects, which is comparably fast. Therefore, each chain at all three levels typically consists of one or more feature-extraction algorithms, followed by a hypothesis algorithm which computes the actual decision of the chain at that level. The caching mechanism of features allows objects which have been reconstructed by one chain to be reused by all other chains.

ATLAS uses an inclusive streaming model [51, 140], where every trigger chain can be associated to one or more data streams. Every event which has been accepted by the trigger is

written to all streams for which at least one of the trigger chains associated with the stream has fired. There are four primary streams for storing events for physics analyses, called `Egamma` (for photon and electron triggers), `Muons` (for muon triggers), `JetTauEtmiss` (for jet, tau-lepton and \cancel{E}_T triggers) and `MinBias` (for minimum-bias and random triggers), according to the type of triggers that feed events into the stream. Muon chains would trigger the event to be written to the `Muons` stream, an event which fires, for example, both a muon and an \cancel{E}_T trigger would be copied to both the `Muons` and `JetTauEtmiss` stream. In addition to writing complete events to a stream, there are also calibration streams, to which only partial information from one or more sub-detectors is written. The inclusive streaming model implies that there will be overlaps between sets of events from different streams, which needs to be taken into account when combining them in physics analyses, and also means that a fraction of the bandwidth is wasted on writing out the same data redundantly.

In ATLAS, prescales can also be used to deactivate a given trigger by assigning a negative prescale to it^a. In contrast to the composition of the trigger chains or their thresholds, the prescales can be changed at any time within a data-taking run at the beginning of a new luminosity block. Primary triggers are usually not prescaled, whereas the other triggers, in particular those for monitoring and calibration purposes, in general are. As said above, the prescales can be changed independently from the trigger menu. In total, three keys are needed to retrieve the configuration of the ATLAS trigger from the trigger database (cf. Section A.10.5), the `SuperMasterKey` (SMK) which selects the trigger menu and the Level 1 and HLT Prescale Keys (L1PSK and HLTPSK) which determine the set of prescales that is applied [110]. Besides, there is the `Bunch Groups` (BGS) Key which specifies the bunch configuration. The keys are unique integers by which the complete configuration for all recorded runs can be accessed.

Dead time of the trigger, during which the accept signal is vetoed, is introduced by the CTP after every Level 1 Accept to allow the read-out systems to finish processing and to return to normal operation, and to protect the front-end read-out buffers from overflowing due to short-scale rate fluctuations (preventive dead time) [111]. This type of dead time mainly affects the first trigger level. At the subsequent levels, the trigger rates can be assumed to be low enough so that dead time of detector components does not play a role under normal running conditions. There are other sources of trigger vetoes, like dead time of the data-acquisition system itself coming from backpressure, i. e. if the trigger rate is higher than the rate at which the events can be processed or written out, or when the data acquisition is paused by the operator.

To study and understand the trigger performance, it is important to know how the trigger algorithms work. Therefore, in the following a summary of the missing transverse energy and jet triggers will be given, complemented by a brief overview of other important triggers.

4.3.2 \cancel{E}_T Trigger

The algorithms which compute \cancel{E}_T in the trigger [139, 141] run unseeded because the \cancel{E}_T is a global quantity and therefore no Regions of Interest can be defined. In fact, the \cancel{E}_T in the trigger is computed for all events at all levels if any trigger has accepted the event at the respective lower level. This makes it possible to emulate the trigger decision for all recorded

^a The absolute value of a negative prescale does not matter, but different negative values may be found when a trigger is deactivated which is normally prescaled. The absolute value then usually corresponds to the prescale the trigger would have if it were active.

events, whether or not they have been accepted by an \cancel{E}_T trigger chain (cf. Section A.1) and to study combined triggers $\cancel{E}_T + X$. In order to compute \cancel{E}_T (and $\sum E_T$), the full detector needs to be read out, which currently is only possible at Level 1 and at the Event Filter, but not at Level 2. The \cancel{E}_T at Level 2 therefore is the same as \cancel{E}_T at Level 1, plus the muon correction (see below), which in turn is not possible to include at Level 1.

At Level 1, \cancel{E}_T and $\sum E_T$ are computed from trigger towers, which are comprised of cells occupying the same area in η - ϕ space and spanning all of the different calorimeter samplings (both from electromagnetic (EM) and hadronic calorimeters). The approximately 7200 trigger towers with a granularity of approximately 0.1×0.1 in η - ϕ space cover the region up to $|\eta| \leq 4.9$ in pseudorapidity. The analog signals from the trigger towers are transmitted to the Level 1 system, where they are digitised and associated with a bunch crossing and the pedestal subtraction is performed. The Level 1 system also applies a noise suppression threshold per trigger tower and extracts the transverse components. Actually, \cancel{E}_T is not computed at Level 1, but instead a look-up table mapped by the vector sums of the components \cancel{E}_x and \cancel{E}_y is used to check whether any of the defined thresholds have been exceeded. If the measurement saturates the look-up table, a flag is set to indicate an overflow so that the event is passed to Level 2 regardless of the configured \cancel{E}_T thresholds.

If any chain accepts the event at Level 1, the \cancel{E}_x and \cancel{E}_y values computed at Level 1 are passed to Level 2, where \cancel{E}_T is computed as the square root of the sum of the squares of \cancel{E}_x and \cancel{E}_y , rather than using a look-up table. The decision of Level 1 and Level 2 may therefore differ slightly if \cancel{E}_T lies close to a trigger threshold. Additionally, the p_T of muons reconstructed at Level 2 are included in the computation of \cancel{E}_T and $\sum E_T$ at this trigger level, but both computations, with and without the muon contribution, are stored separately. The decision of the \cancel{E}_T and $\sum E_T$ triggers at Level 2 and Event Filter, however, is based on the computation of \cancel{E}_T and $\sum E_T$ without the muon contributions. Note that muons in general do not deposit much of their energy in the calorimeter, so that the contributions from muons measured in the muon spectrometer to \cancel{E}_T are not negligible.

At Event Filter, a complete recomputation of \cancel{E}_T and $\sum E_T$ is done by looping over all calorimeter cells, which gives a better resolution than possible by using the trigger tower results. Only cells with energies above a given threshold which corresponds to three times the noise Root Mean Square (RMS) of the cell are included in the summation. This one-sided cut leads to a small bias in the $\sum E_T$ computation of about 10 to 20 GeV (cf. Figure 6.5). For the muon contribution, the result of the Level 2 computation is reused. All measurements of \cancel{E}_T in the trigger are done at the uncalibrated electromagnetic scale.

4.3.3 Jet Trigger

The jet triggers start at Level 1 by constructing jet elements which are comprised of 2×2 blocks of the trigger-towers mentioned above. These are then combined into proto-jets using a sliding-window algorithm, which selects high-energy depositions in a square of programmable size of 0.4×0.4 , 0.6×0.6 or 0.8×0.8 in $\eta \times \phi$. If the multiplicity and energy of the Level 1 trigger jets which have been constructed in this way is sufficient to pass the trigger thresholds, the geometrical positions of the Regions of Interest of the Level 1 jets are passed to Level 2. The Level 1 jet algorithm identifies jets within $|\eta| < 3.2$. Forward jet trigger chains, which are independent from the normal trigger chains, use the forward calorimeter outside this region with a coarser granularity, but they are not used in this thesis. The Level 2 algorithm is a cone clustering algorithm, which runs a limited number of

iterations on calorimeter clusters in a rectangular region centered around the RoIs identified at Level 1 with higher calorimeter granularity.

The Event Filter basically uses the same algorithms to reconstruct jets that are also used in the offline processing of data. Throughout 2010, all of the jet chains have been run with the Event Filter level in pass-through mode, until at the beginning of 2011, the new trigger algorithms based on anti- k_t were commissioned. An important point is that the jet reconstruction algorithms in the trigger at Event Filter level, starting from 2011, run in an unseeded full-scan mode, i. e. they are executed when Level 2 has accepted the event, but they do not rely on the Regions of Interest or jets identified by the lower chains. Moreover, another change in 2011 with respect to 2010 is that the jet trigger chains with the lowest thresholds at Event Filter are seeded by random triggers and not by jet triggers at Level 1 and Level 2, which can be seen in Table 5.7 and will be discussed later. Note that the description of the Level 1 and Level 2 jet algorithms is based on [43] and may differ from the current implementation.

4.3.4 Other Types of Triggers

Besides the jet and \cancel{E}_T triggers, there are many more types of triggers among the approximately 500 triggers which are defined in the ATLAS trigger menu. The ones which are relevant in the context of this thesis are described below. Other groups of triggers that are not used and therefore not discussed here are electron, photon, tau lepton and b -physics triggers.

Muon Triggers

Muons in ATLAS can be triggered on within the range $|\eta| < 2.4$. Two of the four muon subsystems have a timing resolution that is good enough to be used for triggering, at the price of lower precision. In the barrel region, $|\eta| < 1.05$, these are the Resistive Plate Chambers (RPC), and in the endcap region, $1.05 < |\eta| < 2.4$, the Thin Gap Chambers (TGC). While the trigger coverage in the endcap is close to 100%, in the barrel it is significantly lower because of the lower geometrical acceptance due to a crack around the mirror symmetry axis of the detector at $\eta = 0$, the magnet support structure, and two small elevators in the bottom part of the spectrometer [43]. The Level 1 trigger works by finding hit coincidences between the different muon detector layers. Six p_T thresholds can be configured using look-up tables. The Muon CTP Interface (MuCTPI) collects the information from the muon trigger chambers (RPC and TGC) and does overlap removal. Level 2 starts from the Regions of Interest defined at Level 1 and runs three algorithms successively. The first algorithm fits tracks using the MDT measurements, and for the determination of p_T uses look-up tables. The next algorithm combines the previously found tracks from the muon spectrometer with Inner Detector tracks. In the final step, a discrimination between isolated and non-isolated muon candidates from the previous steps is done, taking into account the energy depositions in the calorimeter using two concentric cones around the muon. The inner, smaller cone is used to exclude the energy depositions by the muon itself, and the energy content within the outer cone defines the isolation. The Event Filter can again basically use the same algorithms as offline, seeded by muons found at Level 2 and combining reconstructed tracks from the muon spectrometer and from the Inner Detector.

Minimum-Bias and Random Triggers

Minimum-bias and random triggers are triggers which are supposed to give an unbiased sample of all possible collision events taking place in the detector. In particular the random triggers, if they are correctly implemented, give a truly unbiased sample. They trigger, as can be easily guessed from their name, randomly on any bunch crossing without taking any event properties into account. There are different types of random triggers, which either only trigger on crossings of filled bunches, where proton-proton interactions are expected to take place (e.g. EF_rd0_filled_NoAlg), or only on empty bunches, where no collisions are expected (e.g. EF_rd0_empty_NoAlg). As the prescales implemented within the trigger steering are sufficient for randomly triggering events, no additional algorithm needs to be implemented at Level 2 and Event Filter. At Level 1, the random triggers require the RNDM0 item to accept an event, which in normal running it does at an unprescaled rate of $f_{\text{LHC}}/8$, $f_{\text{LHC}} \approx 40$ MHz being the beam crossing frequency at the interaction points. Furthermore, the items BGRP0 and BGRP1 or BGRP3 are required, which select non-vetoed filled or empty bunch crossings, respectively.

The minimum-bias triggers differ from the random triggers in that they are based on the output from some detector component and not purely random. There are many different triggers which are all subsumed as minimum-bias triggers, having in common that they trigger on events in which some minimum activity of any kind has been detected. The minimum-bias trigger used in Section 5.4 is the EF_mbMbts_1_eff chain which is based on the counts which have been detected in the minimum-bias trigger scintillators (cf. Section 3.4.6).

Combined Triggers

The single triggers from above can be combined in many ways to give combined triggers. The reconstruction of the trigger objects, which has been described above, stays the same, only the hypothesis algorithms, which act on the reconstructed trigger objects, are replaced by new ones which implement a conjunction (logical “AND” condition) of several signatures. The most important example in the context of this thesis are the combined jet + \cancel{E}_T triggers which are studied in Sections 5.1.2 and 5.2. They combine requirements on the online measurement of jets and \cancel{E}_T in the trigger and exist in many variants. They have undergone a number of changes between 2010 and 2011 as explained in the sections referenced above. These changes are partly due to how they combine the jet and \cancel{E}_T trigger objects and what thresholds are applied, but also due to changes in how the underlying trigger objects are reconstructed.

Closely related to the combined triggers are multi-object triggers, which require a given multiplicity of objects of the same type, for example n jets above given thresholds, and special triggers, which combine existing trigger objects in a more sophisticated fashion. An example of the latter are $\Delta\phi$ triggers, which require a certain angular separation between two or more trigger objects that have been reconstructed online. Special triggers are designed to trigger on very specific physics signatures, and as such will become more important when first indications for new physics have been found which allow to narrow down the search, but for the first years of data taking, presumably most of the trigger bandwidth will be spent on inclusive trigger chains rather than exotic signatures.

4.3.5 Trigger Nomenclature

In this thesis, trigger chains will be referred to by giving their official names as they appear in the trigger menu. These names are composed according to the following conventions. The general structure of the trigger names is

$$(\text{Trigger Level})_{\text{-(Multiplicity if } > 1\text{)}}(\text{Threshold Type})(\text{Threshold Value})_{\text{-(Postfix)}}$$

The prefix indicating the trigger level is always L1, L2 or EF, followed by an underscore. For triggers which require a multiplicity of trigger objects larger than 1, the multiplicity is then given in front of the threshold type. The threshold type is an abbreviation stating the type of the trigger. The most important abbreviations for object triggers are e for electrons, g for photons, mu for muons, tau for tau leptons, j for jets and fj for forward jets. Triggers which use global quantities have mb for minimum bias, rd for random triggers, te for the sum of the transverse energy, je for the total energy in jets, xe for the missing transverse energy and, introduced in the middle of 2011, xs for the significance of missing transverse energy (cf. Section 6.6.4). In addition, at Level 1, there is EM, standing for the item triggering on electron/photon objects, which cannot be distinguished at L1. The threshold value is usually given in GeV. Threshold type and value may appear several times for combined triggers.

The names for L1 items always consist only of uppercase letters, whereas for L2 and EF usually lowercase letters are used, but here the postfixes may vary. Common postfixes specify additional trigger features, for example isolation, “filled” or “empty” for the random triggers or “loose” or “tight” for variants of the same chain differing by the thresholds at lower levels. Important postfixes of the chains studied here are “noMu”, which is common to all configured \cancel{E}_T trigger chains in 2010 and 2011 at L2 and EF and says that no muon contributions are taken into account in the trigger decision. Furthermore, “jetNoEF” or “jetNoCut” appear in jet and combined jet + \cancel{E}_T trigger chains in the 2010 trigger menu, to make clear that no cuts are made at Event Filter level on trigger jets. “EFFS” marks jet and combined jet + \cancel{E}_T trigger chains which use the full-scan algorithm at EF.

Some examples for illustration are:

- L1_MU20 = L1 muon trigger item, requiring a muon with a p_T of at least 20 GeV.
- L2_j45_xe20_noMu = jet + \cancel{E}_T trigger at L2, requiring a jet with at least 45 GeV and \cancel{E}_T of at least 20 GeV excluding muon contributions.
- EF_xe30_loose_noMu = \cancel{E}_T trigger, requiring at least 30 GeV at EF, which has looser thresholds at L1 and / or L2.

In the last example, the difference with respect to EF_xe30_noMu becomes clear only from the full specification of the chain, including the thresholds at L1 and L2. Therefore, in the trigger efficiency studies presented in this thesis, for the relevant chains the thresholds are listed for all trigger levels explicitly. Note also that giving a trigger name like “EF...”, unless otherwise specified, usually refers to the whole trigger chain, not only to the cut at Event Filter level.

4.4 Event Reconstruction

In this section, the reconstruction of events from the raw detector output is described. The event reconstruction is the act of turning the pattern of signals from the detector into physics

objects, thereby reducing the large amounts of raw data into a form which is more suitable for physics analysis. Due to the large dimensions of the ATLAS detector and the high interaction rate, combined with the finite velocity of the particles produced in the proton-proton collisions, reconstructed particles which are detected at the same time but at different distances from the collision point at the center of the detector may have been produced in different collisions. This needs to be taken into account when reconstructing event records from the detector output. The same reconstruction algorithms run on real data and simulated Monte Carlo pseudodata, which is both a test for the reconstruction algorithms as well as a device to study the signatures left in the detector by specific particles or reactions and thereby to develop and refine physics analyses.

The reconstruction is done centrally in a post-processing step of the data which has been stored during data taking and uses algorithms which are implemented within the ATHENA framework. In the following, the reconstruction of the different types of physics objects will be discussed separately. The physics objects relevant in the context of this thesis with respect to the trigger studies and the search for Supersymmetry are leptons, in particular electrons and muons, jets, the missing transverse energy and sum of transverse energy, and the primary vertices. Not discussed are e.g. photons or the tagging of jets arising from bottom quarks. At the end of the section, also the measurement of the luminosity and the computation of derived quantities which are used in the search for Supersymmetry such as the effective mass are described. There are two important figures of merit for the quality of the reconstruction of a certain type of physics object. The probability that a given object, which has been produced in a particle collision or subsequent decays, is actually reconstructed is called the efficiency of the reconstruction. In the common nomenclature, the fake rate gives the fraction in which a certain object is reconstructed from the detector data where in reality no such object has been. The purity gives the fraction of correctly identified objects in all reconstructed objects.

The description of the reconstruction algorithms is limited to the information relevant for the scope of this thesis. Like for the ATLAS detector itself, a detailed documentation can be found in [129] and [43]. Note that the reconstruction algorithms are subject to a constant process of evolution and refinement, and therefore the description here may differ from the algorithms in the current release of the ATHENA software.

4.4.1 Leptons

Charged leptons from the first two generations, electrons and muons, are relatively easy to identify, because the muons in particular leave a unique signature in the detector. Identification of tau leptons is more complicated because due to their high mass they have a lot of different possibilities to decay and their lifetime is so short that they will decay close to the interactions point even at large relativistic factors. Neutrinos cannot be reconstructed as individual objects and only be detected by the energy they carry away, which becomes apparent as a violation of the conservation of momentum in the transverse plane in terms of a non-zero value of \cancel{E}_T .

Electrons (and Photons)

The reconstruction of electrons can be subdivided into two tasks: The reconstruction of electron candidates and the subsequent identification of electrons by rejecting fake electrons.

One of the challenges when reconstructing electrons from clustered energy deposits in the calorimeter is to decide whether the cluster is due to an electron, which would leave a track in the Inner Detector due to its electric charge, or a photon. (In fact, in the trigger such a distinction is not even attempted nor possible at Level 1, thus there are only $e\gamma$ objects in the L1 trigger.)

To reconstruct electromagnetic (EM) clusters, a sliding-window algorithm is used, which builds rectangular clusters with a fixed size which are positioned so that the amount of energy within the cluster is maximised. Alternatively, clusters can be formed by connecting neighbouring cells until the cell energy falls below a threshold as is done for jets, but this is not used by the default electron reconstruction algorithms. The optimal cluster size depends on the calorimeter region and is different for electrons and photons, thus different collections of clusters are reconstructed with different window sizes. The calibrated cluster energy is the sum of four contributions: The energy deposited in the material in front of the EM calorimeter, in the calorimeter inside the cluster, outside the cluster (lateral leakage) and behind the EM calorimeter (longitudinal leakage).

An electron is reconstructed from the sliding-window clusters if it matches a track with $p_T > 0.5 \text{ GeV}$. If several tracks are available, the one lying closest in η - ϕ space to the (weighted) center of the cluster is chosen. If there is no track or if there is a reconstructed conversion vertex matching the cluster, a photon is reconstructed. Electron candidates with $|\eta| > 2.5$ are called forward electrons. They lie outside the range of the tracking systems and therefore need to be reconstructed using an alternative algorithm. Forward electrons are not used in the following. To give a feeling for the reconstruction efficiencies and fake rates for electrons, about 0.9% of electrons with $E_T > 20 \text{ GeV}$ and $|\eta| < 2.5$ from $Z \rightarrow ee$ decays are not reconstructed at all and 2.1% are reconstructed as photons [153]. The other way round 2.1% of photons from Higgs decays $H \rightarrow \gamma\gamma$ are reconstructed as electrons.

Actually, three algorithms are run in parallel to reconstruct electrons. The default algorithm, which is seeded by the EM clusters, is dedicated to high p_T isolated electrons. For low p_T electrons and electrons in jets a second algorithm is available, which is seeded by tracks in the Inner Detector. The third algorithm reconstructs forward electrons. For the analysis in this thesis, electrons reconstructed by only the cluster-based `AuthorElectron` algorithm starting from calorimeter seeds or by both the track-based `AuthorSoftE` algorithm and the `AuthorElectron` algorithm are used (cf. Section 7.3.2). An overlap removal to avoid that the same electron enters the collection of reconstructed electrons multiply is included in ATHENA. The identification of electrons (and photons) in ATLAS relies on independent rectangular cuts on variables which have been found to give a good separation between isolated electrons (or photons) and fake signatures from QCD background. These variables include information from the calorimeter and, in the case of electrons, the tracker, and different sets of optimized cuts are provided by the Electron/Gamma combined performance group [43, 153]. In addition to the identification of electrons, which aims at rejecting fake electrons, a procedure has been established to reject electrons the reconstruction of which will suffer from detector problems such as regions in the calorimeter with dead power supplies, faulty read-out electronics (e.g. Front End Boards with dead optical links [142]) or channels with high noise. It relies on so-called Object Quality maps, which for a certain range of data-taking runs specify geometrical regions of the detector which were affected by known detector problems, so that electrons reconstructed in these regions can be treated appropriately in the physics analysis.

Muons

There is a number of different algorithms to reconstruct muons from Inner Detector, muon spectrometer and calorimeter information [150, 162]. Depending on the method which is used to identify and reconstruct them, four kinds of muon candidates are distinguished: The trajectory of *standalone muons* is reconstructed from the muon spectrometer only, correcting the measured momentum for energy loss in the calorimeter and extrapolating the track back to the interaction region. *Combined muons* are muons which are identified by successfully combining a full Inner Detector track and a full track reconstructed in the muon spectrometer. The trajectory in the Inner Detector is used to determine the impact parameter of the muon trajectory with respect to the primary vertex. The other two types start from trajectories in the Inner Detector. If the trajectory in the Inner Detector can be associated with an energy deposition in the calorimeters that is compatible with the hypothesis of a minimum-ionizing particle, it is used to reconstruct a *calorimeter-tagged muon*. If its extrapolation to the muon spectrometer can be associated with a straight-track segment there, it is used to reconstruct a *segment-tagged muon*.

There are two muon algorithm chains in ATLAS, which are both considered suitable for physics analyses, and each is comprised of several algorithms covering all of the above strategies using different pattern recognition strategies. They are called `MuID` and `Staco` after the name of their respective combined reconstruction algorithms. In this thesis, muons from the muon collection created by the `Staco` chain are used. Three different algorithms contribute to this collection: `Staco` attempts to statistically merge the Inner Detector tracks with tracks in the muon spectrometer to produce combined muons. `Muonboy` is a segment-finding code, building a track starting from the outer and middle stations in the barrel and inner and middle MDT stations in the end-cap, i. e. where second coordinate chambers are available, then iteratively adding segments in the other layers until the full track is obtained. `Mutag` starts with Inner Detector tracks above a 3 GeV cutoff and attempts to associate them with `Muonboy` segments, based on the $\eta - \phi$ location and angle where possible. It is used to complement `Staco` in difficult areas and produces segment-tagged muons.

Combined muons are the baseline recommendation [196]. To recover efficiency at $|\eta| \sim 0$ and $|\eta| \sim 1.2$ and at low $p_T < 6$ GeV, in addition to this, segment-tagged muons can be used in physics analyses. Isolation is an efficient way to discriminate between primary muons, i. e. muons which have been produced in the hard scattering process at the interaction point, and muons which have been produced in a jet. To this aim, the energy deposition in the calorimeter within a cone around the muons (excluding the energy deposited by the muon itself), the number of tracks or the sum of momenta of reconstructed objects within such a cone can be used. In studies with cosmic rays [41], a resolution of the relative transverse momentum of about 2% at $p_T = 10$ GeV and 20% at $p_T = 1$ TeV has been measured. The measurement also shows that the momentum resolution of Inner Detector and muon spectrometer complement each other, because the momentum resolution in the muon spectrometer for small p_T is worse due to uncertainties on the energy loss corrections associated with the extrapolation of the muon spectrometer track parameters to the perigee, so that in this range the Inner Detector gives a more precise measurement of the momentum, which improves the performance of combined muons compared to standalone muons. Also, first performance studies on data show that the resolution of the fractional momentum resolution being better than 10% at $p_T = 1$ TeV expected from the simulation is not yet fully achieved [168], the fractional resolution in data being worse by a factor of about 1.5 compared to the simulation.

The ATLAS Muon Combined Performance group therefore suggests to smear the resolution of muons in events from Monte Carlo simulations to achieve a better agreement with the reconstruction of real data, but this is not yet implemented in the analysis of 2010 data presented in this thesis. The reconstruction efficiency has been found to be above 98% for muons from J/ψ decays with $p_T > 6 \text{ GeV}$ [166].

Taus

Tau leptons are much heavier than electron and muon, have a short lifetime and decay mostly hadronically, which makes their identification more difficult than that of the other charged leptons. If they decay leptonically, they are difficult to distinguish from the prompt production of electrons or muons, therefore often only hadronic tau decays are considered. Tau leptons are produced e. g. in decays of W or Z bosons, but are also expected to appear in the decay of neutral Standard Model or MSSM Higgs bosons [171] if these particles exist. Their reconstruction is done starting from jets with nearby tracks as candidates, with a subsequent identification step to distinguish tau from quark jets. In the zero-lepton analysis presented in this thesis, tau identification is not used.

4.4.2 Jets

Jets can in principle be reconstructed at three different levels. From data, jets are always reconstructed as calorimeter jets (detector level), i. e. the jet constituents are groups of calorimeter cells with energy deposits induced by particles. In Monte Carlo, also truth jets can be reconstructed, which are either built from the final stable particles (particle level) in the Monte Carlo event generation, excluding neutrinos and usually also muons, or from the partons in the final state of the hard scattering process (parton level). This is not directly possible in data, because the information about the original partons gets washed out by the process of hadronization. In the following, only calorimeter jets will be considered.

Different inputs can be used to reconstruct jets. The two main concepts are calorimeter towers and topoclusters [143]. Topological energy clusters or short topoclusters are three-dimensional objects designed to follow the shower development, taking advantage of the calorimeter segmentation of the ATLAS detector. Topoclusters are built around seeds, which are calorimeter cells with a signal-to-noise ratio above a threshold of 4, $|E_{\text{cell}}| > 4\sigma_{\text{cell}}$, where σ_{cell} refers to the RMS of the energy distribution for random events and is dependent on both the sampling layer in which the cell resides and the position along the calorimeter in η [157]. In the next step, neighboring cells are iteratively included if they have a signal-to-noise ratio of at least 2. In the final step, one additional surrounding layer of all nearest-neighbor cells is added to the cluster, regardless of their signal-to-noise ratio. This approach was shown to improve energy resolution in single pion test beam studies. It may result in clusters covering large areas of the detector if sufficient energy is present between incident particles. Therefore, in a following step a cluster splitting algorithm splits individual clusters based on local maxima in terms of energy content. Energy from cells may then be shared between some of the new clusters [143]. Due to the subtraction of a pedestal term from the measured energy, some clusters may exhibit a negative energy. These clusters are rejected from the jet reconstruction. The other possible input to jet reconstruction algorithms are calorimeter towers. They use a fixed geometrical grid of calorimeter cells with size $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$, where a tower is created from the energy in all cells within one grid cell in radial direction. Noise-

suppressed towers use only cells belonging to topological clusters, thus applying the same type of noise suppression. Also, towers used as inputs for jet reconstruction are required to have an energy above zero.

Energy scale

The baseline measurement of the energy deposited by a jet is done at the electromagnetic energy scale (EM scale), which is the energy scale measured directly by the calorimeters, i. e. the scale or calibration established using test-beam measurements for electrons and muons in the electromagnetic and hadronic calorimeters. This is the correct scale for the energy of photons and electrons, but does not account for detector effects. As explained in Section 3.2.2, there is a fundamental difference between hadronic and electromagnetic showers for non-compensating calorimeters so that the calorimeter signal from hadrons is lower than that from electrons. Other effects, which concern both hadronic and electromagnetic jets in the same way, are:

- **Dead material:** Energy which is deposited by particles from the jet in inactive regions of the detector is simply lost because it is not measured. These uninstrumented regions of the detector are called dead material.
- **Leakage:** If particles from a jet escape from the calorimeter (e. g. a “punch-through”), the energy of the jet is also underestimated, because its energy is not completely absorbed within the calorimeter.
- **Out-of-acceptance particles:** Jets at high η may lose energy because they are partly out of the region covered by the calorimeter.
- **Out-of-cone particles:** Jets have a fixed size set by the reconstruction algorithm. Particles outside this cone are not included in the energy sum.
- **Differences between truth and calorimeter jets:** Particles may fall out of the reconstructed calorimeter jet, but be included in the truth jet. This may lead to a mismatch between truth jets and reconstructed jets that affects Monte Carlo based calibrations.

To obtain a better resolution in the jet energy measurement, a Jet-Energy Scale (JES) calibration is applied to the measurement at the EM scale. Different calibration schemes are available, all aiming at correcting the energy and momentum of the jets measured in the calorimeter. The calibration that will be used in the analysis presented in this thesis (cf. Section 7) to convert the energy measurements in the calorimeters from electromagnetic calibration (EM scale) to the calibrated hadronic scale is a simple p_T - and η -dependent calibration scheme called EM+JES calibration. It corrects for the non-linear correlation between the energy reconstructed in the calorimeter and the energy of the particles forming jets, and can be calculated from Monte Carlo simulations or from data using γ + jet and dijet balance techniques [157]. The current calibration is based on Monte Carlo simulations and has been validated with collision data [159]. It is also known as jet numerical inversion correction [146]. Other calibration schemes developed in ATLAS are the Global Sequential (GS), Global Cell energy-density Weighting (GCW), and Local Cluster Weighting (LCW calibration) calibration scheme (see below). These use additional information such as the longitudinal and transverse properties of the jet structure (GW) or weights that compensate for

differences between hadronic and electromagnetic jet responses and energy losses in dead material (GCW and LCW) in order to reduce fluctuations in the jet energy measurement, thereby improving the jet resolution between 10 % and 30 % [138].

Resolution

For the computation of the jet resolution, two different methods have been applied to data and Monte Carlo in [158], using the dijet balance [48] and the bi-sector techniques [30]. Both methods give consistent results, and the resolution on Monte Carlo is found to agree with the resolution on data within 14 % for jets with $20 \text{ GeV} < p_T < 80 \text{ GeV}$ and rapidities $|y| < 2.8$. In the updated version, a better agreement within 10 % of the resolution in Monte Carlo and data for jets with $30 \text{ GeV} < p_T < 500 \text{ GeV}$ and rapidities $|y| < 2.8$ is stated [138], the resolution in data being higher than in Monte Carlo by an average 4 %, but consistent within the systematic uncertainties. The parametrization used to fit the relative momentum resolution is the same functional form as typically used for the parametrization of the calorimeter resolution (cf. Equation (3.4)). The uncertainty on the resolution of jets as well as the jet-energy scale uncertainty are input to physics analyses in terms of systematic uncertainties.

Jet Algorithms

The most important algorithm for jet reconstruction in ATLAS is the anti- k_t algorithm [32], which has superseded the previously used Cone algorithm [43] as default algorithm in ATLAS, starting with 2011 data taking also in the trigger at Event Filter level^b. Here, only the anti- k_t algorithm will be considered.

Several theoretical guidelines can be formulated to make the reconstructed jets meaningful with respect to the underlying scattering process in which the particles were produced which deposit energy in the calorimeter. According to these, a good jet algorithm should be infrared safe and collinear safe, i. e. additional soft particles should not disturb the correct reconstruction of the jet and the jet should be reconstructed independent of whether a certain amount of transverse momentum is carried by one particle or this particle is split into two collinear particles. Finally, the same hard scattering should be reconstructed independently at parton, particle and detector level. Moreover, the jet reconstruction should be independent of detector properties, stable against changes in the rest of the event and fast with respect to its implementation.

Algorithms for the reconstruction of jets are typically iterative procedures, which start from seeds that can be partons, particles or reconstructed detector objects with four-momentum representations such as calorimeter clusters, and then repeat the same step a number of times, iteratively combining seeds to form larger entities, until a predefined stopping criterion is met and a jet has been formed. The anti- k_t algorithm belongs to a more general class of sequential recombination algorithms that differ in the exponent p of the energy scale in the distance measure between two entities i and j ,

$$d_{ij} = \min\{k_{ti}^{2p}, k_{tj}^{2p}\} \frac{\Delta_{ij}^2}{R^2}, \quad (4.1)$$

which is used to determine which entities are to be merged next, and is a function of the radius parameter R , the transverse momentum k_t and the geometrical distance Δ_{ij}^2 . For $p = 1$

^b At Level 1 and 2 of the trigger much simpler algorithms are used, cf. Section 4.3.3.

this gives the inclusive k_t algorithm, $p = 0$ leads to the inclusive Cambridge / Aachen algorithm, and $p = -1$ gives the anti- k_t algorithm introduced and discussed in [32]. The advantage of the anti- k_t algorithm is that it shares the positive features of certain iterative cone algorithms, in particular that the reconstructed jets have a regular shape and their boundaries are resilient with respect to soft radiation, but without suffering from collinear unsafety found for iterative cone algorithms.

Relevant Jet Collections

The jets which are used in this thesis are taken from the jet collections `AntiKt4H1Topo` and `AntiKt4TopoNewEM`, which contain jets reconstructed with the anti- k_t algorithm with distance parameter $R = 0.4$ in η - ϕ space, using topological clusters as inputs. The energy of the jets is calibrated by applying a jet-energy scale factor to the energy measurement at the electromagnetic scale, which depends on the transverse jet momentum and the rapidity (EM+JES scale calibration). An implicit minimum p_T of around 7 GeV at the GCW calibrated scale is induced by the reconstruction algorithm (cf. Sections 5.4.3 and A.9), and only central offline jets with $|\eta| < 2.8$ (or smaller) are used here.

4.4.3 Missing Transverse Energy and Sum of Transverse Energy

The missing transverse energy in ATLAS is calculated from the measurements of energy deposits in the calorimeter, supplemented by additional information from other detector components, in particular from the measurement of muons in the muon spectrometers [80, 160]. The missing transverse energy basically is the sum of two terms from the calorimeter and the muon contribution,

$$\cancel{E}_{x,y} = \cancel{E}_{x,y}^{\text{calo}} + \cancel{E}_{x,y}^{\text{muon}}. \quad (4.2)$$

The two components in the transverse plane are obtained as the sum of the projected energies E_i over all N_{cells} cells in the range $|\eta| < 4.5$,

$$\cancel{E}_x = - \sum_{i=1}^{N_{\text{cells}}} E_i \sin \theta_i \cos \phi_i, \quad (4.3)$$

$$\cancel{E}_y = - \sum_{i=1}^{N_{\text{cells}}} E_i \sin \theta_i \sin \phi_i, \quad (4.4)$$

where θ_i and ϕ_i are the polar and azimuthal angle. The sum of the transverse energy $\sum E_T$ is computed via

$$\sum E_T = \sum_{i=1}^{N_{\text{cells}}} E_i \sin \theta_i, \quad (4.5)$$

using the same calibration scheme which is described for \cancel{E}_T in the following.

The ATLAS calorimeters are non-compensating, meaning that the response to hadrons is lower than the response to electromagnetic energy deposits. Energy is also lost in dead material in front of and between the calorimeters, so that dedicated calibration schemes for the hadronic energy deposits in calorimeters need to be developed which include corrections to account for the different calorimeter response in order to improve the resolution of the \cancel{E}_T measurement. The muon term can be computed from muons reconstructed with different

algorithms in two different ways, either from the momentum reconstructed by the external spectrometer only or from the combined measurement of the momentum in the spectrometer and the Inner Detector, where then the energy deposited by the muon in the calorimeters has to be subtracted to avoid double counting.

The high granularity of the calorimeter necessitates to apply a noise suppression which reduces the number of cells which are used in the computation. This electronic noise suppression can be done either by applying an energy threshold to all cells, or by using only cells within topological clusters constructed as described for the jet reconstruction above. The cell energies can be calibrated in three different ways: In the global cell energy-density weighting calibration scheme (global calibration or GCW), the different calorimeter response to hadronic and electromagnetic energy deposits and the energy losses in dead material are compensated for by applying cell-level signal weights obtained from Monte Carlo simulations. In the local cluster weighting calibration scheme (local hadronic calibration or LCW), properties of topological clusters are used to calibrate them individually based on their classification as being of electromagnetic or hadronic nature, using weights which are determined from Monte Carlo simulations of charged and neutral pions. These weights also account for lost energy.

The third way is the refined calibration, which applies different weights to calorimeter cells based on reconstructed physics objects to which the cells can be associated. As the response of the calorimeter differs between physics objects, knowing which object is responsible for the energy deposited in a given calorimeter cell allows to improve the \cancel{E}_T reconstruction because an individual calibration for a particular physics object naturally gives a more precise description and thus a better accuracy than it is possible with only the averaged global calibration. To find the mapping between calorimeter cells and physics objects, calorimeter cells are associated with reconstructed and identified high- p_T objects in the following order: electrons, photons, hadronically decaying tau leptons, jets, muons. Navigating back from the reconstructed object to its component clusters and back again to the constituent calorimeter cells, the calibration weights depending on the type of the reconstructed object are used to replace to initial global calibration of the associated cells. If a cell belongs to several types of reconstructed objects, only the first association is used, avoiding double counting of cells. Cells which cannot be associated with any reconstructed object are included in the \cancel{E}_T calculation using the initial global calibration scheme.

The full composition of the refined \cancel{E}_T is given by the following sum:

$$\cancel{E}_{x,y}^{\text{calo}} = \cancel{E}_{x,y}^e + \cancel{E}_{x,y}^\gamma + \cancel{E}_{x,y}^\tau + \cancel{E}_{x,y}^{\text{jets}} + \cancel{E}_{x,y}^{\text{soft jets}} + \cancel{E}_{x,y}^{\text{calo},\mu} + \cancel{E}_{x,y}^{\text{CellOut}} \quad (4.6)$$

Each term is calibrated independently [80] and computed as the negative sum of the calibrated cluster energies inside the corresponding objects. The upper index specifies the type of object: electrons, photons, tau jets from hadronically decaying tau leptons, jets with calibrated $p_T > 20 \text{ GeV}$ and soft jets with $7 \text{ GeV} < p_T < 20 \text{ GeV}$. $\cancel{E}_{x,y}^{\text{calo},\mu}$ is the contribution from energy lost by muons in the calorimeter and the $\cancel{E}_{x,y}^{\text{CellOut}}$ term is calculated from the cells in topological clusters which are not included in any reconstructed objects. Different choices for the calibrations and selections of the various terms can be done and adopted to the needs of the analysis. A large number of different “flavors” (computation schemes) of the \cancel{E}_T calibration are thus available. In the following, the flavors that are used in this thesis will be further explained.

MET_Topo is used as offline reference for the computation of trigger efficiencies in data

2010 in Section 5.4. It is a very simple computation of \cancel{E}_T , not using the refined calibration described above, but the uncalibrated \cancel{E}_T from cells in topological clusters [156] and comparable to what is done in the trigger (cf. Section 4.3.2), which makes it a useful offline reference for trigger studies. For the trigger studies in 2011 data in Section 5.6, it has been replaced by MET_LocHadTopo, which is also computed from topological clusters, but applying the local hadronic calibration. The missing transverse energy that is used in the analysis presented in Section 7, and also in the trigger efficiency studies on 2010 data, is the \cancel{E}_T defined by the ATLAS Supersymmetry group based on a simplified refined calibration, which is called MET_EMJES_RefFinal_CellOutEM [146]. It is a simplified version that differs from the full refined \cancel{E}_T shown above in that only for a subset of the objects listed above refined calibrations are used. Contributions that do not enter this version of the refined \cancel{E}_T are those from photons and tau leptons. As cells are not available in the Analysis Object Data (AOD) file format, from which the SUSY \cancel{E}_T ^ε is computed, topological clusters (TopoClusters and egClusters) are used instead. All electrons and muons that pass a given selection are included in the \cancel{E}_T , applying no additional overlap removal. (As said above there is an overlap removal done at cell level when computing the refined \cancel{E}_T .) All electrons passing medium purity criteria are used, but as in the analysis an updated, different set of purity criteria (RobustMedium, cf. Section 7.3.2) is used, this will necessitate to correct the \cancel{E}_T for the different definition of electrons. Contributions from all jets above 20 GeV at EM+JES scale are included without any cut on η . Jets below 20 GeV enter the \cancel{E}_T via the CellOut term at electromagnetic scale. For the trigger studies in 2011 data, the updated definition of \cancel{E}_T used by the SUSY group [149] is called MET_Simplified20_RefFinal. It is basically identical to the previous SUSY \cancel{E}_T computation.

Note that the names like MET_Topo are names of data structures (containers), which indicate the calibration or reconstruction variant of the measurements of the energy. These containers do not only contain the \cancel{E}_T values, but e. g. also its components, \cancel{E}_x and \cancel{E}_y , and the sum of the transverse energy. Nevertheless, if it is not specified which quantity is used, the use of the container name usually refers to the \cancel{E}_T measurement. When using \cancel{E}_T as offline reference, e. g. in trigger efficiency plots, the offline \cancel{E}_T variable is written as e. g. MET_Topo_et in this thesis, which refers to the missing transverse energy values taken from the MET_Topo container.

4.4.4 Primary Vertices

The reconstruction of the primary vertex or vertices is part of the track reconstruction [96, 154]. It is divided in two steps, although these steps often may be mixed in the actual implementation of vertex finding algorithms. In the first step, the primary vertex finding, reconstructed tracks are associated to a particular vertex candidate. In the second step, the vertex fitting, the actual vertex position is reconstructed and its quality estimated. Candidates for primary vertices can for example be created from a pre-selection of tracks with a given minimum p_T , which are compatible with the expected bunch-crossing region, and grouping these into clusters. These clusters are then cleaned iteratively for outlying tracks based e. g. on the χ^2 of the fit of the track and the vertex.

^ε To avoid repeating the long name of the \cancel{E}_T variant defined by the ATLAS Supersymmetry group, it will be denoted “SUSY \cancel{E}_T ” in the following.

4.4.5 Effective and Stransverse Mass

Two additional global event variables are employed in the analysis of ATLAS data in the search for Supersymmetry presented in this thesis to define the signal regions: the effective mass m_{eff} and the "stransverse" mass m_{T2} , which is a generalization of the transverse mass to pair decays.

The effective mass is a straightforward sum of energies and momenta obtained from the reconstructed objects passing the baseline selection,

$$m_{\text{eff}} := \cancel{E}_T + \sum_i p_T^{(i)}, \quad (4.7)$$

where the sum with the symbolical index i usually includes a given number of leading jets and all identified leptons, but the concrete implementation may vary from analysis to analysis. In the case of the zero-lepton channel which is relevant for this thesis, no leptons are expected in the selected events and the sum runs only over the n leading jets ($n = 2, 3$) defining the analysis channel. The second term in Equation (4.7) is sometimes denoted H_T . The effective mass provides an estimate of the overall energy scale of the event at hand and can be used to discriminate Standard Model and supersymmetric events. Selecting high m_{eff} values will prefer events in which heavy particles are produced and stand at the beginning of the decay chain, which aims at giving sensitivity to events in which the final state of the proton-proton interaction includes heavy supersymmetric particles.

The transverse mass, which was used for the W boson discovery, is defined as [25]

$$m_T^2 := m_v^2 + m_i^2 + 2(e_v e_i - \vec{v}_T \cdot \vec{q}_T), \quad (4.8)$$

with the masses m_v and m_i and the transverse momenta \vec{v}_T and \vec{q}_T of the visible and invisible particles. The transverse energies are defined to be

$$e_v^2 := m_v^2 + \vec{v}_T^2 \quad \text{and} \quad e_i^2 := m_i^2 + \vec{q}_T^2. \quad (4.9)$$

The transverse momentum of the invisible particle is assumed to be equal to the missing transverse momentum in the event. In decays of a single parent particle with mass m_0 to a visible and an invisible daughter like in decays of W bosons, it holds $m_T \leq m_0$, and the maximum value of the continuous spectrum of m_T over many events can be taken as an estimation for the parent particle mass.

While the transverse mass m_T is a function of the momentum of one visible particle and the missing transverse momentum, the stransverse mass is a function of the momenta of two visible particles and the missing transverse momentum, defined in the following way [71]:

$$m_{T2}^2 := \min_{\sum \vec{q}_T = \vec{\cancel{p}}_T} \left[\max \{ m_T(1), m_T(2) \} \right], \quad (4.10)$$

where $m_T(n)$ is the transverse mass computed from the momenta of the first and second visible particle and the hypothesized unobserved momenta, respectively. Taking the minimum accounts for the fact that the transverse momenta of the invisible particles is not known individually. Only their sum $\sum \vec{q}_T = \vec{q}_T(1) + \vec{q}_T(2)$ is constrained by the measurement of the missing transverse momentum $\vec{\cancel{p}}_T$. In events where two parent particles of mass m_0 decay, m_{T2} gives the lowest upper bound on m_0 consistent with the observed momenta and the hypothesized daughter masses. The maximum of m_{T2} over events can again be taken as

an estimate of the mass of the parent particle. Studies which use a cut on m_{T2} have shown that it is a good compromise between a combination of single-variable cuts with least computational effort but worst signal efficiency, and computing the likelihood for all events for every signal and background hypothesis which is computationally very challenging [25]. It can be used to suppress both physics and detector background [146]. In that, its application is, of course, not limited to searches for signals from supersymmetric processes. The actual computation of m_{T2} in the analysis code is done using the `Basic_Mt2_332_Calculator` class from the “Oxbridge MT2 / Stransverse Mass Library” [24, 71, 195].

4.4.6 Runs and Good Runs Lists

The chronological subdivision of the data taken with the ATLAS detector has several levels: The largest sub-blocks of data which is taken within one year are the data-taking periods. Every data-taking period spans a certain interval of time, or, equivalently, a certain range of run numbers, and is given a letter starting with “A” at the beginning of the year. There may also be subperiods which are numbered with arabic numbers, i. e. period 2011A consists of A1 and A2. The data from proton-proton collisions taken in 2010 is subdivided into nine periods, labelled with letters A – I. The declaration of the a new data-taking period usually indicates that there has been a major change with non-negligible impact on the detector performance.

Each period may consist of one or more physics runs. A physics run usually starts after the beams in the LHC have been accelerated to the nominal center-of-mass energy and declared stable and ends several hours later shortly before a scheduled beam dump, unless there are problems with the detector or the data-taking infrastructure which necessitate an abort of the current data-taking run. Each run is comprised of luminosity blocks (LBs) with a nominal length of one (for 2011 data taking) or two (for 2010) minutes, which in turn consist of several thousand events each. A luminosity block is the smallest entity over which the detector and beam conditions are assumed to be stable. Its length is somewhat arbitrary, but when the data-taking conditions are changed by manual intervention, in most cases due to a change in the trigger prescales (cf. Section 4.3), the beginning of a new luminosity block is enforced.

Not all of the data which has been recorded with the ATLAS detector is suitable for physics analyses. Problems with the detector may compromise the data and render part of the data unusable. The requirements of individual physics analyses with respect to the quality of the data will differ depending on the type of analysis. If, for example, there is a problem with the muon spectrometer deteriorating the muon momentum resolution, analyses which do not use muons can still use the data from the run in which this problem occurred. To accommodate for the needs of different analyses, while providing a unified approach to assess the data quality, the status of the detector is monitored throughout data taking and observed problems with the data are stored in a database. From this information, Good Runs Lists (GRLs) can be created [136]. They are part of the input to every analysis and define which subset of the data taken with the ATLAS detector is suitable to be considered for physics or performance analyses.

The data-quality status flags are stored in the ATLAS conditions database (COOL). They either directly describe the status of various subdetectors or at a more abstract level the quality of some aspect of the data, e. g. some reconstructed physics quantity or the performance of some reconstruction algorithm. The flags can either be green, yellow or red. Green means that no significant problem was spotted with the data taken by this subdetector or

subsystem. Yellow in general means that the data is flawed and not fulfilling all quality requirements, but may be usable for physics analysis after some reprocessing or action taken to fix the problem. Red means that the data is probably lost and cannot be used for physics analysis. (Additionally, they could be gray if the data quality has not been assessed or black if some component has been taken out for this particular run.) It has been found that this “traffic light” categorization in many cases is not flexible enough to consistently reflect the data quality and possible problems, and from the beginning of 2011 has been supplemented by a defect database.

GRLs are stored as text files in the Extensible Markup Language (XML) file format and contain lists of runs, and for each run a list of the luminosity blocks for which all quality requirements have been met. The ATLAS Supersymmetry group has defined their own set of criteria and provides GRLs based on these, which are used in the analysis described in chapter 7. Looser criteria might be sufficient for the other studies which are presented in this thesis, but the gain in statistics is expected to be negligible, and thus for consistency the same GRLs are used to select data for these studies, too.

4.4.7 Luminosity Measurements

The accurate determination of the luminosity is crucial for precision measurements of cross sections of particular physics processes, which are an important part of the physics program of the LHC, because the result for the cross section scales directly with the integrated luminosity. For online measurements of the luminosity, ATLAS can make use of six subdetectors^d, giving a total of 16 independent methods for the luminosity determinations, which differ in performance, sensitivity to background, acceptance and efficiency [173]. The comparison between the measurements of the different subdetectors allows to do cross-checks and helps estimating the systematic uncertainties. The luminosity is measured online, i. e. in real-time, using multiple detectors and multiple techniques, and stored together with the ATLAS data quality information in the COOL database for each luminosity block (LB). In addition to the online measurements, the results from offline algorithms are also available. Offline algorithms can access more detailed information that is not easily accessible online, but can operate only on the small fraction of events which is recorded and are thus less powerful [74]. The luminosity information stored in the database is given in terms of delivered luminosity, without loss corrections. This allows to compute the integrated luminosity for a subset of the data based on the association of LB number and amount of integrated luminosity for the LB, taking into account the trigger prescales, dead time of the data-acquisition system, LBs marked as bad during the data quality assessment and other sources of data loss, and also allows to use and compare different algorithms and calibrations. Each time the calibration is updated, a new database tag is created.

Measurements of the total inelastic pp cross section that rely on Monte Carlo have been found to have relatively large systematic uncertainties. An alternative is the determination of the absolute luminosity using van-der-Meer scans, also called beam-separation scans, which does not require knowledge of the cross section and gives smaller uncertainties, dominated by the measurement of the LHC beam-current [161]. In the first measurements in 2010, the

^d LUCID, a Čerenkov detector specifically designed for luminosity measurements in ATLAS, FCAL, the forward calorimeter, ZDC, the zero-degree calorimeter, BCM, the beam conditions monitors placed at the Inner Detector, MBTS, the minimum-bias trigger scintillators located at the endcaps of the LAr calorimeter, and vertex counting in the HLT. ALFA is not used according to [74].

systematic uncertainty on the absolute calibration of the luminosity via beam-separation scans was found to be 11 % for the pp collision runs taken at the LHC at a center-of-mass energy of 900 GeV and 7 TeV [74]. An update is presented in [165]. It states a 3.7 % larger visible cross section for pp collisions in 2010, compared to [161], with a reduced relative uncertainty of 3.4 %. This is the luminosity calibration that will be used in this thesis to compute the total integrated luminosity (cf. Section 7.2), which is then used for the normalization of the Monte Carlo and setting exclusion limits. Due to a number of changes in the ATLAS luminosity detectors during the shutdown in winter 2010 / 2011, the 2010 calibration cannot be applied directly to 2011 data. A new calibration was recently made available for the 2011 data [170] which has a relative uncertainty of 3.7 %.

In ATLAS, the official tool to compute the integrated luminosity is `ILUMICALC`, which is also available as an online application [187]. It relies on the online and offline information about the instantaneous luminosity that is stored in the ATLAS Conditions Database (COOL) database, which strictly speaking is the instantaneous luminosity integrated over a very short time interval, over which the instantaneous luminosity is assumed to be constant. The calculation of the integrated luminosity for a given run or set of runs is based on GRLs (see above). The same GRL that was used to select the events for the analysis processing is fed into the `ILUMICALC` tool, which then basically sums up the luminosity of all luminosity blocks that are declared good by the GRL.

4.5 Monte Carlo Generation

Monte Carlo (MC) methods are a class of computational algorithms which use random sampling as an essential part of their execution flow. They only yield approximate solutions, but can often provide sufficient accuracy for problems for which exact algorithms would need prohibitively long running times. Another important field of application are computations which include uncertainties on the input data or simulations of physics processes which are of inherently stochastic nature.

In high-energy physics, the simulation of particle collisions and the subsequent interaction of the produced particles with the detector material are the primary example. Monte Carlo methods are used here to generate event data in as much detail and complexity as they would be observed at a real collider and detector. The computer codes simulate the physics processes taking place in the interaction of the colliding particles and produce a simulated final state of an interaction, with particle identities and momenta. Following the notion of a physics event, these codes are called *event generators*. They try to yield results as close as possible to real events, within the bounds of the current understanding of the underlying physics processes, and allow to translate theoretical physics models into simulated physics events. The simulated events can then be compared to the actual measurements to test predictions of the Standard Model and new physics models. Event generators as the first step of the Monte Carlo production chain are independent of the detector. Their output is used as input to the following step, in which the propagation of the produced particles and their interaction with the detector material is simulated. Only in this step, the exact geometry and behavior of the detector and its subsystems enter, and the simulated output of the virtual detector is created. To make a clear distinction from real detector data, events which are produced in Monte Carlo simulations are sometimes called (Monte Carlo) pseudodata.

A fundamental difference between real data and Monte Carlo pseudodata is that in Monte

Carlo events exact information about all particles is available, including the identity and kinematics of particles, positions of decay vertices and the relationship between parent and daughter particles. This information is called (Monte Carlo) *truth information* and is vital for systematic studies of the response of the detector to certain classes of particles or event types and of how signals of new physics processes can be identified. The Particle Data Book [87] gives a numbering scheme for Monte Carlo particles, which was introduced in 1988 to facilitate interfacing between event generators, detector simulation, and analysis packages, and is used for example by the widely used event generators PYTHIA and HERWIG. It not only covers particles which are known to exist, but also a range of hypothetical particles from extensions of the Standard Model like Supersymmetry or technicolor.

Ideally, real data and Monte Carlo pseudodata are fully interchangeable from a structural point of view so that the same analysis code can be run on one or the other without any changes. This makes the data processing transparent and thus less error-prone. It is not always possible and desirable to achieve this though, either because the Monte Carlo simulation may lack some important detail, or because truth information is needed in the analysis which is not available in real data.

The generation of Monte Carlo pseudodata is usually done separately for groups of closely related physics processes rather than simulating inclusive samples which contain all conceivable processes. In addition to the splitting at process levels, there can be another splitting based on event kinematics. The separation has a number of advantages. There are huge differences in the cross sections of different physics processes (cf. Figure 3.6) so that even in large Monte Carlo samples with millions of events, for rare processes only few events would be included in the sample. However, after applying the cuts implementing the analysis selection precisely these rare events may become dominant. Therefore, to achieve a better coverage of the available phase space and not to run out of statistics after harsh cuts, the phase space can be divided into bins, as is done e. g. in dijet samples, using the momentum exchanged between the two initial quarks, and simulate a certain number of events for each bin. To obtain physically meaningful results, all subsamples need to be combined again later, weighting events from the subsamples according to the corresponding cross sections. Having different physics processes split up into different samples also makes it possible to compare different Monte Carlo generators or to study systematic variations, e. g. using different top masses as templates in an analysis aiming at a measurement of the top quark mass [109]. When generating Monte Carlo samples which are binned in some kinematic property, this is usually done by applying a filter which retains only the desired events. This filter efficiency is a multiplicative factor which needs to be taken into account when computing the cross section of the Monte Carlo sample.

4.5.1 Hadronization and Jet Formation

In the discussion of the parton model in Section 2.1.3 and of the measurement of hadronic jets in Section 3.2.2, it was already explained that the partons hadronize and give particle showers before being detected in the calorimeter as clustered energy depositions. The simulation of jets in Monte Carlo event generators can be divided into several steps which reflect this evolution.

The first step is the simulation of the hard process by evaluating the corresponding matrix elements. The hard process takes place at high energy scales, which allows to use perturbative QCD calculations and compute the respective matrix elements exactly. This, how-

ever, can only be done to a limited order in the strong coupling constant α_s , because the complexity increases roughly factorially with the order [6]. A simple example of a higher order process is the emission of a gluon, which eventually gives rise to an additional hadron jet. This corresponds to electromagnetic bremsstrahlung, but due to the larger strong coupling constant is more frequent. Instead of doing an exact calculation of higher orders, a parton shower algorithm can be employed, which does an approximate perturbative treatment of the QCD dynamics above some given threshold for the square of the momentum transfer. It is thereby possible to include higher order terms which can be important in certain regions of phase space. The parton shower proceeds in a sequential and probabilistic manner according to splitting functions, which describe the emission of gluons and creation of quark-antiquark pairs. These branchings in the parton shower can be efficiently implemented in terms of the so-called Sudakov form factor, using Monte Carlo techniques. The parton shower uses as input the partons from the initial and final state of the hard process. Some care has to be taken to avoid double counting of phase-space regions for partons produced in the hard process at higher orders and partons created in the following shower process. The parton shower output, still consisting of partons, is then used as input to the hadronization which converts the partons into hadrons.

The hadronization of partons is an intrinsically non-perturbative process in the low-momentum transfer, long-distance regime, which can only be modeled phenomenologically at present. There are various models for the hadronization [6]. The simplest model assumes the independent fragmentation of every parton individually by iteratively combining quarks with quark-antiquark pairs until the energy falls below a threshold, and splitting gluons into quark-antiquark pairs. An improvement on this model is the string fragmentation model, where the colored partons are connected by extending flux tubes while they move apart, which break when enough energy is stored in them, creating $q\bar{q}$ pairs and forming colorless objects. In cluster hadronization models, after the parton branching colorless clusters of partons form and then decay into hadrons [11]. Finally, short-lived hadrons may undergo decays before reaching detector material. What is finally observed are their decay products. The string model is used e. g. in PYTHIA, the cluster hadronization e. g. in HERWIG.

4.5.2 The k -factor

It may happen that Monte Carlo samples have been produced by a leading order generator, which yields the Leading Order (LO) cross section, but the Next-to-Leading Order (NLO) cross section may be available from dedicated cross section calculations, too. The difference between LO and NLO calculations is quantified in the k -factor, which is defined as the ratio of the cross section at NLO and LO,

$$k = \frac{\sigma_{\text{NLO}}}{\sigma_{\text{LO}}}. \quad (4.11)$$

The k -factor can be used to correct the LO cross section produced by MC generators to compensate for missing higher-order terms. It can become significantly larger than one, especially when new subprocesses appear at next-to-leading order. Scaling all kinematic distributions by the same factor assumes that these distributions are invariant under the inclusion of higher-order terms. This is not necessarily always true, and studies of this can be done by comparing the predictions for fundamental distributions such as that of the p_T of the leading jets from LO generators to those of higher order where available.

4.5.3 Minimum Bias

Here, a few basic terms will be introduced that are used to classify hadron interactions. A distinction is made between *hard processes*, in which a large momentum is transferred between the interacting partons and which are well described by perturbative QCD, and *soft interactions* with a small momentum transfer, which require phenomenological models as they cannot be treated perturbatively due to the large values of the strong coupling constant in this regime. Soft collisions will usually dominate, whereas the production of final states with high- p_T particles and the creation of massive new particles can only happen in hard collisions.

The notion of a *minimum-bias event* is commonly used. Minimum bias is an experimental definition and can be described as everything that would be triggered by a totally inclusive trigger, including any single inelastic collision of two protons, therefore including very rare high- p_T scatters and very common low- p_T scatters. The name comes from the fact that this type of trigger introduces the least possible bias by its selection. Minimum-bias triggers attempt to be as inclusive as possible by making very loose selections (cf. Section 4.3.4).

A related term is the *underlying event*, mostly defined to include all particles from a single particle collision except the process of interest in terms of the hard physics process or all high-energy particles. Still, the underlying event may contain high- p_T particles. It is different from minimum-bias events due to correlations with the hard process. Both a minimum-bias event and the underlying event are dominated by soft partonic interactions.

5 Trigger Studies

The trigger is a vital part of any detector operating at high interaction rates, and trigger studies thus are a vital part of every analysis which evaluates data taken at such a detector. Two major aspects concerning triggers are covered in this chapter: the trigger rates and the trigger efficiencies of the ATLAS trigger system. Concerning the rates, the bandwidth limits of the Trigger and Data-Acquisition System (TDAQ) define how much data can be taken in terms of the amount of data that can be stored per unit of time. Within the resulting maximum total trigger rate, the relative rates of the trigger chains determine the composition of the sample of events that is obtained. The efficiencies account for the fact that the trigger decides in real-time for every event whether it should be kept or discarded. This online decision has to be made within a time span which is much shorter than the time available for the offline reprocessing of the data. Therefore, only a fraction of the full information can be evaluated, and the decision will in some cases differ from the decision which would have been made based on the full information and offline processing so that events which should have been kept are discarded. The trigger mistag rate induced by the necessary compromise between speed and accuracy of the online data-taking process is expressed in terms of the efficiency of the trigger. To know these efficiencies is crucial for the correct interpretation of the recorded data.

The context and the motivation for the studies presented in the following is the search for Supersymmetry at the ATLAS detector, specifically in zero-lepton final states, for which the combined jet + \cancel{E}_T trigger is the primary trigger. The measurements of the trigger rates will allow to demonstrate the advantages of using a combined trigger for this analysis rather than the easier to utilize single object triggers. This necessitates to understand the combined jet + \cancel{E}_T trigger and to measure its efficiencies, important studies to which the largest part of this chapter is dedicated.

5.1 Trigger Rates

5.1.1 Measurements

The measurements and plots of trigger rates presented in this thesis make use of the trigger counts which are stored in the COOL database. These counts are online measurements, which are recorded at the time of data taking and are available for each run and luminosity block. They are written out by the trigger system for all triggers which are configured according to the corresponding trigger menu. On the one hand, the advantage of using this database is that it allows to access a very large range of runs and different triggers easily, without extracting the trigger decision run by run from the data in the different streams. This would be a very laborious process and also occupy a lot of data-processing resources. On the other hand, using these trigger counts it is neither possible to apply selection cuts based on physics objects nor to determine the correlation between triggers or the relative rate of triggers with respect to other triggers because the trigger counts are already summed

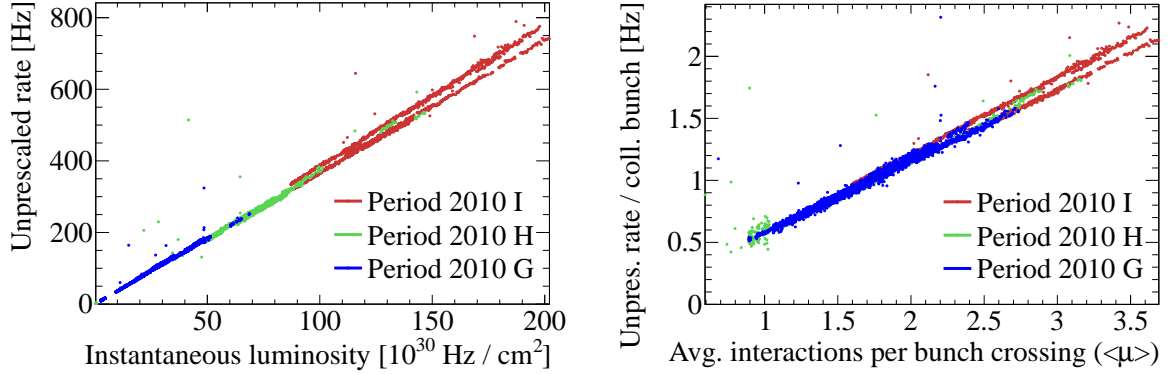


Figure 5.1: Unprescaled rate of the L1_J30 trigger in the last three periods of 2010 data taking (periods G, H and I) as function of the instantaneous luminosity (left) or the average number of interactions per bunch crossing $\langle\mu\rangle$ (right).

over luminosity blocks and therefore an average over a large number of events. However, at least the GRL can be enforced because it is a simple selection based on luminosity blocks, and thus data suffering from bad detector conditions or configurations which differ from normal data taking can be masked out.

The trigger rates can be investigated as function of different variables. Here, the average number of interactions per bunch crossing $\langle\mu\rangle$ (as defined in Section 3.2.5) and the instantaneous luminosity $\mathcal{L}_{\text{inst}}$ are of interest to study the behavior of the rates with changing beam conditions. These quantities are also available per luminosity block in the COOL database. How the plots from COOL are produced is explained in detail in Section A.2 in the Appendix. The rates shown in the following plots are in general corrected for the trigger prescales which have been applied, i. e. the counts which are read from the database are multiplied by the prescale factors at all three trigger levels for the respective run and luminosity block. One dot in the rate plots corresponds to a single luminosity block or an average over several luminosity blocks.

The trigger rates can also be obtained by other methods. An important alternative to predict trigger rates is the measurement of trigger rates relative to another, looser trigger. This allows to study the relative rate reduction introduced by raising trigger thresholds or additional cuts. Measuring absolute rates of triggers is more complicated. As with increasing instantaneous luminosity, the triggers evolve to harsher cuts and not the other way round, the study of evolving trigger cuts relative to existing triggers should be always possible. There are also special runs dedicated to measuring trigger rates for chains with low thresholds, which are called enhanced bias runs. In these runs, triggers with relatively low thresholds are run, and the rates of other triggers can then be measured relative to these. In contrast to the rate measurements stored in COOL, this also allows to study correlations and overlap of triggers and unique rates as well as inclusive trigger rates.

Example Plot: Rate of a Jet Trigger

An example of a trigger rate extracted from the COOL database is given in Figure 5.1. The plot shows the rate of the L1_J30 trigger item before the application of prescales in the periods with the highest instantaneous luminosity in 2010. Every point in the plot corresponds

to one luminosity block from one run so that each point is an average of the trigger rate over approximately two minutes. The GRL of the SUSY group has been applied to mask out data from bad luminosity blocks. The rate is plotted as a function of the instantaneous luminosity and shows a linear behavior over a range of almost two orders of magnitude in luminosity, going from about $5 \cdot 10^{30}$ to $2 \cdot 10^{32}$ Hz/cm². This linear behavior is advantageous because it makes extrapolations of the rates easy. It is an indication for the stability of this trigger, in particular with respect to pile-up. As the center-of-mass energy is constant, and therefore the cross sections are constant, too, the instantaneous luminosity is directly proportional to the number of events taking place in the detector in a given interval of time. A direct proportionality between the trigger rates and the instantaneous luminosity is therefore consistent with the naive expectation. For the \cancel{E}_T triggers, which are discussed in Section 6.1, this will turn out not to be true.

5.1.2 Benefits of Combined Triggers

This section explains the motivation to use combined triggers by demonstrating their benefits over single object triggers. This is done first in terms of general considerations, and in the following part by giving a concrete example for jet + \cancel{E}_T triggers with numbers from ATLAS data taking in 2010.

At the first glance, using a combined trigger entails a number of disadvantages. The measurement of the trigger efficiencies becomes more complicated (cf. Section 5.2.2). At least a two-dimensional binning is needed to accommodate for the two (or more) offline variables, and thus the turn-on curves have a higher dimensionality and more statistics is needed. There may also be fewer possibilities to measure the efficiencies because there are less options for orthogonal triggers, and correlations between the components may give rise to very subtle and hard to understand effects. The advantage of using a combined trigger is that for single object triggers the thresholds will typically be much higher than for combined triggers. Combined triggers therefore allow a higher signal yield than single object triggers within the constraints imposed by the trigger system.

The maximum trigger rate and bandwidth are fixed, and the trigger thresholds need to be optimized within this affordable total rate, i. e. which events are recorded needs to be optimized. The choice of the trigger can be optimized for a specific physics signal by matching the trigger signature to the signature of the searched for signal. Specifically in the context of this thesis discussing the search for Supersymmetry in zero-lepton final states (cf. Section 7), this leads to a trigger which combines a jet signature with an \cancel{E}_T signature, considering the following prerequisite: For a signal where only \cancel{E}_T and none or only soft jets, or where only jets and little \cancel{E}_T are expected, only single object triggers can be used. But many Supersymmetry models have long decay chains, giving hard jets due to the high masses of the supersymmetric partners of the Standard Model particles and lots of \cancel{E}_T due to the heavy LSPs, which evades detection and carries away a lot of energy. The fundamental signature is thus a combination of both large \cancel{E}_T and hard jets. This is illustrated in the exemplary plot to the left in Figure 5.2, which shows a two-dimensional histogram with the p_T of the leading jet at EM+JES scale on the horizontal axis and \cancel{E}_T on the vertical axis, for a Monte Carlo sample implementing a mSUGRA scenario with $m_0 = 360$ GeV and $m_{1/2} = 280$ GeV.

The right plot in Figure 5.2 shows a schematic view of the coverage of different triggers in the plane spanned by leading jet p_T and \cancel{E}_T . The red shaded region represents the single jet trigger efficiency, the blue shaded region represents the \cancel{E}_T trigger efficiency. Events which

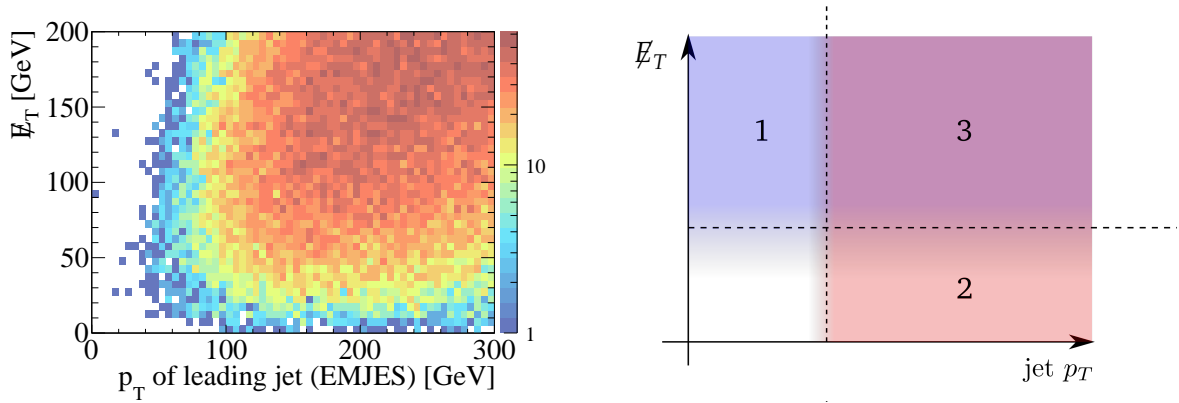


Figure 5.2: Left: Two-dimensional spectrum of the p_T of the leading jet and E_T reconstructed offline as color-coded plot, for one point of the mSUGRA grid used in Section 7 with $m_0 = 360$ GeV and $m_{1/2} = 280$ GeV. Every entry corresponds to one Monte Carlo event. Right: Defining three regions with different types of events in the phase space relevant for jet + E_T triggers.

have E_T and jet p_T in region 1 will only appear in a sample taken with the E_T trigger, events in regions 2 will be triggered by the single jet trigger only, and events in region 3 will be triggered by both the E_T , the jet and a combined jet + E_T trigger. A general guideline is that using an combined trigger is advantageous if the signal has few events in regions 1 and 2 and is concentrated in region 3. The backgrounds, which drive the trigger rates, usually have a steeply falling spectrum and populate in particular regions 1 and 2. By rejecting events in regions 1 and 2, a combined trigger gives a much better signal and background separation than single object triggers, which helps to save bandwidth and enrich the recorded event sample with signal events. To substantiate these considerations, the rates and signal yield of the combined jet + E_T triggers in the ATLAS trigger menu from 2010 will be discussed in the following.

Benefits of Jet + E_T Triggers: Rates

Figure 5.3 gives an overview of the unscaled rates of several jet, E_T and combined jet + E_T triggers in 2010 as function of the instantaneous luminosity. In steps of $25 \cdot 10^{30}$ Hz/cm² in the instantaneous luminosity, all rate measurements from an interval of $\pm 5 \cdot 10^{30}$ Hz/cm² around the central value are collected and fitted with a normal distribution. The fit result gives the rates and their uncertainties, which are shown as vertical error bars in this plot. The horizontal error bars indicate the interval of instantaneous luminosity, from which the rate measurements for the fit are taken. Some points are missing for triggers which were deactivated at the respective instantaneous luminosity. The plot demonstrates that the rates of these triggers cover several orders of magnitude, from the lowest rate for the E_T trigger EF_xe45_noMu with a threshold of 45 GeV up to the highest rate for the jet trigger EF_j20_jetNoEF with a threshold of 20 GeV. In particular the single-object jet and E_T triggers with low thresholds have very high rates, but going to higher thresholds strongly reduces the rates. For low E_T thresholds, already an increase of 5 GeV may reduce the rate by a factor of 10.




Figure 5.3: Unprescaled rates of several jet, \cancel{E}_T and combined jet + \cancel{E}_T triggers in 2010 as function of the instantaneous luminosity.

The primary physics trigger for the analysis of 2010 data in the search for Supersymmetry in zero-lepton final states, which is described in this thesis, is the combined jet + \cancel{E}_T trigger EF_j75_jetNoEF_EFxe25_noMu (cf. Section 7 and Table 7.1). From the plot in Figure 5.3 it can be read off that this trigger has a rate of about 20 Hz at the highest instantaneous luminosity included in the plot, which is $1.75 \cdot 10^{32}$ Hz/cm². This provides an estimate of the rate allotted to this analysis. If, instead of this combined trigger, the single object triggers with the same thresholds were to be used, this would result in an (unprescaled) rate of about 70 Hz for the jet trigger EF_j75_jetNoEF and about 200 Hz for the \cancel{E}_T trigger EFxe25_noMu, resulting in a total rate of 200 to 270 Hz, depending on the overlap^a of the single object triggers. Note that the rates given here are inclusive rates, i. e. they include events which are triggered by other triggers, too. Therefore, switching off these chains does not reduce the total trigger rate by the full rate given here, but having these triggers in the menu means that the total trigger rate will be at least that high. These rates are unaffordable and thus these two triggers have to be prescaled. Their prescales in one of the latest runs in 2010 (run 167776 from period 2010 I2) vary between 36 and 144 for EF_j75_jetNoEF and between 750 and 3000 for EFxe25_noMu, depending on the instantaneous luminosity within the run. In contrast, EF_j75_jetNoEF_EFxe25_noMu ran unprescaled until the end of 2010 *pp* data taking.

The bottom line is that using the combined jet + \cancel{E}_T trigger can effect considerable savings of trigger bandwidth compared to the single object triggers. It reduces the rate roughly by a factor of 10 in this example. Of course, this argument does not take into account how many signal events can be collected with the different triggers. Using a trigger with TeV thresholds would give much lower rates, but it would probably also cut away most of the signal. Therefore, in the following the signal yield of the triggers will be included in the considerations.

^a A total rate of 270 Hz would imply that there is no overlap between these triggers, and then the combined trigger would never fire. This is not the case, of course.




Figure 5.4: Signal efficiency ϵ_T for a number of triggers for Monte Carlo events at the supersymmetric benchmark point SU4. Given is the fraction of events in this sample in which the respective trigger fires. (In the interest of better readability, histogram lines are drawn instead of markers.)

Benefits of Jet + \cancel{E}_T Triggers: Signal Yield

To estimate the impact of using a combined trigger instead of the single object triggers on the signal yield, the fraction of events which pass the trigger selection can be counted on Monte Carlo samples with supersymmetric physics processes. This trigger-dependent signal efficiency is shown in Figure 5.4 for a sample of Monte Carlo events generated for the benchmark point SU4. SU4 is a supersymmetric benchmark scenario implementing mSUGRA (cf. Section 2.3.7), which has a large cross section of 402.19 pb at NLO and $\sqrt{s} = 14$ TeV [43]. The mSUGRA parameter values are $m_0 = 200$ GeV, $m_{1/2} = 160$ GeV, $A_0 = -400$ GeV and $\tan\beta = 10$, and μ is positive. In the m_0 - $m_{1/2}$ plane, SU4 lies in the low-mass region, close to the exclusion bounds set by the Tevatron experiments.

The plot in Figure 5.4 shows several triggers and the fraction of events within the sample for which the respective trigger fires. This fraction will depend on the spectrum of \cancel{E}_T and p_T of the jets which are produced in the supersymmetric decay chains, and therefore also on the masses of the supersymmetric particles. Due to the low masses of the SUSY particles at SU4, it is to be expected that the signal yield will strongly depend on the trigger thresholds, and decrease rapidly for triggers with higher thresholds. For comparison, in Figure 5.5 the same plot for a different point in parameter space is given, using events from one of the Monte Carlo samples of the mSUGRA grid that is introduced in Section 7.6. This point has much larger masses for the supersymmetric particles than the SU4 benchmark point, the relevant mass parameters at the unification scale being $m_{1/2} = 340$ GeV and $m_0 = 1160$ GeV. (Mass spectra for points with similar values of $m_{1/2}$ and m_0 can be seen in Figure 7.5.) Figure 5.5 clearly shows, as expected, that the loss in the signal yield when going to higher trigger thresholds is much lower in this case.




Figure 5.5: Signal efficiency ϵ_T for a number of triggers for Monte Carlo events at one of the points of the mSUGRA grid introduced for the analysis of 2010 data in Section 7.6. This is a point in parameter space with very large masses, $m_{1/2} = 340$ GeV and $m_0 = 1160$ GeV. The numbers give the fraction of events in this sample in which the trigger fires. Note the different range on the vertical axis with respect to Figure 5.4.

A comparison of different types of triggers will serve here as a second motivation why the combined trigger EF_j75_jetNoEF_EFxe25_noMu is indeed a good choice within the bandwidth limits and for the order of magnitude of the instantaneous luminosity in 2010. Background efficiencies and offline cuts will be neglected for this, and an in-depth study would consider a large number of grid points rather than only single points, but the basic argument in favor of combined jet + \cancel{E}_T triggers will be the same. The number of signal events which can be collected at an instantaneous luminosity $\mathcal{L}_{\text{inst}}$ within a time Δt with a particular trigger T is given by

$$\begin{aligned}
 N_T &= \mathcal{L}_{\text{inst}} \cdot \Delta t \cdot \sigma_{\text{SUSY}} \cdot \frac{\epsilon_T}{f_T^{\text{pre}}} \\
 &= c \cdot \frac{\epsilon_T}{f_T^{\text{pre}}},
 \end{aligned}
 \tag{5.1}$$

where ϵ_T is the signal efficiency of the trigger and σ_{SUSY} is the production cross section for the events in the Monte Carlo sample. f_T^{pre} is the prescale of the trigger introduced to limit its rate if it is too high. For simplicity, all quantities are assumed to be constant over Δt . c collects factors which are common to all triggers. N_T/c will be referred to as the signal yield of the trigger T . This is the quantity to be optimized.

The instantaneous luminosity, which determines the trigger rates, will be assumed to be 10^{32} Hz/cm² for this example, about half the peak instantaneous luminosity in 2010. The allotted rate is set to 11 Hz, which is the approximate rate of the trigger EF_j75_jetNoEF_EFxe25_noMu at this instantaneous luminosity. Table 5.1 lists the rates of the triggers shown in Figure 5.3, together with the prescale f_T^{pre} which would be needed to bring the rate within the limit of 11 Hz. The resulting signal yield is N_T/c . The numbers in the table demonstrate

Trigger name	Rate [Hz]	f_T^{pre}	$N_T/c = \epsilon_T/f_T^{\text{pre}}$
EF_j20_jetNoEF	19 043.4	1 731.22	0.001
EF_j30_jetNoEF	2 505.3	227.76	0.004
EF_j35_jetNoEF	1 265.1	115.01	0.009
EF_xe20_noMu	532.2	48.38	0.020
EF_j50_jetNoEF	257.7	23.43	0.043
EF_xe25_noMu	85.1	7.73	0.127
EF_j75_jetNoEF	41.6	3.78	0.263
EF_j75_jetNoEF_EFxe20_noMu	17.2	1.56	0.629
EF_j95_jetNoEF	13.9	1.26	0.774
EF_xe45_noMu	1.5	0.14	0.940
EF_j50_jetNoEF_xe40_noMu	2.5	0.22	0.952
EF_xe40_noMu	3.0	0.27	0.952
EF_j75_jetNoEF_EFxe40_noMu	2.2	0.20	0.956
EF_j75_jetNoEF_EFxe25_noMu	11.2	1.02	0.959
EF_xe35_noMu	6.6	0.60	0.961
EF_j75_jetNoEF_EFxe35_noMu	3.9	0.35	0.963
EF_j75_jetNoEF_EFxe30_noMu	6.8	0.61	0.969
EF_j50_jetNoEF_xe30_noMu	10.5	0.96	0.971

Table 5.1: The rates of different triggers, measured at an instantaneous luminosity of 10^{32} Hz/cm², together with the resulting prescales needed to limit the trigger rate to an assumed allotted rate of 11 Hz and the signal yield N_T/c as defined by Equation (5.1). The row with the primary trigger of the zero-lepton Supersymmetry analysis is highlighted.

that indeed the combined triggers give the best signal yields. It should be noted that the yield depends strongly on the allotted rate, and a lower rate limit will usually prefer triggers with higher thresholds, indicating that the loss in signal yield from higher prescales is more severe than the loss from harsher online cuts, a behavior which is expected for signals with high-energetic objects. Even if triggers with higher thresholds seem capable of higher yields, in comparison to a trigger with lower thresholds which can also be run unprescaled, they would not exhaust the allotted rate and thus give a smaller number of events in the end. EF_j75_jetNoEF_EFxe25_noMu has the best signal yield among the triggers which fully exploit the rate, i. e. have prescales larger than or equal to one, and therefore seems to perform best. This result is biased to some extent because the rate limit has been adjusted to the rate of this particular trigger. It is interesting to see that pure \cancel{E}_T triggers perform quite well, but their rates do not scale linearly with increasing instantaneous luminosity, whereas those of the jet + \cancel{E}_T triggers do. The results of the \cancel{E}_T triggers can therefore not be directly extrapolated to higher luminosities. This is discussed in detail in Section 6.

Conclusion

Considering the savings in terms of trigger rate together with the signal yield, Figure 5.4 shows that even for a low-mass point the signal efficiency of the combined jet + \cancel{E}_T trigger EF_j75_jetNoEF_EFxe25_noMu is still 74 %. Thus in the worst case, only a quarter of events is lost, having on the other hand a factor 10 of reduction in trigger rate compared to single object triggers. To sum up, the search for Supersymmetry in the zero-lepton channel can benefit considerably from using jet + \cancel{E}_T triggers instead of single object triggers, even if these triggers are more complex, and more work is required to understand their behavior. Therefore, the next chapter will be dedicated to the study of the efficiencies of combined jet + \cancel{E}_T triggers.

5.2 Trigger Efficiencies: Methodology

5.2.1 Introduction

The efficiency of a trigger, or trigger efficiency for short, is the probability of a trigger to accept a given event^b. It depends on the detector hardware and implementation details of the trigger system, and needs to be determined individually for every trigger chain that is defined for the online data taking. The trigger efficiency will, of course, also depend on the properties of the events themselves. This dependency may be of arbitrary complexity, although a well-defined trigger should be designed such that these dependencies are reduced to the least possible.

The total number of interactions N in a collider experiment is per definition given by the instantaneous luminosity $\mathcal{L}_{\text{inst}}$ and the cross section σ of the process under study, integrated over time,

$$N = \int \mathcal{L}_{\text{inst}}(t) \sigma dt. \quad (5.2)$$

^b Note that in the following, the term trigger can have two different meanings: It can refer either to a specific trigger chain or to the trigger system as a whole. Which meaning is intended can be inferred from the context.

The number of recorded events n , however, may be smaller,

$$n = \epsilon N, \tag{5.3}$$

due to the trigger efficiency $\epsilon \leq 1$, which includes the acceptance and online reconstruction efficiency. To retrieve the true number of events $N = n/\epsilon$, the efficiency of the trigger needs to be known.

Taking an \cancel{E}_T trigger in ATLAS as an illustration, the efficiency of this trigger will depend on the calorimeter, the calibration and the details of the online measurement of \cancel{E}_T , and needs to be measured specifically for the ATLAS detector. The efficiency in terms of the probability for a given event to fire a given trigger depends on the \cancel{E}_T threshold of the trigger and the actual amount of \cancel{E}_T in the event. As the actual amount of \cancel{E}_T is unknown, the \cancel{E}_T value which is reconstructed offline is taken as the best available estimate to study the efficiency of the trigger. Plotting the probability of the trigger to fire as function of the online \cancel{E}_T measurement would yield a step function: If the measured value is below the trigger threshold, the probability is zero; if it is above, the probability is one. In the ideal case, the relation between the online and the offline measurement of \cancel{E}_T is a monotonous function, and the correlation is one. But as the two measurements are not exactly the same, the probability as function of the offline \cancel{E}_T variable will be a convolution of the step function and the relation between the online and the offline measurement. As said above, there may also be additional dependencies on other event properties. Indeed, for the \cancel{E}_T trigger it can be seen on Monte Carlo that the efficiency strongly depends on the type of \cancel{E}_T and is different for events with fake \cancel{E}_T and events with real \cancel{E}_T [Ⓢ]. In case of the \cancel{E}_T trigger, it is challenging to distinguish between fake and real \cancel{E}_T , and the trigger efficiencies therefore can only be measured as an average over events with both sources of \cancel{E}_T , which will result in larger uncertainties of the trigger efficiency. But there are other examples, in which the influence of the additional dependencies on the trigger efficiency can be determined, e. g. a potential dependence on the isolation of a muon which fires a muon trigger, or the geometrical distance of two jets in an event accepted by a jet trigger.

As there sometimes seems to be some confusion about what is a trigger efficiency, the following is intended to clarify the interpretation which shall be adopted here: The trigger efficiency by itself is a figure which is the same for all analyses because the trigger does not know about the offline selection. The decision which is made by the trigger system online therefore is independent of what happens later in the analysis. Still, care has to be taken because folding in the offline selection will lead to a seemingly different trigger efficiency. If, for example, the offline selection requires many hard jets, this will yield a sample of events which receives large contributions to \cancel{E}_T from mismeasurements of the jets (which is proportional to the jet p_T), so that the measured efficiency of the \cancel{E}_T trigger on this sample will differ from its efficiency on an unbiased sample. This is related to the fact that the trigger by definition acts as a filter, and thus does preferentially select certain event topologies. Thus, if the efficiency of the trigger seems to depend on the offline selection, this is due to the impact of the offline selection on the sample on which the efficiency is measured, rather than the trigger efficiency being different itself. The trigger efficiency is also independent of the type of events in terms of the underlying physics process, in the sense that it only depends on the objects which are reconstructed by the trigger and used in the trigger decision. Along the same lines, it should be noted furthermore that if the offline variable, as function of which

[Ⓢ] This can be seen in Figure 5.15, for example.




Figure 5.6: Typical example of the efficiency of a trigger as function of some offline reference, explaining some of the terms which will be used in this text.

the trigger efficiencies are presented, depends on the event topology or type of event in a different way than the online measurement, this can lead to a distorted shape of the turn-on curves or to dependencies on the offline event selection^d. Such a distortion can be avoided by using an offline variable which comes as close to the online measurement as possible. It is therefore important when interpreting trigger efficiencies to take into account on which sample of events and as function of which variable they are presented.

5.2.2 Terminology of Trigger Efficiencies

In the following, the vocabulary needed for the discussion of trigger efficiencies will be introduced. The plot of a trigger efficiency as function of the offline reference is referred to as the *turn-on curve* of a trigger. A sketch of a one-dimensional turn-on curve is shown in Figure 5.6. Turn-on curves often resemble a Gauss error function in shape and have three distinct regions: one where the efficiency is flat and close to zero for small values of the offline reference, an intermediate region, in which the slope is large, and the region at high values, where the efficiency flattens out again. The latter is called the *plateau (region)* of the efficiency curve. The average efficiency in the plateau region is called the *plateau efficiency*. For jet triggers, the plateau efficiency is usually close to one because jets are hard to miss if they have sufficiently high energy^e. For similar reasons, the plateau efficiency of \cancel{E}_T triggers is close to one, whereas for e. g. muon triggers it may well be below one because muons may escape through cracks in the muon spectrometer or, being minimum ionizing particles, leave too weak a signal. In any case, the efficiency for a well-defined trigger should be constant in the plateau region. The region in which the trigger efficiency changes from zero to its plateau value is called the *turn-on region*. If the turn-on region is broad, the turn-on is slow,

^d This effect can be seen, for example, in Figure 5.19.

^e Even in case parts of the calorimeter are malfunctioning, if this affects both online and offline measurement in the same way, the effect cancels out and is not visible in the trigger efficiency (cf. the impact of the so-called LAr hole on the efficiencies in Section 5.6.3).




Figure 5.7: Schematic behavior of a typical trigger turn-on curve, showing the trigger efficiency ϵ_{sim} as function of the offline quantity Q_{offline} for different functional forms of the resolution $\sigma(Q_{\text{offline}})$ of the online measurement with respect to offline.

if it is narrow, the turn-on is steep. A steep or sharp turn-on is another indicator of a well-defined trigger and of a high resolution of the online measurement with respect to offline. The shaded regions in Figure 5.6 highlight two different, undesired effects, coming from the fact that the trigger efficiency is not a step function as function of the offline reference. The “barren” rate is made up by events in which the trigger fires, but which are not used in the analysis because they fall below the offline cut. This increases the trigger rate, but does not have a direct impact on the analysis and may be affordable or unavoidable. “Lost” events are events for which the trigger does not fire, although they would have exceeded the offline threshold. These events will be missing later in the analysis and will lead to wrong conclusions if they are not accounted for.

If the online measurement of the quantity the trigger cuts on can be described as a smearing of the offline measurement with a Gaussian of constant width, the trigger efficiency as function of the offline reference can be described by a Gauss error function. In general, the resolution of the online measurement with respect to offline may change as function of the offline variable, and then the trigger efficiency deviates from a pure Gauss error function. This is demonstrated in the plot in Figure 5.7, where the trigger efficiency is computed for a resolution of the online measurement which depends in different ways on the value of the offline reference. Assuming that the resolution gets worse for higher values, this makes the turn-on curve asymmetric and its upper part stretch more to the right.

5.2.3 Estimation of Efficiencies

The basic principle of the computation of a trigger efficiency is counting the number of occurrences in which the trigger could have fired (the denominator of the efficiency estimate n), and the number of occurrences in which it actually did (the numerator m), both as a function of the offline variables which are used to parametrize the trigger efficiency. The

ratio of the counts is used as the estimator of the trigger efficiency,

$$\hat{\epsilon} = \frac{m}{n} \xrightarrow{n \rightarrow \infty} \epsilon, \quad (5.4)$$

where m obviously depends on n . The hat, which is often used to denote an estimator, is omitted in the following, identifying for practical applications the estimate of the trigger efficiency with the true efficiency. The counting is usually done using two one- or higher dimensional histograms with an appropriate binning in the relevant offline variables. The difficult part is to determine which offline variables are relevant, i. e. in which variables the efficiency is not flat, and how to reliably do the actual counting, in particular with respect to the choice of the event sample on which the counting is done.

So far, trigger efficiencies have been discussed in terms of probabilities which are assigned to every event as a whole. In some cases, the probability for an event to fire a trigger can be broken down into object probabilities which, appropriately combined, give an approximation of the total event probability. If the object probabilities are independent of the rest of the event, this approximation is exact. In many cases, the dependence is so small that the approximation works well, at least for low object multiplicities. If the trigger efficiency of an object described by a set of parameters $\vec{\sigma}$ is given by $\epsilon_{\text{obj}}(\vec{\sigma})$, the trigger efficiency of an event with n such objects, i. e. the probability of such an event to fire the trigger, is given by

$$\epsilon_{\text{ev}}(\vec{\sigma}_1, \dots, \vec{\sigma}_n) = 1 - \prod_{i=1}^n (1 - \epsilon_{\text{obj}}(\vec{\sigma}_i)), \quad (5.5)$$

under the assumption of independent object probabilities.

To be able to compute the object-wise trigger efficiencies, a matching needs to be done between the objects reconstructed in the processing of all online data by the trigger system, which are called trigger objects, and the respective objects which are reconstructed later in the offline processing of the recorded data. In this matching, corresponding trigger objects and offline objects are identified, usually based on their geometrical distance in $\eta \times \phi$ space. From the properties of a trigger object, it can be inferred whether this object has fired the trigger or not. Alternatively, the information whether a given trigger object has issued the trigger signal is usually stored in the event record. For the binning of the trigger efficiency, not the value measured for the trigger object is relevant, but the value of the respective property of the corresponding offline object. This is made explicit in the description of the concrete implementation of the computation of trigger efficiencies (cf. Section 5.3).

5.2.4 Use Cases and Trigger Strategies

Taking into account the trigger efficiency is an important aspect of every physics analysis and needs to be done with care because the large rate reduction accomplished by the trigger system, which stands at the beginning of every analysis, may have a considerable influence on the analysis results. Generally speaking, the trigger efficiencies are needed to estimate the impact of the trigger selection on the composition of the event sample on which the analysis is carried out.

To this end, the trigger efficiency can be exploited in different ways: The simplest method, for a trigger which attains near 100 % efficiency in its plateau region, is to adjust the offline selection so that only events in the plateau region pass, and then to assume the trigger to

be fully efficient. This approach is adopted in the official search for Supersymmetry in the zero-lepton channel and the analysis presented in Section 7. The jet + \cancel{E}_T triggers used in this analysis fulfill the condition of having a plateau efficiency near 100 % (cf. Section 5.3.2). If changing the offline selection is undesirable or not possible for some reason, the trigger efficiency can be taken into account by using the result of the trigger simulation in Monte Carlo to accept and reject events in the same way as it is done in data. This relies on the correctness of the trigger simulation, and has the disadvantage that for low trigger efficiencies a considerable part of the Monte Carlo is lost, namely all events which are rejected based on the (binary) simulated trigger decision. For triggers with a maximum efficiency which is considerably below one, like it is found e. g. for the muon triggers in the barrel which only reach around 75 % [111], the fraction of discarded Monte Carlo events is sizeable. In this case, it is better not to use the trigger simulation to reject these events, but to assign a weight to each Monte Carlo event, which reflects the trigger efficiency for this particular event and thus the probability that such an event contributes to the data sample collected online. This reweighting of events allows to extract the maximum possible information from the Monte Carlo. Other analyses which, for example, aim at the determination of production cross sections, need precise estimations of trigger efficiencies from data only to be able to account for inefficiencies of the trigger and obtain the true number of selected events.

5.2.5 Trigger Efficiencies and the β -Distribution

The trigger efficiency is defined as the average probability of the trigger to fire on a sample containing an infinite number of events. For the finite case, the outcome of a series of statistically independent experiments with success probability p is in general described by the Binomial distribution. The probability of seeing k successes in n identical trials according to this distribution is given by

$$\text{Bi}(k|n,p) = \binom{n}{k} p^k (1-p)^{n-k}, \quad (5.6)$$

where $\binom{n}{k}$ is the binomial coefficient [2]. The measurement of trigger efficiencies yields two numbers for each bin i : the total number of events or trials n_i and the number of successes in which the trigger fired $k_i \leq n_i$. From these two numbers, the probability distribution $P(\epsilon_i|k_i, n_i)$ for the true efficiency ϵ_i is to be computed.

In order to find the distribution $P(\epsilon_i|k_i, n_i)$ starting from the assumption that the distribution of k_i , given n_i and p_i , is given by $\text{Bi}(k_i|n_i, p_i)$, Bayes' theorem can be employed. In terms of (conditional) probabilities on a discrete set of events[‡] which includes the possible outcomes A and B , it can be written as

$$P(A|B) = \frac{P(B|A) \cdot P(A)}{P(B)}. \quad (5.7)$$

For continuous probability distributions, this translates into $P(\epsilon_i|k_i, n_i) \propto P(k_i|\epsilon_i, n_i)P(\epsilon_i|n_i)$. Inserting the Binomial distribution and requiring the prior $P(\epsilon_i|n_i)$ distribution for the efficiency to be independent of the sample size n_i yields

$$P(\epsilon_i|k_i, n_i) \propto \text{Bi}(k_i|\epsilon_i, n_i)P(\epsilon_i). \quad (5.8)$$

[‡] "Event" being meant here in the general stochastic sense, not as the record of a particle collision.

The conjugate prior distribution of the Binomial distribution is the β -distribution, meaning that if a β -distribution is inserted as the prior probability $P(\epsilon_i)$ in Equation (5.8), the resulting distribution $P(\epsilon_i|k_i, n_i)$ belongs again to the family of β -distributions.

The β -distributions $\beta(x|a, b)$ are a family of continuous probability distributions over the interval $x \in (0, 1)$ with two positive parameters a and b . They are given by

$$\beta(x|a, b) = \frac{x^{a-1}(1-x)^{b-1}}{\int_0^1 u^{a-1}(1-u)^{b-1} du}, \quad (5.9)$$

and the following of their properties are relevant here:

$$\text{mean: } E(X) = \frac{a}{a+b} \quad (5.10)$$

$$\text{variance: } V(X) = \frac{ab}{(a+b)^2(a+b+1)} \quad (5.11)$$

$$\text{mode: } M(X) = \frac{a-1}{a+b-2} \quad (5.12)$$

From setting $P(\epsilon_i) = \beta(\epsilon_i|a, b)$ in Equation (5.8), it follows that

$$P(\epsilon_i|k_i, n_i) = \beta(\epsilon_i|a', b'), \quad (5.13)$$

with $a' = a + k_i$ and $b' = b + n_i - k_i$ [34]. In this thesis, a uniform (“flat”) prior $P(\epsilon_i)$ will be used, because this makes the mode, the most likely value of the β -distribution, fall together with the ratio of the number of passed events k_i over all events n_i . This is the intuitive expectation for the trigger efficiency, and it can also be motivated *a priori* as the assumption of no value for the efficiency being more likely than any other. The uniform distribution corresponds to parameters $a_1 = b_1 = 1$, so that the posterior probability of the trigger efficiency is given by

$$P(\epsilon_i|k_i, n_i) = \beta(\epsilon_i|k_i + 1, n_i - k_i + 1). \quad (5.14)$$

From Equation (5.12), it follows that the mode then is k_i/n_i as expected. The error bars which are given in the trigger efficiency plots in this thesis indicate the smallest interval covering 68.3% of the β -distribution obtained in the way described above.

5.2.6 Overview of Methods

In this section, an overview of methods is given that can be used to determine trigger efficiencies. One way is to use Monte Carlo, where the trigger efficiencies can be determined by direct counting. This is sometimes referred to as *Monte Carlo counting*. Monte Carlo counting relies on the correctness of the trigger simulation and can be used to validate the other methods described below. The efficiencies obtained from Monte Carlo are always specific to the type of events in the sample and thus biased. This can be an advantage, for example to optimize the trigger on certain event topologies. However, it is not possible to measure the trigger efficiency for an unbiased sample in Monte Carlo because statistics in Monte Carlo are too small to obtain meaningful results for triggers which select rare events. A general advantage of data-driven techniques is that they do not rely on the trigger simulation in Monte Carlo.

Therefore, the common objective of the methods presented here to measure trigger efficiencies on data is to obtain an unbiased sample of events and thus an unbiased estimate of the trigger efficiency:

- The *tag & probe* method exploits the fact that for a given two-body decay with known kinematics, from the observation of one of the decay products it can be inferred that also the other particle should be present in the event. Decays with two identical particles, usually leptons, in the final state are used. One of the particles is selected as the tag which is used to collect the event sample, the other particle is the probe which is used to perform the actual measurement of the trigger efficiency. The efficiency is computed as the number of probes firing the trigger divided by the total number of probes. The tag & probe method can be used to study e. g. electron [122] or muon [166] trigger efficiencies in decays of Z bosons, J/ψ mesons or Y mesons. A cut on the invariant mass of the reconstructed pair of particles and the back-to-back topology can be used to obtain a well-defined sample of such decays. Note that in this method a sample dependence remains, in particular with respect to the kinematic distributions of the decay products. An important assumption is furthermore that the trigger efficiencies of the objects are uncorrelated, i. e. the trigger efficiency for the probe is assumed to be independent of the fact that there is always another object of the same kind in the event which has fired the trigger. The tag & probe method is also employed to compute (offline) reconstruction efficiencies [169].

- An event sample to compute an unbiased estimate of the trigger efficiency can also be collected using an *orthogonal trigger*, which is a trigger that, in the ideal case, is completely independent of the trigger under study. It is not always possible to identify such a trigger and the orthogonality will often be only approximate. For example, for the computation of \cancel{E}_T trigger efficiencies, muon triggers could be used to collect the event sample because the \cancel{E}_T triggers in ATLAS at the time of writing are independent of the measurements in the muon spectrometer. However, the muon selection will enhance the fraction of events with $W \rightarrow \mu\nu_\mu$ decays, which have real \cancel{E}_T from the neutrino, thereby changing the composition of \cancel{E}_T in the event sample.

To obtain a sample with very little or no bias at all, *minimum-bias* or *random triggers* can be used. For most studies, these triggers yield far too small statistics, which is obvious considering the magnitude of the rate reduction in the trigger system. To get down from the design collision rate of 40 MHz to the rates of a typical physics trigger with a rate of a few Hertz, a rate reduction of the order of at least 10^6 is necessary. This means that to collect a sample of a few thousands of events in which the target trigger actually fires, billions of events would have to be collected with random or minimum-bias triggers, and these triggers typically are run at rather low rates, so that even for the lowest thresholds direct studies on these samples are not feasible.

- *Bootstrapping* overcomes the problem of insufficient statistics in samples taken with unbiased triggers by using a biased event sample instead, on which the trigger efficiencies are computed. The bias which is introduced by the sample trigger can then be corrected under certain conditions, using Bayes' theorem and the turn-on curve of the sample trigger, which must be known to perform the bias correction.

Bootstrapping is the main data-driven method used in this thesis to compute trigger efficiencies, and is therefore explained in detail in Section 5.2.7. In this section and the following

Section 5.2.8, it is also described how to perform the actual computation of the efficiency estimates from the event counts and how to combine the trigger efficiencies from several weighted samples, which is needed to correctly propagate the uncertainties of efficiencies computed on weighted Monte Carlo samples.

5.2.7 Bootstrapping

In this section, the notion of bootstrapping in the context of measuring trigger efficiencies will be explained. The trigger under study will be referred to as the target trigger in the following. The trigger that is used to collect the sample of events on which the efficiency of the target trigger is measured will be referred to as the sample trigger. The basic idea has already been sketched in Section 5.2.6. Instead of insisting to have an unbiased sample of events, on which the efficiency of the target trigger is measured, this sample may now also be collected using a sample trigger which introduces a bias. This solves the problem that for triggers that are tailored towards selecting rare events, often it is simply impossible to acquire an unbiased sample, comprised of minimum-bias or random events, which contains enough of these rare events so that it is possible to determine the efficiency of the target trigger. Of course, measuring the trigger efficiency on a biased sample yields a biased measurement, but this bias can be removed by applying the bootstrapping method.

Bootstrapping is based on Bayes' theorem, which is given by (cf. Equation (5.7))

$$P(T) = \frac{P(T|S) \cdot P(S)}{P(S|T)}. \quad (5.15)$$

The probabilities $P(S)$ and $P(T)$ can now be identified with the probability that the sample trigger chain issues a trigger signal in a given collision event, $P(S)$, and the probability that the target trigger issues a trigger signal in this collision event, $P(T)$. This makes $P(T|S)$ the probability of the target trigger to fire in an event where the sample trigger has fired, which is true for all events from the sample which have been collected using the sample trigger.

If the probability $P(S|T)$ in the denominator were known, the unconditional (unbiased) probability $P(T)$ of the target trigger to fire could be computed from $P(T|S)$ and $P(S)$. In practice, the sample and target trigger are chosen such that

$$P(S|T) = 1, \quad (5.16)$$

i. e. if in a given event the target trigger fires, the sample trigger fires with probability one, too. It is then possible to calculate the unbiased target trigger efficiency $P(T)$ from the trigger efficiency of the sample trigger $P(S)$ and the biased efficiency of the target trigger on the sample of events collected with the sample trigger:

$$P(T) = P(T|S) \cdot P(S). \quad (5.17)$$

This is the central equation of the bootstrapping method. Although in this simple derivation, the necessity of Equation (5.16) for this method to work is obvious, it should be stressed that it is vital to make sure that this condition really holds when applying Equation (5.17). If both the sample trigger and the target trigger make their decision based on the same online trigger quantity, e.g. on the E_T reconstructed by the trigger, and if the sample trigger has a threshold lower or equal to the target trigger, then Equation (5.16) holds by definition.

In short, the following three conditions must be met to apply the bootstrapping method:

- The unbiased sample trigger efficiency needs to be known, either from another method or by bootstrapping the sample trigger efficiency from a third (even more inclusive) trigger.
- The target trigger efficiency needs to be measured with respect to the sample trigger, i. e. all events which fire the sample trigger go into the denominator and all events which fire the target trigger go into the numerator.
- The sample trigger must be chosen such that Equation (5.16) holds. This will be discussed in Section 5.4.2.

Propagation of Uncertainties

Equation (5.17) is written in terms of discrete probabilities, rather than the probability distribution which appears in Equation (5.14) as the result of the computation of the trigger efficiency. As such, it only holds in the case of infinite statistics, i. e. if the probabilities were known exactly. The form of Equation (5.14) expresses the fact that this is not the case, but only estimates of the efficiencies can be derived from the measurement. It accounts for the uncertainty on the efficiency estimates by describing the efficiency in terms of a probability distribution. The generalization of Bayes' theorem and Equation (5.17) to probability distributions is straightforward, and then implies the multiplication of two probability distributions. Note that this multiplication does not depend in any way on offline or online trigger quantities, but solely and independently on the counts in the respective bins of the numerator and denominator histograms.

This multiplication of two probability distributions suffers from a fundamental problem. In case of low event counts, both distributions are relatively broad so that, even if the estimator for both efficiencies is one and the probability distribution for both efficiency estimates is peaked at one, the multiplication will yield a distribution that no longer is peaked at one due to the tails. This leads to a systematic underestimation of the efficiency, and in particular makes the efficiency appear to decrease for the bins at high values of the offline reference, which typically have less statistics. It is then not clear whether this decrease stems from the decrease in statistics or is due to an inherent deficiency of the trigger algorithm. To amend this problem, a reweighting of the product probability distribution can be done, as it is explained in the next section. The actual multiplication of the probability distributions (which are β -distributions) may be done analytically [31], but for simplicity a numerical approach is preferred here. It consists of randomly sampling numbers from the two distributions and filling a histogram with the product. From this histogram the mode and the smallest interval covering 68.3% are computed as the estimate of the efficiency and its lower and upper confidence intervals, covering quantiles equivalent to one standard deviation for a Gaussian distribution.

Accounting for Bias from Bootstrapping: Logarithmic Reweighting

The multiplication of two β -distributions, as is necessary in the propagation of uncertainties in the bootstrapping method, will in general not leave the mode invariant. In particular, the




Figure 5.8: Turn-on of the jet trigger $EF_{j75_jetNoEF}$ measured on events taken with $EF_{j30_jetNoEF}$ in 2010 G – I with different methods: Biased turn-on before bootstrapping (upper left), after bootstrap correction of bias (upper right), after bootstrapping without error propagation (lower left) and with bootstrapping and reweighting (lower right). The offline reference is the p_T of jets from the `AntiKt4H1Topo` collection at EM+JES scale (cf Section 4.4.2).

mode of the product distribution is lower than the product of the modes, if the two distributions are peaked at one. This will lead to an underestimation of the trigger efficiency, if the probability distributions are broad as is the case for low statistics. This is demonstrated in Figure 5.8, using the example of the efficiency of the jet trigger `EF_j75_jetNoEF` (target trigger), which is measured on events taken with `EF_j30_jetNoEF` as sample trigger in 2010 G – I. The four plots compare four different ways of computing the efficiency. The upper left plot shows for comparison the biased efficiency of the target trigger measured on the sample of events taken with `EF_j30_jetNoEF` before bootstrapping. In the other three plots, bootstrapping is applied. For the lower left plot, it was assumed that the estimate of the efficiency of the sample trigger efficiency is exact, i. e. the uncertainties of the estimate of the sample trigger efficiency are not propagated, effectively multiplying with the mode instead of the full distribution of the sample trigger efficiency[§]. This plot serves as a reference of what the bootstrapping method ought to return as the estimate of the target trigger efficiency (but not for the uncertainties, which naturally are too small). Comparing this to the upper right plot, which shows normal bootstrapping using the full sample trigger efficiency distribution, the underestimation of the efficiencies is clearly visible in the plateau.

The solution to this problem is a phenomenological approach, in which the resulting product distribution is reweighted. This approach has been introduced and described in [172], and is motivated by considering the product $Y = X_1 \cdot X_2$ of two independent random variables X_1 and X_2 with uniform distribution over the interval $[0, 1]$. Its probability distribution is given by

$$f_Y(y) = \mathcal{N} \iint f_{X_1}(x_1) f_{X_2}(x_2) \delta(x_1 \cdot x_2 - y) dx_2 dx_1, \quad (5.18)$$

where \mathcal{N} is a suitably chosen normalization factor, f denotes the probability density function of the respective random variable, and the integrals are over the support of the probability distributions. For $0 \leq y \leq 1$, the integrals evaluate to $f_Y(y) = -\ln y$. This factor can be understood from considering the extreme cases: There is only one possibility to combine x_1 and x_2 to obtain $y = 1$ ($= 1 \cdot 1$), but infinitely many for 0 ($= 0 \cdot x \forall x \in [0, 1]$). Identifying X_1 and X_2 with the trigger efficiencies which are multiplied in Equation (5.17), a factor $-1/\ln(y)$ is used to reweight the resulting distribution. This corrects for the combinatoric effect involved in the efficiency multiplication, which would otherwise lead to the underestimation of the bootstrapped trigger efficiency discussed above. This is not a stringent derivation, but is shown in [172] and in examples in the following to work reasonably well, in the sense that it can remove the bias without distorting the propagated uncertainties.

First evidence is given in the lower right plot in Figure 5.8, in which the bootstrapping is combined with a subsequent reweighting. It shows indeed that the efficiencies are consistent with the left plots, and an improvement over the upper right plot which does not have the reweighting. Moreover, the error bars appear to have the correct magnitude. The plot in Figure 5.9 shows the result of a study on toy Monte Carlo. It compares the distribution of the efficiency estimate obtained from bootstrapping, with and without the reweighting factor applied, to the true distribution of the efficiency and demonstrates that here the reweighted

[§] Note that the two plots to the left are the same, except that in the lower version some points are missing. Bootstrapping does not change the mode of the efficiency estimate of the `EF_j75_jetNoEF` here, because `EF_j30_jetNoEF` reaches its plateau below the onset of region of `EF_j75_jetNoEF`. This is typical for jet triggers which usually have comparably steep turn-on curves. The missing points stem from empty bins in the denominator of the sample trigger efficiency, where the efficiency of the sample trigger is thus undefined.

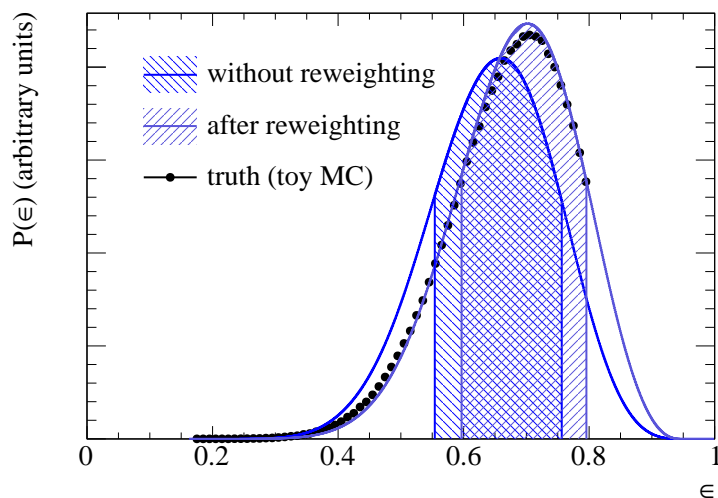


Figure 5.9: Example of the impact of the reweighting correction for bootstrapped efficiencies, in a toy Monte Carlo study for a specific choice of parameters (see text), comparing the bootstrapping results with and without reweighting (blue lines) against the distribution of the true efficiency (black dots, the error bars are smaller than the marker size). The shaded regions give the smallest intervals covering 68.3%, the curves are normalized to have the same area in the range $[0, 0.8]$.

distribution matches the true distribution much better. For this plot, the efficiency of the sample trigger was chosen to be $16/20 = 0.8$, and the efficiency of the target trigger chosen to be 0.7 so that the relative efficiency of the target trigger on the sample is $14/16 = 0.875$, the fraction giving the event counts which dictate the uncertainties. Note that Equation (5.16) limits the true efficiency to values below 0.8 . To create the histogram of the true efficiency, care has to be taken not to do the toy Monte Carlo the wrong way round: In bootstrapping, the efficiencies are estimated from the event counts in the numerator and denominator histograms for the sample trigger and for the target trigger. Thus, the toy Monte Carlo must yield the true distribution of the target trigger efficiency for a given observed number of the event counts, not for a given target trigger efficiency.

Systematic Studies and Improved Reweighting Factor

In the course of the trigger studies, with the increase of statistics available for the computation of the combined trigger efficiencies, it became clear that the performance of the reweighting approach is not optimal under certain circumstances. This affects in particular the computation of efficiencies of the combined jet + \cancel{E}_T trigger using a jet trigger as sample trigger, where a one-dimensional turn-on curve describing the jet trigger efficiency as function of jet p_T is employed for the bootstrapping. In that case, it may happen for a bin with high \cancel{E}_T in the plateau region that the event counts for the efficiency of the target trigger are much smaller than for the sample trigger at the same offline jet p_T . In particular, the denominator for the target trigger may be much smaller than the numerator for the sample trigger, which would otherwise be equal if no additional binning (here: in \cancel{E}_T) were done. This imbalance makes the probability distribution of the efficiency of the sample trigger very




Figure 5.10: Comparison of the estimates of the target trigger efficiency (eff. trig.) from bootstrapping when applying the new or old reweighting (new / old rew.), against the true target trigger efficiency (true eff.) in a toy Monte Carlo study. The plot also shows the sample trigger efficiency (eff. samp.) and the bootstrapped efficiency without reweighting (no rew.). The numbers in the labels on the horizontal axis are the event counts in the denominator for the efficiency measurement of the sample (n_1) and target trigger (n_2), the observed number of events in which the target trigger fires on the event sample taken with the sample trigger (k_{2r}) and the assumed sample trigger efficiency (ϵ). ϵ_X on the vertical axis stands for the different efficiencies explained in the legend.

narrow, whereas the probability distribution for the target trigger efficiency is broad, and in this edge case, the reweighting factor $-1/\ln(y)$ leads to an overestimation.

It turns out that an improvement of the reweighting factor can be achieved by restricting the integral in Equation (5.18) to the interval over which the numerical multiplication is done instead of the full interval $[0, 1] \times [0, 1]$ [178]. The numerical multiplication effectively uses only a subset of $[0, 1]$. Due to discretization effects, bins outside this subset do not contribute to the product in the sampling process, because their weight is so small that they will never be hit in the random sampling. Denoting these limits with $a_1 \leq x_1 \leq b_1$ and $a_2 \leq x_2 \leq b_2$,

Equation (5.18) turns into

$$f_Y(y) = \mathcal{N} \int_{a_2}^{b_2} \int_{a_1}^{b_1} f_{X_1}(x_1) f_{X_2}(x_2) \delta(x_1 \cdot x_2 - y) dx_2 dx_1 \quad (5.19)$$

$$= \mathcal{N} \int_{a_2}^{b_2} \frac{1}{x_2} [\Theta(b_1 x_2 - y) - \Theta(a_1 x_2 - y)] dx_2, \quad (5.20)$$

writing $\Theta(x)$ for the Heaviside step function. The expression in square brackets translates into the condition

$$\frac{y}{b_1} \leq x_2 \leq \frac{y}{a_1}, \quad (5.21)$$

so that, writing $I := [\min(b_2, y/a_1), \max(a_2, y/b_1)]$ for the domain of integration,

$$\Rightarrow f_Y(y) = \mathcal{N} \int_I \frac{1}{x_2} dx_2 \quad (5.22)$$

$$= \mathcal{N} \ln \left(\frac{\min(b_2, y/a_1)}{\max(a_2, y/b_1)} \right). \quad (5.23)$$

For $a_1 = a_2 = 0, b_1 = b_2 = 1$, this goes over into $f_Y(y) = -\ln y$ again.

Figure 5.10 is the result of a toy Monte Carlo study of the same type as described above. It shows a systematic comparison of the estimates of the target trigger efficiency from bootstrapping when applying the new or old reweighting, and the true target trigger efficiency. Each test is defined by the event counts in the denominator for the efficiency measurement of the sample (n_1) and target trigger (n_2), the observed number of events in which the target trigger fires on the event sample taken with the sample trigger (k_{2r}) and the assumed sample trigger efficiency (ϵ). The plot conveys three messages:

- It confirms the underestimation of the efficiency without the reweighting.
- It demonstrates the improvement achieved by reweighting, where the old reweighting overestimates the efficiency under the conditions described above.
- It shows the better agreement of the new reweighting with the true value compared to the old reweighting in most cases.

The histograms in Figure 5.11 have been created by repeating the toy Monte Carlo on a larger set of parameters. They show the distribution of the discrepancies between the true and the non-reweighted efficiencies, and between the true efficiencies and the efficiencies with the old and new reweighting, all scaled by the inverse of the symmetrized uncertainties. It demonstrates the three items listed above in less detailed representation. To quantify the differences between the three curves, a fit with a Gaussian function is done, which yields the fit parameters in Table 5.2. The mean x_0 moves closer to zero, which shows the improvement of new reweighting over the old reweighting. Note that the composition of the test parameters is arbitrary to a certain degree, so that the slight downward bias of the new reweighting may stem from the dominance of high efficiencies. Finally, Figure 5.12 compares the distribution of the error of the computed uncertainties on the efficiency estimates between the

Reweighting	Fit parameter		
	\mathcal{N}	x_0	σ
none	112(7)	-0.105(6)	0.116(4)
old	119(7)	0.036(6)	0.110(4)
new	116(7)	-0.021(6)	0.115(4)

Table 5.2: Fit parameters from a fit of the three distributions in Figure 5.11 with a Gaussian function $f(x) = \mathcal{N} \exp\left(-\frac{(x-x_0)^2}{2\sigma^2}\right)$.




Figure 5.11: Pull distribution of the estimates of the target trigger efficiency from bootstrapping when applying no, the old or the new reweighting. Figure 5.10 shows a subset of the data used to fill these histograms in detail.

three approaches. It shows that basically the reweighting does not change the distribution of the uncertainties, which is important because the reweighting would be pointless if it changed the uncertainties: bootstrapping without error propagation can be achieved without multiplying two distributions (see above). In conclusion, for the efficiency plots which rely on bootstrapping in 2011 and for which the edge cases causing a slight overestimation have been observed for the first time, the new reweighting is used by default.

Bootstrapping and Projections

The efficiencies of combined triggers like the jet + \cancel{E}_T triggers require at least a two-dimensional binning to account for the two different offline quantities which are involved. To study the turn-on behavior of one of the two components of the trigger in detail, a projection onto one of the two axes is convenient. This is done in many examples in Sections 5.3.1 and Section 5.3.2. If a one-dimensional turn-on of the sample trigger is used, which is the case




Figure 5.12: Absolute error on the uncertainty intervals of the efficiency estimate. The plots show the differences of the lower (left plot) and upper (right plot) uncertainty to the true uncertainty from the toy MC experiments, for the new and the old reweighting method and without reweighting.

throughout all of the studies in these sections^h. The projection however can only be combined with bootstrapping if the projection axis coincides with the axis in which the binning of the sample trigger efficiency is done: In this case, it is possible to add up all entries from the numerator and denominator bins above a given threshold on the orthogonal axis, and multiply the resulting probability distribution of the efficiency with the known probability distribution of the sample trigger efficiency for this bin. In the opposite case, when the axis in which the sample trigger efficiency is binned and the projection axis are orthogonal, it is not possible to perform bootstrapping. Only a weighted projection can be done, which by multiplying with the efficiency of the sample trigger in the respective bin yields the correct efficiency estimate, but cannot propagate the errors correctly. This corresponds to the assumption that the error on the sample trigger efficiency can be neglected, which in general is a valid assumption due to the cut being applied, which, if it brings the trigger into its plateau, also ensures that the sample trigger has reached the plateau of its efficiency.

5.2.8 Computation of Trigger Efficiencies on Weighted Samples

Often when generating Monte Carlo pseudodata, a physics process is split into several subsamples to achieve a better coverage of the phase space. This splitting can for example be done in the number of additional partons or the transferred momentum in the hard scattering process. To obtain physically meaningful results, the events from all subsamples have to be added up, using the corresponding cross sections as weights. The same needs to be done when computing trigger efficiencies, so the number of events passing a given trigger k and the total number of events n need to be scaled properly. For the expectation value, the weighting is simple, the estimated efficiency being the ratio of the weighted number of events passing the trigger and the weighed total number of events (cf. Equation (5.24) below).

The uncertainty computation is more involved, as the error bars need to be computed from the β -distribution representing the probability distribution of the efficiencies, which is

^h The one exception being Figure 5.23.

not known *a priori*. For a sample that has ten times as many events as another sample but only a tenth of the weight, even if the values for k and n are the same, the uncertainty on the efficiency estimate should be smaller due to the higher statistics, but by just scaling k and n this is not accounted for. Instead, one solution is to use the method of moments to obtain an approximation of the probability distribution of the efficiency ϵ in terms of a β -distribution based on the values of the expectation value $E(\epsilon)$ and the variance $V(\epsilon)$ [34], from which the asymmetric uncertainties can be determined, as will be shown in the followingⁱ.

Fixing the index i of the binning in the offline variable and using the index j to run over the samples which are averaged, the best estimate of the efficiency for the complete sample is given by the expectation value

$$E(\epsilon_i) = \frac{\sum_j \omega_j k_i^j}{\sum_j \omega_j n_i^j}, \quad (5.24)$$

where ω_j is the weight of the sample j . The variance scales with the square of the weights,

$$V(\epsilon_i) = \frac{\sum_j \omega_j^2 V(\epsilon_i^j)}{(\sum_j \omega_j)^2}. \quad (5.25)$$

In the following, the parameter ϵ_i will be left out to simplify the notation, i. e. $E \equiv E(\epsilon_i)$ and $V \equiv V(\epsilon_i)$. In the same way, all other quantities are understood to refer to a given bin i . The estimates of E and V obtained from $\{k_i^j\}$ and $\{n_i^j\}$ in Equations (5.24) and (5.25) can be used to find a β -distribution with this mean and variance. Solving Equations (5.10) and (5.11) for a and b yields the method-of-moments estimates of the parameters [197]

$$\begin{aligned} a &= E \left(\frac{E(1-E)}{V} - 1 \right) \quad \text{and} \\ b &= (1-E) \left(\frac{E(1-E)}{V} - 1 \right). \end{aligned} \quad (5.26)$$

There is a problem though. From Equations (5.26), it can be seen that this does not work when $E = 0$ or $E = 1$ because then a and b take values (negative or zero) for which the β -distribution is not defined. This problem stems from the fact that there is no β -distribution with mean 0 or 1. Therefore, a similar approach is used here, but instead of the mean E the mode M is used. Solving Equations (5.11) and (5.12) for a and b yields the slightly more complicated expression

$$b = (a + 2) + \frac{a - 1}{M}, \quad (5.27)$$

where a is given by the solutions of a cubic equation,

$$\begin{aligned} & a^3 + \\ & \left(M^3/V - M^2/V + 7M - 3 \right) a^2 + \\ & \left(-2M^3/V + 16M^2 + M^2/V - 14M + 3 \right) a^1 + \\ & \left(12M^3 - 16M^2 + 7M - 1 \right) = 0. \end{aligned} \quad (5.28)$$

ⁱ Note that the expectation value $E(\epsilon)$ and variance $V(\epsilon)$ of course depend on the probability distribution $P(\epsilon)$ rather than on ϵ , although this is not made explicit in the notation.

In the implementation, the solutions are found numerically using `gsl_poly_solve_cubic` from the GNU SCIENTIFIC LIBRARY. The case $E = 1/2$, in which the coefficients of the monomials a^1 and a^0 vanish, is treated specially. If a is found to be 1, it follows $M = 0$, and the result of Equation (5.27) is undefined. In that case, Equation (5.11) simplifies to

$$V = \frac{b}{(b^2 + 2b + 1)(b + 2)}, \quad (5.29)$$

yielding

$$\Rightarrow b^3 + 4b^2 + (5 - 1/V)b + 2 = 0, \quad (5.30)$$

which is solved to find b using `gsl_poly_solve_cubic` again. A closure test has been performed to ensure that the solutions for a and b using the largest root of the cubic equations yield a β -distribution with the desired values of M and V .

5.3 Measurements of Trigger Efficiencies

The methods which have been introduced in Section 5.2.6 will now be applied to measure trigger efficiencies both in Monte Carlo and data. The ultimate goal of the trigger efficiency studies is the measurement of the efficiencies of the combined jet + \cancel{E}_T triggers, which are used by the ATLAS Supersymmetry group as the primary trigger for the zero-lepton analysis and in the search for Supersymmetry which is presented in this thesis in Section 7. Before the measurements of trigger efficiencies on data are presented, on which these analyses are based, the results of complementary studies of the efficiencies of jet + \cancel{E}_T triggers on Monte Carlo samples are discussed in the next subsection. The determination of the efficiencies on data follows afterwards in Section 5.3.2.

Which method is employed to compute the efficiencies is specified for each turn-on curve individually. The efficiencies of triggers involving jets are in general determined object-wise, i. e. a matching between trigger jets and jets which were reconstructed offline is done based on the geometrical distance in (η, ϕ) space with the usual metric $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$. For this matching, the trigger jet at the highest active trigger level, Level 2 in 2010 or Event Filter in 2011, is used. The p_T of all jets from the respective offline selection is filled into the denominator histograms, the numerator is filled using only the p_T of offline jets which have been matched to a trigger jet that has actually fired the trigger. There may be more than one trigger jet in the same event firing a given trigger chain. For combined jet + \cancel{E}_T chains, an additional condition for filling the numerator is, of course, that the amount of \cancel{E}_T computed by the trigger at all levels exceeds the respective thresholds.

5.3.1 Measurements of Jet + \cancel{E}_T Trigger Efficiencies in Monte Carlo

The efficiencies of the combined jet + \cancel{E}_T triggers are evaluated on two different Monte Carlo samples which differ strongly in the composition of \cancel{E}_T . The first sample is the PYTHIA dijet sample used later in the analysis to model the QCD background, and is intended to obtain an idea of the performance of the jet + \cancel{E}_T trigger on the majority of events the trigger will run on, because the QCD background is the dominant background at the LHC. (A measurement on minimum-bias Monte Carlo is not possible because of insufficient statistics. On the other hand it is safe to assume that minimum-bias events indeed contain only very soft jets and low \cancel{E}_T and will therefore always be rejected by the combined jet + \cancel{E}_T triggers. This is in




Figure 5.13: Spectrum of \cancel{E}_T coming from non-interacting particles in PYTHIA QCD (left) and ALPGEN $Z \rightarrow \nu\nu + \text{jets}$ (right). The error bars indicate the statistical uncertainties due to limited Monte Carlo statistics. The number of events is normalized to an integrated luminosity of 1 pb^{-1} .

contrast to the dijet samples which contain events with hard jets which have a non-negligible cross section.) Whilst the QCD sample is expected to contain jets with high p_T , the main contribution to \cancel{E}_T comes from jet mismeasurements and does not include real \cancel{E}_T , which in QCD processes could only arise from neutrinos produced within jets.

The second sample contains $Z \rightarrow \nu\nu$ decays and jets from additional partons produced in the hard scattering process. If the neutrinos are produced back-to-back in the laboratory frame, this again does not give \cancel{E}_T , because their transverse momenta add up to zero. However, additional jets recoiling against the neutrino system can give rise to significant \cancel{E}_T , which can be seen in the right plot in Figure 5.13. This figure shows the spectrum of \cancel{E}_T coming from non-interacting particles for both samples, QCD and $Z \rightarrow \nu\nu$, where the subsamples have been weighted according to their cross section as given in Section 7.5, and the \cancel{E}_T spectrum is plotted from Monte Carlo Truth information. As expected from the above explanations, the real \cancel{E}_T in the vast majority of the QCD events is negligible, whereas most events in the $Z \rightarrow \nu\nu + \text{jets}$ sample contain real \cancel{E}_T of around 40 GeV.

As outlined in the previous section on methods to compute trigger efficiencies, in Monte Carlo the trigger efficiencies can be determined using simple event counting, i. e. relying on the results of the simulation of the trigger. Both of the Monte Carlo samples from above are split into several subsamples with different production cross sections and therefore different weights. The subsamples of the $Z \rightarrow \nu\nu$ sample differ in the number of additional partons, while the QCD sample is binned in the \hat{p}_T exchanged in the hard scattering (cf. Sections 7.5.2 and A.8), the subsamples being denoted by J0 through J6[‡]. This means that the results from Monte Carlo counting on each individual sample need to be combined using the method presented in Section 5.2.8 in order to obtain a physically meaningful efficiency for the full sample. This combination is shown in the two plots in Figure 5.14, where the trigger efficiency is presented using a two-dimensional, color-coded histogram, with the offline p_T of the jets at EM+JES scale (cf. Section 4.4.2) on the horizontal and the offline \cancel{E}_T on the vertical axis. Note that in this thesis, offline \cancel{E}_T variants reconstructed by different

[‡] J7 and J8 are not used, because J0 through J6 suffice to fill the phase space region covered by the plots, and J7 and J8 are not expected to contribute to the combination due to their low cross section.

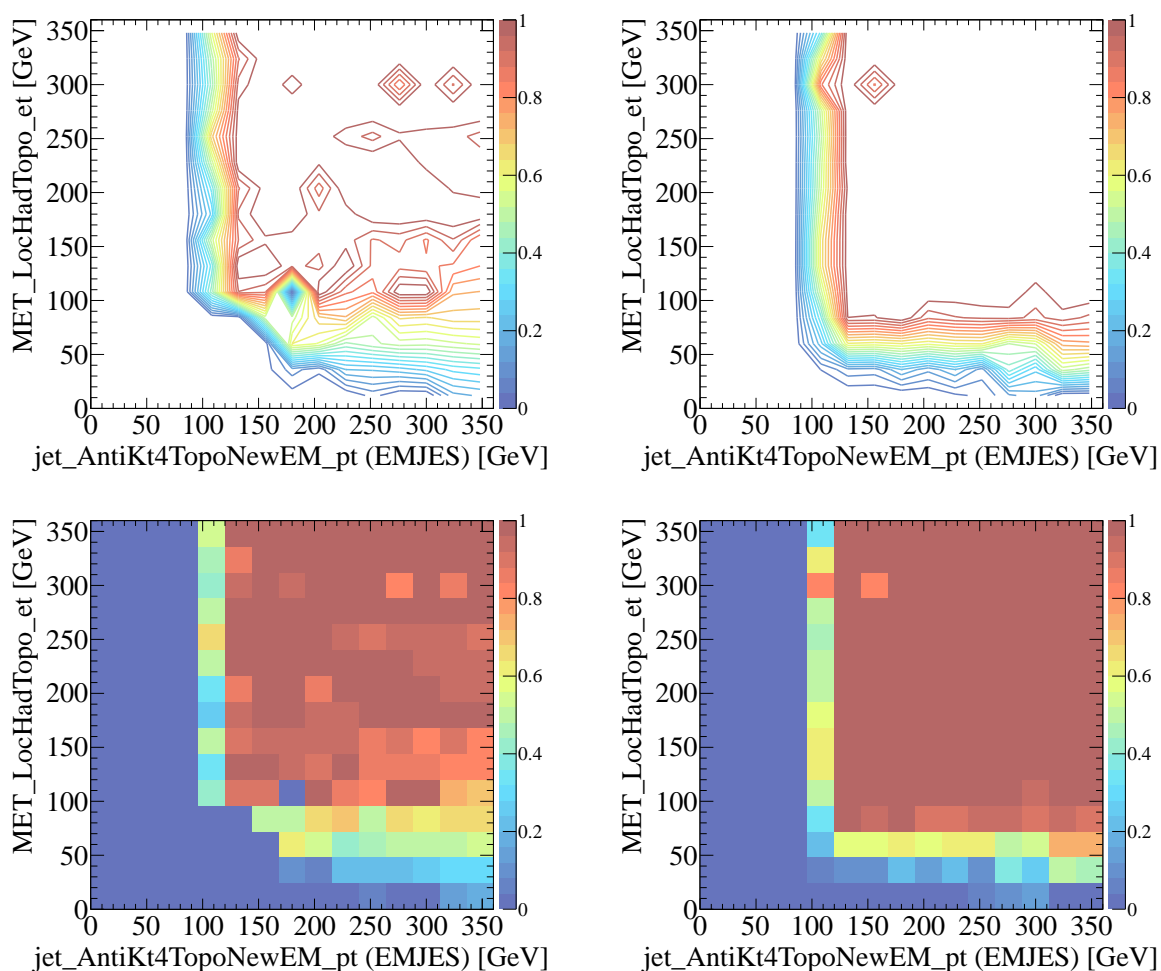


Figure 5.14: Two-dimensional efficiency plots of the combined jet + \cancel{E}_T trigger EF_j75_a4tc_EFFS_xe45_loose_noMu in Monte Carlo events, for a QCD sample (left) and for a $Z \rightarrow \nu\nu + \text{jets}$ sample (right). The two upper plots show the efficiencies as contours, which highlights details in the regions where the efficiencies change, in particular in the turn-on regions. The two lower plots use a color-coding of the efficiencies, thereby making it easier to distinguish regions with low and high efficiencies.

algorithms are used. The flavor of the \cancel{E}_T reconstruction can be read off from the axis title and is explained in Section 4.4.3. In general, the standard for the offline \cancel{E}_T reconstruction used in the computation of trigger efficiencies will be the \cancel{E}_T from topological cluster cells calibrated using weights from the local hadronic calibration (LocHadTopo) as used in Figure 5.14. This \cancel{E}_T comes very close to how the \cancel{E}_T calculation is done by the trigger system, and therefore gives the best offline reference. The analog of the plots which are shown in this section with the offline \cancel{E}_T variant used for the SUSY zero-lepton analysis instead can be found in the Appendix in Section A.7. The binning is chosen such as to give quadratic bins with sizes that are a compromise between smaller statistical fluctuations for larger bins and better sensitivity to features of the turn-on region for smaller bins.

The two plots to the left in Figure 5.14 show the trigger efficiency of the combined jet + \cancel{E}_T trigger with the lowest thresholds running in 2011, EF_j75_a4tc_EFFS_xe45_loose_noMu, for the QCD sample, the two plots to the right the efficiency for the same trigger, but for $Z \rightarrow \nu\nu$. This trigger applies a threshold of 75 GeV for the jet part and 45 GeV for the \cancel{E}_T part. The plots with $Z \rightarrow \nu\nu$ events and real \cancel{E}_T exhibit much sharper turn-on curves, in particular as function of \cancel{E}_T , and less fluctuations in the plateau, where the efficiency is almost everywhere close to one. In these plots, it can therefore be seen best that the \cancel{E}_T turn-on of the jet + \cancel{E}_T trigger actually lies around the nominal value of 45 GeV, whereas the jet turn-on only begins at values around 100 GeV, much higher than the nominal 75 GeV. This is due to the fact that the offline and trigger jet calibration is different. The trigger jets are calibrated at EM scale, whereas for the offline jets, the EM+JES calibrated p_T is used, so that the same jet will be reconstructed with higher p_T offline than within the trigger. This shifts the turn-on to higher values, which can also be seen in the projections in Figure 5.16 discussed below^k. In the lower left plot for the QCD sample, the blue region at small \cancel{E}_T and jet p_T values cuts into the lower left corner of the plateau region. This is due to the dominance of the J0 and J1 subsample in this region, and is revealed by the two-dimensional efficiency plots per subsample. For the J3 and higher subsamples, this region looks more like in the right plot for the $Z \rightarrow \nu\nu$ sample. This matches the observation that in the high \hat{p}_T subsamples of the QCD sample, the contribution of real \cancel{E}_T to the total \cancel{E}_T should be comparable to the real \cancel{E}_T in $Z \rightarrow \nu\nu$.

As in the two-dimensional plots it is difficult to visualize error bars, often projections of the two-dimensional histograms onto one of their axes will be employed. These projections are created from all bins above a minimum value on the respective orthogonal axis, i. e. when projecting onto the \cancel{E}_T axis, all bins above a given jet p_T threshold are used and vice versa. The threshold is chosen such that the efficiency of the jet part of the combined trigger has reached its plateau value. Otherwise, the projection would yield an average over the full range of jet p_T values and, regarded as function of \cancel{E}_T , the efficiency would not reach an efficiency of one in the plateau, because of the bins with an efficiency of zero for low jet p_T . As can be seen in particular in the plots to the left in Figure 5.14, tracing the turn-on

^k Note that there is no general problem in using offline jets at EM scale instead. Indeed, the p_T at EM scale would be a more natural variable to plot the trigger efficiency as function of, but as most of the jet + \cancel{E}_T trigger studies have been done for the official zero-lepton Supersymmetry analysis, here it has been conceded to use the same offline variable as in the analysis. As a different calibration of the jets does not induce such a fundamental difference in the offline variable with respect to the trigger variable as it is seen when using a different offline \cancel{E}_T , the conclusions stay the same apart from the shift in the turn-on to higher values. The last fact makes it possible to directly read off from the turn-on curves the threshold for the offline jet p_T cut which is needed to only select events in the plateau region of the trigger.




Figure 5.15: Projections of the efficiency of the combined jet + \cancel{E}_T trigger EF_j75_a4tc_EFFS_xe45_loose_noMu in Monte Carlo events for the combined QCD sample and for the combined $Z \rightarrow \nu\nu$ + jets sample. Left: onto offline \cancel{E}_T , right: onto jet p_T , after cuts on the respective orthogonal variable at 130 GeV.

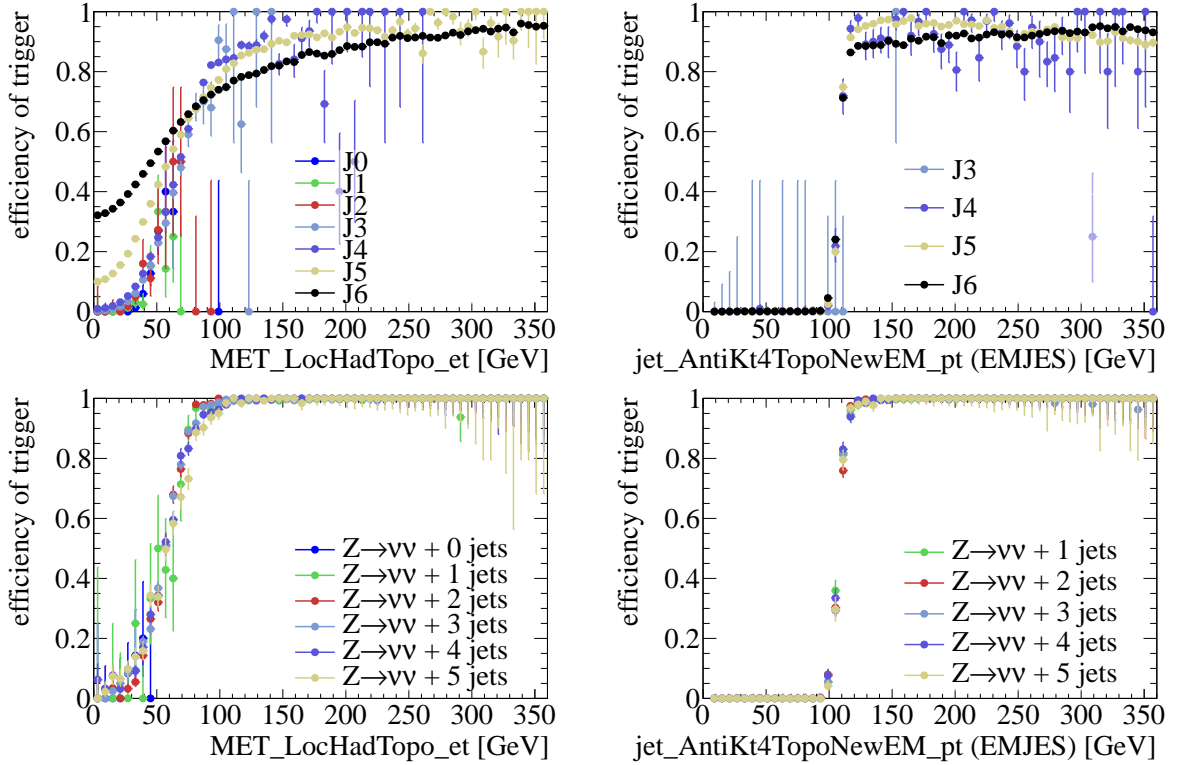


Figure 5.16: Projections of the efficiency of the combined jet + \cancel{E}_T trigger EF_j75_a4tc_EFFS_xe45_loose_noMu in Monte Carlo events for each subsample of the QCD sample (top) and the $Z \rightarrow \nu\nu$ + jets sample (bottom) individually. Left: onto offline \cancel{E}_T , right: onto jet p_T , after cuts on the respective orthogonal variable at 130 GeV.

region around \cancel{E}_T values of 100 GeV along the horizontal axis, the behavior of the turn-on may change as function of the orthogonal variable. Here, as function of jet p_T , the turn-on region of the \cancel{E}_T part of the combined trigger gets broader. Features like this of the two-dimensional structure of the trigger efficiency are naturally lost in projections, so that it is always advisable to consider the two-dimensional plots, too. Figure 5.15 displays the projections onto the \cancel{E}_T and jet p_T axes for both samples, created using the combination of all subsamples. The efficiencies per subsample are shown in Figure 5.16, which demonstrates the large differences between the QCD subsamples, affecting in particular the \cancel{E}_T part, but much less so the jet part. The $Z \rightarrow \nu\nu$ subsamples all behave alike with respect to the trigger turn-on. The projections also show that the width of the turn-on region is very different, the turn-on of the jet part being much sharper than the one of \cancel{E}_T . This is a typical difference between jet and \cancel{E}_T triggers. Note that some of the lower subsamples do not appear in the right plots in Figure 5.16 because no events remain after the cut on \cancel{E}_T .

The results of the studies on Monte Carlo in this section demonstrate a strong dependence of the outcome of the efficiency measurement on the event sample which is used as input. They thereby underline the importance of finding a way to measure the efficiencies in data on an unbiased, representative sample of events. It may be tempting to say that in an analysis which is looking for Supersymmetry, the selection is designed to accept only events which have real \cancel{E}_T , but that is not what is actually happening in the evaluation of the event selection, where event counts from Monte Carlo and data are compared without taking into account e. g. the origin of the \cancel{E}_T which makes an event pass the selection cuts.

5.3.2 Measurements of Trigger Efficiencies in Data

As was mentioned in the introduction, the goal of the efficiency studies in this thesis is the measurement of the efficiencies of the combined jet + \cancel{E}_T triggers, which are employed by the ATLAS Supersymmetry group in the zero-lepton analysis as well as in the analysis which is presented in Section 7. The trigger menus have undergone a fast-paced evolution in particular in 2010, following the requirements of the increasing instantaneous luminosity. In addition, they are different between 2010 and 2011 due to changes in the jet trigger algorithms. Altogether, there have been a number of changes in the ATLAS trigger between 2010 and 2011 which necessitate changes in the measurement of the efficiencies with respect to which triggers are of interest and which methods can be employed. Therefore, the measurements of the trigger efficiencies are presented and discussed separately for 2010 and 2011. The first part of this section deals with the trigger conditions and efficiency measurements in 2010, which were coined by the rapid increase in the instantaneous luminosity, entailing constant changes in the structure of the trigger menu, and the necessity to test which methods are suitable and feasible to determine the trigger efficiencies. This part is also documented in an internal ATLAS note [133]. In the second part, the changes with respect to 2010 and efficiencies of newly introduced triggers are discussed, the steadier beam conditions allowing for consistency checks between different periods. The studies presented in this section have become part of the supplementary documentation of the official zero-lepton analyses performed by the ATLAS Supersymmetry group in 2010 and 2011 [146, 148, 149].

Name of chain	Cut on jet E_T [GeV]	Cut on \cancel{E}_T [GeV]
EF_j50_jetNoEF_xe20, 30, 40_noMu	L1: 30	L1: 10, 15, 25
	L2: 45	L2: 12, 20, 30
	EF: —	EF: 20, 30, 40
EF_j75_jetNoEF_EFxe20, 25, ..., 40_noMu	L1: 55	L1: —
	L2: 70	L2: —
	EF: —	EF: 20, 25, ..., 40

Table 5.3: Jet + \cancel{E}_T chains in the trigger menu `Physics_pp_v1`, which was used at the end of pp data taking in 2010. The table shows the cuts which are applied on the online measurement of \cancel{E}_T and jet E_T at Level 1, Level 2 and Event Filter for two different groups of jet + \cancel{E}_T triggers.

5.4 Measurements of Trigger Efficiencies in Data Taken in 2010

Before coming to the actual measurements, a brief summary of the jet + \cancel{E}_T triggers used by ATLAS in 2010 will be presented. Two different groups of jet + \cancel{E}_T trigger chains have been defined in the trigger menus in 2010, which not only differ in the value of their thresholds at the three trigger levels, but also by which trigger objects cuts are applied to. As at Event Filter level event rejection based on jet cuts has not been activated in 2010, all jet triggers only cut at Level 1 and Level 2. The EF hypothesis algorithms are set to forced accept. This also holds for the combined jet + \cancel{E}_T triggers. Unless otherwise stated, all of the rest of this section refers to 2010.

5.4.1 Combined Jet + \cancel{E}_T Triggers in the Trigger Menu

Table 5.3 gives an overview of the two groups of jet + \cancel{E}_T triggers included in the trigger menu of 2010 in terms of the values of the thresholds applied to the transverse energy of the jets and the missing transverse energy at each trigger level. The first group in the upper block of the table contains the “traditional” jet + \cancel{E}_T trigger chains, seeded by trigger items at L1 which combine cuts both on jets and \cancel{E}_T . There are three chains in this group, which comprise the same set of jet cuts, but differ by their thresholds on \cancel{E}_T , which are 20, 30 and 40 GeV at EF level and correspondingly lower cuts on \cancel{E}_T at L1 and L2. The other group of triggers is seeded by the L1 item `L1_J55`, i. e. L1 jets with $E_T \geq 55$ GeV, includes a 70 GeV cut on jets at L2, but only applies cuts on \cancel{E}_T at EF. In this group, there are five triggers, which again only differ in their cuts on \cancel{E}_T , whereas their cuts on jets are identical. The motivation for having chains which are seeded by jets at L1 and cut at \cancel{E}_T only at EF is to profit from the higher resolution of the \cancel{E}_T measurement at EF owing to to the higher latency affordable at this trigger level. `L1_J55` could be kept unrescaled throughout the whole data-taking period of 2010, whereas the chains seeded by the lowest jet + \cancel{E}_T trigger item at L1 with a L1 \cancel{E}_T cut of 10 GeV had to be deactivated towards the end of the data-taking period. Note that the (first) number in the name of the jet + \cancel{E}_T chains refers to the EF jet threshold (50 GeV or 75 GeV, respectively), which for these chains is in general 5 GeV higher than the L2 cut. As the names of the jet + \cancel{E}_T chains are quite long and easy to confuse, in the following the

Target trigger	Potential sample triggers	
	Jet type	\cancel{E}_T type
EF_j50_jetNoEF_xe20,30,40_noMu	L1_J30	L1_XE10,15,25
	L2_j45	L2_xe12,20,30_noMu
	EF_j50_jetNoEF	EF_xe20,30,40_noMu
EF_j75_jetNoEF_EFxe20,25,...,40_noMu	L1_J55	—
	L2_j70	—
	EF_j75_jetNoEF	(EF_xe15_unbiased_noMu)

Table 5.4: Overview of the sample triggers that can be used for bootstrapping jet + \cancel{E}_T triggers in 2010 taking Equation (5.16) into account.

first group of triggers in Table 5.3 will be referred to as “full-chain jet + \cancel{E}_T triggers” and the second group as “jet-seeded jet + \cancel{E}_T triggers”.

5.4.2 Sample Triggers for Measuring Jet + \cancel{E}_T Trigger Efficiencies

For jet + \cancel{E}_T triggers, two types of triggers fulfill Equation (5.16), provided the thresholds are low enough: single jet triggers and \cancel{E}_T triggers. This allows to bootstrap jet + \cancel{E}_T triggers on two different samples, one taken with a jet trigger, the other with an \cancel{E}_T trigger, and to cross-check the results for consistency. Taking Equation (5.16) as guideline, the set of sample triggers given in Table 5.4 is identified, which can be used to measure the efficiencies of the combined jet + \cancel{E}_T triggers. The table shows for both groups of triggers, jet-seeded and full-chain, all possible jet and \cancel{E}_T triggers which can be used as sample trigger, together with their thresholds at Level 1 and Level 2 in terms of the items and chain names at these levels. Consistency checks for the jet-seeded triggers using \cancel{E}_T sample triggers in addition to jet sample triggers are not possible, because the jet-seeded triggers only cut on \cancel{E}_T at EF level, and therefore a sample would be needed taken with an EF-only \cancel{E}_T trigger. There is only one such trigger, EF_xe15_unbiased_noMu, in the menu, which cannot provide sufficient statistics.

All event samples for the measurements of trigger efficiencies are selected by requiring a positive EF decision of a single trigger chain. In the following, it is explained how the selection of this sample trigger is optimized to obtain the maximal possible sample sizes. The fastest way to find out which trigger provides the largest statistics is to read this information from the ATLAS COOL database, where the number of events per trigger, run and luminosity block are stored. This allows to compare the absolute number of events after prescales in luminosity blocks marked as good in the corresponding GRL which were triggered by e. g. the various \cancel{E}_T triggers with different thresholds, and to select the trigger which has the highest count of accepted events. Note that COOL only holds the trigger counts for periods G and later. As the prescales evolve over time, the trigger chains with highest counts vary between the different periods. The trigger with the highest count summed over all three periods is chosen as sample trigger.

Table 5.5 shows the resulting optimal choices of the sample triggers for the respective target triggers and the size of the event samples. For EF_j50_jetNoEF_xe30_noMu, both

Target trigger	EF sample triggers with highest count	
	Jet type	\cancel{E}_T type
EF_j50_jetNoEF_xe20_noMu	EF_j50_jetNoEF (616 911)	EF_xe20_noMu (817 636)
EF_j50_jetNoEF_xe30_noMu	EF_j50_jetNoEF (616 911)	EF_xe20_noMu (817 636) EF_xe30_noMu (892 808)
EF_j50_jetNoEF_xe40_noMu	EF_j50_jetNoEF (616 911)	EF_xe30_tight_noMu (1 309 130)
All jet-seeded triggers	EF_j75_jetNoEF (2 266 903)	—

Table 5.5: Overview of the triggers with highest counts which can be used as sample triggers for the respective target triggers in periods G – I. The number of events selected by each trigger is given in brackets to give a feeling for the number of events needed to generate turn-on curves such as shown in this section.

EF_xe20_noMu and EF_xe30_noMu yield comparable sample sizes. In this case, using the trigger with the higher threshold gives better results, because it has a better coverage of the plateau region of the target trigger for high \cancel{E}_T values. Note that EF_xe30_tight_noMu cannot be used for EF_j50_jetNoEF_xe30_noMu due to its higher thresholds at L1 and L2, which would violate Equation (5.16). The cross-checks on the Muons stream require a muon trigger as sample trigger. EF_mu13 has the highest trigger count of all single muon triggers.

It turns out that the sample trigger with the highest event count is always the one with the highest threshold that is allowed by condition Equation (5.16). The conclusion from this is that probably not much could be gained by including triggers with lower thresholds in a logical OR with the optimal sample trigger because events are scarce especially in regions with high \cancel{E}_T or jet p_T , and including chains with lower thresholds will not help there. This, of course, only holds if the chosen sample triggers with the highest possible thresholds are unprescaled.

5.4.3 Data Selection

Unless stated otherwise, data from periods 2010 G, H, and I will be used in the following, as the trigger menu has been relatively stable in these periods. Moreover, they comprise the latest and largest part (about 92 %) of the data collected in 2010 in terms of luminosity (about 41 pb^{-1} of recorded luminosity [185]). All of this data has been reconstructed with ATHENA in release 15. Mostly, the NTUP_SUSY data format has been used, which is produced from all relevant data streams apart from the MinBias stream. The event selection relies on basic cuts that are common to most physics analyses in order to keep the results general and to discard events which suffer from bad detector conditions, while retaining as many events as possible to have an event sample as large as possible. The following event-level cuts have been applied: Events must belong to luminosity blocks that have been marked as suitable for physics analysis in the GRL generated for the ATLAS SUSY group. Events are rejected if they contain a bad jet arising from hardware problems, cosmic-ray showers, or general problems with the beam quality, as defined by the ATLAS Jet/Etmiss combined performance group in their updated recommendation for loose cleaning of release 15 data




Figure 5.17: Efficiency curves for a set of jet triggers, some of which will be used as sample triggers for the computation of jet + \cancel{E}_T efficiencies. Left: Direct measurement on MinBias stream, periods F – I, using EF_mbMlbs_1_eff as sample trigger. Right: Bootstrapping on JetTauEtmiss stream with EF_j30_jetNoEF as sample trigger, periods G – I.

after the PISA workshop in 2010 [155, 189] (cf. Section 7.3.1). A selection based on the number of reconstructed primary vertices, as is made in physics analyses to reject events with no collisions, is not applied.

As offline reference for jet triggers the jet collection used by the ATLAS Supersymmetry group is chosen, which contains offline jets reconstructed using an anti- k_t jet finding algorithm [32] with a cone size of 0.4 based on topological clusters (AntiKt4H1Topo, cf. also Sections 4.4.2 and 7.3.2). The energy is calibrated by applying a pseudorapidity and p_T dependent jet-energy-scale factor to the energy measurement at the electromagnetic scale, denoted by EM+JES in the labels of the axes. For the turn-on curves, only offline jets within $|\eta| < 2.8$ are used, and no explicit minimum p_T requirement is made.

Two different variants of offline \cancel{E}_T are considered as offline reference: MET_Topo as primary reference, and MET_EMJES_RefFinal_CellOutEM as alternative (cf. Section 4.4.3), which is the definition of \cancel{E}_T the ATLAS Supersymmetry group has agreed on to be used in publications analysing 2010 data, and which will be used to facilitate conclusions from the analyst’s point of view. To avoid the long name, it will be referred to as the “ \cancel{E}_T definition of the SUSY group” or “SUSY \cancel{E}_T ” for short.

5.4.4 Jet Triggers

Bootstrapping the jet + \cancel{E}_T turn-on curves from events collected with jet sample triggers requires to compute the jet-trigger efficiencies first. There are several possibilities to measure the jet-trigger turn-on curves:

- A direct measurement on events from the MinBias stream taken with a minimum-bias trigger.
- Bootstrapping, using a jet trigger with some even lower threshold as sample trigger on the JetTauEtmiss stream.
- Using a muon trigger as an orthogonal trigger on the Muons stream.

Note that the muon-triggered events do not yield a completely unbiased sample because requiring a muon enhances the admixture of jets from heavy flavor decays. Figure 5.17 shows the turn-on curves of the jet triggers $EF_j30_jetNoEF$, $EF_j50_jetNoEF$ and $EF_j75_jetNoEF$. The left plot shows a direct measurement on the `MinBias` stream using $EF_mbMbts_1_eff$ as sample trigger, with data from periods F through I. The turn-on curves from the four different periods agree among each other within the statistical uncertainties, i.e. the trigger is stable enough to allow a combination of these periods. It can be seen in Figure 5.17 that the efficiency of $EF_j50_jetNoEF$ has large statistical uncertainties when measured on the `MinBias` stream, and the uncertainties for $EF_j75_jetNoEF$ are even larger, of course. As these triggers are to be used as sample triggers for bootstrapping the jet + \cancel{E}_T trigger efficiencies, the efficiencies of these two jet triggers need to be determined with higher statistics on the `JetTauEtmiss` stream with the bootstrap method, otherwise the very large uncertainties from the measurement on the `MinBias` stream would deteriorate the turn-on curves obtained for the combined jet + \cancel{E}_T triggers. The resulting turn-on curves for $EF_j50_jetNoEF$ and $EF_j75_jetNoEF$ are shown in the right plot in Figure 5.17, using bootstrapping and $EF_j30_jetNoEF$ as sample trigger on the `JetTauEtmiss` stream. The data is from periods G – I. It has been checked that the efficiencies in these periods are consistent so that it is justifiable to combine them. Note that comparing the left and right plot, within the uncertainties, the efficiencies from the `MinBias` and the `JetTauEtmiss` stream agree. This is an important cross-check because the potential non-execution of the L2 jet reconstruction algorithms in events triggered by minimum-bias triggers might bias the jet trigger efficiencies measured on the `MinBias` stream, a problem that will be explained in more detail at the end of Section 5.4.6.

All turn-on curves shown in the plots above and the following plots for jet and combined jet + \cancel{E}_T triggers are integrated over pseudorapidity η and azimuthal angle ϕ . It should be mentioned that there is a drop in efficiency of the jet triggers as a function of η , limited to the transition region around $|\eta| = 1.5$ between the barrel and the endcap calorimeters [163]. The efficiency is lower because in this region only a part of the calorimeter is used in the L1 trigger in 2010. The drop affects the jet + \cancel{E}_T triggers in the same manner. In order to collect sufficient statistics for the jet + \cancel{E}_T trigger turn-on curves, which require additional binning in \cancel{E}_T , it was decided to relinquish binning in η here.

5.4.5 Missing Transverse Energy Triggers

It has been explained above that for bootstrapping jet + \cancel{E}_T triggers, both parts of the combined trigger, jets and \cancel{E}_T , can be used as sample triggers. Actually, this is only true for the full-chain triggers because there is no \cancel{E}_T trigger cutting only at EF (cf. Section 5.4.2). For the employment of \cancel{E}_T triggers as sample trigger however, the problem lies in the measurement of the turn-on curves for the \cancel{E}_T triggers themselves. The main difficulty is to define a good sample of events to measure the \cancel{E}_T trigger efficiency on. Using an unbiased sample from the `MinBias` stream is not possible because it runs out of statistics before reaching the plateau of even the lowest useful \cancel{E}_T trigger which is EF_xe20_noMu . Bootstrapping EF_xe20_noMu from a lower \cancel{E}_T trigger is not feasible, because there is only one lower trigger, EF_xe15_noMu , which is so close that it suffers from the problem of insufficient statistics, too. Hence, only orthogonal triggers remain as a viable alternative for the sample selection. Here, muon triggers will be used, which can be considered as orthogonal triggers with respect to the \cancel{E}_T and jet + \cancel{E}_T triggers because both jets and \cancel{E}_T are quantities which in the trigger are calculated from calorimeter measurements only (cf. Section 4.3.2).




Figure 5.18: Comparison of the turn-on curves for the \cancel{E}_T triggers EF_xe20_noMu and EF_xe30_noMu, using the minimum-bias trigger EF_mbMbts_1_eff (MinBias) and the single muon trigger EF_mu13 (Muons) as sample triggers, respectively.

Note that turn-on curves for \cancel{E}_T triggers can also be measured on event samples with decays of W bosons to electrons and muons [111, 179]. Alas, these are not unbiased in the general sense discussed above.

In the following, muon triggers will be used as orthogonal triggers with respect to \cancel{E}_T . This allows to determine the \cancel{E}_T turn-on, and use \cancel{E}_T triggers for bootstrapping jet + \cancel{E}_T triggers. Later on, in Section 5.4.6, muon triggers will again prove useful as an overall cross-check of the jet + \cancel{E}_T efficiencies from bootstrapping. Figure 5.18 shows the turn-on curves of two \cancel{E}_T triggers with thresholds of 20 and 30 GeV to be used later on for bootstrapping of jet + \cancel{E}_T trigger efficiencies. It compares the efficiency estimates on events from the MinBias stream, using the minimum-bias trigger EF_mbMbts_1_eff as sample trigger, with those on events from the Muons stream, using the single muon trigger EF_mu13 as sample trigger. The agreement between the two samples is far from perfect, but it is clear that only the Muons stream can provide enough statistics to measure the full turn-on curve. A potential overestimation of the efficiencies in the turn-on region when using the Muons stream has to be kept in mind. It should be noted that the efficiencies in the turn-on region also vary slightly with the threshold of the muon trigger which is chosen as sample trigger, decreasing with increasing muon trigger threshold, as can be seen in the left plot in Figure 5.19, where the turn-on curves of EF_xe30_noMu on event samples taken with single muon triggers with different thresholds are compared.

When computing the turn-on curves of \cancel{E}_T triggers with respect to the SUSY definition of offline \cancel{E}_T , using muon triggers to obtain an orthogonal event selection has proven to exhibit undesired effects. The reason is that muons are included in the SUSY \cancel{E}_T , whereas in the trigger \cancel{E}_T they are not. Therefore, the \cancel{E}_T turn-on as function of SUSY \cancel{E}_T does not reach a plateau for high values of offline \cancel{E}_T , but instead drops again with increasing \cancel{E}_T , as can be seen in the right plot in Figure 5.19 for four \cancel{E}_T triggers with different thresholds. This behavior is presumably due to events where muons and neutrinos are produced back-to-back. In these events, there is real \cancel{E}_T coming from the neutrino. This real \cancel{E}_T enters

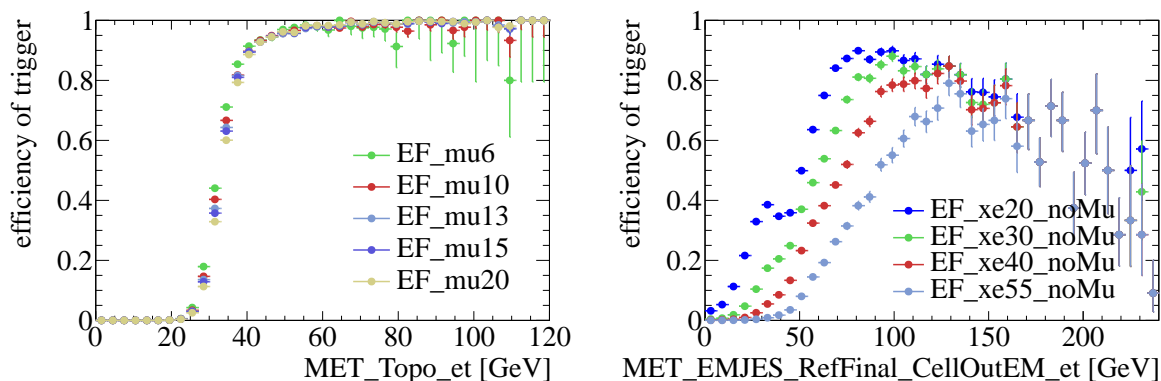


Figure 5.19: Left: Comparison of the turn-on curve for EF_xe30_noMu, on event samples taken with single muon triggers with different thresholds (given in the name in GeV), using MET_Topo as offline reference.

Right: Turn-on curves for several \cancel{E}_T triggers using SUSY \cancel{E}_T as offline reference, measured on events taken with EF_mu13. Note that for high offline \cancel{E}_T values the points for all four triggers fall together.

the computation of SUSY \cancel{E}_T , giving high values for SUSY \cancel{E}_T via the measurement of the transverse momentum of the muon. But as the trigger \cancel{E}_T does not include the muon, it cannot measure the imbalance in E_T due to the energy taken away by the neutrino so that it measures only little or no \cancel{E}_T , and the probability of the \cancel{E}_T trigger chains to fire is low.

Figure 5.20 shows two comparisons between MET_Topo and the SUSY definition of \cancel{E}_T on events sampled with EF_mu13 from the Muons stream in period I. In the left plot, showing the correlation between the two different \cancel{E}_T definitions in events with exactly one reconstructed muon with p_T above 10 GeV, note the vertical structure which can be seen at the left border of the plot, corresponding to events with vanishing values of MET_Topo but high SUSY \cancel{E}_T . It can be seen that, taking horizontal slices of this plot, the events from this vertical feature become predominant in number for high \cancel{E}_T , which explains the decrease of the \cancel{E}_T trigger efficiency when measured as function of the SUSY \cancel{E}_T , but not when using MET_Topo as offline reference. The right plot in Figure 5.20 demonstrates the difference between the two \cancel{E}_T definitions arising from the muon correction. The distribution of the angle between the \cancel{E}_T direction and the direction of the leading muon in the transverse plane peaks at zero for MET_Topo and EF \cancel{E}_T , whereas for SUSY \cancel{E}_T most events have \cancel{E}_T and the leading muon going into anti-parallel directions. This corroborates the argument that events with muons and neutrinos produced back-to-back are responsible for the problems arising when measuring the \cancel{E}_T trigger efficiency with SUSY \cancel{E}_T as offline reference, which is enhanced on the Muons stream due to the selection of only events in which a muon was identified by the trigger.

5.4.6 Combined Jet + \cancel{E}_T Triggers

The measurements of the efficiencies of the jet and \cancel{E}_T triggers which have been discussed above can now be used to determine the efficiencies of the combined jet + \cancel{E}_T triggers with the bootstrapping method. The two different types of jet + \cancel{E}_T triggers are first discussed separately, first the jet-seeded and then the full chain triggers, and then compared against each other. As offline reference for \cancel{E}_T , both MET_Topo and, where possible and expedient,




Figure 5.20: Comparisons between MET_Topo E_T and the SUSY definition of E_T on events sampled with EF_mu13 from the Muons stream, period I. Left: Correlation of the two different E_T definitions in events with exactly one reconstructed muon with p_T above 10 GeV. Right: Distribution of the angle between the E_T direction and the direction of the leading muon in the transverse plane, showing a completely different behavior for MET_Topo and SUSY E_T , whereas the distributions of E_T for EF and MET_Topo are very similar.

also the SUSY definition are used. Often, the turn-on curves are shown as contour plots. Note that this plotting style is not intended to allow to read off exact efficiency values. Instead, it ought to provide an overview of the behavior of the jet + E_T trigger at one glance, in particular with regards to where the plateau region begins, how broad the turn-on is and how it behaves as function of the respective orthogonal offline variable. Projections will be used to compare triggers with different thresholds or of different type.

Jet-seeded Triggers

As discussed in Section 5.4.2, no E_T trigger can be used to bootstrap the jet-seeded combined jet + E_T triggers, because all E_T triggers include cuts at L1 and thus violate condition (5.16). Therefore, only jet triggers will be considered for this type of jet + E_T trigger. According to Table 5.5, EF_j75_jetNoEF is the best sample trigger here. It was shown above that its efficiency cannot be measured on a minimum-bias sample due to lack of statistics, but that it has to be taken from the JetTauEtmis stream. These events can be sampled using the jet trigger EF_j30_jetNoEF, which reaches its plateau for offline jet p_T values below the onset of the turn-on region of EF_j75_jetNoEF so that in principle no bootstrap is needed, but for consistency with the EF_j50_jetNoEF_xe*_noMu turn-on curves that are shown later and for which using the bootstrap method is essential, it is done here as well. The complete sequence used to compute jet or jet + E_T trigger efficiencies, respectively, thus is the following:

$$\begin{aligned} \text{Jet} & : \text{EF_mbMbts_1_eff} \xrightarrow[\text{MinBias}]{\text{counting}} \text{EF_j30_jetNoEF} \xrightarrow[\text{JetTauEtmis}]{\text{bootstrapping}} \text{EF_j75_jetNoEF}, \\ \text{Jet} + E_T & : \text{EF_j75_jetNoEF} \xrightarrow[\text{JetTauEtmis}]{\text{bootstrapping}} \text{EF_j75_jetNoEF_EFxe*_noMu}. \end{aligned}$$

The method used to compute the turn-on is given above the arrows and the name of the stream from which the events are sampled is given below. The sample trigger is written to the left and the target trigger to the right of the arrows.

The upper plots in Figure 5.21 show the two-dimensional turn-on curve of the jet-seeded

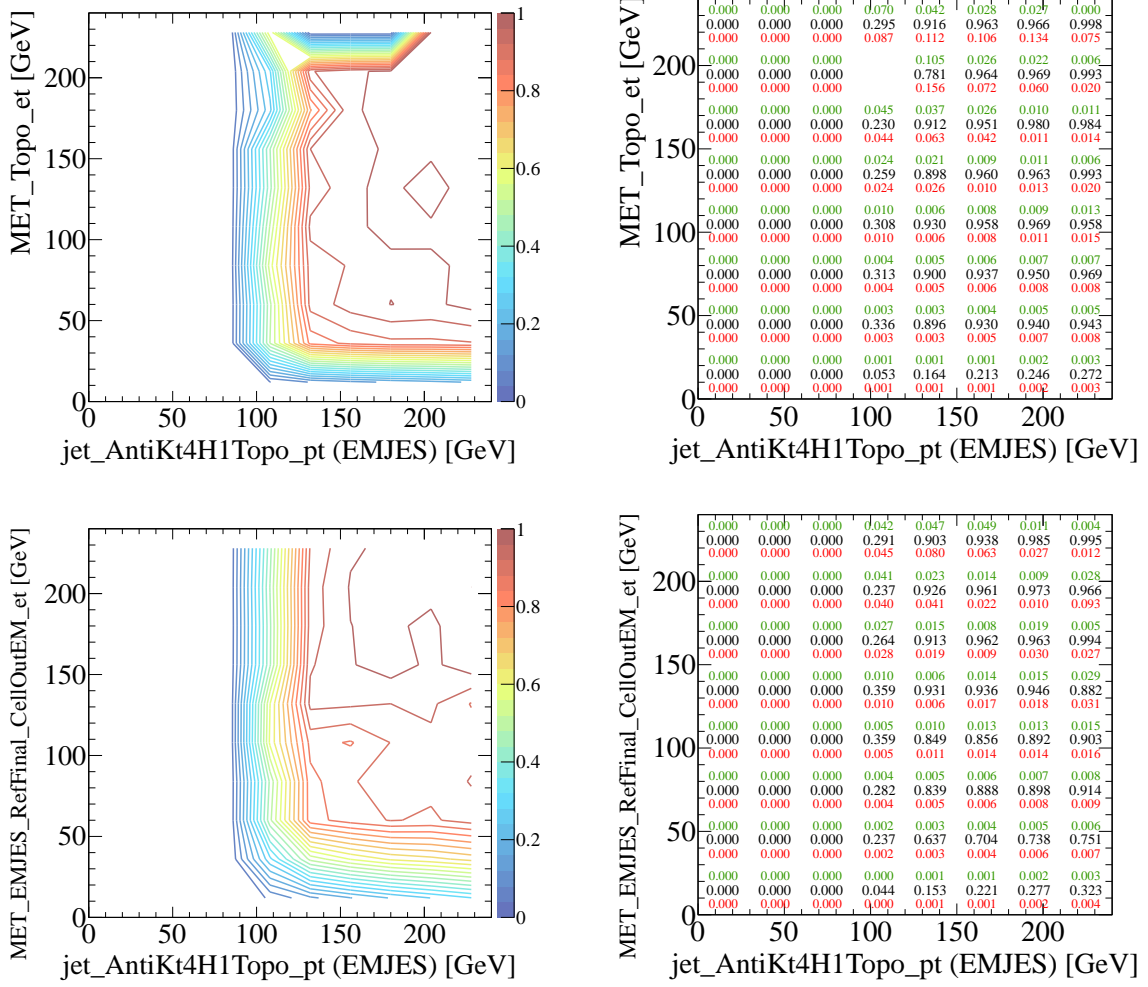


Figure 5.21: Efficiency of jet-seeded jet + \cancel{E}_T trigger EF_j75_jetNoEF_EFxe25_noMu from bootstrapping using MET_Topo (top) and the SUSY \cancel{E}_T (bottom) as offline reference and EF_j75_jetNoEF as sample trigger. Left: contour plot, right: efficiency estimates together with upper (green / above) and lower (red / below) errors for each bin.

combined jet + \cancel{E}_T trigger EF_j75_jetNoEF_EFxe25_noMu as function of MET_Topo and jet p_T . This trigger is the most important amongst the jet + \cancel{E}_T triggers, because it has been used as primary trigger to select events for the search for Supersymmetry in zero-lepton final states in 2010 data. To the left, a contour plot of the efficiencies is shown, the right plot gives the estimates of the efficiencies as numbers printed in black, adding the upper and lower uncertainties above and below the efficiency values in green and red, respectively. The empty spot in the right plot is one bin which happens to have no entries in the denominator. It can be seen that the plateau has a well-distinguished rectangular shape with an efficiency of about 90%, going to values above 95% for $\cancel{E}_T > 90$ GeV and jet $p_T > 150$ GeV. A slight increase of the trigger efficiency with jet p_T can be seen outside the plateau region for events with little or no \cancel{E}_T . This can be explained by the fact that due to the lower resolution of the \cancel{E}_T calculation in the trigger with respect to offline in events which contain high-energetic jets, it is more likely that a larger amount of (fake) \cancel{E}_T is reconstructed online, and these events are therefore more likely to also pass the \cancel{E}_T threshold of the trigger. The bump in the horizontally running turn-on region of the jet part of the combined trigger as well as the seeming drop of efficiency for very high \cancel{E}_T are due to a lack of statistics.

The lower plots in Figure 5.21 show the efficiencies of the same trigger, but using the \cancel{E}_T definition of the ATLAS SUSY group instead of MET_Topo. The events from which all four plots were produced are the same, they only differ in the way of computing the offline \cancel{E}_T . The lower plots, in particular the left one, facilitate a comparison of the trigger efficiency and the offline cuts used in the zero-lepton analysis of the ATLAS SUSY group and in this thesis in Section 7, where cuts on the missing transverse energy at 100 GeV and on the transverse momentum of the leading jet p_T at 120 GeV were applied, using exactly the same offline variables as plotted on the axes of the turn-on curve shown in these plots. Because of the calibration factors, the values of the SUSY \cancel{E}_T are usually larger than those of MET_Topo for the same event. Neglecting the muon contributions, the vertical axis is thus roughly a stretched version of what is shown in the upper plots, so that the vertically running turn-on region of the \cancel{E}_T part gets broadened, whereas the position and width of the turn-on of the jet part remain the same. Note that in these turn-on curves no correlation between \cancel{E}_T and jets can be seen, whereas, when looking at the denominator and numerator histograms separately, the correlation is visible as a diagonal feature.

Full-Chain Triggers

The jet + \cancel{E}_T triggers including the full \cancel{E}_T chain all have lower jet thresholds than the one studied above (cf. Table 5.3), which prohibits the use of EF_j75_jetNoEF as sample trigger. The next lower jet trigger which can be used as sample trigger is EF_j50_jetNoEF. The full chain of computations is then

$$\begin{aligned} \text{Jet} & : \text{EF_mbMbts_1_eff} \xrightarrow[\text{MinBias}]{\text{counting}} \text{EF_j30_jetNoEF} \xrightarrow[\text{JetTauEtmis}]{\text{bootstrapping}} \text{EF_j50_jetNoEF}, \\ \text{Jet} + \cancel{E}_T & : \text{EF_j50_jetNoEF} \xrightarrow[\text{JetTauEtmis}]{\text{bootstrapping}} \text{EF_j50_jetNoEF_xe*_noMu}. \end{aligned}$$

Figure 5.22 shows the two-dimensional turn-on curves of the full-chain jet + \cancel{E}_T trigger EF_j50_jetNoEF_xe30_noMu, which has a threshold of 30 GeV on \cancel{E}_T . The efficiencies are computed on an event sample taken with EF_j50_jetNoEF via bootstrapping. In the left plot, MET_Topo \cancel{E}_T is used as offline variable, in the right plot, the SUSY definition of \cancel{E}_T is used. It is obvious that in particular the high \cancel{E}_T range needs more statistics. While this was not




Figure 5.22: Efficiency of the full-chain jet + \cancel{E}_T trigger EF_j50_jetNoEF_xe30_noMu from bootstrapping, using the jet trigger EF_j50_jetNoEF as sample trigger. Left: using MET_Topo as offline reference, right: using the SUSY \cancel{E}_T definition.

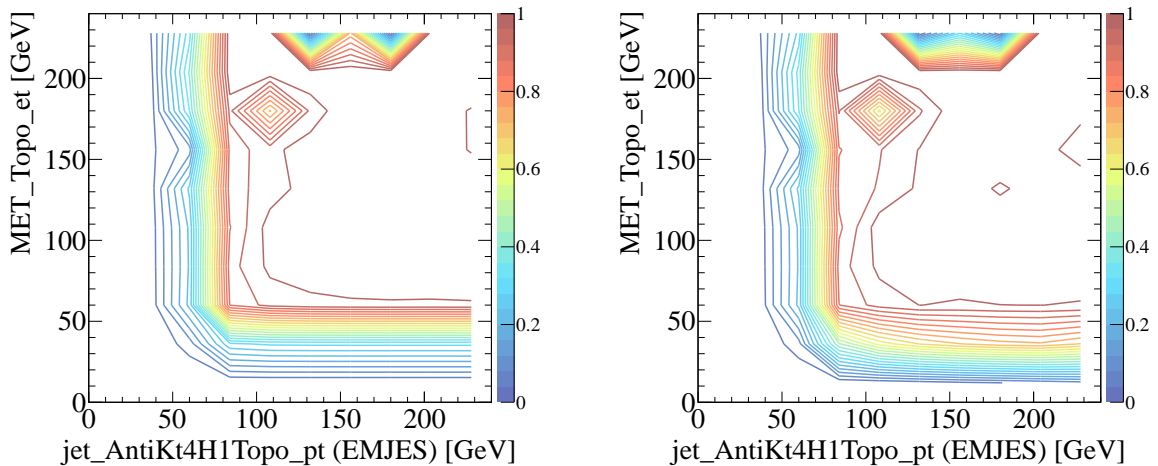


Figure 5.23: Efficiency of the full-chain jet + \cancel{E}_T trigger EF_j50_jetNoEF_xe30_noMu from bootstrapping, using EF_xe30_noMu as sample trigger. Both plots use MET_Topo as offline reference. Left: using a one-dimensional turn-on for bootstrapping, right: using a two-dimensional turn-on which is also binned in jet p_T .

possible for the jet-seeded triggers, it is possible to determine the turn-on of the full-chain triggers also on event samples taken with \cancel{E}_T triggers. The result is shown in the left plot in Figure 5.23, where EF_xe30_noMu has been used as sample trigger. The turn-on of the sample trigger has been computed on the Muons stream, as was discussed in Section 5.4.5. The computation thus comprises two steps:

$$\begin{aligned} \text{Jet} & : \text{EF_mu13} \xrightarrow[\text{Muons}]{\text{counting}} \text{EF_xe30_noMu} \\ \text{Jet} + \cancel{E}_T & : \text{EF_xe30_noMu} \xrightarrow[\text{Jet Tau Etmiss}]{\text{bootstrapping}} \text{EF_j50_jetNoEF_xe30_noMu}. \end{aligned}$$

As can be seen in Figure 5.23, using the \cancel{E}_T trigger yields higher statistics and a better coverage of the plateau for high \cancel{E}_T values than the jet sample trigger in Figure 5.22. An interesting observation is that while the position of the plateau is unchanged, the turn-on region of the \cancel{E}_T part in the left plot of Figure 5.23 does not show the variation with jet p_T that is clearly visible in Figure 5.22. The reason for this is that the *relative* efficiency of the combined jet + \cancel{E}_T trigger on the sample taken with EF_xe30_noMu is close to one in this range so that the shape visible in the plot is completely dominated by the turn-on behavior of the sample trigger EF_xe30_noMu. When using a one-dimensional turn-on of EF_xe30_noMu for bootstrapping, as was done in this plot, the p_T dependence of the turn-on therefore is lost. In a similar spirit, the slow turn-on of the jet part, which can be seen in the left plot in Figure 5.23 in the vertical slice between 40 and 60 GeV of offline jet p_T , is not reproduced in Figure 5.22 due to the steep turn-on of the jet sample trigger.

The jet- p_T dependence of the \cancel{E}_T turn-on of the combined jet + \cancel{E}_T trigger can be recovered using a two-dimensional turn-on curve of the \cancel{E}_T sample trigger, where in addition to the binning in offline \cancel{E}_T on the first axis, the second axis is binned in jet p_T . This is possible thanks to the high statistics available from the Muons stream. The result of bootstrapping using a two-dimensional turn-on of the sample trigger is shown in the right plot of Figure 5.23. In comparison to the left plot from the same figure and to the left plot in Figure 5.22, it becomes clear that the two-dimensional turn-on better reproduces the shape of the turn-on of the \cancel{E}_T part.

Corresponding plots using an \cancel{E}_T sample trigger with SUSY \cancel{E}_T as offline reference are not possible, because the \cancel{E}_T trigger turn-ons cannot be measured properly on the Muons stream as discussed above. Finally, it needs to be said that the caveats concerning missing L2 jet information given in Section 5.4.6 also hold here when using \cancel{E}_T triggers as sample triggers.

Comparisons of Different Jet + \cancel{E}_T Triggers

In the remaining part of this section, one-dimensional projections of the two-dimensional turn-on curves of the combined jet + \cancel{E}_T triggers will be presented, making possible a direct comparison between different thresholds and trigger types in one plot. For consistency among the different types of triggers, only jet sample triggers will be used here, because these are the only ones that can be used for bootstrapping the jet-seeded chains, too.

Figure 5.24 shows the turn-on behavior of the \cancel{E}_T part of the two groups of jet-seeded and full-chain jet + \cancel{E}_T triggers, presenting from each group three triggers with thresholds on \cancel{E}_T of 20, 30, and 40 GeV and using EF_j75_jetNoEF and EF_j50_jetNoEF as sample trigger for the jet-seeded (with a threshold of 75 GeV at EF) and full-chain triggers (with a threshold of 50 GeV at EF), respectively. The left plot uses MET_Topo as offline \cancel{E}_T , the right plot the SUSY \cancel{E}_T definition. In both cases, the projection is done after a cut on offline jet p_T at 140 GeV. The

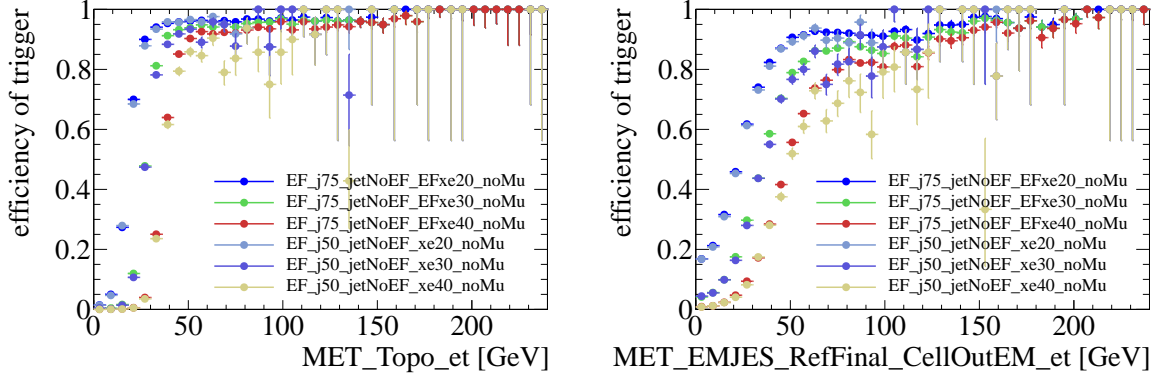


Figure 5.24: Efficiency of the \cancel{E}_T part of both types of jet + \cancel{E}_T triggers with thresholds on \cancel{E}_T at 20, 30 and 40 GeV, selecting events with offline jet p_T of at least 140 GeV. Left plot: as function of MET_{Topo} , right plot: as function of SUSY \cancel{E}_T . The sample triggers are $\text{EF}_{j75_jetNoEF}$ and $\text{EF}_{j50_jetNoEF}$ for the jet-seeded and full-chain triggers, respectively.

left plot shows a plateau efficiency around 95 %, whereas it appears to be little bit lower in the right plot. In both plots, the efficiency of the jet-seeded triggers is slightly higher in the turn-on region with respect to the full-chain triggers, whereas the plateau value is consistent within error bars.

A slight dependence of the plateau value on the \cancel{E}_T threshold value is visible, too, and becomes even more distinguished in the left plot in Figure 5.25, where the turn-on behavior of the \cancel{E}_T part of only the jet-seeded triggers is shown as function of SUSY \cancel{E}_T . Between 100 and 200 GeV, the plateau behavior of these triggers is not stable, but all thresholds show a wavy structure as function of SUSY \cancel{E}_T . Raising the offline cut on jet p_T seems to slightly dampen the variations, but a lack of statistics impedes a final conclusion here. In the right plot of the same figure, the turn-on behavior of the jet part of the jet + \cancel{E}_T triggers is compared, again showing efficiencies for both groups of jet + \cancel{E}_T triggers. As the two types of jet + \cancel{E}_T triggers differ by their EF jet threshold and the dependence on the \cancel{E}_T threshold seems to be low after the offline \cancel{E}_T cut on MET_{Topo} which is applied for the projection, the curves fall into two groups, each exhibiting approximately the same shape. It can be seen that for $\text{EF}_{j50_jetNoEF_xe40_noMu}$ (in yellow), the \cancel{E}_T cut of 70 GeV is too low because this trigger has not yet reached its plateau efficiency at that value of \cancel{E}_T . The efficiencies therefore appear to be lower than those of the other two full-chain jet + \cancel{E}_T triggers with thresholds at 20 and 30 GeV, which fall exactly together.

Cross-checks Using Muon Triggers as Orthogonal Sample Triggers

The Muons stream has been assigned approximately the same bandwidth as the JetTau-Etmiss stream and therefore offers a promising alternative to cross-check the bootstrap results. Again, $\text{EF}_{\mu13}$ is used to select the event sample, and Figure 5.26 shows the turn-on curves of the jet-seeded trigger $\text{EF}_{j75_jetNoEF_EFxe25_noMu}$ and the full-chain trigger $\text{EF}_{j50_jetNoEF_xe30_noMu}$, in both plots using MET_{Topo} as offline \cancel{E}_T variable. The two plots need to be compared to the upper plots in Figure 5.21 and Figures 5.22 and 5.23, respectively, and agree quite well with the results obtained from bootstrapping. Note however that the turn-on curves of the jet + \cancel{E}_T triggers computed from the Muons stream are reliable

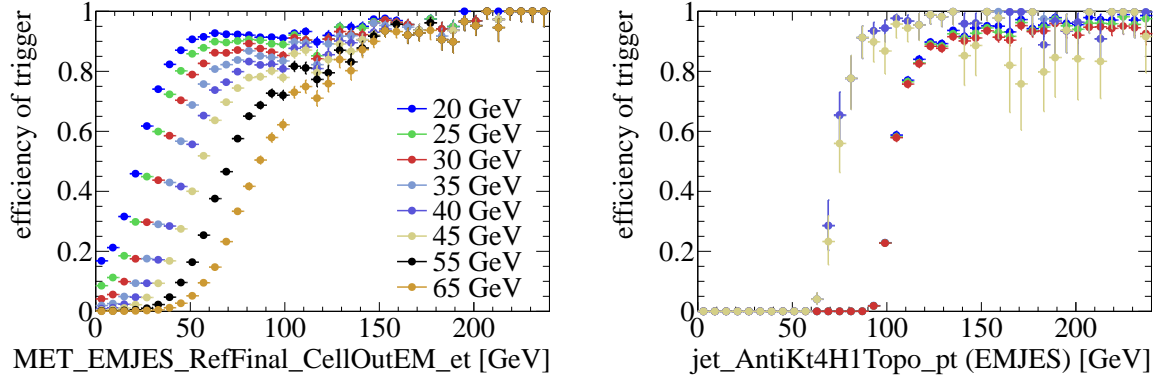


Figure 5.25: Left plot: Efficiency of the \cancel{E}_T part of the jet-seeded triggers EF_j75_jetNoEF_EFxe*_noMu, after a cut on offline jet p_T at 140 GeV as function of SUSY \cancel{E}_T . The legend gives the cut on \cancel{E}_T at EF. This plot includes three of the curves shown in the right plot of Figure 5.24.

Right plot: Efficiency of the jet part of both types of jet + \cancel{E}_T triggers, after a cut on MET_Topo at 70 GeV as function of offline jet p_T . The two types of jet + \cancel{E}_T triggers have different EF jet thresholds of 50 and 75 GeV, therefore the curves fall into two groups with the same shape. The legend for the right plot is the same as in the plots in Figure 5.24.




Figure 5.26: Efficiency of jet + \cancel{E}_T triggers measured on the Muons stream under the assumption of orthogonality of the muon triggers with respect to jets and \cancel{E}_T . Left: EF_j75_jetNoEF_EFxe25_noMu, right: EF_j50_jetNoEF_xe30_noMu. Both plots use \cancel{E}_T from the MET_Topo container as offline reference.

only as long as the assumption of orthogonality of the muon triggers holds. The corresponding plots with the SUSY definition of \cancel{E}_T instead of MET_Topo suffer from the same problem which was explained for the \cancel{E}_T turn-on curves in Section 5.4.5 and are therefore not suitable as cross-check.

It must be stressed that in those events for which the feature-extraction algorithms which reconstruct the L2 jet triggers objects have not been run, because no L1 jet trigger has accepted the event due to prescales and thereby seeded the execution of the feature-extraction algorithms at L2, the trigger efficiency cannot be correctly evaluated from the trigger objects as it is done here using the trigger emulation. The trigger emulation would always reject the event because it does not find any L2 jets, even though the event possibly may have jets with sufficient energy which would have fired the emulated trigger if the L1 accept signal had not been suppressed by prescales. For the jet-seeded jet + \cancel{E}_T triggers, this is not a problem because they require L1_J55, which has been running unprescaled throughout 2010 so that in all events where any of the jet-seeded jet + \cancel{E}_T triggers could possibly have fired, at least L1_J55 must have fired, too, and seeded the reconstruction of trigger jets at L2. For the full-chain jet + \cancel{E}_T triggers, the argumentation is not as straightforward, because they require only L1_J30, which has become heavily prescaled since period 2010 F. This means that in events which have L1 jets between 30 and 55 GeV, the prescale of L1_J30 will inhibit the creation of L2 jets in most events. Events containing only L1 jets with smaller energy can be ignored because those events would not fire the full-chain jet + \cancel{E}_T triggers anyway. Events containing L1 jets with larger energy also trigger the unprescaled L1_J55 item, and L2 jets become again available. In any case, the agreement of the turn-on curves using the Muons stream or \cancel{E}_T sample triggers on the JetTauEtmiss stream, in which case the L2 jets may also be missing, and those using jet sample triggers on the JetTauEtmiss stream (cf. in particular Figures 5.22 and 5.23), where L2 jets are always available due to the sample trigger requirement, suggests that the impact of the potential lack of jet information at L2 can be safely neglected for 2010. The same conclusion can be drawn from Figure 5.17 for the jet trigger efficiencies.

5.5 Conclusions for 2010

In the above, the application of the bootstrapping method to compute efficiency estimates of combined jet + \cancel{E}_T triggers has been described, and its application to the two different groups of jet + \cancel{E}_T chains that were defined in the 2010 data-taking menu for proton-proton collisions in the ATLAS experiment was demonstrated. It was shown how bootstrapping can be used to obtain unbiased turn-on efficiencies, and the performance of the two different groups of jet + \cancel{E}_T trigger chains was compared. Making use of different types of sample triggers, based on jets, \cancel{E}_T , and muons, allowed to compare their coverage of the relevant part of phase space and to do several cross-checks of the results yielding satisfying agreement.

The studies have been carried out using two different offline \cancel{E}_T definitions, MET_Topo, which in terms of calibration and composition is similar to the computation of \cancel{E}_T in the trigger, and a refined calibration used by the ATLAS Supersymmetry group for their physics analyses. It has become clear that care has to be taken not only when choosing the thresholds of the physics triggers, but also to define sample triggers that can be used to collect event samples of sufficient size, on which the efficiencies and the turn-on behavior of the primary triggers employed in physics analyses can be determined. It should be stressed that, while

Name of chain	Cut on jet E_T [GeV]	Cut on \cancel{E}_T [GeV]
EF_j75_a4tc_EFFS_xe45,55_loose_noMu ¹	L1: 50	L1: 20, 25/30 ³
	L2: 70	L2: 20, 25
	EF: 75	EF: 45, 55
EF_j75_a4tc_EFFS_xe55_noMu ²	L1: 50	L1: 35
	L2: 70	L2: 35
	EF: 75	EF: 55
EF_j80_a4tc_EFFS_xe60_noMu ²	L1: 50	L1: 35
	L2: 75	L2: 40
	EF: 80	EF: 60

¹ There is also the chain EF_j75_a4tc_EFFS_xe40_loose_noMu, but it has never been running.

² This chain has been introduced in 2011 I.

³ The definition of EF_j75_a4tc_EFFS_xe55_loose_noMu has changed starting with period 2011 I in that the L1 threshold on \cancel{E}_T was raised to 30 GeV.

Table 5.6: Overview of the relevant set of combined jet + \cancel{E}_T trigger chains which were defined in the trigger menu `Physics_pp_v2` for data taking in 2011. The table shows the cuts which are applied to the values of \cancel{E}_T and jet E_T measured online at Level 1, Level 2 and Event Filter.

being a good cross-check for the MET_Topo variant of \cancel{E}_T , the muon triggers cannot be used as sample triggers for turn-ons when using the SUSY \cancel{E}_T definition (or probably any \cancel{E}_T which includes muon contributions) as offline reference, as was shown here. For computing the turn-on curves of jet + \cancel{E}_T triggers with respect to the offline SUSY \cancel{E}_T definition \cancel{E}_T triggers could therefore not be used as sample trigger, because no suitable turn-on of these to perform bootstrapping could be obtained.

5.6 Measurements of Trigger Efficiencies in Data Taken in 2011

Combined jet + \cancel{E}_T triggers have remained the primary physics triggers for the zero-lepton studies of the ATLAS Supersymmetry group in 2011. This section documents the continuation of the studies which were begun in 2010 and presented above. The efficiency studies in 2011 profit a lot from the experience gained in 2010 in terms of which methods and sample triggers are suitable to measure the efficiencies of the combined jet + \cancel{E}_T triggers. Although the general type of the primary triggers has remained the same, there were a number of changes in the trigger menu and in the trigger system itself, which are explained in detail in the next section.

5.6.1 Overview of Changes and New Jet + \cancel{E}_T Triggers

The most important change between 2010 and 2011 with respect to the combined jet + \cancel{E}_T triggers concerns the jet trigger algorithm at Event Filter level, which has been replaced by a full-scan algorithm (indicated by “EFFS” in the trigger name) and has been used for rejecting

events for the first time in 2011 (cf. Section 4.3.3). The efficiencies of triggers involving jets are therefore expected to potentially exhibit fundamental changes.

Table 5.6.1 shows the relevant combined jet + \cancel{E}_T triggers which are defined in the trigger menu in 2011. The thresholds lie in the same order of magnitude as the highest thresholds which were defined in 2010. Jet-seeded chains are no longer included in the menu, since the increased instantaneous luminosity prohibits keeping low threshold jet items at Level 1 unprescaled because the input rate to Level 2 would be too high. As a too-close spacing of the cuts at Level 1 and Event Filter on \cancel{E}_T deteriorates the \cancel{E}_T trigger performance due to the very broad turn-on of the \cancel{E}_T at Level 1, the spacing of the thresholds at Level 1 and Level 2 and Event Filter is kept relatively large, as can be seen from the table. The combined trigger EF_j75_a4tc_EFFS_xe45_loose_noMu has been the primary trigger for the first analyses of 2011 data by the ATLAS Supersymmetry group [148, 149].

Due to the difficulties with the measurements of the \cancel{E}_T turn-on curves, in particular when using \cancel{E}_T with muon corrections as offline reference, which were discussed above in the context of the trigger efficiencies in 2010, and because of the larger statistics coming through jet triggers, whereas the low threshold \cancel{E}_T triggers had to be deactivated or heavily prescaled, in 2011, \cancel{E}_T triggers have not been employed for bootstrapping. At the end of this section, a brief cross-check is given though.

The event selection is done in the same way as in 2010, applying the SUSY GRL and the jet cleaning in its updated version for ATHENA release 16 [189]. Which periods were used to generate the turn-on curves is specified below for each plot individually. If not stated otherwise, data from all available periods in 2011 is used, i. e. up to period 2011 K, except for the first three periods A, B and C. Periods A and B have been mostly for testing after the winter shutdown of the LHC and do not contain much usable data, and 2011 C has a reduced center-of-mass energy of 2.76 TeV. The object definitions have been updated following the recommendations of the performance groups and the changes in the zero-lepton analysis of the ATLAS SUSY group. Instead of MET_Topo, MET_LocHadTopo is used, as recommended by the ATLAS Jet/Etmiss combined performance group for 2011 [190]. The SUSY \cancel{E}_T definition in 2011 is called MET_Simplified20_RefFinal¹, and the new jet collection is called AntiKt4TopoNewEM, but still contains jets calibrated at EM+JES scale (cf. Sections 4.4.3 and 4.4.2).

5.6.2 Jet Triggers

The jet triggers will play an even more important role as sample triggers for the measurement of the combined jet + \cancel{E}_T trigger efficiencies in 2011 than they did in 2010. In particular, the efficiencies of the jet trigger EF_j75_a4tc_EFFS will need to be studied because this trigger will be the sample trigger for all efficiencies of jet + \cancel{E}_T triggers presented in Section 5.6.3. Table 5.7 gives an overview of the jet trigger chains which were defined in the trigger menu Physics_pp_v2 for data taking in 2011. They fall into two groups: The lowest chains with thresholds of 10, 15 and 20 GeV at Event Filter level are seeded by random triggers at Level 1 and Level 2 (L1_RD0_FILLED and L2_rd0_filled_NoAlg), and only work with trigger jet objects at Event Filter level. The chains with higher thresholds at Event Filter are full chains again in the sense that they require trigger jets to have been reconstructed at Level 1 and Level 2, too.

¹ Note that the term “SUSY \cancel{E}_T ” in the following therefore denotes a different version of offline \cancel{E}_T than in the plots for 2010, which in most plots is obvious from the axis title.

EF chain name	L2 chain name	L1 item name
EF_j10_a4tc_EFFS	L2_rd0_filled_NoAlg	L1_RD0_FILLED
EF_j15_a4tc_EFFS	L2_rd0_filled_NoAlg	L1_RD0_FILLED
EF_j20_a4tc_EFFS	L2_rd0_filled_NoAlg	L1_RD0_FILLED
EF_j30_a4tc_EFFS	L2_j25	L1_J10
EF_j40_a4tc_EFFS	L2_j35	L1_J15
EF_j55_a4tc_EFFS	L2_j50	L1_J30
EF_j75_a4tc_EFFS	L2_j70	L1_J50

Table 5.7: Overview of the single jet trigger chains which were defined in the trigger menu `Physics_pp_v2` for data taking in 2011. The table shows the chain or item names at Level 1, Level 2 and Event Filter, the two-digit numbers specify the thresholds on trigger jets in GeV (cf. Section 4.3.5).

The efficiency of `EF_j75_a4tc_EFFS` needs to be determined using a lower jet trigger as sample trigger, in the same way as was done in 2010 for the jet triggers. The lowest jet trigger fulfilling the bootstrapping condition (5.16) is `EF_j30_a4tc_EFFS`, which is the lowest jet trigger chain not seeded by random triggers (cf. Table 5.7). The first step, therefore, is to obtain the efficiency of `EF_j30_a4tc_EFFS`. A direct measurement of the efficiency of `EF_j30_a4tc_EFFS` in terms of the full chain of Level 1, Level 2 and Event Filter, i. e. the product of the efficiencies at all three trigger levels in ATLAS, cannot be done on an event sample selected with the random trigger `EF_rd0_filled_NoAlg`, because of the high prescales of the respective jet triggers at Level 1 and Level 2, leading to the already explained problem that trigger objects are missing. This means that the trigger efficiencies would be strongly underestimated and biased by events which are actually only recorded, because they have also fired a jet trigger with a higher threshold, which led to the execution of the feature-extraction algorithms at the higher trigger levels. However, it is possible to compute the efficiencies for the three random-seeded jet triggers `EF_j10,15,20_a4tc_EFFS` using such a sample, because if the random trigger chain (`EF_rd0_filled_NoAlg` seeded by `L2_rd0_filled_NoAlg`) fires at L2, the feature-extraction algorithms for jets at Event Filter level are executed according to Table 5.7. The efficiency of `EF_j20_a4tc_EFFS` will therefore be measured like this and employed below.

As a consequence, in order to obtain the efficiency of `EF_j30_a4tc_EFFS`, the efficiencies of all three levels need to be measured separately:

- At Level 1, the efficiency for `L1_J10` can be measured on an event sample taken with a random trigger.
- At Level 2, the efficiency of `L2_j25` relative to `L1_J10` can be measured by selecting a sample of events in which both `EF_rd0_filled_NoAlg` and `L1_J10` fire. This will, of course, give very low statistics. Note that it is not enough to require `L1_J10` only because not all events passing `L1_J10` will be recorded due to prescales, and the sample obtained by requiring `L1_J10` is thus not an unbiased sample.
- The efficiency of `EF_j30_a4tc_EFFS` with respect to Level 2 can be measured either by again requiring both `EF_rd0_filled_NoAlg` and `L2_j25` to fire, or on a sample of




Figure 5.27: Left: Efficiencies of the three levels of the trigger chain EF_j30_a4tc_EFFS, measuring L1 and L2 on a sample of events taken with the random trigger EF_rd0_filled_NoAlg from the MinBias stream, and EF with bootstrapping on EF_j20_a4tc_EFFS on the JetTauEtmisss stream. Note that for L2, this is the relative efficiency with respect to L1, and that the curve for EF_j30_a4tc_EFFS shows the efficiency of the EF level only and not the efficiency of the full chain.

Right: Comparison of the efficiency of the jet trigger EF_j75_a4tc_EFFS in periods 2011 D and K, on a sample of events taken with EF_j30_a4tc_EFFS. No bootstrapping is applied.

events triggered by L2_j25 and recorded due to the active pass-through factor at EF level for this chain, or — as is done here — using bootstrapping on events taken with EF_j20_a4tc_EFFS, the efficiency of which can be measured on random-triggered events.

This is obviously more complicated than a direct measurement of the turn-on curve for EF_j30_a4tc_EFFS, but for the following it will be sufficient to convince oneself that this trigger reaches its plateau before the onset of the turn-on region of EF_j75_a4tc_EFFS, which is around 100 GeV. The correction of the sample trigger bias in the measurement of the efficiency of EF_j75_a4tc_EFFS in a sample taken with EF_j30_a4tc_EFFS can then be relinquished. The left plot in Figure 5.27 shows the efficiencies of L1_J10, L2_j25 and EF_j30_a4tc_EFFS, which comprise the full chain EF_j30_a4tc_EFFS^m. The efficiency for Level 2 is measured relative to its lower chain L1_J10, i. e. the denominator of the efficiency of L2_j25 is filled with the p_T of jets from events which fire both EF_rd0_filled_NoAlg and L1_J10. The efficiency for EF_j30_a4tc_EFFS is measured on events taken with EF_j20_a4tc_EFFS, by applying bootstrapping. Note again that the efficiency here refers only to the Event Filter level of the trigger chain, not to the combined efficiency of all trigger levels. As expected, statistics is very low for L2 due to the conjunction of two trigger requirements. Still, the plot allows to estimate the position of the turn-on and the beginning of the plateau region of the full chain EF_j30_a4tc_EFFS quite accurately. It shows that Level 1 and Level 2 should not effect a large shift of the efficiency of the full chain to the right, and it can be seen that the full turn-on, i. e. the product of the efficiencies of all three trigger levels, should indeed reach 100 % safely

^m The convention by which the full chain bears the same name as the Event Filter part of the chain is a relic from the early commissioning of the trigger when event rejection was not activated for all trigger levels and the prefix was needed to make clear up to which level the performance of the chain is studied. Recently, it has become common to refer to the full chain by leaving out the prefix, e. g. j30_a4tc_EFFS.




Figure 5.28: Efficiency of jet triggers with EF thresholds of 55 and 75 GeV measured on the JetTauEtmisss stream on an event sample taken with EF_j30_a4tc_EFFS. No bootstrapping is applied.

below 100 GeV, and even below 70 GeV, which would be the condition for measuring also EF_j55_a4tc_EFFS on a sample taken by EF_j30_a4tc_EFFS directly.

The efficiency of EF_j75_a4tc_EFFS is plotted in Figure 5.28, together with the efficiency of EF_j55_a4tc_EFFS. No bootstrapping has been done to correct for a bias of EF_j30_a4tc_EFFS in the event sample. Instead, the plot relies on the assumption that EF_j30_a4tc_EFFS reaches its efficiency plateau well below the onset of the turn-on region of EF_j75_a4tc_EFFS, which has been demonstrated above. Only data from period 2011 D has been used to produce this plot, but this is sufficient to cover the full range up to 360 GeV in offline jet p_T , which will be the upper limit chosen for the measurement of the jet + \cancel{E}_T trigger efficiencies in the following. The increase in statistics which could be achieved by summing over periods E – G in addition to period D is very small. To make sure that the efficiency of this sample trigger is stable and valid also for data taken in later periods, a cross-check has been done by measuring the efficiency of EF_j55_a4tc_EFFS again in period 2011 K. The result is shown in the right plot in Figure 5.27, where no changes with respect to the horizontal position of the turn-on region can be observed. The slight decrease of the efficiency in the turn-on region may be attributed to the pile-up noise suppression which was introduced between periods D and K.

Recent trigger performance plots published by ATLAS for 2011 data [134] confirm that the efficiency of the jet trigger EF_j40_a4tc_EFFS with Event Filter threshold of 40 GeV reaches a plateau efficiency near 100% at a calibrated offline jet E_T of around 70 to 75 GeV. Using the jet trigger EF_j30_a4tc_EFFS with a lower Event Filter threshold of 30 GeV as sample trigger to determine the efficiency of the jet trigger EF_j75_a4tc_EFFS with an Event Filter threshold of 75 GeV therefore will give correct results even without the need to apply the bootstrap method.




Figure 5.29: Efficiency of the jet + E_T trigger EF_j75_a4tc_EFFS_xe45_loose_noMu from bootstrapping, using the jet trigger EF_j75_a4tc_EFFS as sample trigger and data from periods 2011 D – K. The efficiencies are shown as contour and as color-coded plots, using MET_LocHadTopo as offline reference (left column) and using the SUSY E_T definition (right column). White spots arise from empty bins in the denominator histograms due to insufficient statistics.

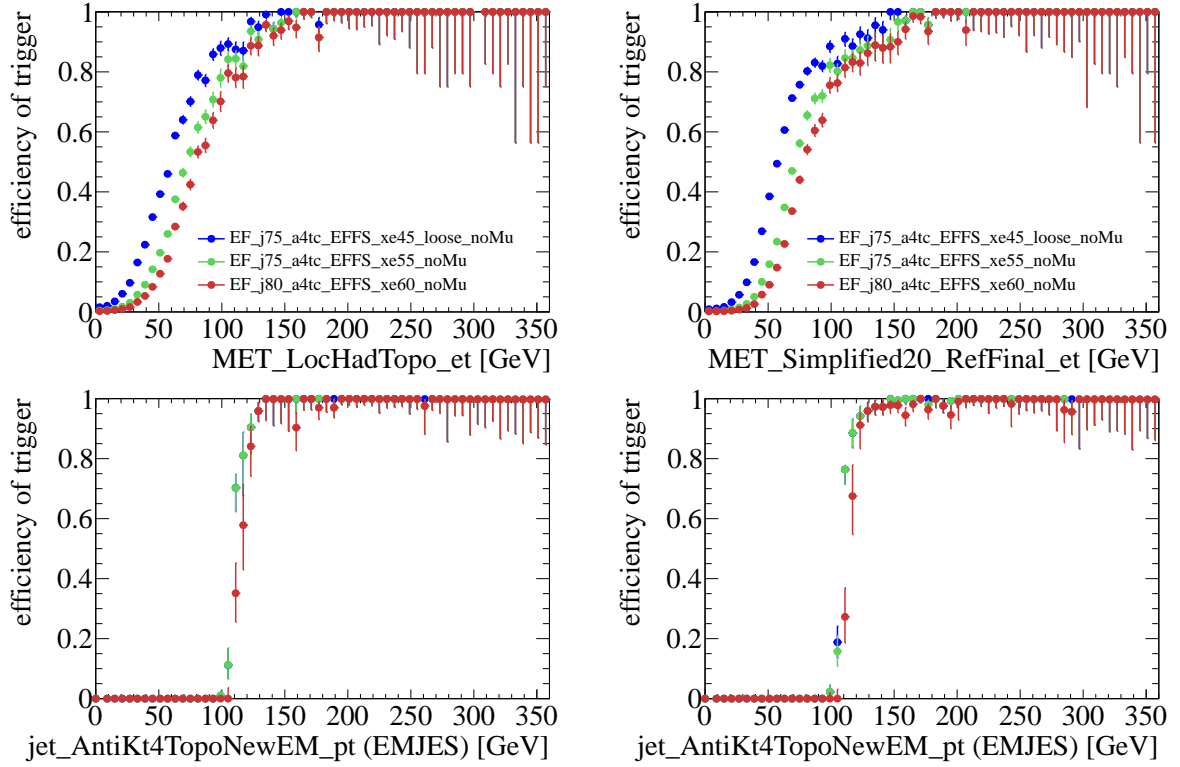


Figure 5.30: Efficiencies of three jet + \cancel{E}_T triggers from bootstrapping, using the jet trigger EF_j75_a4tc_EFFS as sample trigger and data from periods 2011 D – K. Projections onto the \cancel{E}_T axis (top) and jet p_T axis (bottom) are shown, after cuts on the respective orthogonal variable at 150 GeV. In the left column, MET_LocHadTopo is the offline reference, in the right column the SUSY \cancel{E}_T definition. (The legend in the upper plots also applies to the lower plots.)

5.6.3 Combined Jet + \cancel{E}_T Triggers

The efficiency of the combined jet + \cancel{E}_T trigger EF_j75_a4tc_EFFS_xe45_loose_noMu has been studied right from the beginning of data-taking in 2011, updating the plots with new data whenever it became available. Here, only the most recent plots including all so-far available data are shown, which includes periods D – K. The four plots in Figure 5.29 show the efficiency of this trigger with two different offline \cancel{E}_T references, MET_LocHadTopo and the SUSY \cancel{E}_T definition, and two different plot styles, as contour plots and as color-coded plots.

Figure 5.30 shows projections of the efficiencies onto the \cancel{E}_T and jet p_T axes, after cuts on the respective orthogonal variable at 150 GeV. In these plots, it is possible to compare several triggers. In addition to EF_j75_a4tc_EFFS_xe45_loose_noMu the primary physics trigger at the beginning of 2011, two other triggers are included which were introduced in period 2011 I, when the former primary physics trigger with an \cancel{E}_T threshold of 45 GeV had to be prescaled because its rate had become too high due to the increased instantaneous luminosity.

These three jet + \cancel{E}_T triggers all differ in the thresholds on \cancel{E}_T (45, 55 and 60 GeV at EF), but only one of them has a higher threshold also on jet p_T (80 instead of 75 GeV at EF). Correspondingly, in the projection onto jet p_T two of the three curves coincide, and even the turn-on of the trigger with the higher jet p_T threshold is very close, as the increase in the threshold is only 5 GeV. The spacing of the triggers with respect to their \cancel{E}_T turn-ons also agrees with the expectations, their thresholds at Event Filter being different by 10 GeV and 5 GeV, respectively. In the upper right plot in Figure 5.30, it can be seen that the offline cut of 150 GeV on SUSY \cancel{E}_T is too low to bring the jet + \cancel{E}_T trigger with the highest threshold of 60 GeV on \cancel{E}_T in its plateau, with the consequence that in the lower right plot, the efficiency of this trigger seems to be a little lower than 100%. For the projection onto \cancel{E}_T , in particular the upper right plot, which uses the SUSY \cancel{E}_T definition, seems to hint at a slight decrease in efficiency in the region between roughly 100 and 150 GeV of offline \cancel{E}_T . It is much less pronounced for MET_LocHadTopo (if present at all), and may be connected to details of the definition of SUSY \cancel{E}_T .

Note that although the new combined jet + \cancel{E}_T triggers were only defined in periods 2011 I and later, their performance can still be studied in earlier periods thanks to the emulation of the trigger decision for periods where these triggers were not yet defined (cf. Section A.1). This gives a significant gain in statistics.

Consistency of Periods

As there have been a number of changes and incidents with potential impact on the jet + \cancel{E}_T trigger efficiencies, a cross-check has been done of the kind that the efficiencies of the triggers are plotted separately for four small sets of data, each combining only data from between one and three periods. The result is shown in Figure 5.31, again only for the primary trigger EF_j75_a4tc_EFFS_xe45_loose_noMu, because only this trigger has actually been used for taking data up to period 2011 I^a. The four sets contain data from periods D, E – F, G – H and I – K. This division has been chosen because period 2011 E is the first period in which calorimeter measurements were affected by the LAr hole (cf. Section 3.4.3). In period G, the noise suppression cuts in the trigger have been adjusted to account for pile-up contributions, and beginning with period I, the read-out of 4 of the 6 front-end boards constituting the

^a The plots for the other two triggers are very similar to the ones shown in Figure 5.31.




Figure 5.31: Efficiency of the jet + \cancel{E}_T trigger EF_j75_a4tc_EFFS_xe45_loose_noMu from bootstrapping, using the jet trigger EF_j75_a4tc_EFFS as sample trigger. Here, the efficiencies are shown for small sets of periods separately, using the usual projections after cuts onto the respective orthogonal variable at 150 GeV. The offline reference for \cancel{E}_T is MET_LocHadTopo.

LAr hole could be re-established. In the two projection plots, MET_LocHadTopo is used as offline reference, but the plots for SUSY \cancel{E}_T basically show the same features. The left plot, with the projection onto \cancel{E}_T , indicates a small difference between periods D – F and G – K, the trigger efficiency for small values of \cancel{E}_T being higher in periods G – K. This is, however, the opposite of what would be expected from the change (i. e. increase) of the noise thresholds which should lead to a decrease in the \cancel{E}_T seen by the trigger and thus a lower efficiency. Note that the projection on jet p_T is actually not expected to exhibit a strong dependency on the period, because the turn-on behavior it shows is likely dominated by the turn-on that is used for bootstrapping, i. e. the one of the sample trigger EF_j75_a4tc_EFFS, because it has the same cuts on jet p_T as the combined trigger itself. As the same turn-on from period D is used for all periods, the projection onto jet p_T is expected to agree for all periods. The stability of the sample trigger and thus the jet part of the combined trigger is demonstrated by the right plot in Figure 5.27 from above.

In general, the efficiencies from all four sets of data are in good agreement. The onset of the plateau region suffers from low statistics and apparent fluctuations in the efficiency measurements as mentioned above. From the contour plots in Figure 5.29, this behavior seems to be coming dominantly from the region of high jet p_T , and thus really be due to the limited statistics rather than being a genuine trigger issue. In any case, it should be investigated further when more data becomes available.

Cross-checks Using Muon Triggers as Orthogonal Sample Triggers

The idea of doing a cross-check of the efficiencies measured on the JetTauEtmiss stream under the assumption of muon triggers being orthogonal to the jet + \cancel{E}_T trigger definition, thus allowing for the use of a muon-triggered event sample, has been picked up for the efficiencies in 2011 again. Figure 5.32 shows two-dimensional turn-on curves for EF_j75_a4tc_EFFS_xe45_loose_noMu on the sample of events selected with the single muon trigger EF_mu18 from the Muons stream in periods 2011 I – K. A synopsis together with the two other jet + \cancel{E}_T triggers introduced in 2011 I is shown in the projection plots in Figure 5.33. All plots have very high statistics and therefore give a very good coverage of the plateau regions. Noticably,




Figure 5.32: Efficiency of the jet + \cancel{E}_T trigger EF_j75_a4tc_EFFS_xe45_loose_noMu on the Muons stream, using EF_mu18 as sample trigger under the assumption of orthogonality. The offline reference for \cancel{E}_T is MET_LocHadTopo. The two plots are produced from the same data and only differ in style.

Name of chain	Sample trigger	Lower chains
EF_j75_a4tc_EFFS_xe45,55_loose_noMu	EF_xe30_noMu	L2_xe20_noMu L1_XE20
EF_j75_a4tc_EFFS_xe55_noMu	EF_xe40_noMu	L2_xe30_noMu L1_XE30
EF_j80_a4tc_EFFS_xe60_noMu	EF_xe50_noMu	L2_xe35_noMu L1_XE35

Table 5.8: Overview of \cancel{E}_T triggers with highest thresholds that respect the bootstrapping condition (5.16) and can thus be used as sample triggers for the combined jet + \cancel{E}_T triggers. The lower chain names make clear the thresholds applied at Level 1 and Level 2.




Figure 5.33: Efficiency of three jet + \cancel{E}_T triggers on the Muons stream, using EF_mu18 as sample trigger under the assumption of orthogonality and projecting onto the \cancel{E}_T (left) and jet p_T (right) axis. (The legend in the left plot also applies for the right one.)

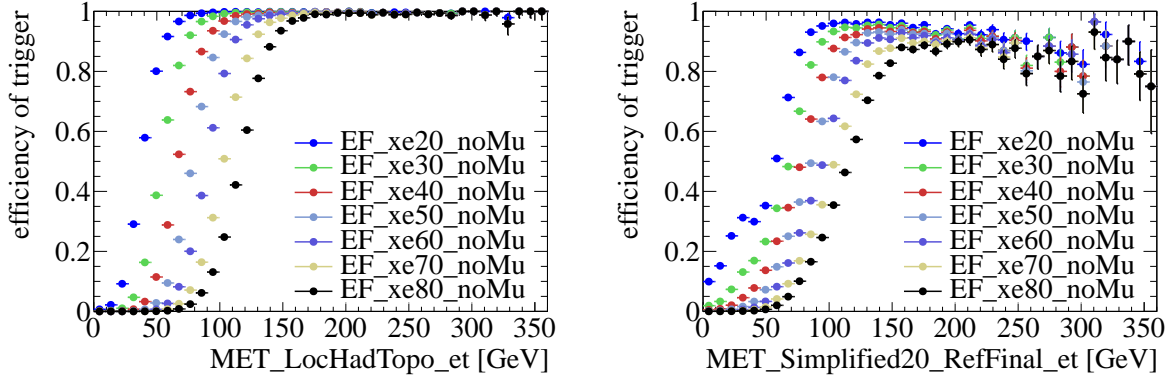


Figure 5.34: Efficiency of \cancel{E}_T triggers with thresholds between 20 and 80 GeV on the Muons stream, using EF_mu18 as sample trigger under the assumption of orthogonality, as function of MET_LocHadTopo (left) and SUSY \cancel{E}_T (right).

the dip in the efficiency which can be observed in Figure 5.30 is not reproduced in these plots at all. This indicates that it is either only a problem of statistics, or it is related to the composition of \cancel{E}_T . In the muon-triggered event sample, as explained above, the \cancel{E}_T composition is dominated by real \cancel{E}_T so that, if the dip is due to fake \cancel{E}_T , e. g. from jet mismeasurements, it will not show up in the efficiency plots from the Muons stream.

Figure 5.34 shows the turn-on curves for several triggers with different thresholds on the missing transverse energy, as they are computed on events taken with the muon trigger EF_mu18 as sample trigger, under the assumption of orthogonality. In the left plot, MET_LocHadTopo is the offline reference for \cancel{E}_T , in the right plot the SUSY \cancel{E}_T definition. The comparison of the efficiencies using the SUSY \cancel{E}_T definition in the right plot in Figure 5.34 to the right plot in Figure 5.19 shows that, although the instability of the plateau when using SUSY \cancel{E}_T as offline reference appears to be less pronounced than it is in 2010 data, it definitely still can be seen. Note also the hump in the efficiency curve clearly visible for EF_xe20_noMu at the typical position of the peak of the transverse energy spectrum of muons produced in W decays. The improvement over 2010 may be explained by the updated \cancel{E}_T definition as well as by the different composition of \cancel{E}_T in muon-triggered events, which in 2011 will have a higher level of fake \cancel{E}_T due to the increased number of overlaid in-time pile-up events.

In terms of event counts, the \cancel{E}_T trigger EF_xe20_noMu even gives slightly better statistics than the jet trigger EF_j75_a4tc_EFFS, but it cannot compete with this jet trigger, which has been used as sample trigger above. Due to the large spacing of its \cancel{E}_T threshold at Event Filter level with respect to the combined triggers, there are too few events left at high \cancel{E}_T to cover the plateau region. For the combined jet + \cancel{E}_T triggers with higher thresholds on \cancel{E}_T , which already have or will supersede the current primary triggers, \cancel{E}_T triggers as sample triggers may become an appealing option again if the problem of how to obtain the sample trigger turn-on can be solved. For MET_LocHadTopo as offline reference, the turn-on curves for the \cancel{E}_T trigger from the Muons stream can be used for bootstrapping. An example is given in Figure 5.35, where the efficiency of the combined jet + \cancel{E}_T trigger EF_j75_a4tc_EFFS_xe55_noMu, which was introduced in 2011 I, is shown on an event sample taken with EF_xe40_noMu after bootstrapping. The plot demonstrates that this \cancel{E}_T trigger gives a very good coverage of the plateau region, better in particular at high \cancel{E}_T values

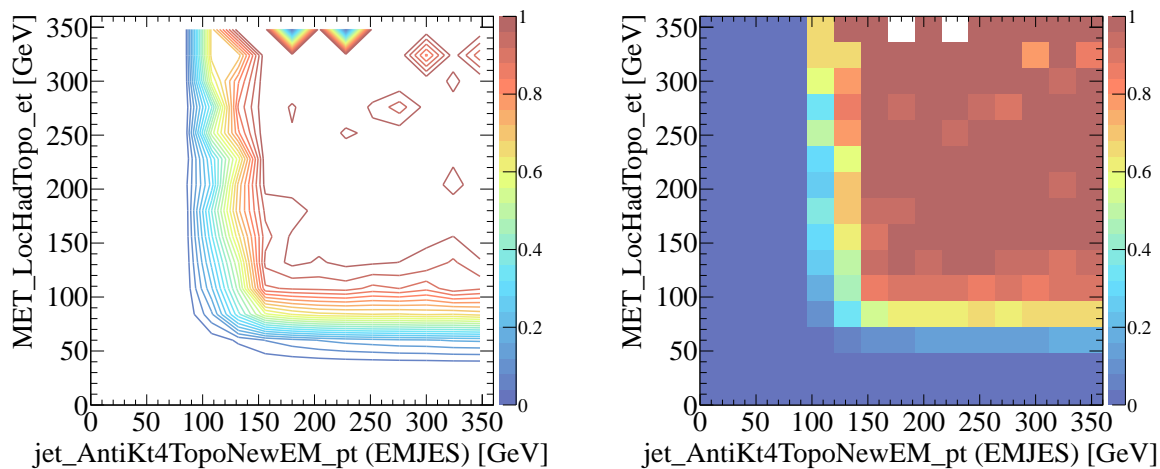


Figure 5.35: Efficiency of the combined jet + \cancel{E}_T triggers EF_j75_a4tc_EFFS_xe55_noMu bootstrapped from an event sample taken with EF_xe40_noMu.

in comparison to Figure 5.29. Noticeably, the vertical jet turn-on region in this plot is much broader and extends to higher offline values than it does in Figure 5.29 using a jet sample trigger. This suggests that in the \cancel{E}_T triggered events used to produce the plot, the problem of missing trigger jet objects again deteriorates the measurement of the jet turn-on. The \cancel{E}_T turn-on on the other hand is much steeper. This can be explained by the fact that the sample trigger turn-on is measured on an event sample collected with a muon trigger, which is dominated by real \cancel{E}_T rather than the fake \cancel{E}_T found in most events in the JetTauEtmiss stream, and thus has a sharper turn-on, as could be seen e. g. in Figure 5.15. Through the bootstrapping which is applied, the turn-on of the \cancel{E}_T trigger will shape the turn-on of the jet part of the combined trigger. This is not a problem with the bootstrapping procedure itself, but highlights again the fact that the sample trigger turn-on used for bootstrapping needs to be unbiased, too.

5.7 Conclusions for 2011

In Section 5.6, the efficiencies of the combined jet + \cancel{E}_T trigger in the 2011 trigger menu have been studied. The focus has been on the physics triggers which are used by the ATLAS Supersymmetry group. The computation of the turn-on curves largely relies on the jet trigger EF_j75_a4tc_EFFS as sample trigger, and employs the bootstrap method to correct for its bias in the event sample. The results built upon experience from the studies in 2010 and constitute a continuation of these studies. A number of additional cross-checks have been presented, showing an overall good agreement and demonstrating the consistency of the efficiencies in different periods, despite many changes in the trigger.

With respect to the official Supersymmetry analysis, the offline cuts follow the trigger requirements to be in the plateau of the trigger efficiency: It has been decided to use an offline cut of 130 GeV on both offline \cancel{E}_T and jet p_T for the analysis using EF_j75_a4tc_EFFS_xe45_loose_noMu in the zero-lepton channel [149]. There is an apparent tension between this cut and the results on the latest periods I – K because the beginning of the plateau seems

to be affected by a drop in efficiency, which could not be seen previously due to the smaller statistics. There are several possible explanations. It is conceivable that the change in the turn-on is caused by the increasing level of in-time pile-up and changing composition of \cancel{E}_T in the events. It could be due to a crossing of the \cancel{E}_T turn-on curves of the different trigger levels constituting the jet + \cancel{E}_T chain, but it also cannot be ruled out that it is an artefact that will vanish with increasing statistics. In any case, this drop cannot be seen in events collected with muon triggers so that it can be assumed that in events dominated by real \cancel{E}_T the trigger is not afflicted by this decrease of efficiency. This is also consistent with the argument of crossing turn-ons, because if the turn-ons of the different trigger levels are steeper in this sample due to the different composition of \cancel{E}_T , they would possibly not cross for this type of events.

The efficiencies of the two new jet + \cancel{E}_T triggers introduced in period I could be measured on the full dataset thanks to trigger emulation. For the next update of the Supersymmetry analysis^o, the trigger EF_j75_a4tc_EFFS_xe55_noMu will become the primary physics trigger, which necessitates to increase the offline \cancel{E}_T cut to presumably 150 GeV. Using \cancel{E}_T triggers as sample triggers might become interesting due to the higher statistics, because the currently used jet sample trigger is already heavily prescaled. Nevertheless, there are two challenges that would have to be met: First, the \cancel{E}_T turn-on needs to be measured, which is hindered by finding an unbiased event sample and by the difference between SUSY \cancel{E}_T and trigger \cancel{E}_T . Second, there is the problem of missing trigger jets, which by using jet sample triggers can be circumvented elegantly.

An approach that would be interesting to try and adopt in future analyses relying on this type of trigger would be to find the value of the offline variable at which the trigger efficiency exceeds a given threshold and an uncertainty on this value. This could then be converted into a systematic uncertainty in the analysis. One way to find this value is to fit the turn-on curve with e. g. a linear combination or a product of two Gauss error functions, which would be able to account for a crossing of two turn-on curves. Such an approach could also fully exploit the computation of the uncertainties on the trigger efficiencies, which would then enter the analysis in terms of an additional systematic uncertainty. So far, systematic uncertainties coming from the trigger efficiency have been neglected in the zero-lepton analyses, because the offline cuts were chosen such that the efficiency of the primary trigger is close to 100%, and the systematic uncertainties arising from the trigger efficiency expected to be small compared to other sources of systematic errors. With a better control of the other uncertainties, which are dominating the total uncertainty at the moment, and an on-going optimization of the analysis for sensitivity, a refined treatment of the trigger uncertainties will become more important. Besides, the progression of the instantaneous luminosity might render it unfavorable to adapt the offline cuts to the trigger requirements to be in the plateau. In particular, when fitting of kinematic distributions instead of event counting will be introduced, the impact of the variation of the trigger efficiency between different bins needs to be evaluated. Furthermore, it would be interesting to study the correlation of jet p_T and \cancel{E}_T , which is found in particular in events where jet mismeasurements make up for a large part of the \cancel{E}_T , and to estimate its influence on the trigger efficiencies.

^o Moriond 2012

6 A Model for \cancel{E}_T Trigger Rates

Missing transverse energy (\cancel{E}_T) is one of the most involved measurements at a particle detector, because it is basically a sum over all particles that are produced in the proton-proton collisions. Its resolution will therefore suffer from all possible sources of mismeasurements of all subdetectors combined in the sum. What concerns the trigger, the measurement of \cancel{E}_T is additionally complicated by the fact that due to time constraints it is prohibitive, at least at Level 1, to do a sum over actually all detector components. In practice this means that measurements from the muon spectrometers cannot be included in the \cancel{E}_T sum at Level 1 (cf. Section 4.3.2). Currently, also at Level 2 and Event Filter no muon correction is employed to make the trigger decision, although here the muon contribution is computed and could in principle be used. Furthermore, the \cancel{E}_T measurements in general (and even more $\sum E_T$ measurements) will directly receive contributions from pile-up, which leads to a strong dependence of the trigger rates on the activity in the event and the level of pile-up.

At the Large Hadron Collider, the trigger rates are dominated by QCD processes (cf. Figure 3.6). In these processes, events with a large real \cancel{E}_T are very rare because no invisible particles like neutrinos are produced in the hard scattering process. Neutrinos may still be produced in jets originating from heavy quarks, but these neutrinos mostly carry away only comparably small energies. The largest contribution to the \cancel{E}_T measured by the detector will thus be fake \cancel{E}_T due to resolution effects in the calorimeter (cf. Section 3.2.2). As \cancel{E}_T , being the magnitude of a vector quantity, is by definition a positive quantity, fluctuations due to noise do not cancel, and \cancel{E}_T will attain a positive mean even if there is no real \cancel{E}_T .

6.1 Motivation

As was said in the introduction, the \cancel{E}_T measurement depends on a lot of different inputs. It will in particular also be directly dependent on the activity in the event. The higher the activity and the more collisions take place at the same time and the more particles are produced, the larger the sum of the energy deposited in the calorimeter systems of the detector will be. The impact of this dependence can be seen by comparing the rate of different types of triggers as function of the instantaneous luminosity $\mathcal{L}_{\text{inst}}$ and the average number of interactions^a per bunch-crossing, which is denoted as $\langle\mu\rangle$ (cf. Section 3.2.5). The trigger rate plots in this section which show rate measurements from the ATLAS trigger system have all been produced from the data stored in the COOL database, as described in Section A.2.

Figure 6.1 shows the rates of one of the Level 1 \cancel{E}_T triggers in ATLAS, from periods G – I of 2010 data taking, as a function of the instantaneous luminosity. The plot exhibits two important features:

- Unlike the rate of other triggers, for example the jet trigger for which the rates are shown in Figure 5.1, the rate of this trigger is not a linear function of the instantaneous

^a Interactions means proton-proton collisions here, happening when two bunches of protons are brought to collision.




Figure 6.1: Rate of the \cancel{E}_T trigger at Level 1 with a threshold of 20 GeV, L1_XE20, in runs from periods G – I of 2010, drawn as function of the instantaneous luminosity. The colors correspond to different numbers of colliding bunches n_{coll} , which are given in brackets together with the period. The intervals in the legend are the numbers of the runs, which have been grouped according to the number of colliding bunches.

luminosity, but depends on the luminosity with an exponent larger than one. This means that the rate of this trigger grows stronger than the instantaneous luminosity, and it therefore needs to be watched carefully in order not to exceed the bandwidth capacities of the TDAQ system. This is a major motivation for trying to achieve a better understanding of what causes this behavior.

- The rate shows a discontinuous behavior, as if parts of the curve had not been properly scaled to match the others. This does not come from prescales. These have been taken into account and corrected for when drawing the rates for the different runs. (Which is confirmed by the fact that this trigger has been prescaled for the first time in run 167607 from period I.) The steps which can be observed in this plot do not only occur at the boundaries of periods, where trigger rates might be expected to change due to changing detector conditions, but also within periods. Their origin can be explained by looking at the bunch structure of the LHC fills in the different runs.

Table 6.1 shows the number and structure of the proton bunches which are brought to collision at the interaction point in ATLAS. Given are the number of colliding bunches n_{coll} and the number of trains n_{train} in which these colliding bunches are organized (cf. Section 3.3.2). Note that these numbers only refer to the period of time in which the LHC has declared stable beams. In particular before that, in the phase of bunch injection into the collider ring, the number of bunches is lower, but these luminosity blocks are cut away by the GRL. Moreover, the trigger system is not activated (i. e. not issuing trigger signals) before stable beams have been declared. The spacing of bunches within the bunch trains is always (at least) five empty buckets between two colliding proton bunches, corresponding to a spacing of (at least) 150 ns.

Comparing the positions of the steps in the trigger rate plot in Figure 6.1 to the changes

Run no.	n_{coll}	n_{train}	Run no.	n_{coll}	n_{train}	Run no.	n_{coll}	n_{train}
165591	16	2	166097	140	18	166925	295	38
165632	47	6	166142	140	18	166927	295	38
165703	47	6	166143	140	18	166964	295	38
165732	93	12	166198	186	24	167575	295	38
165767	93	12	166305	186	24	167576	295	38
165815	93	12	166383	186	24	167607	348	46
165817	93	12	166466	233	30	167661	348	46
165818	93	12	166658	233	30	167680	348	46
165821	93	12	166786	233	30	167776	348	46
165954	140	18	166850	3	3	167844	348	46
165956	140	18	166856	295	38			
166094	140	18	166924	295	38			

Table 6.1: Number and structure of the proton bunches colliding in ATLAS in the runs from period G – I in 2010. The table lists the run number together with the number of colliding bunches n_{coll} and the number of bunch trains n_{train} in which the colliding bunches are organized.

in the bunch structure from Table 6.1 shows a correlation between the two: whenever the number of bunches is increased for a fixed instantaneous luminosity, the rate of the \cancel{E}_T trigger becomes smaller. This is made obvious by the coloring that has been used in Figure 6.1. The runs have been grouped by the number of colliding bunches in ATLAS, and the color changes whenever the number of colliding bunches changes. All the different, apparently disconnected pieces of the rate curves have different colors, i. e. correspond to different numbers of colliding bunches.

The origin of the effect is a combination of two factors: the sensitivity of the \cancel{E}_T rates to the total activity in the event and the relation between the instantaneous luminosity and the number of bunches, which is simply

$$\mathcal{L}_{\text{inst}} = \sum_{i=1}^{n_{\text{bunch}}} \mathcal{L}_{\text{inst}i}^1 \quad (6.1)$$

assuming in the following for simplicity that the instantaneous luminosity per bunch $\mathcal{L}_{\text{inst}i}^1$ is the same for all colliding bunches. From this equation, it follows that, for a larger n_{bunch} at a fixed instantaneous luminosity, the luminosity per bunch becomes smaller, and therefore the activity within this bunch crossing becomes smaller.

It should be stressed again that this behavior is typical for the \cancel{E}_T and also for the $\sum E_T$ triggers, but different for triggers like the aforementioned jet triggers, which also may have a dependence on the event activity, but this dependence is small enough to be neglected so that the jet trigger rate only depends on the total instantaneous luminosity, independent of whether this total instantaneous luminosity is generated by a few or by many colliding bunches (cf. Figure 5.1).

Figure 6.2 shows the rate of the L1_XE20 \cancel{E}_T trigger for the same periods as in Figure 6.1, but plotted as function of the average number of interactions per bunch crossing, which is related to the instantaneous luminosity by $\langle \mu \rangle \sim \mathcal{L}_{\text{inst}}/n_{\text{coll}}$ (cf. Section A.2). Indeed, after




Figure 6.2: Rate of the \cancel{E}_T trigger L1_XE20 in runs from periods G – I of 2010, drawn as function of the average number of concurrent interactions $\langle\mu\rangle$. Again the colors correspond to different numbers of colliding bunches n_{coll} which are given in brackets together with the period. The intervals in the legend are the numbers of the runs which have been grouped according to the number of colliding bunches.

rescaling with the number of colliding bunches, the curves for all runs show a consistent behavior. The conclusion from this for the model building is that, instead of regarding rates as function of instantaneous luminosity, for the \cancel{E}_T triggers the event activity, i. e. the average number of interactions per bunch crossing $\langle\mu\rangle$, is a more suitable choice because it does not depend on the number of bunches, which would otherwise have to be included as an additional parameter in the model.

6.2 Overview

In this section, a model will be presented which is aimed at helping to understand and explain the behavior of the missing transverse energy (\cancel{E}_T) and sum of transverse energy ($\sum E_T$) triggers. Although the model has been developed using data from the ATLAS experiment, it is deliberately kept very general. This allows to transfer the conclusions made here to trigger systems at other collider experiments. Besides, another strong motivation for keeping the model as simple as possible is that this simplicity allows to understand the essential features which drive the trigger behavior. The model does not seek to replace a full Monte Carlo detector simulation and cannot possibly describe the \cancel{E}_T distribution found in data in all detail. Instead, it is designed to be sufficiently precise to allow predictions of the pile-up dependence of the trigger rates, while still exhibiting the connections between the underlying ideas and the resulting impact on the rate, and to thereby complement existing Monte Carlo simulations.

The goal of the model is to describe the fake \cancel{E}_T contribution coming from the limited detector resolution and to reproduce the $\sum E_T$ and \cancel{E}_T spectra, from which the trigger rates can then be computed. A key aspect is to account for in-time pile-up to be able to predict




Figure 6.3: Distribution of EF \cancel{E}_x , \cancel{E}_y and \cancel{E}_T for minimum-bias events from run 167776 (period 2010 I) which have $53 \text{ GeV} < \sum E_T < 75 \text{ GeV}$. The solid lines are fits to the data points with a Gaussian (for \cancel{E}_x and \cancel{E}_y) or \cancel{E}_T distribution, for which the relevant parameters are given in the legend.

changes as function of the average number of interactions per bunch-crossing $\langle \mu \rangle$. Another important ingredient is that the \cancel{E}_T resolution is independent of $\langle \mu \rangle$ and only depends on $\sum E_T$, as will be shown below.

The model presented here has been developed in collaboration with the ATLAS \cancel{E}_T trigger group. The documentation of the model is also published in an ATLAS note [137].

6.3 Model Building

6.3.1 \cancel{E}_T Shape

Figure 6.3 shows the distributions of \cancel{E}_x , \cancel{E}_y and \cancel{E}_T measured at Event Filter, for a sample of events with EF $\sum E_T$ in a given range, $53 \text{ GeV} < \sum E_T < 75 \text{ GeV}$. The events are taken from run 167776 in period 2010 I after applying the GRL. They are selected by the minimum-bias trigger EF_mbMbts_1_eff and can therefore be assumed to consist mostly of events which do not contain real \cancel{E}_T from particles invisible to the detector, carrying away part of the energy unnoticed. All of the \cancel{E}_T is therefore fake \cancel{E}_T coming from the limited detector resolution. The measurement of \cancel{E}_T or its components in the transverse plane, \cancel{E}_x and \cancel{E}_y , is thus, in fact, a measurement of the detector resolution. As can be seen from the fit in Figure 6.3, the bulk

of the distribution of \cancel{E}_x and \cancel{E}_y is well described by a Gaussian^b,

$$N_0 P(\cancel{E}_x = x | \sigma_x, x_0) = \frac{N_0}{\sqrt{2\pi\sigma_x^2}} \exp\left(-\frac{(x - x_0)^2}{2\sigma_x^2}\right), \quad (6.2)$$

and correspondingly for \cancel{E}_y . The values for the width parameter $\sigma_{x,y}$ of the Gaussians, together with their uncertainties as they are obtained from the fit, are given in the legend. The normalization N_0 is not of importance here, and also the small shifts x_0 and y_0 will be neglected in the following and in the model in general because this allows to derive an analytical form of the distribution of \cancel{E}_T . The distribution of \cancel{E}_T evidently does not follow a Gaussian, but is described by

$$N_0 P(\cancel{E}_T = x | \sigma_T) = \frac{N_0}{\sigma_T^2} x \exp\left(-\frac{x^2}{2\sigma_T^2}\right). \quad (6.3)$$

The derivation of this form is given in Section A.5.2 in the Appendix, and follows from the geometrical relation $\cancel{E}_T = \sqrt{(\cancel{E}_x)^2 + (\cancel{E}_y)^2}$ under the assumption that \cancel{E}_x and \cancel{E}_y can be described by Gaussian distributions with zero mean. In particular, if $\sigma_x = \sigma_y =: \sigma$, then also $\sigma_T = \sigma$. A comparison of Equations (6.2) and (6.3) shows that even if the mean of both \cancel{E}_x and \cancel{E}_y is zero, the measurement of \cancel{E}_T will still yield a non-vanishing value leading to fake \cancel{E}_T directly correlated to the resolution.

The value of σ_T obtained from the fit in Figure 6.3 is also given in the legend. It can be seen that σ_x and σ_y do not match perfectly, but agree within less than two standard deviations. In the same spirit, also σ_T , which ought to be equal to σ_x and σ_y , does not match perfectly, but is slightly larger due to the small shift of \cancel{E}_y . Still, the overall agreement shows that the \cancel{E}_T and its components can be reasonably well described by the given forms, which shows that the underlying assumption of the dominance of fake \cancel{E}_T due to the limited detector resolution is correct.

The reason that the event sample used to produce the plot in Figure 6.3 has been selected from a narrow slice of $\sum E_T$ is that the resolution of the x and y components of the missing energy does depend on $\sum E_T$. But apart from that, the resolution is a detector parameter that is independent from pile-up, i. e. from the average number of interactions per bunch crossing. One important step for the model building is therefore to find a parametrization of the resolution of \cancel{E}_T as function of $\sum E_T$ and to measure this on data. How this can be done is demonstrated below, and the result can be seen in Figure 6.9, which also shows that the parametrization of the resolution is independent from the activity in the event due to pile-up. On the other hand, the dependence of the \cancel{E}_T resolution on $\sum E_T$ means that in order to model the \cancel{E}_T distribution, a model for the underlying $\sum E_T$ distribution is needed as well. The next step is therefore to analyse the distribution of $\sum E_T$.




Figure 6.4: Distribution of $EF \sum E_T$ in minimum-bias Monte Carlo (left) and data taken with a minimum-bias trigger (right) from an early run with low luminosity. The solid line in both plots is a fit with an exponentially falling function $f(x) = N \exp(-\lambda x)$, yielding $\lambda_{MC} = 0.0482(1)$ in the left and $\lambda_{data} = 0.0454(1)$ in the right plot.

6.3.2 $\sum E_T$ Shape

The baseline of the model of the $\sum E_T$ shape is an exponential distribution describing the energy measured in the calorimeter in events with exactly one interaction and without significant contributions of real \cancel{E}_T . The plots in Figure 6.4 motivate this assumption by showing the distribution of $\sum E_T$ measured in the Event Filter in minimum-bias Monte Carlo (left) and in data (right). The minimum-bias Monte Carlo generated with PYTHIA in the left plot is without pile-up. Only events with exactly one interaction are simulated, and no event selection is done. The data used in the right plot is triggered by a minimum-bias trigger EF_mbMbts_1 in data from early 2010, where the instantaneous luminosity was so low that only rarely more than one interaction would take place at the same time, and the contamination by pile-up events should be low. The average instantaneous luminosity in this run was $2 \cdot 10^{28}$ Hz/cm². The average number of interactions per bunch crossing averaged over the whole run was about 0.1. In addition, in the data sample only events with exactly one reconstructed primary vertex were included. Fits with exponential functions $f(x) = N \exp(-\lambda x)$ have been done in both plots with $x = \sum E_T/\text{GeV}$, yielding a slope of $\lambda_{MC} = 0.0480(1)$ for Monte Carlo and $\lambda_{data} = 0.0454(1)$ for data. (Only the range between 20 and 200 GeV has been used for the fit as indicated by the range over which the solid line is drawn.) The fitted slopes in data and Monte Carlo thus do not agree within their uncertainties, but the modelling of $\sum E_T$ in minimum-bias events is known not to be ideal, and besides there may still be a small contamination of the data sample by events with more than one concurrent interaction, leading to a slightly harder spectrum and thus a slope which is a bit smaller.

Evidently, the assumption of $\sum E_T$ following an exponential distribution is not perfect for the full range of $\sum E_T$, but it has the important advantage that this distribution can be

^b In the following, a lot of probability density functions for quantities like \cancel{E}_T and $\sum E_T$ will be given, often decorated with additional Poisson distributions. A compromise between an abstract and an explicit notation has been found in that it was decided to abandon a lean notation in favor of a more explicit notation, which will hopefully allow for a fluent reading. The expression $P(Q = x|\vec{a})$ should be understood to denote the density of probabilities of the values x which can be attained by a physical quantity Q (e. g. $\sum E_T$ or \cancel{E}_T) for a given set of parameters \vec{a} .




Figure 6.5: Distribution of $EF \sum E_T$ (left) and offline $\sum E_T$ (right, Topo calibration) in 2011 data (period 2011 D) taken with a random trigger on empty bunch crossings with no proton-proton collisions.

treated analytically and still adequately reflects the overall behavior of the spectrum. Two deviations can be seen: For values below approximately 20 GeV, the spectrum of $\sum E_T$ drops very steeply. This is not an effect of the trigger being a minimum-bias trigger, but also found using a random trigger. It is a result of the one-sided cut applied to the measurement of cell energies in the Event Filter (cf. Section 4.3.2) and discussed below. For high values, in the tail of the distribution, the slope gets steeper. As this effect can be seen both in data and in Monte Carlo which has exactly one simulated minimum-bias interaction per event, it is unlikely to be caused by a mixture of events with different slopes due to different numbers of concurrent interactions.

Figure 6.5 shows in the left plot the $\sum E_T$ spectrum computed online in the trigger at Event Filter level and in the right plot the offline \cancel{E}_T obtained from the Topo calibration (cf. Section 4.4.3). Both plots use data from period 2011 D, where the minimum-bias triggers have been replaced by random triggers. The events have been selected using the random trigger `EF_rd0_empty_NoAlg`, which triggers on empty bunch crossings^c and the probability is thus very high that no interaction has taken place in these events. The comparison between online and offline $\sum E_T$ shows the impact of the one-sided noise cuts which are implemented in the trigger at Event Filter level. They lead to a bias which appears as a shift of the mean of the Gaussian distribution^d of the online $\sum E_T$ measurements from zero to about 18 GeV. This one-sided cut is unique to the trigger and cannot be seen in the offline $\sum E_T$ spectrum, which peaks at zero as would be expected.

The shift arising from this measurement bias is included in the model by subtracting it from the $\sum E_T$ values which enter the fits when extracting the parametrization. The argument of the analytic form of the $\sum E_T$ distribution, cf. Equation (6.4), is replaced by $x - x_0$,

^c The term empty bunch crossing refers to LHC time buckets (cf. Section 3.3.2) into which no proton bunch has been injected in this particular LHC fill. They therefore contain no (or at least only negligibly few) protons so that no proton-proton interactions are expected when these “empty bunches” cross.

^d It turns out that using a log-normal distribution gives an even better fit. This can be motivated by applying the central limit theorem not directly to the sum of energy measurements over a large number of calorimeter cells, taking the measurements as independent and identically distributed random variables (i.i.d.), but taking the order of magnitude of the measurements as i.i.d. random variables, such that the logarithm of the measurements is i.i.d., i. e. $P_0(\sum E_T = x | x_0, \sigma) = \frac{1}{x\sigma\sqrt{2\pi}} \exp\left(-\frac{(\ln x - \ln x_0)^2}{2\sigma^2}\right)$, $x > 0$.

where $x_0 = 11.5 \text{ GeV}$ for 2010 and $x_0 = 17.6 \text{ GeV}$ for 2011 is the mean of a fit to the peak of $\sum E_T$ in empty events. That the $\sum E_T$ in empty events in 2011 is higher than in 2010 can be attributed to the higher instantaneous luminosity, leading to an overall higher activity in the detector. This higher activity can be transferred to empty bunch crossings either by instrumental effects, e. g. due to the lag in the response to energy depositions in the calorimeter, i. e. the time the detector needs to reach its zero state again, or by physical effects, e. g. higher remaining activity from previous collisions and, in general, a higher background level in the cavern. Note that subtracting the full shift probably tends to be an overestimation of the pedestal, as part of the sources of $\sum E_T$ contributing to the shift in empty events will also contribute to $\sum E_T$ in collision events that the model seeks to describe. This would be interesting to study in more detail. However, it would require to develop a method to discriminate between the different contributions of fake \cancel{E}_T .

As said above, the benefit of using an exponential distribution to model the shape of $\sum E_T$ is that the resulting shape of $\sum E_T$ for μ concurrent interactions under these assumptions can be derived analytically. This derivation is shown in Section A.5.1 and yields (cf. Equation (A.21))

$$P\left(\sum E_T = x|\mu\right) = \frac{\lambda^\mu}{(\mu-1)!} x^{\mu-1} \exp(-\lambda x). \quad (6.4)$$

The domain of this probability distribution is restricted to non-negative values. In the end, a model for $\sum E_T$ as function of the average number of interactions per bunch crossing $\langle\mu\rangle$ is required. The distribution of $\sum E_T$ for a given $\langle\mu\rangle$ is simply the Poisson-weighted sum of the distributions of $\sum E_T$ for a fixed number of interactions per bunch crossing given by Equation (6.4),

$$P\left(\sum E_T = x|\langle\mu\rangle\right) = \mathcal{N}_0 \sum_{\mu=1}^{\infty} P\left(\sum E_T = x|\mu\right) \text{Pois}(\mu|\langle\mu\rangle). \quad (6.5)$$

Note that in the sum the term $\mu = 0$ has been left out, although it has non-zero weight in the Poisson distribution. The distribution therefore needs an additional normalization factor \mathcal{N}_0 . This is done because the model does not include events without any interaction, which will contribute only a negligible amount of energy in the calorimeters anyway. This can be seen in the left plot in Figure 6.5, where it should be remembered that the actual sum of the energy depositions needs to be corrected for the shift from the measurement bias. Concerning the $\sum E_T$ triggers, even the lowest $\sum E_T$ thresholds are of the order of a few hundred GeV at Event Filter and therefore lie safely above the peak stemming from the events with no interactions. The normalization factor \mathcal{N}_0 is simply given by the Poisson weight of the $\mu = 0$ bin for the given value of the model parameter $\langle\mu\rangle$,

$$\mathcal{N}_0 = \frac{1}{\sum_{\mu=1}^{\infty} \text{Pois}(\mu|\langle\mu\rangle)} = \frac{1}{1 - \text{Pois}(0|\langle\mu\rangle)}. \quad (6.6)$$

Equation (6.5), giving the shape of $\sum E_T$ with $\langle\mu\rangle$ as parameter, completes the model. Using the distribution of \cancel{E}_T as described by Equation (6.3) and the knowledge that the resolution σ_T is dependent on $\sum E_T$, the complete analytic form for the \cancel{E}_T distribution can now be

written down:

$$P(\cancel{E}_T = x|\langle\mu\rangle) = \frac{1}{\sigma_T(\sum E_T)^2} x \exp\left(-\frac{x^2}{2\sigma_T(\sum E_T)^2}\right) \quad (6.7)$$

$$= \int_0^\infty \frac{1}{\sigma_T(y)^2} x \exp\left(-\frac{x^2}{2\sigma_T(y)^2}\right) P(\sum E_T = y|\langle\mu\rangle) dy \quad (6.8)$$

$$= \mathcal{N}_0 \sum_{\mu=1}^\infty \text{Pois}(\mu|\langle\mu\rangle) \int_0^\infty \frac{P(\sum E_T = y|\mu)}{\sigma_T(y)^2} x \exp\left(-\frac{x^2}{2\sigma_T(y)^2}\right) dy, \quad (6.9)$$

with $P(\sum E_T = y|\mu)$ given by Equation (6.4).

6.3.3 Trigger Rates

For any modelled trigger quantity Ω the average probability P_{ev} of an event to issue a trigger is the probability of the trigger measurement to exceed the given trigger threshold Ω_i ,

$$P_{\text{ev}}(\Omega_i) = \int_{\Omega_i}^\infty P(\Omega = x) dx, \quad (6.10)$$

not writing dependencies on additional parameters explicitly. Note that this only holds for one trigger level and does not take correlations of the measurements at different levels into account. To compute the absolute trigger rate $R(\Omega_i)$, this probability needs to be multiplied with the total cross section σ_{tot} and the instantaneous luminosity $\mathcal{L}_{\text{inst}}$,

$$R(\Omega_i) = P_{\text{ev}}(\Omega_i) \sigma_{\text{tot}} \mathcal{L}_{\text{inst}}. \quad (6.11)$$

This normalization is difficult to obtain, but when computing relative rates, assuming without loss of generality $\Omega_i < \Omega_j$,

$$\frac{R(\Omega_i)}{R(\Omega_j)} = \frac{P_{\text{ev}}(\Omega_i)}{P_{\text{ev}}(\Omega_j)} = \frac{\int_{\Omega_i}^\infty P(\Omega = x) dx}{\int_{\Omega_j}^\infty P(\Omega = x) dx} = 1 + \frac{\int_{\Omega_i}^{\Omega_j} P(\Omega = x) dx}{\int_{\Omega_j}^\infty P(\Omega = x) dx}, \quad (6.12)$$

it cancels out. In this context it should be stressed that the low tail of the trigger measurements, below the lowest trigger threshold, neither affects the relative nor the absolute trigger rates, and thus discrepancies between the model and the real spectrum in this range do not deteriorate the quality of the model predictions.

6.3.4 Summary

The full model for the prediction of the \cancel{E}_T trigger rates for a given value of the average number of interactions per bunch crossing $\langle\mu\rangle$ is thus comprised of three steps:

1. Compute the $\sum E_T$ spectrum. This step uses $\langle\mu\rangle$ and the slope of the exponential distribution of the energy found for events with one proton-proton interaction and the $\sum E_T$ offset as input, and employs Equation (6.5) to obtain the distribution of $\sum E_T$. The

assumptions used in this step are that the $\sum E_T$ for one interaction per bunch crossing, $\mu = 1$, can be described by an exponential distribution, and that $\sum E_T$ from μ concurrent interactions can be described as an overlay of independent $\sum E_T$ contributions from μ individual interactions^e.

2. Compute the \cancel{E}_T spectrum. This step uses the parametrization of the detector resolution as a function of $\sum E_T$ and the $\sum E_T$ distribution obtained in the previous step as input. It employs Equation (6.8) to obtain the distribution of \cancel{E}_T . The assumptions used in this step are that the detector resolution is independent of $\langle\mu\rangle$, which will be shown for ATLAS in Section 6.4.3, and that all \cancel{E}_T is fake \cancel{E}_T coming from the limited detector resolution. Note that the model assumptions also imply that the resolution parameter for \cancel{E}_x , \cancel{E}_y and \cancel{E}_T is the same.
3. Compute the trigger rates. This step uses the \cancel{E}_T and the $\sum E_T$ spectrum from the previous steps as inputs, together with a set of trigger thresholds as parameters. The trigger rate is then given by the integral over the spectrum above the given threshold, as in Equations (6.11) and (6.12). Here, only relative trigger rates are computed.

The number of input parameters of the model is, following the guiding principle of simplicity, very small: There are the slope λ of $\sum E_T$ for events with a single interaction, the offset of $\sum E_T$ estimated from the $\sum E_T$ distribution in empty bunch crossings and two parameters describing the resolution of \cancel{E}_T as function of $\sum E_T$. The target average number of concurrent interactions $\langle\mu\rangle$ is an external parameter. In the following, the extraction of the input parameters from data is described. Afterwards, the model is used to produce trigger rate predictions.

6.3.5 Implementation

The actual implementation of the model closely follows the summary given above. Some care needs to be taken when choosing the cut-off for the integrals and summations which have upper limits sent to infinity. The function evaluating Equation (6.5) for given $\sum E_T$ includes a loop over μ . This loop runs from 1 to $\min(\{20, \langle\mu\rangle + 5\sqrt{\langle\mu\rangle}\})$, where the last term comes from an approximation of the Poisson distribution by a Gaussian, which is cut off at 5 standard deviations to the positive side of its mean. The $\sum E_T$ distribution is obtained by filling a histogram, calling this function once for every bin. The \cancel{E}_T distribution is then obtained by looping over the previously filled histogram holding the $\sum E_T$ distribution, and for each bin computing the corresponding resolution of \cancel{E}_T . A histogram with the \cancel{E}_T distribution is filled by summing up the contributions from all $\sum E_T$ bins with the appropriate weight, which is given by the content of the respective $\sum E_T$ bin. It is thus important to choose the range of the $\sum E_T$ histogram suitably because in particular the tail of the \cancel{E}_T distribution will receive significant contributions from relatively high $\sum E_T$ values. A fixed value of 2 TeV has been used here. This could be improved by dynamically adjusting the upper bounds automatically during the computation as needed.

^e That this part of the computation can be done in a first independent step is somewhat concealed in Equation (6.9), but clear from the fact that the model $\sum E_T$ shape does not depend on \cancel{E}_T .

6.4 Determination of Input Parameters

One of the two important inputs to the model which need to be measured on data is the slope of the exponential distribution of the sum of the transverse energy in events which have one proton-proton interaction. As the actual number of interactions which have taken place in a data event is unknown, this number has to be estimated from some suitable quantity which can be reconstructed. The average number of interactions per bunch-crossing $\langle\mu\rangle$ is the expectation value of a Poisson distribution and as such only meaningful for a sample of events. Currently, the best per-event estimate for the number of concurrent proton-proton interactions is the number of reconstructed primary vertices. The reconstruction of primary vertices has an intrinsic efficiency $\epsilon_{\text{vx, reco}} \leq 1$, so that the number of reconstructed primary vertices $n_{\text{vx, reco}}$ is expected to be smaller than $\langle\mu\rangle$. However, assuming that the reconstruction efficiency is independent of the number of primary vertices, it is easy to show that the distribution of $n_{\text{vx, reco}}$ for a sample of many events also follows a Poisson distribution and that its expectation value $\langle n_{\text{vx, reco}} \rangle$ is proportional to $\langle\mu\rangle$, with a coefficient given by $\epsilon_{\text{vx, reco}}$. The quality of this estimation will be discussed in the first part of this section.

6.4.1 Number of Concurrent Collisions

The primary vertices are reconstructed by finding intersections of tracks reconstructed in the Inner Detector and applying a suitable set of cuts (cf. Section 4.4.4). In physics analyses, often a cut is applied on the number of tracks associated to a primary vertex. For instance, in the analysis presented in Section 7 at least 5 tracks are required to select only events which have a reliably reconstructed primary vertex. The algorithm which creates the collection of reconstructed primary vertices is written in such a way that one dummy vertex with no associated tracks is always included in addition to the reconstructed primary vertices. This has to be taken into account when no cut on the number of associated tracks is made and can be seen in the following figure.

The plots in Figure 6.6 shows the correlation of the number of reconstructed primary vertices and the true number of interactions for MC09 Monte Carlo samples which contain minimum-bias events generated with PYTHIA and have an average number of interactions per bunch crossing of $\langle\mu\rangle = 5 + 1^{\ddagger}$. The four plots differ in the way the number of reconstructed primary vertices is counted: Only reconstructed primary vertices which have at least 0 (upper left), 2 (upper right), 3 (lower left) or 8 (lower right) associated tracks (n_{track}) are counted. Comparing the two plots in the upper row makes obvious that requiring two tracks only removes the dummy vertex with no associated tracks. Apart from that, the distribution stays the same. A large fraction (54% for $n_{\text{track}} \geq 2$) of events lies on the diagonal marked by the black line. Here, the reconstructed number of primary vertices corresponds to the true number of interactions. For most of the rest of the events, the number of primary vertices is smaller than the true number of interactions, mostly by one, which comes from the inefficiency of the reconstruction algorithm for primary vertices. In contrast to this, an overestimation of the true number of interactions is very rare because this could only be

[‡] Writing the average number of interactions per bunch crossing as $\langle\mu\rangle = n + 1$ shall indicate that every event contains one interaction according to the physics process which defines the Monte Carlo sample, and on top several minimum-bias interactions from pile-up. The number of additional interactions is Poisson distributed with an expectation value which here is 5. Note that the resulting total number of interactions per bunch crossing is different from a Poisson distribution with an expectation value of 6.




Figure 6.6: Two-dimensional correlation plots of the number of reconstructed primary vertices ($n_{pv, reco}$, vertical axis), counting only vertices with a given minimum number of associated tracks, and the true number of interactions ($n_{pv, true}$, horizontal axis) for Monte Carlo minimum-bias events. Four criteria have been used to count the number of reconstructed primary vertices, requiring at least 0, 2, 3 or 8 associated tracks. The resulting two plots in the top row are identical, the left one being shifted up by 1.




Figure 6.7: Distribution of the number of reconstructed primary vertices $n_{\text{pv, reco}}$ in data from 2010 for three small event samples with a narrow distribution of $\langle\mu\rangle$. The data points are fitted with a Poisson distribution (solid lines) with free rate parameter μ and the fit result $\langle\mu\rangle_{\text{fit}}$ is given in the legend together with the ratio of $\langle\mu\rangle_{\text{fit}}$ over the nominal $\langle\mu\rangle$ expressed as a percentage value in brackets. The black points have been scaled up by a factor 10. (A corresponding plot for 2011 is shown in Figure A.3 in the Appendix.)

due to the accidental crossing of two or more tracks. For higher numbers of true primary vertices, the distribution around the diagonal gets broader as can be expected. Comparing the plots with $n_{\text{track}} \geq 2$ and $n_{\text{track}} \geq 3$ shows that already here requiring more tracks starts to make the agreement between the reconstructed number of primary vertices and the true number of interactions worse. For $n_{\text{track}} \geq 3$, the number of events on the diagonal is about 49%. The number of interactions is underestimated from the number of reconstructed primary vertices. Going to $n_{\text{track}} \geq 8$ makes this even worse, of course. The conclusion is that in the following, in order to obtain an estimate of the number of interactions per event, either no selection based on the number of tracks is done, keeping in mind to subtract one from the number of reconstructed primary vertices to account for the dummy vertex, or, equivalently as shown above, only reconstructed primary vertices with at least two associated tracks are counted.

From the plots on Monte Carlo in Figure 6.6, it can be seen that the number of reconstructed vertices tends to underestimate the number of interactions. It is therefore natural to expect this also to be the case for data, and indeed this effect can also be made visible on data as will be discussed now. The number of interactions per event in data is unknown and therefore cannot be compared directly to the number of reconstructed vertices, but for ensembles of events the average number of interactions can be estimated from the instantaneous luminosity, which is available per luminosity block. Again, this number is taken from the COOL database, as was done to produce the trigger rate plots. For events from a small set of luminosity blocks over which the average number of interactions $\langle\mu\rangle$ does not vary too much, the distribution of the number of reconstructed primary vertices is plotted in Figure 6.7. Three different sets of luminosity blocks from run 167776 in period 2010 I have been used, corresponding to an $\langle\mu\rangle$ of approximately 1.6 (luminosity blocks 541–546), 2.2

(320–326) and 3.3 (120–125), respectively. The relative spread of the nominal value of $\langle\mu\rangle$ over the selected luminosity blocks is smaller than half a percent. The events have been selected using the minimum-bias trigger `EF_mbMbts_1_eff`, and the SUSY GRL has been applied. A Poisson distribution has been fitted to the three distributions, and the fit results for the expectation value of the Poisson distribution, $\langle\mu\rangle_{\text{fit}}$, are listed in the legend. For all three sets, $\langle\mu\rangle_{\text{fit}}$ is about one quarter smaller than the expected average number of interactions $\langle\mu\rangle$. The number of interactions is thus again underestimated by the reconstructed number of primary vertices, due to the inefficiency in the vertex reconstruction. The underestimation tends to become larger when going to higher numbers of concurrent interactions. This can be explained as a result of vertex merging, which occurs when in busier events with a higher number of tracks the probability increases that two primary vertices are merged and mistaken to be the same vertex. Figure 6.7 suggests that correction factors can be found to improve the estimator of the number of concurrent interactions. However, this only allows to correct the average number of interactions per bunch crossing $\langle\mu\rangle$ for event samples, but does not help with the selection of events with a fixed number value of μ , for which a reliable event-by-event estimate of μ rather than $\langle\mu\rangle$ is needed.

6.4.2 $\sum E_T$ Slope

Coming back now to the original question, the extraction of the slope of $\sum E_T$ for events with one proton-proton interaction, the conclusion from the above is that it is not easy to reliably determine the number of interactions that have taken place in a given event. When selecting a sample of events with exactly one interaction, a major problem will be the contamination of this sample by events in which actually more than one interaction has taken place, but one or more of the primary vertices have been missed by the reconstruction algorithm. This will lead to an underestimation of the slope because events with a higher number of interactions have a harder $\sum E_T$ spectrum and thus enhance the tail of the distribution, leading to a seemingly smaller slope.

The solution to this is not to try and fit the $\sum E_T$ distribution for exactly one interaction, but instead to use the full model for $\sum E_T$, as given in Equation (6.5), and to fit $\sum E_T$ in a sample of events in which the average number of interactions follows, as usual, a Poisson distribution. The data for this fit has been taken with the random trigger `EF_rd0_filled_NoAlg` in data from period 2011 D. For each of the runs used to produce the plot, only data from luminosity blocks was used which have an average number of interactions per bunch crossing $\langle\mu\rangle$ around 5.0 (selecting events with $4.9 < \langle\mu\rangle < 5.1$), where $\langle\mu\rangle$ is taken from the COOL database. This selection is needed to ensure that the distribution of $\langle\mu\rangle$ in the event sample closely follows a Poisson distribution as is assumed in the fit.

The plot in Figure 6.8 shows a comparison of several distributions of Event Filter $\sum E_T$ on different subsamples. The data points in red and gray correspond to subsamples in which the number of reconstructed vertices $n_{\text{pv, reco}}$ is 1 or 2 up to 7, respectively. The red points with $n_{\text{pv, reco}} = 1$ make obvious that selecting one reconstructed vertex no longer gives a pure sample of events with one concurrent interaction, as was the case in Figure 6.4, because it is no longer an exponential distribution, which would follow a straight line in logarithmic scaling. The red and gray data points are fitted using the simple form of $\sum E_T$ in Equation (6.4), where the slope λ and the parameter μ are left free in the fit. This allows the fitted curves to closely follow the data points, but the good agreement is misleading because the fit yields values for λ which vary strongly with $n_{\text{pv, reco}}$. This contradicts the model assumption




Figure 6.8: Comparison of fits to the Event Filter $\sum E_T$ spectrum in events from 2011 data with $\langle\mu\rangle \approx 5.0$. Shown are the distributions from selecting subsamples by the number of reconstructed vertices $n_{\text{pv, reco}}$ being 1 (red) and 2 through 7 (gray) as well as the full sample (blue). The data points represent the distributions in data, the curves are fits which are explained in the text. The intervals over which the lines are drawn indicate the fit ranges.

of $\sum E_T$ being the result of an overlay of μ independent spectra with the same fixed slope λ for all values of μ . This is, of course, again a result of the contamination of the samples by events with μ larger than the selected $n_{\text{pv, reco}}$, which is also reflected in the fact that the value of $\langle\mu\rangle_{\text{fit}}$ obtained from the fit is larger than $n_{\text{pv, reco}}$ for small $n_{\text{pv, reco}}$, but less so for high $n_{\text{pv, reco}}$.

All this confirms that it is not possible to extract the slope λ from a subsample with given $n_{\text{pv, reco}}$. In contrast, the blue data points in Figure 6.8 show the distribution of $\sum E_T$ in the full sample, with no selection on $n_{\text{pv, reco}}$ and fitted with the full $\sum E_T$ model from Equation (6.5). In the fit, both $\langle\mu\rangle$ and the slope λ of $\sum E_T$ for $\mu = 1$ are left free. The fit yields $\langle\mu\rangle_{\text{fit}} = 4.96(5)$ and a slope of $\lambda_{\text{fit}} = 0.0410(4)$, which is compatible with the selected $\langle\mu\rangle$ and reasonably close to the slope obtained above for data from the fit in Figure 6.4. The same procedure has been carried out on data from 2010 taken with a minimum-bias trigger. The result is shown in the Appendix in Figure A.4 and is in very good agreement with the value from 2011 data. This approach is therefore found suitable to determine the slope used as input to the model for the shape and rate predictions.

6.4.3 Missing Energy Resolution

The other important input to the model is the resolution of the \cancel{E}_T measurement in the detector. Figure 6.9 shows the resolution of \cancel{E}_x and \cancel{E}_y measured by the trigger at Event Filter level as function of Event Filter $\sum E_T$ for run 167776 from period 2010 I. The resolution is computed by doing Gaussian fits of the EF \cancel{E}_x and \cancel{E}_y distribution for small slices of $\sum E_T$ as shown in Figure 6.3. The fit result for σ is used as resolution parameter. Dots in three different colors mark three event samples, which differ in the average numbers of concurrent interactions $\langle\mu\rangle$, as shown in the legend of the plot. The error bars on the dots are the




Figure 6.9: Resolution of EF \cancel{E}_x (left) and \cancel{E}_y (right) plotted as function of EF $\sum E_T$ for run 167776 from period 2010 I. Three event samples with different average numbers of concurrent interactions $\langle\mu\rangle$ are compared. The black dotted line is a fit with a square root function.

Resolution [GeV]	a	b
$\sigma_x(\sum E_T/\text{GeV})$	-0.584 ± 0.008	0.421 ± 0.002
$\sigma_y(\sum E_T/\text{GeV})$	-0.658 ± 0.008	0.425 ± 0.002

Table 6.2: Parametrization of the resolution of \cancel{E}_x and \cancel{E}_y , denoted σ_x and σ_y , as function of $\sum E_T$. The parameter values are obtained from the square root fit $f(x) = a + b\sqrt{x}$ shown as the black dotted line in the resolution plot in Figure 6.9 as function of $x = \sum E_T/\text{GeV}$.

uncertainties on the resolution parameter obtained from the Gaussian fit. The comparison between the three samples shows that the resolution in very good approximation is independent from $\langle\mu\rangle$. It can therefore be assumed to be a detector parameter that is independent of the pile-up level.

To derive a parametrization which can later be used as input to the model, a square root function $f(x) = a + b\sqrt{x}$ is fitted to the combined set of dots from all three event samples. The functional form of the fit is inspired by the general parametrization of the energy resolution of the calorimeter in Equation (3.4). The parameters obtained from the fit marked by the black dotted line in the figure are given in Table 6.2. Concerning the seemingly large value of the constant term a , it should be noted that the shape of the curve is dominated by the square root term with prefactor b , as is obvious from the plot. The fitted function gives negative values for the resolution for small values of $\sum E_T$ below about 2 GeV. These negative values for the resolution are unphysical, of course, but will not affect the model as such small $\sum E_T$ values never occur in the model due to the inclusion of the pedestal of $\sum E_T$ for empty events as described above. What is furthermore apparent from the values for b is that σ_y is slightly larger than σ_x for large $\sum E_T$. This is a result of the shift of the mean of the y -component of the missing energy. In the model, the mean of the fit parameters obtained for σ_x and σ_y is used for the parametrization of the \cancel{E}_T resolution because both \cancel{E}_x and \cancel{E}_y are assumed to have the same resolution.

Figure 6.10 shows the distributions of the components of the missing energy along the x - and y -direction as well as its transverse component, for the same slice of $\sum E_T$ as in Figure 6.3,




Figure 6.10: Distribution of EF \cancel{E}_x , \cancel{E}_y and \cancel{E}_T for events from period 2011 D which have $53 \text{ GeV} < \sum E_T < 75 \text{ GeV}$, taken with a random trigger. The solid lines are fits to the data points with a Gaussian (for \cancel{E}_x and \cancel{E}_y) or \cancel{E}_T distribution, for which the relevant parameters are given in the legend.

but for data taken with a random trigger in period D from 2011 instead of 2010. All three distributions are fitted with the appropriate functions. The fact that the same resolution parameter σ_T is found as in the fit from 2010 confirms the assumption that the resolution parameters from 2010 can also be applied in 2011.

6.5 Model Predictions

In this section, the model will be applied and compared to data taken in 2010 and 2011. The data taken in each of these two years will be treated separately in the following two sections, due to considerable changes in the trigger and the instantaneous luminosity. In addition, within 2011 there have been further changes in the trigger, so that for 2011 only data up to and including period 2011 F will be used.

6.5.1 Performance in 2010

Shape Predictions

Figure 6.11 compares the predictions of the model for Event Filter $\sum E_T$ and \cancel{E}_T to the respective spectrum measured in data from 2010. The points represent data which has been collected from several runs in periods F – I, using the minimum-bias trigger EF_mbMbits_1_eff[§] after applying the GRL defined by the ATLAS Supersymmetry group. The range of luminosity blocks contributing to the plot has been chosen for each run such that $\langle \mu \rangle_{\text{LB}}$, com-

[§] In 2010, the random triggers were prescaled so heavily that they cannot be used for these studies. On the other hand, in 2011 the minimum-bias triggers are prescaled heavily because due to the high activity they would probably trigger on any bunch crossing anyway, and thus they have been replaced by purely random triggers. Therefore, in 2010 a minimum-bias trigger is used, whereas in 2011 a random trigger is used.




Figure 6.11: Comparison of the distribution of Event Filter ΣE_T (top) and E_T (bottom), from the model prediction (lines) and 2010 data (points) for several values of $\langle\mu\rangle$. The normalization of the data points is chosen such that the tails are clearly visible.




Figure 6.12: Comparison of predicted and observed rates for missing transverse energy triggers with different thresholds (15, 20, ..., 35 GeV) in 2010, as function of the average number of interactions per bunch crossing $\langle\mu\rangle$. The filled circles are the model predictions, the clouds of small dots represent the trigger rate measurements.

puted from the instantaneous luminosity for a given luminosity block, lies within a range $|\langle\mu\rangle_{\text{LB}} - \langle\mu\rangle| < 0.1$, with $\langle\mu\rangle$ ranging between 1 and 3.5 in steps of 0.5 for the six sets of data points in the plot. This covers most of the range of $\langle\mu\rangle$ values that is available in 2010 data, the peak value being 3.8.

The curves in the plot give the model predictions for the respective value of $\langle\mu\rangle$. They have been generated using as input the detector resolution measured in 2010 data and the slope of $\sum E_T$ for one concurrent interaction obtained as described above^h. The normalization of the data points and curves has been chosen such that the tails of the distributions are clearly visible, by shifting the distributions so that they do not cross. Comparing the Event Filter $\sum E_T$ between data and model, the distributions in data fit much better the predictions of the model for values of $\langle\mu\rangle$ which are smaller by about 0.5, in particular for $\langle\mu\rangle \geq 2.0$. Indeed, this effect can also be captured by fitting the model $\sum E_T$ function to the measured $\sum E_T$ distribution in data (cf. Figure A.3 in the Appendix) and seems to affect primarily data from 2010, where $\langle\mu\rangle$ is in general smaller than in 2011. Normalizing data and the model prediction in the $\sum E_T$ range above a given threshold of e.g. 50 GeV gives a better agreement, confirming that it is especially difficult to model the low part of the $\sum E_T$ distribution. The \cancel{E}_T distribution is stable against this effect and the agreement between data and model prediction is satisfactory for all $\langle\mu\rangle$, as can be seen in the lower plot in Figure 6.11.

Rate Predictions

The spectra of Event Filter \cancel{E}_T in Figure 6.11 for different values of $\langle\mu\rangle$ show good agreement between prediction and observation. From the predicted spectra, the \cancel{E}_T trigger rates can

^h Note that for the fractional values of μ , the same analytic form from Equation (6.5) is used for $\sum E_T$, extending (without proof) its definition from integer μ to fractional μ by using the Gamma function $\Gamma(n) = (n-1)!$, which generalizes the factorial function to real (and complex) arguments n .

now be computed by integration according to Equation (6.12). As the model cannot predict absolute rates, but only relative rates, the rate predictions are normalized to the measured trigger rate of the lowest \cancel{E}_T trigger. The comparison of predicted and measured rates is shown in Figure 6.12. The clouds of points represent the measured trigger rates, which are obtained from the COOL database and normalized to the number of colliding bunches, as described in detail in Section 6.1. To reduce the spread, each point is an average over two luminosity blocks in this plot. The predicted rates are computed for integer values of $\langle\mu\rangle$ and marked by filled circles in the color matching the corresponding measured rates. The green points and markers correspond to the lowest threshold of 15 GeV, and the measured rates for $\langle\mu\rangle = 2$ for this threshold are used to normalize the model predictions.

Note that the prediction of the rate from the \cancel{E}_T spectrum measured by the Event Filter only encompasses one level of the three-level trigger system of ATLAS. A simulation of more than one level is not possible with the current approach because correlations between the measurements at the different levels are not included in the model. However, the prediction works as long as the overlap between the different levels is small enough. To put it the other way round, if the spacing between Level 2 and Event Filter is large enough, the Event Filter sees an undistorted spectrum of \cancel{E}_T , and therefore its trigger rate is consistent with the (properly normalized) integral of the \cancel{E}_T spectrum above the Event Filter \cancel{E}_T threshold.

It is no surprise that the rates for the lowest threshold of 15 GeV agree well with the prediction, as this is the rate which the prediction is normalized to. Note, however, that this normalization only uses the point with $\langle\mu\rangle = 2$. The agreement of the predictions for $\langle\mu\rangle = 1$ and 3 with the measured rates is not automatic, but shows that the dependence of the rates on $\langle\mu\rangle$ is correctly described by the model, at least over the small range where rate measurements are available in 2010. On the other hand, the plot also shows that already for the next higher threshold of 20 GeV, the prediction underestimates the measured rate by a factor of about 2, and for higher thresholds the discrepancy gets larger. Thus, even the spacing of the thresholds in terms of rates is apparently not well described by the model. The reason for the large discrepancies for higher thresholds lies in the composition of \cancel{E}_T , which arises from several sources. Only one of these, namely fake \cancel{E}_T from the limited detector resolution, is included in the model. This source of \cancel{E}_T is no longer dominant for high trigger thresholds, which leads to an underestimation of the measured trigger rates by the model prediction. This is explained in more detail below, but first the results for 2011 are presented.

6.5.2 Performance in 2011

One of the noticeable differences between data taken in 2010 and 2011 is that, as part of the progression in the instantaneous luminosity, in 2011 also the bunch spacing was reduced, going from 150 ns in 2010 I down to 50 ns in 2011 D, so that out-of-time pile-up possibly can no longer be neglected¹. It is not *a priori* clear that the model would still give sensible results. As will become clear from the plots in this section, however, the agreement for the shapes in 2011 is even better than in 2010.




Figure 6.13: Comparison of the distribution of Event Filter ΣE_T (top) and \cancel{E}_T (bottom), from the model prediction (lines) and 2011 data (points) for several values of $\langle\mu\rangle$. The normalization of the data points is chosen such that the tails are clearly visible.




Figure 6.14: Comparison of predicted and observed rates for missing transverse energy triggers with different thresholds (20, 30, 40, 50 GeV) in 2011, as function of the average number of interactions per bunch crossing $\langle\mu\rangle$. The filled circles are the model predictions, the clouds of small dots represent the trigger rates measurements.

Shape Predictions

In Figure 6.13, the predictions of the model and the spectra measured in data from period 2011 D for the Event Filter $\sum E_T$ and \cancel{E}_T are shown. The data has been collected from several runs, this time using a random trigger, EF_rd0_filled_NoAlg, and applying the SUSY GRL. Again, different subsamples are shown, which are selected by requiring $\langle\mu\rangle_{\text{LB}}$ to lie within a range $|\langle\mu\rangle_{\text{LB}} - \langle\mu\rangle| < 0.1$, with $\langle\mu\rangle \in \{4, 5, 6, 7\}$ for the four different sets of data points. The four curves compare the model predictions for the respective value of $\langle\mu\rangle$ to the observed spectra from data, which are plotted as data points. The model predictions use the detector resolution from 2010 as input, and the slope of $\sum E_T$ for one concurrent interaction, which has been extracted from the fit of the full model to the $\sum E_T$ spectrum for $\langle\mu\rangle \sim 5$. The normalization of the data points and curves has been chosen such that the tails of the distributions are clearly visible.

It can be seen from these plots that the model gives a very good description for both \cancel{E}_T and $\sum E_T$. The good description of the $\sum E_T$ spectrum for $\langle\mu\rangle = 5$ may be attributed to the fact that the slope λ which is used as input to the model was extracted from this subsample of events, but still the value of $\langle\mu\rangle$ was left free in the fit from which λ was taken, and is set to 5 here, so the good agreement is not completely artificial. The other three $\sum E_T$ distributions and all distributions for \cancel{E}_T follow from Equations (6.5) and (6.9) by adjusting only $\langle\mu\rangle$.

Rate Predictions

Figure 6.14 shows the predictions of the missing transverse energy trigger rates compared to the rates measured in 2011, periods 2011 D – F. Runs from later periods have higher $\langle\mu\rangle$

ⁱ The ATLAS Technical Design Report (TDR) states a peaking time of the bipolar shapers in the read-out of the liquid-argon Calorimeters of approximately 35 ns [129]. In fact, a (small) destructive interference due to out-of-time pile-up can be expected even for large bunch spacings up to 500 ns [21].

(almost up to 10, cf. e. g. Figure 6.18), but cannot be compared to previous periods nor to the model predictions due to changes in the trigger which are discussed below. Like for 2010, also here the predicted rates are normalized to the lowest threshold, which is 20 GeV for 2011. Again, the dependence on $\langle\mu\rangle$ of the rate of this trigger threshold seems to be well described by the model, but unfortunately less so for the next higher threshold, which exhibits a slower increase of its rate with increasing $\langle\mu\rangle$ than is predicted by the model. Note that here, the spacing of the two lowest thresholds is twice as large as in 2010, but the disagreement for the second-lowest threshold is not much worse than in 2010. For higher thresholds, the predicted rates stay behind the observed rates because again for the higher thresholds other components of \cancel{E}_T become important which are not included in the model.

6.6 Discussion and Outlook

In this section, a model has been presented to predict the \cancel{E}_T trigger rates as function of the average number of interactions per bunch-crossing $\langle\mu\rangle$. The basic idea is that the rates are dominated by fake \cancel{E}_T from the limited detector resolution. That this assumption is indeed a good description is confirmed by the fact that the functional forms which are implied by this assumption yields consistent resolution parameters for \cancel{E}_x , \cancel{E}_y and \cancel{E}_T when fitted to data. The resolution depends strongly on $\sum E_T$ and increases (i. e. becomes worse) with higher $\sum E_T$. The baseline for the explanation of the stronger than linear increase of the rates is therefore that higher pile-up means more activity in the event, which implies higher calorimeter occupancy and higher $\sum E_T$, and thus an increased resolution from larger fluctuations. This in turn leads to higher trigger rates. It has been shown that the \cancel{E}_T resolution is, however, independent of $\langle\mu\rangle$, i. e. its pile-up dependence is swallowed by the parametrization as function of $\sum E_T$.

Using only three input parameters, the model already provides a good description of both the \cancel{E}_T and $\sum E_T$ shapes as function of $\langle\mu\rangle$. The predicted rates are only in good agreement with the actual trigger rates for the lowest thresholds. Their dependence on $\langle\mu\rangle$ seems to well described, underlining the predictive power of the model for high $\langle\mu\rangle$. The good agreement of the shapes, together with the deviations in the rates, means that there must be a change in the \cancel{E}_T and $\sum E_T$ spectra outside the range over which the model and the observed spectra are compared. This is indeed found when plotting \cancel{E}_T and $\sum E_T$ over an extended range, which is only possible by using a combination of \cancel{E}_T triggers in addition to the minimum-bias and random triggers employed so far.

6.6.1 Extended \cancel{E}_T and $\sum E_T$ Spectra

To make the change in the slope of \cancel{E}_T visible, a large range of \cancel{E}_T needs to be covered. Due to the steeply falling nature of \cancel{E}_T , this range cannot be covered by a single trigger alone. In an example \cancel{E}_T spectrum taken with a random trigger in 2011, $3 \cdot 10^6$ raw events ($1.1 \cdot 10^6$ after applying the GRL and trigger selection) cover the range up to roughly 35 GeV. (This is the distribution in red in Figure 6.15.) The slope indicates that a factor of 60 times more events is needed to reach higher in \cancel{E}_T by 10 GeV. Extrapolating this and assuming no change in the slope[†], e. g. $1.4 \cdot 10^9$ raw events would be needed to cover the \cancel{E}_T spectrum up to 50 GeV,

[†] The assumption that the slope is constant up to much larger value of \cancel{E}_T is only correct for the contribution of fake \cancel{E}_T , as will become clear below.




Figure 6.15: Spectrum of \cancel{E}_T constructed by overlaying the spectra from several samples of events taken with different triggers in period 2011 D. Left: \cancel{E}_T at Event Filter, right: offline \cancel{E}_T (Topo calibration). The normalization of the different sets was done in the left plot, as described in the text, and then applied to the right plot without changes.

which would in turn correspond to a raw integrated luminosity of 35 fb^{-1} with the prescales of the existing random trigger dataset, many times more than the LHC has delivered in 2011.

What is done, therefore, in the plots in Figure 6.15, is to use spectra taken with several different triggers, which cover different ranges of \cancel{E}_T , and to normalize the event samples by matching them in the overlapping regions. A first justification that this normalization indeed is correct can be drawn from the fact that the shapes in the overlapping regions fit well. Below, another and more objective justification will be given. The triggers used to collect the different data samples are a random trigger on filled bunch crossings (with colliding bunches) for the lowest part of the spectrum (EF_rd0_filled_NoAlg) and four \cancel{E}_T triggers with different thresholds (EF_xe20,30,40,60_noMu). Every \cancel{E}_T trigger covers the EF \cancel{E}_T range above its threshold. For higher \cancel{E}_T , the statistics gets lower, until at the next threshold the next higher \cancel{E}_T trigger sets in. For the offline \cancel{E}_T , the coverage is smeared out, as can be seen in the plot. Both plots cover a range between 0 GeV and 200 GeV of \cancel{E}_T , in which the effective number of events per bin varies between about 10^2 and 10^{11} . Note that this, of course, is not the actual event count per bin, which is far below 10^{11} , but the number of events multiplied by the corresponding normalization factors.

Table 6.3 shows details of the normalization of the \cancel{E}_T spectrum and the data used in Figure 6.15. For all spectra, data from period 2011 D is used, selecting a number of runs from this period arbitrarily. The total number of events after applying the GRL from all selected runs and the integrated luminosity \mathcal{L}_{int} are also shown in the table. The integrated luminosity can be used to compute the relative normalization factors f_{pre} , which are needed to match the spectra taken with the different triggers. They are just the relative prescale factors and these can be computed from the ratio of the integrated luminosity of the data samples taken with the different triggers. The values of the integrated luminosity, which are stated in the table, have been computed using `ILUMICALC` (see Section 4.4.7). The factors f_{pre} in the penultimate column can be compared to the factors f_{match} , which are the normalization factors obtained from the overlap of the different distributions.

For the \cancel{E}_T triggers, the values of f_{pre} and f_{match} agree very well, the values for f_{pre} being slightly larger, which is probably due to some bias when matching the distributions in the

Stream	Period	Trigger	Events	\mathcal{L}_{int} [nb^{-1}]	f_{pre}	f_{match}
MinBias	D (8 runs)	EF_rd0_filled_NoAlg	1 140 870	0.556	$8 \cdot 16.5$	60
JetTauEtmiss	D (4 runs)	EF_xe20_noMu	139 793	9.15	22.2	20
		EF_xe30_noMu	72 313	203.0	15.3	14
		EF_xe40_noMu	105 871	3 098	23.5	23
		EF_xe60_noMu	408 883	72 650		

Table 6.3: Details of the normalization of the \cancel{E}_T spectrum and the data used in Figure 6.15. The left-most five columns hold information about which data was used, including the number of events and the respective integrated luminosity \mathcal{L}_{int} . The last two columns display the relative normalization factors, computed from prescales (f_{pre}) or from a matching in the overlapping regions (f_{match}), where e. g. $f_{\text{match}} = 60$ is the relative normalization of the data samples from the EF_rd0_filled_NoAlg and the EF_xe20_noMu triggers.

overlap regions. (No errors are given on these numbers as these are not vital for the discussion.) The first value for f_{pre} in the table needs some explanation. In addition to the prescales set at the three trigger levels, this trigger has an intrinsic prescale-like selection, which is configured to pick on average one out of eight bunch crossings at random. The raw rate of this trigger therefore is $f_{\text{LHC}}/8 \approx 5$ MHz, where f_{LHC} is the nominal rate of bunch collisions in ATLAS. Therefore, an additional prescale of eight on top of the relative prescale to the lowest \cancel{E}_T trigger needs to be included in the normalization. As can be seen from the table though, this does not give the correct result, but is too large by a factor of about two compared to the overlap normalization. What causes this discrepancy is not clear. Explaining it by the ratio of filled bunches over the total number of buckets, on which the nominal rate of bunch collisions is based, comes to nothing because this ratio is too low (up to 598/3564) for the set of runs that have been used, and it should affect the rates of all triggers in the same way, including the \cancel{E}_T triggers.

Figure 6.16 shows the spectrum of $\sum E_T$ and was produced in a similar way as Figure 6.15, but this time using $\sum E_T$ triggers instead of the \cancel{E}_T triggers. The lowest spectrum is again taken with the same random trigger as before and spans quite a substantial range of $\sum E_T$. To be consistent with the \cancel{E}_T spectrum in Figure 6.15, the spectrum taken with the random trigger has been scaled by an additional factor of four instead of eight. In contrast to the \cancel{E}_T triggers, no manual normalization was possible for $\sum E_T$ because the overlap of the spectra taken with $\sum E_T$ triggers is not sufficient. Instead, only the prescale-based scaling was applied. A striking difference between the spectra for EF \cancel{E}_T and EF $\sum E_T$ in Figures 6.15 and 6.16 is that the individual $\sum E_T$ spectra taken with $\sum E_T$ triggers are not monotonically falling functions, but reach a maximum which lies above the nominal trigger threshold. This is an indication that the respective lower triggers (L1_TE at Level 1 and L2_te at Level 2) have not reached their plateau at the threshold value of the EF cut on $\sum E_T$. In fact, the L1 and L2 $\sum E_T$ triggers have a very slow turn-on, so that, in order not to have to go to too high values at EF and still be able to stay within the rate limits at L2, this kind of selection overlap between the trigger levels is unavoidable. Another difference to the \cancel{E}_T spectrum is that the $\sum E_T$ spectrum seems to be rather well described by a single exponentially falling distribution over the range from 200 GeV up to 1.5 TeV, whereas the \cancel{E}_T spectrum shows a




Figure 6.16: Spectrum of $\sum E_T$ constructed by overlaying the spectra from several samples of events taken with different triggers in period 2011 D. Left: $\sum E_T$ at Event Filter, right: offline $\sum E_T$ (Topo calibration). The normalization of the different sets was done based on prescales for both plots, as described in the text. The nominal trigger thresholds at Event Filter level are as usual given in the chain name in GeV.

clear change in its slope as discussed above. Note that no selection on the number of interactions has been done, so that this $\sum E_T$ spectrum arises from a mixture of events with different numbers of interactions and the slope is not the same as the slope for events with exactly one interaction used as input parameter to the model.

6.6.2 Other Sources of \cancel{E}_T

To allow a direct comparison of the model prediction and the extended \cancel{E}_T spectrum, the plot in Figure 6.17 repeats the \cancel{E}_T spectrum from Figure 6.15, but this time again with an additional selection, which requires events to be from luminosity blocks for which $\langle \mu \rangle \approx 5$. The black curve in the plot is the model prediction for this value of $\langle \mu \rangle$. The change in the slope of the \cancel{E}_T spectrum happens around 40 GeV, but the deviation from the prediction by the model already starts below this value: for \cancel{E}_T values above this value, the observed \cancel{E}_T spectrum lies above the prediction. It is clear that the \cancel{E}_T rates for larger thresholds cannot be expected to be well described by the model, and indeed, the deviation from the plot is reflected in the underestimation of the rate for the trigger with an \cancel{E}_T threshold of 30 GeV in Figure 6.14. The underlying reason for this is that the model includes only one source of \cancel{E}_T , the fake \cancel{E}_T from resolution effects, which is dominant at low \cancel{E}_T , but the \cancel{E}_T measurement receives additional contributions from other sources at higher values:

- \cancel{E}_T from mismeasured jets leading to an imbalance of the total energy sum is expected to be proportional to the number of jets which are produced per second and thus scales linearly with the instantaneous luminosity. This \cancel{E}_T contribution shall also be called fake \cancel{E}_T , as opposed to real \cancel{E}_T from invisible particles. Yet it is of different origin than the resolution-based fake \cancel{E}_T and not to be confused with it.
- The trigger rate due to real \cancel{E}_T from physics processes is proportional to the number of invisible objects produced per second and thus also scales linearly with instantaneous luminosity.




Figure 6.17: Spectrum of \cancel{E}_T constructed by overlaying the spectra from several samples of events taken with different triggers and restricted to luminosity blocks with $\langle\mu\rangle \approx 5$. In addition the model prediction is plotted (solid line).

- Detector problems like “hot” cells constantly giving large contributions to the energy measurement are independent of $\langle\mu\rangle$ and instantaneous luminosity.

The rates attributed to all of these sources are expected to grow slower with $\langle\mu\rangle$ than the contribution to the trigger rate from fake \cancel{E}_T due to the detector resolution. Indeed, the dependence on $\langle\mu\rangle$ for higher \cancel{E}_T triggers in the range of \cancel{E}_T where the slope in the extended \cancel{E}_T spectrum deviates from the model prediction is linear, confirming that here the above sources drive the trigger rate. This linear dependence can be seen in Figure 6.18, where the rates of \cancel{E}_T triggers with thresholds of 60, 70, 80 and 90 GeV are plotted as function of $\langle\mu\rangle$ in periods I – K of 2011. Note that in this plot, the vertical axis is not on logarithmic scale. The spread of the data points is relatively large, although the values are averaged over 8 luminosity blocks, due to the low rates which have larger statistical fluctuations.

Thus, even though the model only accounts for one source of \cancel{E}_T , it covers the most problematic source of fake \cancel{E}_T , the one which gives rise to non-linear trigger rates, and its predictions can be understood as a lower bound for the trigger rates. With increasing $\langle\mu\rangle$, fake \cancel{E}_T from resolution effects will become the dominant contribution for higher and higher thresholds, and the model is important for estimating the contribution of this type of fake \cancel{E}_T when going to very high $\langle\mu\rangle$. The estimation can continuously be adapted and improved as new data for higher $\langle\mu\rangle$ becomes available.

6.6.3 Pile-up Noise Suppression

One example of a change in the ATLAS trigger system in 2011 that necessitates a recalibration of the model is the introduction of the pile-up noise suppression in the online system with the start of period 2011 G. The noise suppression uses an energy-dependent RMS value to do a one-sided cut for each cell, requiring the energy to lie above 3 RMS. The set of RMS values on which the cuts are based has been replaced by a set which is tuned to a calorimeter




Figure 6.18: Rates of several high-threshold \cancel{E}_T triggers in runs from period I – K of 2011. A linear behavior as function of the average number of interactions per bunch crossing $\langle\mu\rangle$ can be seen.

activity corresponding to eight pile-up interactions, whereas before it has only been accounting for electronic noise.

The rates of the three lowest \cancel{E}_T triggers in 2011 are plotted in Figure 6.19 for periods D – F, before the introduction of the pile-up suppression in the noise cuts, and for periods G – K, which have the pile-up suppression. The impact on the rates is clearly visible for EF_xe20_noMu. EF_xe40_noMu has already such a high threshold that its rate is not strongly affected by the pile-up noise suppression. The intermediate threshold EF_xe30_noMu has been deactivated in period G and later periods, except for one run (185353) with very low instantaneous luminosity (peak value: $1.2 \cdot 10^{30}$ Hz/cm²), so that almost no data for this trigger with the pile-up noise suppression in place is available.

This change at the beginning of period G is the reason why the above plots, in which the trigger rates in 2011 between prediction and measurement are compared, use periods D – F only. (Periods A and B have been mostly for testing after the winter shutdown of the LHC and do not contain much usable data, 2011 C is left out because it has a reduced center-of-mass energy of 2.76 TeV.)

The aim of including the updated noise suppression was to reduce the rates of the \cancel{E}_T triggers. As increasing the thresholds directly reduces the \cancel{E}_T seen by the trigger, this has naturally been achieved. Still, the strong dependence of the rates on $\langle\mu\rangle$ appears to be unchanged, considering the slope of the rates of EF_xe20_noMu in Figure 6.19. With respect to the model, the impact of the new noise suppression scheme on the resolution of the \cancel{E}_T measurement in the trigger and possibly the $\sum E_T$ spectrum needs to be studied, and the input parameters of the model need to be updated accordingly to obtain meaningful comparisons for the new periods. This will be the next step when pursuing the development of the model further. Other possibilities for future developments of the model are discussed in the remainder of this section.

So far, no uncertainties are available for the rates predicted by the model. The extraction of the input parameters yields uncertainties on these parameters which have been stated




Figure 6.19: Rates of the three \cancel{E}_T triggers EF_xe20,30,40_noMu as function of the average number of interactions per bunch-crossing compared between periods 2011 D – F (without) and G – K (with additional pile-up noise suppression in the trigger). For EF_xe40_noMu the points for periods D – F are mostly hidden behind G – K and of comparable spread.

above and which can be propagated to the spectra and trigger rates. This can help to assess the reliability of the extrapolation to higher pile-up levels and the sensitivity of the model predictions to the input parameters. Furthermore, there will also be systematic errors arising from the modeling itself. To quantize these, careful studies are needed. It has been stressed several times that the model is deliberately kept simple. Nevertheless, now that the underlying mechanisms driving the trigger rates seem to be understood quite well, the above discussion invites to do more checks and extend the existing model to incorporate some more of the features found in the observed \cancel{E}_T spectrum, which it currently does not cover. With respect to the implementation, the numerical stability could be checked. In the computation of the shapes, often very small values are used, and it would be good to confirm that the numerical precision at which the computations are carried out does not have a significant impact on the predictions. Predictions of the rates of the $\sum E_T$ triggers have not been attempted so far, as they are not widely used in ATLAS, and the prediction of $\sum E_T$ seems to be less stable than that of \cancel{E}_T . Yet it would be interesting to do some studies of how the predictions of $\sum E_T$ trigger rates compare against observed rates. It may also be interesting to check the influence of varying bunch sizes on the trigger rates: Not all bunches which are injected into the LHC contain the same number of protons. Depending on the magnitude of this variation, there may be a dependence of the trigger rate on the bunch crossing, which at least for triggers with low \cancel{E}_T thresholds will not average out, due to the non-linear dependence of the trigger rate on the average number of interactions within a bunch crossing. As said before, the model cannot predict the behavior of the full trigger chain. Whether it makes sense to try and include correlations between different trigger levels in the model is not obvious. It would, in any case, require a complete change of how the model works, and as long as the spacing of the thresholds at the different trigger levels is sufficiently large, there does not seem to be an urgent need to do so. Another possible extension which would probably promise the largest benefit with respect to the rate predictions is to include the

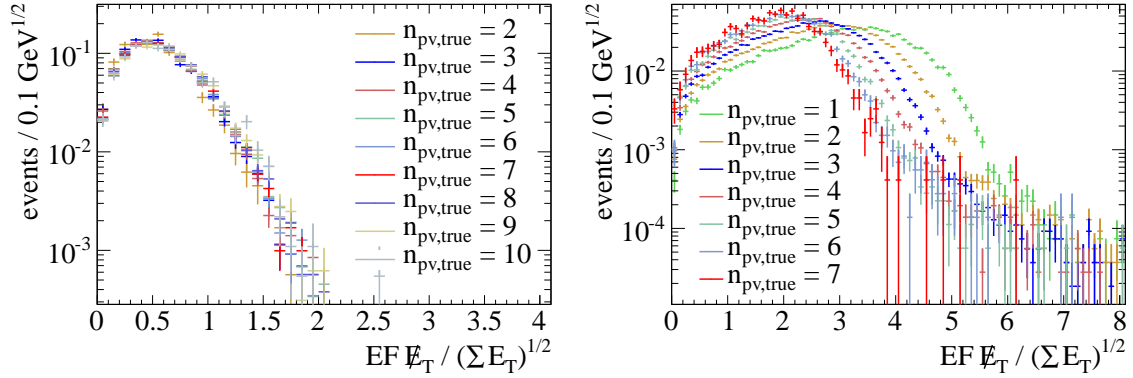


Figure 6.20: Distribution of the ratio $\cancel{E}_T / \sqrt{\sum E_T}$, as measured by the trigger at Event Filter level in simulated Monte Carlo (mc09) events with pile-up, with both \cancel{E}_T and $\sum E_T$ in GeV. The events are split up by the number of true primary vertices $n_{pv,true}$. Left: Minimum-bias events, right: $W \rightarrow e\nu_e$ decays. All distributions are normalized to unity to allow to compare their shapes.

fake \cancel{E}_T coming from jet mismeasurements in the model. This would require to find an analytic description which at least phenomenologically can describe the \cancel{E}_T contribution from jet mismeasurements. Its dependence on $\langle\mu\rangle$ should simply be linear as explained above. In any case, the model has shown that mismeasurements of the energy of jets make up an important part of the trigger rate, so that it might be worthwhile to consider how the \cancel{E}_T trigger itself can be improved to be more stable against this type of fake \cancel{E}_T , e. g. by vetoing events in which the direction of the missing transverse energy and the leading jet or a jet with significant energy are anti-aligned. Another offspring of the \cancel{E}_T trigger studies is the development of \cancel{E}_T significance triggers, which are briefly presented in the last part of this chapter.

6.6.4 \cancel{E}_T Significance Triggers

One of the conclusions drawn from the dominance of fake \cancel{E}_T with respect to the trigger rates is to design a new type of triggers, which use the \cancel{E}_T significance instead of the measurement of \cancel{E}_T alone. The \cancel{E}_T significance XS is given by the ratio of \cancel{E}_T and $\sqrt{\sum E_T}$. The actual implementation in the trigger is

$$XS := \frac{\cancel{E}_T / \text{GeV}}{A + B\sqrt{\sum E_T / \text{GeV}} + C(\sum E_T / \text{GeV})}, \quad (6.13)$$

where A , B and C are constants which allow to account for the offset due to noise suppression and to tune the online implementation of the trigger^k. Roughly speaking, resolution effects are proportional to the square root of the sum of the energy which is deposited in the calorimeter. Thus, the denominator normalizes the measured \cancel{E}_T to the total activity in the event and levels out the impact of in-time pile-up on the \cancel{E}_T measurement, which otherwise gives increasing contributions to $\sum E_T$ as the instantaneous luminosity per bunch increases.

^k The description here focuses on the main concept of the online implementation, leaving out additional accept and reject conditions which are introduced to ensure robustness in edge cases.




Figure 6.21: Rates of several \cancel{E}_T significance triggers with different thresholds as function of the average number of interactions per bunch-crossing $\langle\mu\rangle$ from 2011 data, periods 2011 G – K. The meaning of the threshold values is explained in the text.

Following this argument, the rates of the \cancel{E}_T significance triggers are expected to be much more stable against in-time pile-up than the those of the \cancel{E}_T or $\sum E_T$ triggers. The general behavior of this type of XS triggers is studied here on two types of Monte Carlo events. Both plots in Figure 6.20 show the distribution of the ratio $\cancel{E}_T/\sqrt{\sum E_T}$ measured at Event Filter level. The different distributions within each plot stand for different selections based on the true number of interactions per event, $n_{pv,true}$. The left plot is produced from Monte Carlo with minimum-bias events overlaid with an average number of five additional interactions from pile-up per event. It clearly shows that in this sample, $\cancel{E}_T/\sqrt{\sum E_T}$ does not depend on the number of interactions because indeed all distributions have the same shape, independent of the number of interactions, and fall together when normalized to one, as is done in this plot. For this type of background rate, coming from minimum-bias events with negligible real \cancel{E}_T , the XS triggers are therefore expected to be stable against pile-up contributions. The right plot shows the distribution of XS on a very different Monte Carlo sample containing $W \rightarrow e\nu_e$ events, again overlaid with on average five additional interactions from pile-up per event. In this type of events, a significant contribution to \cancel{E}_T is real \cancel{E}_T from the neutrino. With increasing pile-up, the normalization by $\sqrt{\sum E_T}$ shifts the peak in the spectrum to the left. (Note that the range of the horizontal axis is twice the range in the left plot so that the peak still lies at higher values on the horizontal axis than the peak from the minimum-bias background.) A direct consequence is that the contribution from events with real \cancel{E}_T to the trigger rate will decrease with increasing pile-up, which may be regarded an unusual if not undesired behavior. In any case, this is important to remember when considering to use this type of triggers for selecting events for physics analyses or doing studies of trigger and signal efficiencies on these events. It would also be interesting to study how big the gain of the XS triggers over \cancel{E}_T triggers from the improved stability with respect to pile-up is, e. g. whether the efficiency for certain (signal) types of events is improved considerably.

The XS-based triggers have been running online since the middle of 2011 in ATLAS. In the plot in Figure 6.21, the rates of four XS triggers from periods G – K in 2011 are shown,

in which the \cancel{E}_T significance triggers were running for the first time with the final calibration. The nomenclature for these triggers is such that the number in the trigger name is 10 times the value of XS as defined in Equation (6.13). Indeed, the measured trigger rates show no or only a weak dependence of the average number of interactions per bunch-crossing $\langle\mu\rangle$. EF_xs30_noMu is the lowest XS threshold, which had to be heavily prescaled already in 2011 G, as can be deduced from the rate plot in Figure 6.21 (note that the rate needs to be multiplied by the number of colliding bunches). Nonetheless, from the right plot in Figure 6.20, it already rejects most of this type of events for high levels of pile-up, i. e. has a low signal efficiency for $W \rightarrow e\nu_e$ events, which shows that there are limitations for this new type of trigger.

7 Search for Supersymmetry

In this chapter, a search in data taken with the ATLAS detector in the year 2010 for signatures from decays of supersymmetric particles is presented. First, the search channel is introduced, and the outline of the analysis is given. Afterwards, the data sample and event selection are described, including the object definitions on which the event selection is based and the definition of the signal regions. This is followed by the description of the Monte Carlo samples for the Standard Model backgrounds and the signal grids which implement the supersymmetric model scenarios. The uncertainties on the measurements in data and Monte Carlo are discussed, and the result of the analysis is presented in terms of the obtained event counts. From the event counts and the systematic uncertainties limits are computed on the mass parameters of two different supersymmetric models.

7.1 Introduction

7.1.1 The Zero-Lepton Channel

Supersymmetry can manifest itself through many different signatures in a collider experiment. Based on the type of signature that is searched for, a rough characterization of different analyses can be done, for example by specifying the multiplicity of particles in the final state. In this thesis, the zero-lepton channel is studied, meaning that in the selected events no leptons appear in the final state. Instead, the targeted signature consists of several hard (i. e. high-energetic) jets and a large amount of missing transverse energy, which is carried away by a heavy particle that evades detection because it only weakly interacts with matter.

The left diagram in Figure 7.1 shows an example of a supersymmetric decay chain which yields the detector signature of the signal events searched for in the zero-lepton channel. Due to R -parity conservation, the decay chain ends in the LSP, and every event contains an even number of these decay chains. The missing transverse energy is due to the LSP, the hard jets are produced in the decays of the heavy supersymmetric particles. For comparison, in the right diagram a longer supersymmetric cascade decay is shown, in which an opposite-sign lepton pair is produced. Which decay modes are preferred depends on the mixings of the neutralinos [106]. In the analysis at hand, where leptons are vetoed, the term leptons is understood to only include electrons and muons, but not tau leptons. Tau leptons are excluded because of their decay properties, giving a very different detector signature. They have very short lifetimes and in contrast to electrons and muons have enough mass to decay into hadrons, giving rise to hadronic jets. They are thus harder to identify than electrons and in particular muons, which give clean signatures with high efficiencies and low fake rates. Jets from hadronic tau decays are always interpreted as ordinary jets, no tagging is used.

To trigger on the signature of the supersymmetric signal events, jet triggers, missing transverse energy triggers or combinations of both can be used. The benefits of using combined triggers have been demonstrated in Section 5.1.2. Consequently, the primary triggers for the zero-lepton channel analysis are combined jet + \cancel{E}_T triggers. The efficiencies of these triggers




Figure 7.1: Left: A supersymmetric decay chain which produces the detector signature being searched for in the zero-lepton channel: jets, no leptons and missing transverse energy. Right: For comparison, a longer supersymmetric decay chain in which an opposite-sign lepton pair is produced.

have been discussed in Section 5.2, for 2010 in particular in Section 5.4. The offline cuts of the analysis ensure that for events passing all cuts, the trigger can be assumed to be fully efficient.

7.1.2 Outline of the Analysis

In the general approach for a search for new physics that is pursued here, observed data is compared to theoretical expectations which are predicted by the Standard Model of particle physics. As it is not possible to describe the expected detector response in detail analytically, Monte Carlo simulations are employed. They encode the expectations of detector signals, both from Standard Model processes and from new types of physics processes that are being searched for, in terms of simulated data. The simulated data is made to resemble the actual detector measurements as closely as possible. If deviations are found in event counts, distributions of kinematic variables or other quantities, which cannot be explained by uncertainties in the measurement itself, these discrepancies may be interpreted in terms of physics beyond the Standard model. If the inclusion of the new physics processes can explain the data, it has to be checked whether this model is superior to the wealth of competing models which may also be used to explain the data. If the model is better than all other (of equal simplicity), then it has good chances of superseding parts of previous models, but this is a process at timescales of years or decades. If no deviations are found, the results can still be interpreted in the framework of some concrete model including new physics, by setting limits on the probability for this model to be realized in nature, or at least on its parameters which have impact on the probability to find traces of this model at the detector. Usually, new physics processes are supposed to lead to deviations in terms of an excess of events rather than finding fewer events than expected. An important counterexample are neutrino oscillations (cf. Section 2.1.3), which are the solution to the solar neutrino problem, denoting the fact that when the first measurements of the flow of solar neutrinos to the earth were done, much less neutrinos (of a given flavor) were found than expected. This was resolved by giving up the assumption of massless neutrinos. Assuming instead that neutrinos have a small but non-zero mass, they can oscillate between different flavors. Summing up neutrinos of all flavors, the flux was found to agree with the predictions again [104]. Thus in principle, new physics processes can lead to a smaller number of observed events, but in the following, it will be assumed that Supersymmetry only yields a positive contribution to the event count on top of the Standard Model background.

In counting analyses, the question whether physics beyond the Standard Model exists, is answered by looking at the number of events or the distribution of kinematic quantities after applying a set of cuts which define a signal region. The term region refers to a certain

subset of the phase space of events, or, better, the physics objects the events are composed of. The event count in the signal region is then compared to the expectation from Monte Carlo. The significance of the result depends on how well the contamination of the signal region with background events can be predicted. As the Monte Carlo may suffer from uncertainties and systematic bias, many analyses try to be as little dependent on Monte Carlo as possible, and to constrain the prediction of the background contamination in the signal region by additional measurements in control regions. The control regions are designed to be enriched with background events, in order to allow a better determination of the background level, which then is extrapolated to the signal region. The extrapolation is done using transfer functions, which in the simplest case are linear factors, giving the ratio of the selection efficiencies i. e. the number of events expected to be observed in the control and signal region. In case of several types of backgrounds and several control regions, the transfer functions become a set of linear equations. Transfer functions can be measured on Monte Carlo. This still relies on Monte Carlo, but no longer on the absolute normalization, which is often most affected by uncertainties.

For the present analysis the official definitions of the signal regions which are used within the ATLAS SUSY group for the zero-lepton analysis of 2010 data are adopted. This allows to validate the results of the official results, and on top to take additional aspects into account. An updated calibration for the luminosity is used, giving a slightly smaller value for the integrated luminosity (see Section 7.2). The impact of pile-up, which was neglected in the official analysis, is studied, and two different methods for the interpretation of the results in terms of limits on the mass parameters are compared.

The studies of the ATLAS SUSY group show good agreement between the Monte Carlo predictions for the Standard Model backgrounds and cross-checks from measurements in control regions on real data. Therefore, the Monte Carlo predictions of the backgrounds are used as central values here, and discrepancies taken into account in terms of systematic uncertainties. The background studies using control regions which were done by the ATLAS SUSY group are not repeated here. Instead, for these aspects, the following analysis relies on the results of the background studies given in the internal note [146], the relevant parts of which will be summarized when discussing the normalization of the background Monte Carlo samples in Section 7.5.2.

7.2 Integrated Luminosity, Data Samples and Trigger Selection

The integrated luminosity is a measure for the amount of collision data that has been collected (cf. Section 3.2.4). It is needed for almost every kind of analysis of collision data. In analyses that compare data to Monte Carlo predictions, as is done here, it enters as a normalization factor to scale the Monte Carlo event count to the number of events in the dataset.

For the calculation of the integrated luminosity of the 2010 dataset, which is evaluated in the analysis presented in this thesis, the official `ILUMICALC` tool described in Section 4.4.7 is used. As live-fraction trigger, needed to estimate the dead time of the trigger system, the minimum-bias trigger `L1_MBTS_2` is used, and the database tag for the luminosity calibration is `0f1Lumi-7TeV-003`, following the recommendations for 2010 pp collision data [187].

Table 7.1 gives an overview of the luminosity per data-taking period in 2010, calculated using the `ILUMICALC` tool. Note that the resulting value for the total integrated luminosity in 2010 is about 4% lower than the value which was officially used by the ATLAS SUSY

Period	Run numbers	Trigger requirement for event selection	Online trigger	Integrated luminosity in nb ⁻¹
A	152166 – 153200	L1_J55 + L2_j70 + EF_EFonly_xe25_noMu	L1_J55	0.31
B	153565 – 155160	L1_J55 + L2_j70 + EF_EFonly_xe25_noMu	L1_J55	7.69
C	155228 – 156682	L1_J55 + L2_j70 + EF_EFonly_xe25_noMu	L1_J55	8.14
D	158045 – 159224	L1_J55 + L2_j70 + EF_EFonly_xe25_noMu	L1_J55	278
E	160387 – 161948	L1_J55 + L2_j70 + EF_EFonly_xe25_noMu	EF_j75_jetNoCut	906
F	162347 – 162882	L1_J55 + L2_j70 + EF_EFonly_xe25_noMu	EF_j75_jetNoCut	1 652
G1	165591 – 165632	EF_j75_jetNoEF_EFxe25_noMu		
G2	165703 – 165732	EF_j75_jetNoEF_EFxe25_noMu		
G3a	165767	EF_j75_jetNoEF_EFxe25_noMu		
G3b	165815	EF_j75_jetNoEF + EF_EFonly_xe25_noMu	EF_j75_jetNoEF	536
G4	165817 – 165818	EF_j75_jetNoEF + EF_EFonly_xe25_noMu	EF_j75_jetNoEF	
G5a	165821	EF_j75_jetNoEF + EF_EFonly_xe25_noMu	EF_j75_jetNoEF	
G5b	165954 – 166143	EF_j75_jetNoEF_EFxe25_noMu		
G6	166198 – 166383	EF_j75_jetNoEF_EFxe25_noMu		4 915 [‡]
H	166466 – 166964	EF_j75_jetNoEF_EFxe25_noMu		6 803
I	167575 – 167844	EF_j75_jetNoEF_EFxe25_noMu		18 319
Sum over all periods				33 425

[‡] 4915 nb⁻¹ is the integrated luminosity for period G excluding G3b, G4, and G5a.

Table 7.1: Integrated luminosities and triggers per period for pp collisions taken in 2010. The online trigger used for data collection is given if it is different from the trigger requirement imposed offline to select events.

group for the first analyses of 2010 data [105] because at that time the luminosity calibration implemented in `0f1Lumi-7TeV-003` was not yet available (cf. Section 4.4.7). The run numbers are not consecutive because the pool of run numbers is shared between different ATLAS projects, and not all run numbers are given to runs in which data for physics analysis is taken. In total, an integrated luminosity of

$$\mathcal{L}_{\text{int}}^{2010} := \int_{2010} \mathcal{L}_{\text{inst}}(t) dt = 33.4 \pm 1.2 \text{ pb}^{-1} \quad (7.1)$$

is computed for the dataset taken in 2010, which is used for the search for Supersymmetry in this thesis. The data was taken by the ATLAS collaboration between 30th March 2010 and 27th October 2010.

Table 7.1 also contains details about the trigger requirements for the analysis per period. Due to the constant changes in the trigger menu in the start-up phase of the LHC, the trigger selection is not homogeneous, which is reflected in the table. The trigger used for the event selection in the analysis is a combined jet + \cancel{E}_T trigger. As this trigger was not available in the trigger menu in all periods, a different, looser trigger is used to select events and the jet + \cancel{E}_T trigger decision is emulated offline. Therefore, in the table there is, for some periods, a distinction being made between the trigger which was used online to record selected events and which therefore is relevant to the luminosity calculation in terms of prescales, and the trigger requirement that is later imposed to select events in the analysis. The latter must, of course, select a subset of the previous. Otherwise the event sample would lack events and probably be systematically biased in a very subtle way, because the composition of events would be determined by the configuration of the trigger menu, in particular which triggers were running in parallel, rather than physics properties and the trigger efficiency.

In periods A – D, the HLT was still being commissioned: it was running, but not rejecting events. Every event that passed L1 was recorded, and in an analysis it was enough to require a positive L1 trigger decision to select a valid event sample. From period E onwards, the HLT has commenced normal operation. The jet + \cancel{E}_T trigger was not yet implemented and events therefore selected with the jet trigger `EF_j75_jetNoCut`, the suffix `jetNoCut` implying that still no cut was done in the HLT. From period G onwards, the combined jet + \cancel{E}_T trigger, the primary trigger for the analysis presented in this thesis, became available and the online trigger decision could therefore be used directly to select events. At the same time, event rejection by the HLT was also activated for the jet slice, although the jet algorithms at EF were still in pass-through mode and have never been rejecting events in 2010 (therefore the suffix `jetNoEF`), so that effectively only the L2 jet algorithm was activated beginning with period G. In a subset of runs (165815 – 165821) from this period, the jet + \cancel{E}_T triggers were disabled, but all events can be recovered by using `EF_j75_jetNoEF` which was running unprescaled in these runs. For periods H and I, the jet + \cancel{E}_T trigger could be used again.

The Monte Carlo version used in this thesis is `mc09_7TeV`. The data format used as input for the analysis scripts is the `NTUP_SUSY` ntuple format (see Section A.10.3). This format can be read by ROOT directly and is produced from data and Monte Carlo AOD files in the central production with the SUSY D3PD maker, which is a derivate of the official D3PD maker and maintained by the ATLAS SUSY group. To ensure homogeneity of the data with regard to the code used for the reconstruction and the calibration, all ntuples used in this thesis have been produced with version 06-20-00 of SUSY D3PD maker in ATHENA release

AtlasProduction 16.0.2.2^a. These datasets bear the processing tag p305 and have been produced without skimming and thus contain all events from the respective datasets. The same holds for the data samples. The NTUP_SUSY ntuples with production tag p305 are used, their project name is data10.

7.3 Event Selection

For the analysis of ATLAS 2010 data implementing the search for Supersymmetry in the channel with jets, missing transverse energy and no identified leptons in the final state, the official cuts developed within the ATLAS SUSY group have been implemented. It has been checked that the results on data in terms of event counts after cuts are consistent with the official analysis. In particular, the number of data events reaching each of the signal regions agrees exactly with the official analysis published in [105]. Details about the definition of the signal regions are given below.

The cuts used in the analysis are described in detail in this section. To give a first idea of how the analysis proceeds, the cuts can be grouped and summed up in the following scheme:

1. Selection of collision candidates and event clean-up. This step filters out bad data from runs with known detector problems, but also implements the trigger selection.
2. Reject events with reconstructed leptons to define the zero-lepton search channel. As mentioned in the introduction, only electrons and muons are considered in the lepton veto.
3. Cuts on reconstructed jets to select events of a given jet multiplicity corresponding to the jet multiplicity of the signal regions. This defines the sub-channel.
4. Cuts on \cancel{E}_T which is also part of the baseline selection.
5. Cuts on derived event variables that further refine the signal regions. These are m_{eff} and m_{T2} (cf. Section 4.4.5).

Table 7.2 shows the sequence of cuts, collectively called cutflow, together with the number of events in data left after each cut. The efficiency of the cuts ϵ_{cut} on data is also listed^b. Note that not all cuts are part of the definition of a signal region. A few steps in the cutflow are not evaluated when setting limits. These are the cuts selecting events with a monojet (step 9 and 10) or 4- (or more) jet topology (step 24). They are included for completeness and in order to make sure that the numbering of the cuts is consistent everywhere.

About 270 millions of raw events from the JetTauEtmis stream are processed, leaving about 230 millions events in luminosity blocks which are marked as suitable for physics analysis according to the GRL defined by the SUSY group (cf. Section 4.4.6). The triggers used to select a well-defined sample of events from the stream are listed in Table 7.1 as explained

^aNote that although this is release 16 of ATHENA, the reconstruction of the data has been done in release 15, e.g. in AtlasProduction 15.6.13.7 for datasets with tag f299 such as data10_7TeV.00167776.physics_JetTauEtmis.merge.A0D.f299_m639, from which data10_7TeV.00167776.physics_JetTauEtmis.merge.NTUP_SUSY.f299_m639_p305 is produced.

^bAssuming Poisson statistics and using Gaussian error propagation, rarely more than two digits are significant for the efficiencies, even for large event counts.

Cutflow step	Cut description	Event count	ϵ_{cut} [%]	SR
1	Raw events	268 706 385		
2	GRL	229 004 362	85.2	
3	Main trigger	3 889 200	1.7	
4	Jet cleaning	3 848 197	98.9	
5	Vertex requirement	3 838 202	99.7	
6	Veto on crack electrons	3 833 805	99.9	
	(Object identification and overlap removal)			
7	Veto on leptons	3 787 202	98.8	
8	Leading jet $p_T > 120$ GeV	1 910 736	50.5	
9	Monojet	49 839	2.6	
10	Charge fraction	45 237	90.8	
11	Second jet $p_T > 40$ GeV	1 789 007	93.6	
12	Charge fraction	1 763 609	98.6	
13	$\cancel{E}_T > 100$ GeV	9 121	0.5	
14	$\Delta\phi(\text{jet}, \cancel{E}_T) > 0.4$	934	10.2	
15	$\cancel{E}_T/m_{\text{eff}} > 0.3$	621	66.5	
16	$m_{\text{eff}} > 500$ GeV	87	14.0	A
17	$m_{T2} > 300$ GeV	11	1.2	B
18	Third jet $p_T > 40$ GeV	680 058	38.0	
19	$\cancel{E}_T > 100$ GeV	4 510	0.7	
20	$\Delta\phi(\text{jet}, \cancel{E}_T) > 0.4$	397	8.8	
21	$\cancel{E}_T/m_{\text{eff}} > 0.25$	198	49.9	
22	$m_{\text{eff}} > 500$ GeV	66	33.3	C
23	$m_{\text{eff}} > 1$ TeV	2	3.0	D
24	Fourth jet $p_T > 40$ GeV	152 405	22.4	

Table 7.2: Overview of the complete cutflow of the analysis. A short description of the cuts is given together with the number of events in data left after each cut and the efficiency of the cut ϵ_{cut} . In the last column, the four signal regions are denoted with the letters A through D (cf. Table 7.6). The indentation of the cut description hints at the hierarchy of the cuts and is explained in the text.

Target	Loose	Tight
EM coherent noise	$f_{\text{EM}} > 0.95$ and $ Q > 0.8$	$f_{\text{EM}} > 0.90$ and $ Q > 0.6$
HEC spikes	$f_{\text{HEC}} > 0.8$ and $N_{90} \leq 5$ $f_{\text{HEC}} > 0.5$ and $ Q > 0.5$	$f_{\text{HEC}} > 0.3$ and $ Q > 0.3$ $f_{\text{HEC}} > 1 - Q $
Cosmic rays and beam background	$ t_{\text{jet}} > 25$ ns $f_{\text{EM}} < 0.05$ $f_{\text{max}} > 0.99$ and $ \eta_{\text{EM}} < 2$	$f_{\text{EM}} < 0.10$ $f_{\text{max}} > 0.95$ and $ \eta_{\text{EM}} < 2$

Table 7.3: Defining conditions for a jet to be flagged as bad according to the loose and tight jet cleaning recommendations [155, 175]. Only one of the conditions listed in the two rows has to be fulfilled. The variables appearing in the table are explained in the text.

above. For the Monte Carlo samples, a combination of L2_j70 and EF_EFonly_xe25_noMu is used as trigger. EF_EFonly_xe25_noMu refers to an emulated trigger (cf. Section A.1), which cuts at 25 GeV on the measurement of missing transverse energy obtained by the trigger at Event Filter level.

7.3.1 Jet Cleaning

The jet cleaning procedure implies a veto on events in which the reconstruction of one or more jets is considered to be unreliable due to suspicious values in variables quantifying the reconstruction quality [155]. To this aim, every jet is classified as either good or bad[§], where different sets of criteria can be used to define bad jets. Bad jets in this sense are energy deposits in the calorimeters from various sources other than the true jets from the proton-proton interaction. These include hardware problems, LHC beam conditions and cosmic-ray showers.

Both the loose and the tight definition for bad jets recommended by the Jet/Etmiss combined performance group for detector data reconstructed with ATHENA in release 15 are used [189]. A jet is marked as bad under the loose or tight selection, respectively, if any of the conditions listed in the corresponding column in Table 7.3 is met. Note that jets fulfilling any of the loose conditions also are considered bad when applying a selection with the stricter tight conditions. Here, stricter selection means that more jets are considered bad and thus more events rejected, i. e. the set of bad jets under the tight selection is the union of all bad jets from the loose selection plus all jets fulfilling any of the conditions listed under the tight selection criteria. The cuts listed in the table target different sources of potential mismeasurements or misidentifications of jets and make use of the following variables:

- f_{EM} and f_{HEC} are the fraction of the energy of the jet deposited in the electromagnetic and the hadronic end-cap calorimeter, respectively.
- Q is the fraction of jet energy from LAr cells with a quality factor larger than 4000. The cell quality factor is given by the quadratic difference $\sum_i (a_i - a_i^0)^2$ between the samplings of the measured pulse a_i and the reference pulse shape a_i^0 that is used to

[§] The third category containing “ugly jets” is not used in this analysis.

reconstruct the cell energy. It can be used to detect energy accounted for by coherent noise.

- N_{90} is the minimum number of cells containing at least 90% of the jet energy. An unusually small value for N_{90} is a good indication for single-cell jets which are due to spikes affecting only one or a few cells due to cross-talk.
- The timing variable t_{jet} is the energy-squared-weighted cell time, defined with respect to the event time and used to find jets reconstructed from large out-of-time energy depositions, e. g. due to photons produced by cosmic ray muons.
- f_{max} gives the maximum energy fraction deposited in only one calorimeter layer. In the central region, jets should spread out over more than one layer.
- The η_{EM} value used in this definition is the pseudorapidity of the jet at EM scale^d.

Events containing a bad jet according to the loose definition with $p_T > 20 \text{ GeV}$ at EM+JES scale are discarded. The inefficiency of this cut is small, as can be seen from Table 7.1. The tight definition is applied later on in the 2-jet sub-channel. Events passing the jet cleaning are in the next step required to have at least one reconstructed primary vertex with ≥ 5 tracks associated to it, to filter out empty events in which no hard interaction has taken place. Afterwards, a veto is applied, rejecting events if they contain a reconstructed electron which fulfills the object identification criteria described below and goes into the region $1.37 < |\eta_{\text{cluster}}| < 1.52$. This transition region between the barrel and the end-cap electromagnetic calorimeters is expected to have poorer performance, because of the large amount of passive material in front of the first active layer of the electromagnetic calorimeter.

7.3.2 Object Identification

The next step is the object identification and overlap removal. Here, no filtering of events is done, but a promotion of candidate objects to (selected) reconstructed objects is made, rejecting objects which do not fulfill the selection criteria. This step is called object identification, object definition or object selection. The criteria used for the selection of physics objects are summarized in Tables 7.4 and 7.5 and explained in detail below. This step also includes an overlap removal in which objects are deleted which overlap geometrically, are of different physics type (electron, muon or jet), but are probably reconstructed from different interpretations of detector signals caused by the same real physics object.

Electrons

The selection of electron candidates follows the recommendations of the ElectronGamma combined performance group for the identification of electrons within the region $|\eta| < 2.5$ covered by the Inner Detector [50]. It is based on a number of rectangular cuts (as opposed to multivariate techniques), defined in three reference sets which allow to trade efficiency for purity of the electron sample obtained. In the analysis, electrons reconstructed by either only the cluster-based algorithm, or by both the cluster-based and track-based algorithm are used

^d Although a rescaling of the energy does not change the jet direction η , small discrepancies between the two η variables can be observed which may occur when calculating the direction from reweighted cell contributions.

Cut	Electron	Jet
Reconstruction algorithm	AuthorElectron or AuthorSoft	Anti- k_t , EM+JES scale
Acceptance	$p_T > 10 \text{ GeV}$ and $ \eta < 2.47$	$p_T > 20 \text{ GeV}$ and $ \eta < 2.5$
Isolation	—	—
Overlap	removed if $0.2 < \Delta R_{\min}^{\text{jet}} < 0.4$	removed if $\Delta R_{\min}^e < 0.2$
Quality cuts	RobustMedium OQMaps	—

Table 7.4: Object identification criteria for electrons and jets. ΔR_{\min}^a is the smallest distance in η - ϕ space to an object of type a .

(cf. Section 4.4.1). The quality and purity cuts defined in the `RobustMedium` set are applied. This requires the electron to satisfy both the loose [153] and the updated (robuster) version of the medium cuts [144], which, among other changes, takes into account differences in the shower shape between data and Monte Carlo. The acceptance cuts for electrons are $p_T > 10 \text{ GeV}$ and $|\eta| < 2.47$. In addition, electrons must not fall into a region which suffers from detector problems. This is ensured by using the Object Quality maps (OQMaps) provided by the ElectronGamma group. Selected electrons that fall into the transition region $1.37 < |\eta_{\text{cluster}}| < 1.52$ trigger the rejection of the event as a whole.

Jets

For the jets there are only few requirements on top of the jet cleaning as described above. All jets reconstructed by the anti- k_t algorithm with a radius parameter $\Delta R = 0.4$ from topological clusters are used, if they have $p_T > 20 \text{ GeV}$ and $|\eta| < 2.5$. Another common value for the radius parameter in the jet reconstruction algorithms in ATLAS is $\Delta R = 0.6$, but as many jets are expected in typical signal events, the smaller cone radius is used. The energy of jets is calibrated at EM+JES scale, as described in Section 4.4.2. This follows the recommendations of the ATLAS Jet and \cancel{E}_T Group for physics analyses on 2010 data [190].

Muons

Muons are taken from the container filled by the `Staco` reconstruction algorithms. Following the recommendations from the Muon combined performance group for release 15, combined muons and segment-tagged muons are used (cf. Section 4.4.1), and the respective cuts applied [196]: To make sure that identified muons have traversed the Inner Detector with hits associated to the fitted track from all sub-detectors, the number of pixel hits n_{pixel} is required to be at least one and in addition six or more hits in the SCT on the muon track are required. Within the acceptance of the TRT, a minimum number of associated hits is required, plus a limit on the number of TRT outliers on the muon track $n_{\text{TRT}}^{\text{outliers}}$ is set, where $n = n_{\text{TRT}}^{\text{hits}} + n_{\text{TRT}}^{\text{outliers}}$, and $n_{\text{TRT}}^{\text{hits}}$ is the number of TRT hits on the muon track. For combined muons, in addition to the above cuts, the quality of the combination is tested by requiring that the χ_{match}^2 of the matching of the tracks from the muon spectrometer and the Inner Detector at the perigee must be smaller than 150, and for combined muons with $p_T < 50 \text{ GeV}$, a cut on the difference between the extrapolated momentum measurement from the muon

Cut	Muon
Reconstruction algorithm	Staco (Combined or segment tagged)
Acceptance	$p_T > 10 \text{ GeV}$ and $ \eta < 2.4$
Isolation	$\sum_{\Delta R < 0.2} p_T^{\text{track}} < 1.8 \text{ GeV}$
Overlap	removed if $\Delta R_{\text{min}}^{\text{jet}} < 0.4$
Quality cuts	$n_{\text{pixel}} \geq 1, n_{\text{SCT}} \geq 6$ for $ \eta < 1.9: n > 5$ and $n_{\text{TRT}}^{\text{outliers}} < 0.9n$ for $ \eta \geq 1.9$ and $n > 5: n_{\text{TRT}}^{\text{outliers}} < 0.9n$ for combined muons: $\chi_{\text{match}}^2 \leq 150$ for combined muons with $p_T < 50 \text{ GeV}$: $p_{\text{MS}}^{\text{extrapol.}} - p_{\text{ID}} > -0.4p_{\text{ID}}$

Table 7.5: Object identification criteria for muons. The quality cuts correspond to the recommendation from the Muon combined performance group for release 15 and are explained in the text, together with the variables they are based on.

spectrometer $p_{\text{MS}}^{\text{extrapol.}}$ and the Inner Detector p_{ID} is also applied. The acceptance cuts for all muons are $p_T > 10 \text{ GeV}$ and $|\eta| < 2.4$. To select only isolated muons, a relatively tight cut is applied on the summed p_T of tracks, which must be smaller than 1.8 GeV within a cone of $\Delta R = 0.2$ around the muon.

Missing Transverse Energy

The missing transverse energy that is used in the analysis is called MET_EMJES_RefFinal_Ce110outEM and calculated with an object-based refined \cancel{E}_T algorithm, as described at the end of Section 4.4.3. Two fixes need to be applied. As the tool for the computation of the \cancel{E}_T does not allow for the muon selection criteria from above, the muon term is subtracted from the refined \cancel{E}_T and replaced by adding the negative x and y components of the muon momentum to the \cancel{E}_T components. Similarly, the \cancel{E}_T calculation is corrected by subtracting $p_{x,y}$ from the missing energy for all electrons that pass the RobustMedium, but not the medium purity criteria, and lie within a distance of $\Delta R < 0.2$ to a jet. Then, $p_{x,y}$ of the corresponding jet is added to the respective components of the missing energy. This corrects for the fact that the wrong electron identification is used in the \cancel{E}_T computation in data. In addition to these fixes, changes in \cancel{E}_T arising from the calculation of the jet-energy scale and jet-energy resolution uncertainties described below are taken into account.

Overlap Removal and Lepton Veto

Having thus selected the physics objects for the following analysis, any occurring overlaps need to be resolved. This is done according to the following strategy:

1. Selected jets are removed if they fall within a cone of $\Delta R < 0.2$ around a selected electron. The idea behind this is that if both a jet and an electron are reconstructed so

close geometrically, the real physics object is more likely to be an electron rather than a jet.

2. Selected electrons are removed if they fall within a cone of $\Delta R < 0.4$ around one of the remaining selected jets, because they are then believed to have been produced in a (hadronic) jet.
3. Along the same lines, selected muons are removed if they fall within a cone of $\Delta R < 0.4$ around one of the remaining selected jets.

(The ΔR distance is computed using uncalibrated (EM scale) coordinates of the jet here.)

Events that contain identified leptons (electrons or muons passing the above selection) after the overlap removal are vetoed. This veto prevents events from entering the zero-lepton channel analysis that are also selected in other SUSY search channels like the one-, two- or multi-lepton analyses which include leptons in their selection. It thereby avoids double counting of events, so that it is easier to compute combined limits from the results of search channels with different lepton multiplicities. Note that in contrast to this, the jet selection within the zero-lepton channel is inclusive, so that events passing the three-jet selection may also appear in the signal regions with two jets. This needs to be remembered when setting limits and is discussed in Section 7.11. It is important for the limit-setting procedure, because the signal regions cannot be taken as independent search channels.

7.3.3 Jet Multiplicity and Sub-Channel Selection

The rest of the cuts after the lepton veto no longer involve leptons, but only jets, \cancel{E}_T and derived event quantities, namely m_{eff} and m_{T2} . The cuts are straightforward and lead to four different signal regions described below. The indentations in Table 7.2 mark branching points in the cutflow. Cuts that are indented are not relevant for following cuts on a lower indentation level. The cut in step 11, for example, is applied to all events left after step 8, regardless of the two intermediate steps.

Monojet Channel

All sub-channels require a leading jet with $p_T > 120$ GeV. Steps 9 and 10 define the monojet channel. In step 9, events are rejected if they contain any further jets with $p_T > 30$ GeV or if they contain jets (the leading or further jets with $20 \text{ GeV} < p_T < 30 \text{ GeV}$) that are marked as bad according to the tight selection criteria from Table 7.3. Step 10 reject events which have a jet with $p_T \geq 120$ GeV, $|\eta| \leq 2$ and a charge fraction $f_{\text{charge}} \leq 0.02$, where the charge fraction is defined as the summed p_T of all tracks associated with the jet, divided by the p_T of the jet at EM+JES scale. The monojet channel does not define any signal region, but is kept for completeness.

2-Jet Channel

Steps 11 and 12 reject events if they do not have a second jet with at least $p_T > 40$ GeV, if they have a bad jet according to the tight definition or a jet with $p_T \geq 120$ GeV, $|\eta| \leq 2$ and a charge fraction $f_{\text{charge}} \leq 0.02$. The last step of the baseline selection in the 2-jet channel then requires $\cancel{E}_T > 100$ GeV in order to obtain an event sample for which the trigger can be

Signal region	A	B	C	D
Cutflow step	16	17	22	23
Target final state	squark–squark	squark–squark	gluino–gluino	squark–gluino
Jet multiplicity	≥ 2	≥ 2	≥ 3	≥ 3
$\cancel{E}_T/m_{\text{eff}}$	> 0.3	—	> 0.25	> 0.25
m_{eff} [GeV]	> 500	—	> 500	> 1000
m_{T2} [GeV]	—	> 300	—	—

Table 7.6: Overview of the signal regions defined for the analysis [105]. A cut at $\cancel{E}_T > 100 \text{ GeV}$ and $\Delta\phi(\text{jet}, \cancel{E}_T) > 0.4$ is common to all four signal regions as well as other cuts from the baseline selection which is explained in the text and summarized in Table 7.2.

assumed to be fully efficient. Steps 14 – 17 define the first two signal regions A and B (see below).

3-Jet Channel

Step 18 requires events to have a third jet with at least $p_T > 40 \text{ GeV}$. In contrast to the 2-jet channel, in the 3-jet channel the tight jet cleaning is not enforced, only a simplified additional jet cleaning is done: Events are removed if any of the leading three jets falls within $|\eta| \leq 2$, has a charge fraction $f_{\text{charge}} \leq 0.02$ and the fraction of its energy deposited in the electromagnetic calorimeter is $f_{\text{EM}} \geq 0.98$. Step 19 completes the baseline selection in the 3-jet channel, requiring $\cancel{E}_T > 100 \text{ GeV}$. The following steps define the signal regions C and D, with D being a subset of C.

4-Jet Channel

The 4-jet channel requires a fourth jet with $p_T > 40 \text{ GeV}$ and the same simplified jet cleaning as in the 3-jet channel, but extended to all four leading jets. It is kept for completeness and is not used for limit setting.

7.3.4 Signal Regions

Four signal regions are used in the analysis, optimized for different supersymmetric final states. The definitions of the signal regions are given in Table 7.6, following [105]^e, together with the final state they target. The cutflow for the baseline selection has been explained in the previous section. On top of the baseline selection, all four signal regions have a cut on the minimum angle between the vector of the direction of the missing energy and the direction of the three leading jets with $p_T > 40 \text{ GeV}$ at EM+JES scale projected onto the transverse plane. The minimum angle is denoted as $\Delta\phi(\text{jet}, \cancel{E}_T)$ and required to be larger than 0.4. The signal regions are then defined by cuts on the effective mass m_{eff} , the transverse mass m_{T2} or the ratio of missing transverse energy and effective mass $\cancel{E}_T/m_{\text{eff}}$ as summarized in Table 7.6

^e In some places in the internal note [146], the letters assigned to the signal regions are swapped, $A \leftrightarrow C$ and $B \leftrightarrow D$.

or the cutflow in Table 7.2. The cut on the ratio of missing transverse energy and effective mass $\cancel{E}_T/m_{\text{eff}}$ aims at quantifying the significance of the measured \cancel{E}_T . As said above, the four signal regions are not orthogonal, which is obvious for C and D, where every event from D is also contained in C, but also the other signal regions may overlap.

Signal regions A and B are optimized for squark-squark pair production which, under the assumption that $\tilde{q} \rightarrow q\tilde{g}$ is kinematically forbidden, will give two jets from each of the squarks decaying to a quark plus a neutralino or chargino. If squarks may decay into gluinos, even more jets are produced, but to include all events only two jets are required. Signal regions C and D require three or more jets as they arise in the decay of pair-produced gluinos, $\tilde{g}\tilde{g}$, or from the associated production of a squark together with a gluino, $\tilde{q}\tilde{g}$. The gluino can only decay through a squark, $\tilde{g} \rightarrow q\tilde{q}$, which increases the number of jets with respect to the direct production of squark pairs. Cutting on a higher number of jets, as is done in the three-jet channel, then gives better sensitivity to the SUSY signal over the background. In signal region D, tighter cuts than in signal region C can be used due to the higher total cross section of the associated production compared to gluino pair production.

7.4 Standard Model Backgrounds

This section will discuss the Standard Model backgrounds which are relevant for the search for Supersymmetry in the zero-lepton channel, before in Section 7.5 the Monte Carlo samples that are used to model the Standard Model background will be presented in detail.

In a search for new physics, where “new” refers to any type of physics that goes beyond the current version of the Standard Model, all known processes from the Standard Model comprise the background to the search. Usually background processes dominate over the theorized new physics processes in terms of cross section and hence event counts; background-free searches are the exception. Using suitable cuts on kinematic variables, event properties or derived quantities, the signal is separated from the background, with the goal that in the signal regions, the ratio of signal events over background events should be as high as possible. The ratio of the number of signal events s and background events b is an important figure of merit when optimizing the cuts in the analysis. Searches usually optimize the ratio s/\sqrt{b} , motivated by the fact that this ratio gives the approximate number of standard deviations $\sigma \sim \sqrt{b}$ that the background needs to fluctuate upwards to fake a potential signal. For measurements, typically the ratio s/b is considered instead. Note that s/b is constant as function of the luminosity, whereas s/\sqrt{b} grows as the square-root of the luminosity.

The optimal cuts are a compromise between purity and efficiency in the signal region. The purity measures the fraction of signal events in all events that reach the signal region, and the efficiency denotes the ratio of signal events that remain after the cuts defining the signal region to the original number of signal events. Monte Carlo serves as a test bed to optimize cuts with respect to purity and efficiency. Depending on their similarity to the signal topology, some types of backgrounds can be suppressed efficiently, others cannot. Backgrounds that are indistinguishable from signal events, because they give the same detector signature, are called irreducible.

The number of background events reaching the signal region needs to be estimated both for irreducible and reducible backgrounds. It may prove more difficult to estimate the amount of background contamination for backgrounds which are easy to cut away, but have

very large cross sections, than for backgrounds which are more signal-like, but also have relatively small cross sections. At the LHC, in contrast to lepton colliders, the event rate is completely dominated by QCD processes, giving rise to events with lots of hadronic jets. The cross section for this QCD background is many orders of magnitude larger than for all potential signal processes (cf. Figure 3.6). QCD events are expected to contain only little real \cancel{E}_T , which can arise from neutrinos produced in jets, but most of the \cancel{E}_T in QCD events will be fake \cancel{E}_T from mismeasurement of jets. The QCD background can therefore be strongly reduced by cutting on \cancel{E}_T , the ratio of \cancel{E}_T and m_{eff} , or the angle between \cancel{E}_T and the leading jet, so that in the analysis, it will only be one of the minor backgrounds. Still, the main challenge with respect to the QCD background is that events with high jet multiplicities are difficult to model in Monte Carlo (see e. g. [115]), so that the predictions of the total cross section and the number of QCD events entering the signal region from Monte Carlo suffer from large uncertainties.

Standard Model processes that yield real \cancel{E}_T and jets are very similar to the signal topologies and make up the main background. Three types of events contribute to this background:

- W + jets. W bosons that decay leptonically are an important background due to the \cancel{E}_T from the neutrino created in the decay. If the lepton is misidentified, out of acceptance or is a tau lepton, the event will pass the lepton veto, and additional jets will make it pass the jet selection. The dominating contribution in the signal regions here comes from the decays of W bosons including a tau lepton. In addition to leptonic W decays, W decays to bottom-antibottom quark pairs are considered explicitly.
- Z + jets. Z + jets events are the second important background to the analysis. They can in particular give rise to an irreducible background of the zero-lepton search, if the Z boson decays into a neutrino-antineutrino pair. Together with the additional jets, this gives a signature that exactly matches the topology of events which are looked for in the jet plus \cancel{E}_T based zero-lepton search. Less likely to pass the selection are Z boson decays to electron-positron and muon-antimuon pairs, because it is unlikely that both leptons fail to be identified, and therefore the lepton veto leads to a significant reduction of these two decay channels. For $Z \rightarrow \tau^+ \tau^-$, the lepton veto is not effective, as tau leptons are not regarded as leptons here, still electrons and muons from tau decays may trigger the lepton veto.
- Top quarks. As the top quark is the heaviest of the Standard Model particles, it is not surprising that its production cross section depends strongly on the center-of-mass energy. At the moment, top quarks contribute one of the subleading backgrounds, in between QCD processes and electroweak vector boson production. Two different types of contributions are considered in the analysis, the production of top quark pairs and the less likely single top production. Top quarks can be assumed to always decay via $t \rightarrow Wb$, so that every top quark produces at least one heavy b -jet, plus possibly \cancel{E}_T from the subsequent decay of the W boson. Top quark decays are therefore similar in their signature to W + jets events which also produce jets and \cancel{E}_T , so that most of the above said holds. Due the topology of their decays, there are a number of kinematic constraints which allow to identify top quark events, but this is not exploited in the present analysis.

7.5 Monte Carlo for Standard Model Backgrounds

7.5.1 Generators

In summary, the following classes of background events are considered: Decays of Z and W bosons, decays of pairwise or singly produced top quarks, and inclusive jet events. The latter are modeled via $2 \rightarrow 2$ matrix elements for the hard scattering at leading order, giving two hard jets [40], which make up the main component of the QCD background with a very large cross section. Top quark production is not included in the QCD background sample.

The W and Z samples are generated with ALPGEN [18], with HERWIG [57, 58] for the parton showering and fragmentation and JIMMY [64] for the underlying event. The top quark pair samples were generated with MC@NLO [75, 76], as were the samples with single top quarks [77, 78], both in combination again with HERWIG and JIMMY. The inclusive jet samples were generated with PYTHIA [97].

Different parton distribution functions and tunes of the minimum-bias and the underlying-event description are used in ATLAS in combination with the different Monte Carlo generators [151]. The tunes are modifications of phenomenological model parameters so that the MC predictions describe existing data (from other collider experiments such as those at the Tevatron) as well as possible. For the generation of the inclusive jet samples in the PYTHIA generator at leading order, the modified leading order parton distribution functions MRST 2007 LO* are used [83]. They are an optimized version of the LO parton distributions, introduced to amend the flaws observed for both LO and NLO parton distributions. They also aim at an improved consistency of cross sections and differential distributions with NLO calculations when using this optimized parton distribution with LO generators. For the generation of the events with top quark decays with MC@NLO, the CTEQ6.6 Parton Distribution Function (PDF) set [45, 184] is used for the matrix element, parton shower and underlying event. The W and Z + jets samples were generated using CTEQ6L1 [44].

7.5.2 Cross Sections and Background Normalization

A detailed list of the cross sections for the Monte Carlo samples that have been used is given in the Appendix in Tables A.1 and A.2. The cross sections σ which are given in these tables in most cases are the official cross sections used within the ATLAS Supersymmetry group [202]. They include k -factors and generator efficiencies.

In general, whenever possible, a data-driven normalization of the Monte Carlo samples, using additional measurements in dedicated control regions, is preferred over the normalization using the bare cross section as predicted from theory, because this cross section, being an independent prediction, cannot account for efficiencies of the detector. Furthermore, distributions of kinematic variables may not be correctly reflected in the Monte Carlo simulations. Elaborate studies of the background estimates from Monte Carlo have been done within the ATLAS SUSY group, using different approaches for each of the backgrounds listed above [146]. In general, very good agreement with the Monte Carlo prediction is found. This justifies to use the Monte Carlo predictions as a first approximation and not to repeat similar studies in the scope of this thesis. For completeness, the main results from [146] are briefly summarized in the following, before coming to the description of the normalization of the QCD background in Section 7.5.3.

W + Jets and Z + Jets

Cross sections at next-to-next-to-leading order (NNLO) for the W and Z samples have been computed with the FEWZ code. For the $Z \rightarrow l^\pm l^\mp$ processes, $l \in \{e, \mu, \tau\}$, the cross section at LO given by the generator ALPGEN has been scaled with a k -factor common to all subprocesses, such as to give a total cross section of 1.069 nb for each of the three lepton flavors. In the same way, the cross sections for $W^+ \rightarrow l^+ \nu_l$ and $W^- \rightarrow l^- \bar{\nu}_l$ have been scaled with a common k -factor, so that the total cross section amounts to 10.46 nb per flavor. The same cross section is assumed for all three flavors to enforce lepton universality, which holds in weak interactions up to minor corrections from the mass differences of the produced particles. Finally, for $Z \rightarrow \nu \bar{\nu}$, a NNLO cross section of 5.817 nb is used as total cross section summing over the three lepton flavors. Note that the sum of the cross sections for all $Z \rightarrow \nu \bar{\nu}$ processes in Table A.1 does not give 5.817 nb, because the generator efficiency in this case is significantly smaller than 1.0, and as stated above the cross sections listed in this table are the final cross sections used for normalization, which also include the generator efficiencies.

The production of bottom quark pairs in W boson decays, $W \rightarrow b \bar{b}$, is included via a separate MC sample. This sample has been found to have a reasonably small overlap with the other W + (light) jets samples of around 4 [145] to 10% [203], so that it can be used in combination with these. The difference between the W + light jets samples and the $W \rightarrow b \bar{b}$ sample is that the W + light jets samples include only hard jets from the matrix element produced by gluons or quarks from the first two generations, which are assumed to be massless for this purpose. Jets arising from b quarks can also be found in this sample, then having been produced in the hadron shower, but those jets will predominantly have $p_T < 15$ GeV [145]. For the cross sections of the samples describing this decay channel, the cross section computed by the generators is used and normalized with a k -factor of 1.22 as given in [145].

For the W and Z + jets background, the number of events in the data sample is too small to allow for a purely data-driven estimate of the background normalization. Hence, the prediction from ALPGEN Monte Carlo is used as central value, which can be justified by the good agreement with data. The systematic uncertainties on the central value, such as the jet-energy scale, the jet-energy resolution and luminosity, are derived from Monte Carlo, as explained in Section 7.7. Other uncertainties arise from control measurements, where the control region statistics, being the dominant uncertainty in the comparison between data and Monte Carlo, is taken as a systematic uncertainty of the Monte Carlo predictions. For Z + jets, the control measurement is done on a sample of events with $Z \rightarrow l^+ l^-$ decays and additional jets. Here, the uncertainties are the control region statistics and the acceptance, for which a conservative 100% for the unmeasured phase space region at high η is assumed. The uncertainty from the control region statistics is taken as a replacement for the theoretical uncertainties, which are expected to be smaller. It includes, for the high mass signal regions, for which the m_{eff} or m_{T2} cuts on the control regions had to be relaxed to have any events at all, an uncertainty from the necessary extrapolation from the control region to the signal region. Both acceptance and control region statistics give contributions to the uncertainty which are of the size of the jet-energy scale uncertainty which is the dominating uncertainty otherwise (cf. Table 7.14). For W + jets, the control measurement uses $W \rightarrow l \nu$ + jets. In addition to the control region statistics, the lepton reconstruction efficiency is taken into account as a systematic uncertainty. For both W and Z + jets, alternative methods have been tested and have been found to give consistent results among each other and with the ALPGEN predictions. The uncertainty from the control region statistics is always the dominating

Subsample	MC ID	$\hat{p}_{T,\min}$ [GeV]	$\hat{p}_{T,\max}$ [GeV]
J0	105009	8	17
J1	105010	17	35
J2	105011	35	70
J3	105012	70	140
J4	105013	140	280
J5	105014	280	560
J6	105015	560	1120
J7	105016	1120	2240
J8	105017	2240	∞

Table 7.7: The \hat{p}_T intervals covered by the subsamples of the PYTHIA Monte Carlo inclusive jet samples used in ATLAS.

uncertainty, again together with the jet-energy scale uncertainty.

Top Quarks

For the Monte Carlo sample with decays of pair-produced top quarks, the NNLO cross section of 160.79 pb declared in the Atlas Metadata Interface (AMI) is used, which is consistent with the value recommended by the SUSY group. This cross section is shared between two subsamples, one of which includes only events where both top quarks decay hadronically, $t\bar{t} \rightarrow WWbb \rightarrow bbqqqq$, and the other the rest of events where one or both top quarks decay leptonically, $t\bar{t} \rightarrow WWbb \rightarrow bbqq\nu\nu$ or $bbl\nu\nu$. Top quark decays give a very high jet multiplicity and, in case of leptonic decays, also \cancel{E}_T , and taking into account the comparably large cross section at the center-of-mass energy of the LHC, they are an important background. For the production of single top quarks in the s and t channel, a NLO cross section of 7.15 pb and 0.47 pb, respectively, and 14.6 pb for the associated production of a top quark and a W boson is assumed. This is similar to the values recommended by the SUSY group and follows [145]. A recent computation including higher-order corrections at next-to-next-to-leading-logarithm accuracy gives slightly higher values for s -channel and Wt production [66].

Top events are a sub-leading background in the signal regions. A partially data-driven cross-check has been done, using a control region measurement in the same way as for the other backgrounds. The control region has the same jet requirements as the signal regions, but with an additional muon, a b -tagged jet and $40 \text{ GeV} < \cancel{E}_T < 100 \text{ GeV}$. No additional uncertainties on top of the uncertainties which are discussed in detail in Section 7.7 are assumed in [146]. Also, no theoretical uncertainty is given, because it is found to be small compared to statistical and jet-energy scale uncertainties from several other cross-checks.

Inclusive Jets (QCD)

The inclusive jet sample generated with PYTHIA, modeling the QCD background, is split into several subsamples. The splitting is implemented by specifying a p_T range, setting

the CKIN(3) and CKIN(4) variables of the PYTHIA generator which determine the range of allowed \hat{p}_T values for the hard $2 \rightarrow 2$ process. \hat{p}_T is the transverse momentum defined in the rest frame of the hard interaction [97]. This slicing is advisable because of the steeply falling cross section as function of p_T , which would otherwise make it impossible to cover all of the large p_T range with reasonable statistics. The large differences in the cross sections between the slices that can be found in Table A.2 make obvious the need for slicing to get sufficient statistics. How the slicing translates into \hat{p}_T intervals is shown in Table 7.7. In general, the higher the number in Jn, the higher the \hat{p}_T values and consequently the lower the cross section. Samples J0 through J7 are exclusive samples, J8 is the highest inclusive bin without an upper bound on \hat{p}_T . J8 is not available as NTUP_SUSY and therefore not included in the analysis. Its cross section including k -factor is $8 \cdot 10^{-6}$ pb and thus negligibly small. The cross sections given in Table A.2 for the QCD samples are the generator cross sections provided in the AMI, multiplied by a scale factor of 1.28 that is discussed below. The values used by the ATLAS SUSY group differ from these by about 1 %.

QCD creates an overwhelming background before cuts, which can be efficiently reduced to give a small background after cuts in the signal regions. It is nevertheless difficult to model in Monte Carlo, and therefore needs to be well controlled. QCD events that enter the signal regions must have high jet multiplicity, and are likely not well modeled in $2 \rightarrow 2$ PYTHIA Monte Carlo. The baseline method in [146] is a control region measurement based on an inverted $\Delta\phi$ cut. From the control region measurements, an effective scale factor is derived

$$k_R = \frac{N_{\text{data}}^{\text{CR}} - N_{\text{MC}(W+Z+\text{top})}^{\text{CR}}}{N_{\text{MC}(QCD)}^{\text{CR}}} \quad \text{with} \quad N_{\text{exp}}^{\text{SR}} = k_R N_{\text{MC}(QCD)}^{\text{SR}}. \quad (7.2)$$

An inverted $\cancel{E}_T/m_{\text{eff}}$ cut has also been checked instead of the inverted $\Delta\phi$ cut, and gives results which are well compatible within the uncertainties. The reason to prefer the inverted $\Delta\phi$ cut is that it contains events with large missing momentum due to jet mismeasurements so that more systematics will cancel out. The central value of the QCD background normalization is not the bare Monte Carlo prediction, but its value multiplied by the k_R -factor. The uncertainties on this include generator systematics, heavy flavor systematics and non-QCD background uncertainties. The combined uncertainties on the nominal value for the QCD prediction lie above 100 % for all four control regions, because a conservative 100 % systematic uncertainty from varying detector conditions is assigned to the ratio.

The scale factor used for normalization of the QCD Monte Carlo to data that has been adopted in this thesis is $k_R = 1.28$. It is motivated below and consistent with all values for k_R from [146]. It can be regarded as a simple normalization, making the rate of data and Monte Carlo agree after the \cancel{E}_T cut.

7.5.3 Normalization of the QCD Background

Figure 7.2 shows the factor for the normalization of the QCD background that would be needed to make data to Monte Carlo agree, binned in \cancel{E}_T . It is computed from a direct comparison of the Monte Carlo and data after the cuts up to and including step 12 from the cutflow in Table 7.2, i. e. before the \cancel{E}_T cut in the 2-jet channel. This scale factor is computed for each bin according to Equation (7.2) as the ratio of the difference of the event counts in data and the sum from all background Monte Carlo samples, apart from the PYTHIA QCD samples, over the number of QCD events in Monte Carlo. Obviously, the scale factor is not




Figure 7.2: Ratio of the estimation of the number of QCD events in data and in Monte Carlo, using the missing transverse energy \cancel{E}_T for the binning of the horizontal axis. The red line marks the value 1.28 of the scale factor that is applied to the cross section of the QCD Monte Carlo in the analysis. “MC[WZt]” denotes the sum of events in weak boson and top quark Monte Carlo samples.

independent of \cancel{E}_T . For small \cancel{E}_T it is below one, for large values it settles around a value that is consistent with 1.28, the value derived by the ATLAS SUSY group from different data-driven methods as said above. The good agreement between these values can be taken as a sanity check and motivation to use the value of 1.28 as scale factor for the normalization of the PYTHIA QCD samples.

The dependence of the scale factor on \cancel{E}_T indicates that the distribution of \cancel{E}_T in Monte Carlo does not describe the shape of \cancel{E}_T in data well, at least for small values and especially in the PYTHIA QCD sample which makes up the dominating background at this stage (cf. Figure 7.10). However, this part of phase space is cut away by the baseline selection cuts which cut on $\cancel{E}_T > 100$ GeV. After the cut on \cancel{E}_T at 100 GeV, a normalization factor of 1.28 can be expected to give a good agreement of data and Monte Carlo. Figure 7.3 shows scale factors which are computed in the same way, but binned in another important kinematic variable, the p_T of the leading jet. The two plots show the distributions at different steps in the cutflow (cf. Table 7.2). The left plot in this figure is after step 7, i. e. directly after the lepton veto and without any cuts on the leading jet. The dependence of the scale factor on this variable is much less pronounced. The scale factor is smaller than 1.28 over the whole range, but this is due to the dominance of events with small \cancel{E}_T , which have a smaller scale factor as shown in the previous plot in Figure 7.2. After applying the cuts up to step 13, thereby including the cut $\cancel{E}_T > 100$ GeV, the scale factors again are consistent with 1.28 over the whole range, as can be seen in the right plot in the same figure.

7.5.4 Pile-up Effects in Monte Carlo

The policy for producing Monte Carlo with pile-up contributions has changed from 2010 to 2011, following the evolution of the beam conditions. In 2010, in addition to the sam-




Figure 7.3: Ratio of the estimation of the number of QCD events in data and in Monte Carlo, using the p_T of the leading jet at EM+JES scale for the binning of the horizontal axis. The red line again marks the value 1.28. The left plot is before, the right plot after cuts on \cancel{E}_T and leading jet p_T .

ples without pile-up, samples have been produced in which the events are overlaid with minimum-bias events before the reconstruction step. The average number of additional minimum-bias events was fixed to two or five, with a large bunch spacing of 900 ns and accordingly only in-time pile-up. In 2011, instead of using a fixed average number of overlaid events, this number was varied, with an equal number of events being produced for Poisson expectation values of six to ten, and less for lower and higher values (from zero up to 18). Moreover, the simulated bunch train structure was modified, using now three bunch trains with 225 ns separation and a bunch spacing of 50 ns within the trains, consistent with the online conditions and leading to both in-time and out-of-time pile-up.

The Monte Carlo samples which are used to model the background as described above do not contain pile-up contributions. However, in the analysis which is presented in this chapter, the impact of pile-up is taken into account in terms of a systematic uncertainty, using a different Monte Carlo reconstruction that includes pile-up contributions. The Monte Carlo for all background types described above, except for the QCD samples, are available without pile-up, as well as with pile-up from two overlaid minimum-bias events. The baseline analysis is done using the Monte Carlo samples without pile-up contributions. The difference to this central value when using the pile-up samples instead, is taken to be a systematic uncertainty from pile-up contributions. This approach is chosen, because a reweighting of the Monte Carlo events to match the distribution of pile-up in data is not possible. Such a reweighting would require to have Monte Carlo events which are overlaid with a range of different numbers of additional pile-up interactions.

As for the PYTHIA QCD samples no pile-up version is available, an estimation of the size of the pile-up uncertainty has been tried on ALPGEN QCD samples instead. It has been found that the size of these samples is too small to have sufficient statistics in the signal regions. Therefore, for the QCD samples, the pile-up uncertainty is set to 100%.

For the signal samples described in Section 7.6, no Monte Carlo version with overlaid pile-up is available. Supersymmetry events in general are assumed to give hard (high energy) objects, whereas the overlaid pile-up events are rather soft and contribute little to the variables that are cut on to select events. Therefore, for the SUSY signal, neglecting pile-up ought

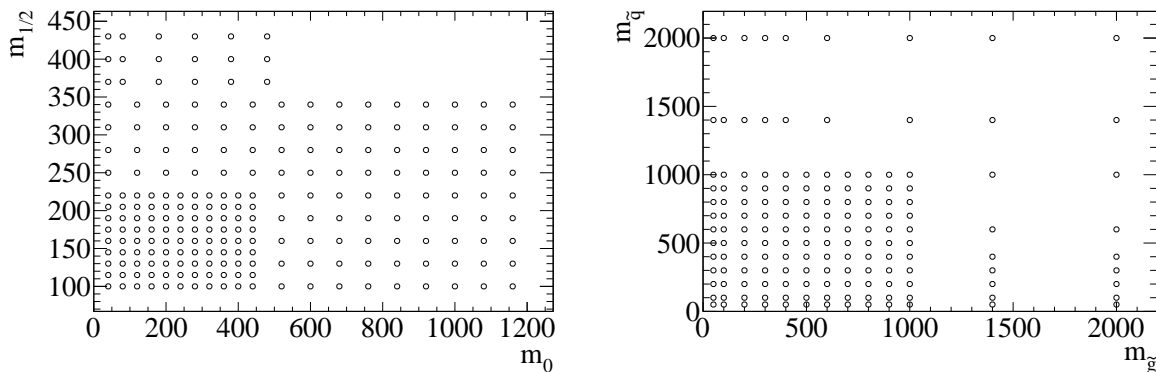


Figure 7.4: Overview of the points defined in the two grids used in this thesis. Left: mSUGRA grid, defined in the m_0 - $m_{1/2}$ plane. Right: MSSM grid, defined in the gluino- ($m_{\tilde{g}}$) and squark-mass ($m_{\tilde{q}}$) plane.

to be an acceptable approximation for 2010, where the average number of interactions per bunch crossing in most runs lay below 3. For 2011, this approximation should be revisited, the pile-up level being several times as high as in 2010.

7.6 Signal Grids

As explained in Section 2.3, the parameter space of Supersymmetry models has a very large dimensionality. A restriction to a specific mechanism of Supersymmetry breaking like the minimal supergravity framework (cf. Section 2.3.7) reduces the number of free parameters considerably. Still, this number often is too large to scan the complete phase space, mainly due to limitations in computing time and storage space needed to generate Monte Carlo for the necessary number of grid points. To ease the presentation of results, for example when visualizing exclusion limits, often two-dimensional subspaces of the parameter space are chosen, which are then scanned using a fixed grid of parameter sets for which Monte Carlo is generated. In this thesis, two such grids of points in parameter space are used. Each point in parameter space corresponds to a specific choice of SUSY scenario, and for each point a sample of Monte Carlo events is generated.

7.6.1 mSUGRA Grid

The first grid is defined in the plane spanned by the m_0 and $m_{1/2}$ parameters of mSUGRA for a relatively low value of $\tan \beta = 3$. The trilinear coupling is set to $A_0 = 0$ GeV, and the sign of the Higgsino mass parameter is positive, $\mu > 0$. The Monte Carlo event generation has been done with HERWIG++ [59, 60] and ISAJET 7.79 [63, 188], using ISASUGRA from this program suite as spectrum calculator to calculate the weak scale parameters from those at the high scale, and the MRST 2007 L_0^* parton distributions [44].

The mass range that is covered by this grid is $40 \text{ GeV} < m_0 < 1160 \text{ GeV}$ and $100 \text{ GeV} < m_{1/2} < 430 \text{ GeV}$, with a varying density of grid points that for low m_0 - $m_{1/2}$ values is four times as high as in the rest of the plane to allow for a finer resolution of the exclusion contour in this region. An overview of the points defined in this grid is shown in the left plot in




Figure 7.6: Mass spectrum of the supersymmetric particles at one example point of the MSSM grid, with $m_{\tilde{q}} = 1200$ GeV, $m_{\tilde{g}} = 1600$ GeV and a massless lightest neutralino $\tilde{\chi}_1^0$. The masses of all other supersymmetric particles including \tilde{b} and \tilde{t} are set to 5 TeV. Arrows indicate decays with a branching ratio of at least 1 %.

Figure 7.4 in which each marker corresponds to one MC dataset. Note that the position of the dots in Figure 7.4 is exact, whereas in the two-dimensional color-coded histograms in the following the grid points may appear to be slightly shifted, but this is only an artefact from the binning of the histograms. The number of events within the MC datasets varies between approximately $5 \cdot 10^3$ and $1.5 \cdot 10^5$. This grid consists of 222 points and will be referred to as the mSUGRA grid in the following.

Figure 7.5 shows four examples for mass spectra from the mSUGRA grid with low and high m_0 and $m_{1/2}$, where low and high refer to the mass range covered by the mSUGRA Monte Carlo grid used for the analysis. The two upper plots show the spectrum for $m_{1/2} = 100$ GeV, the lower plots for $m_{1/2} = 340$ GeV; the two plots to the left have m_0 set to 120 GeV, the two plots to the right have $m_0 = 1000$ GeV. All other parameters, $\tan \beta$, $\text{sign } \mu$ and A_0 are set to the values used for the production of the mSUGRA grid. The dashed lines show the most important decays with a branching ratio of at least 1 %, the thickness and opacity of the line indicating at the branching ratio.

7.6.2 MSSM Grid

The second grid is a non-universal supergravity model with minimal particle content, in which the grid points have a fixed, low value for the LSP and vary in the masses of the squarks and gluinos, both covering the same range of $50 \text{ GeV} < m_{\tilde{g}} < 2000 \text{ GeV}$ and $50 \text{ GeV} < m_{\tilde{q}} < 2000 \text{ GeV}$. This type of grid is available for three different masses of the LSP, 0, 50, and 95 GeV, where for the massive LSPs only points have been simulated for which the lighter of $m_{\tilde{g}}$ and $m_{\tilde{q}}$ is less than $3 \cdot m_{\text{LSP}}$. The masses of all other supersymmetric particles, including the masses of the third generation of scalar quarks, are set to 5 TeV to decouple them. The mass spectra at the different grid points therefore all are very simple, and identical except for the squark and gluino masses. An example is given in Figure 7.6.

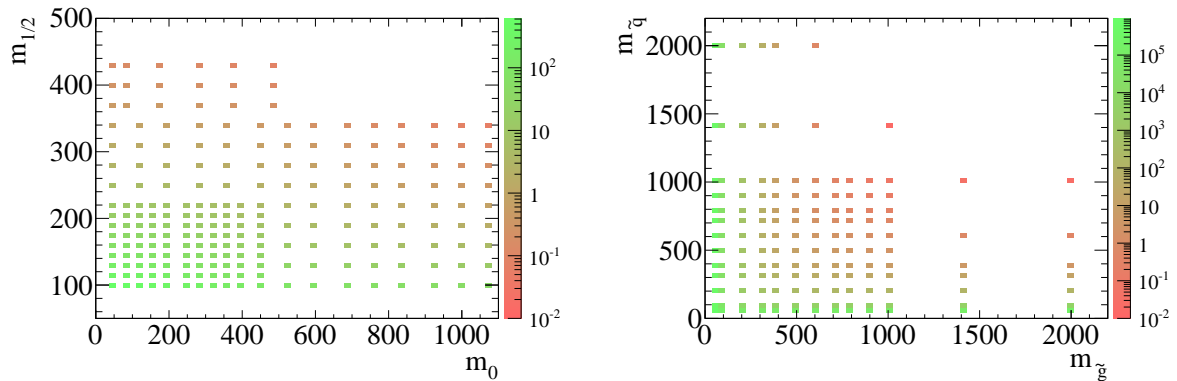


Figure 7.7: Cross sections of the SUSY grid Monte Carlo samples in picobarn, left: mSUGRA grid, right: MSSM grid.

The arrows indicate decays with a branching ratio larger than 1%. Particles without outgoing arrows are not stable, but dominantly decay to non-supersymmetric particles which are not shown in this plot. The H^\pm bosons e.g. decay almost always to $t\bar{b}$ or $\bar{t}b$ pairs, and the gluino decays to all light squark flavors with the same branching ratio of 1/16. The mass spectrum for this grid is generated using ISASUGRA in a non-universal supergravity model, using a set of soft-symmetry breaking parameters that should approximately give the desired masses. The masses returned by ISASUGRA slightly differ from those specified, and are therefore edited by hand, then ISASUSY is used to calculate the branching ratios and decay widths [194].

In the following, only the grid with a massless LSP is considered, because it covers the full plane within the limits given above. This grid consists of 165 points and will be referred to as the MSSM grid in the following. An overview of the points defined in this grid is shown in the right plot in Figure 7.4. Note that this plot only shows 153 points, because the remaining points lie at much larger values of the masses, at around 5 TeV, and are not relevant in the context of the following analysis. The number of events within the MC datasets is approximately 10^4 for all grid points, apart for very few exceptions for which probably some of the reconstruction jobs failed.

In Figure 7.7, the cross sections of the SUSY grid Monte Carlo samples for the mSUGRA grid (left plot) and the MSSM grid (right plot) are shown. In both cases, the highest cross sections are found for low masses, of course. In the mSUGRA grid, the cross section depends strongly on $m_{1/2}$ and less so on m_0 . For the MSSM grid, the dependence on the squark and gluino masses is comparable. Note that the MSSM grid covers a range up to much larger cross sections than the mSUGRA grid which are attained at grid points with very low squark and gluino masses.

7.6.3 Final-State Dependent Cross Sections

There is a large variety of different Monte Carlo samples for Standard Model processes, which together make up the background in a search for new physics. Instead of having one big Monte Carlo sample containing all possible types of events, for several reasons, as outlined in Section 4.5, it makes more sense to simulate different classes of physics processes individually and combine them again later. In the combination, an appropriate weight is

ID	Process	ID	Process
1	squark – gluino	71, 72, 73, 74	$\tilde{\chi}_{1,2,3,4}^0$ – gluino
2	gluino – gluino	75, 76	$\tilde{\chi}_{1,2}^+$ – gluino
3	squark – squark	77, 78	$\tilde{\chi}_{1,2}^-$ – gluino
4	squark – antisquark	81, 82, 83, 84	$\tilde{\chi}_{1,2,3,4}^0$ – squark
51	sbottom 1 pair	85, 86	$\tilde{\chi}_{1,2}^+$ – squark
52	sbottom 2 pair	87, 88	$\tilde{\chi}_{1,2}^-$ – squark
61	stop – antistop 1		
62	stop – antistop 2		

Table 7.8: ID numbers of the final states used to assign events to classes with given cross sections.

applied to each event, which is computed from the cross section of the respective physics process.

For the SUSY signal samples, there is only one sample of Monte Carlo events for each point in parameter space. The distinction between different processes, all of which are of supersymmetric nature, is done afterwards in the analysis. Using the Monte Carlo truth information, every event is identified as belonging to a certain class. The classes are defined by the particles in the final state of the hard scattering process, which contains two supersymmetric, pair-produced particles, e. g. two squarks or a squark and a gluino. (Although these two particles stand at the beginning of the decay cascade, they are the final state from the point-of-view of the generator of the hard scattering process.) Instead of using one overall cross section for the full MC sample to compute an average weight for each event, this procedure allows to use individual weights for each subprocess that produces SUSY particles. These subprocesses do not only have different cross sections, but also the acceptance of the event selection that is done in the analysis will depend on the particles that were produced, and therefore will vary significantly between different final states. In a simplified example, a selection requiring two jets would mostly select squark pairs, because (assuming $\tilde{q} \rightarrow q\tilde{g}$ is kinematically forbidden) the two final state squarks will each decay to a quark giving a jet plus a neutralino or chargino. Requiring four jets would prefer gluino pairs, each of which can only decay through a squark, $\tilde{g} \rightarrow q\tilde{q}$, and therefore likely give two quarks, *ergo* two jets, in the decay chain (see e. g. [106]). Using the overall cross section is not wrong, but yields an averaged result. Different final states may have different k -factors, and for SUSY processes the k -factors tend to be large [27], so that going from the generator cross section to the NLO cross section will enhance different final states by different factors.

Computing the cross sections takes quite a lot of CPU time, because there are many grid points, many final states for each, and the estimation of the systematic uncertainties (cf. Section 7.7.7) again increases the computational effort by a significant factor. The cross sections are therefore computed centrally by the ATLAS SUSY group using PROSPINO and provided as a ROOT file for each of the SUSY grids. Not all possible final states are considered in this computation, but only those that are expected to be dominant. Table 7.8 lists 24 classes of events that are defined by the final state, together with the integer numbers which are used as their ID in the official cross section tables.

For the MSSM grid, cross sections are available only for final states with IDs 1 – 4. For the mSUGRA grid, also the other 20 final states from Table 7.8 are taken into account. Of course, these 24 final states do not cover all possibilities, and events which in the analysis are identified as not belonging to one of the 24 classes are ignored, because their cross section is unknown. Ignoring these unidentified events is equivalent to assuming a cross section of zero for the unidentified processes and means that the expected number of signal events is artificially reduced, leading to exclusion limits that are more conservative than those that would be obtained when including all events. Tables A.3 and A.4 in the Appendix give an overview of all final states that are encountered in the Monte Carlo signal grid samples. From Table A.3, it is clear that in the MSSM samples using final states with IDs 1 – 4 is mostly sufficient, because only these types of final states exist in the Monte Carlo, apart from final states where a neutralino is produced in association with a gluino or squark and some more, negligibly small exceptions. In particular, final states never include scalar bottom nor scalar top quarks. Looking at Figure 7.6 again, this makes clear that the heavy supersymmetric particles are indeed decoupled, because only the much lighter squarks and gluinos are produced, and only their decays are relevant. For the mSUGRA grid, the variety of final states is larger, including scalar bottom and scalar top quarks and also associated production with charginos and neutralinos other than the lightest neutralino.

According to the documentation of PROSPINO, for final states including squarks (ID 1, 3 and 4 in Table 7.8), a summation over both squark chiralities and over all possible squark flavors is understood, excluding stops [26]. The sum over the squark flavors is due to the computation being done in a supergravity-inspired model in which all squarks are assumed to have a common mass. This is not correct for stops, and these are therefore excluded from the final states. The official statement from the ATLAS SUSY group is that scalar bottoms are included, but in the updated version of PROSPINO, which was used to compute the cross sections used by the SUSY group, scalar bottom quarks can be treated separately and are not included in squark–gluino and squark–squark by default, whereas they are included by default in squark–antisquark[‡]. The final states with ID 1 and 3, listed in the table as squark–squark and squark–gluino, also include a summation over charge-conjugate final states [26, 27]. The final states with ID 51 and 52 are taken to refer to sbottom–antisbotttom production corresponding to the respective final state option in PROSPINO[§]. Final states with ID 61 and 62 refer to $\tilde{t}_1\tilde{t}_1$ and $\tilde{t}_2\tilde{t}_2$. Mixed pairs ($\tilde{t}_1\tilde{t}_2$ and $\tilde{t}_2\tilde{t}_1$) can only be produced at higher orders ($\mathcal{O}(\alpha_s^4)$) at strongly suppressed rates [28], which can also be seen from Table A.4, where $\tilde{t}_1\tilde{t}_1$ appears about 6000 times as often as $\tilde{t}_1\tilde{t}_2$.

From the above considerations, the following mapping is derived. Final states that consist of a squark–antisquark pair, one of which or both may be scalar bottom squarks, are all included in ID 4, independently of which of the two mixed mass states the sbottom and / or antisbotttom is. The cross sections computed for ID 51 and 52 are therefore not used in the analysis, because they are a subset of all squark–antisquark final states and overlap with ID 4. Events with final states including scalar bottom quarks that do not fall into the class with ID 4, e. g. sbottom–sbotttom or sbottom–gluino, are ignored, because they are also not

[‡] This can be seen in the output of the official version of the PROSPINO code used by the SUSY group, where an integer value `isquark*` is running from `-4` through `4` for `final_state_in = 'ss'`, which corresponds to the different squark flavors and does not include scalar bottoms. Scalar bottoms would be `-5` and `5` [177], values that only appear when computing the cross section for squark–antisquark final states (`final_state_in = 'sb'`).

[§] `final_state_in = 'bb'`

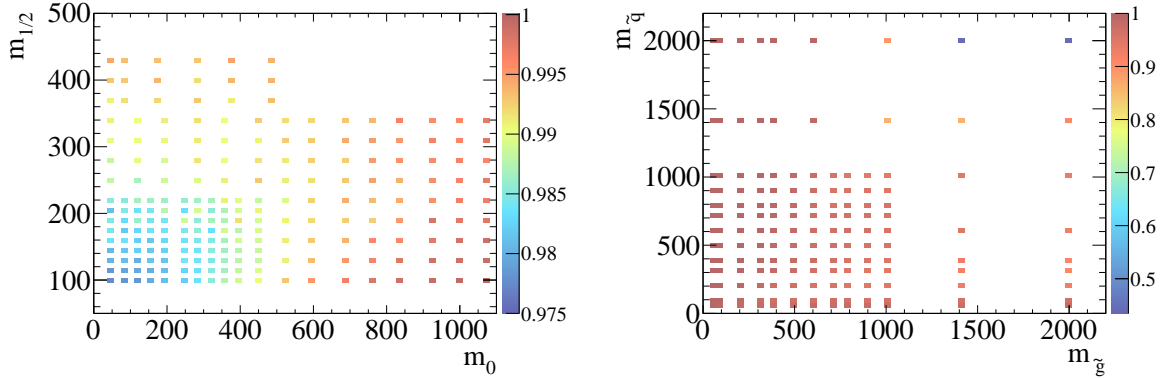


Figure 7.8: Fraction of events in the Monte Carlo samples from the signal grids for which the cross section is known. Left: mSUGRA grid, right: MSSM grid.

included in the cross section computation by PROSPINO for ID 1, nor for ID 3.

With respect to the implementation, the classification of events according to the final state of the hard process, as given in Table 7.8, is done by evaluating the Monte Carlo truth information, searching the Monte Carlo tree for the two SUSY particles that have a Standard Model particle as ancestor. From the identity of these particles, it can be determined to which of the defined classes the event belongs. If it does not match any of the classes, the event is discarded.

The impact of neglecting certain classes of events is discussed briefly now. It can be estimated from the fraction of events that do not fall into one of the classes for which the cross section is known^h. If this fraction is small, this is an indication that the impact should be negligible. Figure 7.8 shows two plots, with the fraction of events in the grid samples for which the cross section is known, i. e. which fall into one of the classes as defined above, for which the PROSPINO NLO cross section is available. The left plot shows the fractions for the mSUGRA grid, the right one for the MSSM grid. For the mSUGRA grid, the fractions increase from about 98 % in the lower left corner to nearly a 100 % close to the right and top border. In the MSSM grid, for low $m_{\tilde{g}}$, the fraction of known events is close to 100 %, and falls off to the right, going down to 94 % in the dense grid region ($m_{\tilde{g}}$ and $m_{\tilde{q}}$ below 1 TeV). In the upper right corner, the fraction is only 45 %. Because of the large masses of the squarks and gluinos, the associated production of squarks and gluinos with neutralinos $\tilde{\chi}_1^0$ becomes dominating here. Note that the yield of the different final states with squarks and gluinos depends in different ways on the squark and gluino masses [27]. For $m_{\tilde{q}} \approx m_{\tilde{g}}$, squark–gluino production dominates at low masses, whereas at higher masses of about 1 TeV squark–pair production has the highest yield. From the plots in Figure 7.8, it can be concluded that the impact of the neglected events should not deteriorate the power of the analysis too much. A complementary cross-check is presented in Section A.8.2. In any case, as said above, neglecting events which are contained in the Monte Carlo will make the limits more conservative, but it is safe in that it cannot lead to false exclusions.

^h The overall magnitude of this fraction of this can also be taken from the counts in Table A.4.

Source	Measured	Signal / Background	Symmetric	Correlated
Statistics	on MC	S / B	yes	no
Integrated luminosity	fixed	S / B	yes	yes
Jet resolution	on MC	S / B	yes	yes
Jet-energy scale	on MC	S / B	no	yes
Pile-up contribution	on MC	- / B	yes	yes
Background normalization	fixed	- / B	yes	no
Signal normalization	fixed	S / -	yes	no

Table 7.9: Overview of systematic uncertainties. The second column describes how the impact of the uncertainty was estimated, the third states which Monte Carlo the uncertainty applies to. The last two columns state whether this uncertainty is symmetric and whether it is correlated between the signal and the different background contributions. Further explanations can be found in the text.

7.7 Statistical and Systematic Uncertainties

The statistical evaluation of the observed data requires a careful treatment of its systematic and statistical uncertainties, because only then it becomes possible to arrive at a conclusion which can be said to hold at a certain probability. Table 7.9 gives an overview of the systematic uncertainties that were considered. Statistical uncertainties are included in this list, too, because they enter the limit setting procedure (see Section 7.11) in a fashion very similar to the systematic uncertainties. It is common to assume the number of observed events in the data to be a fixed number without uncertainties and to assign the statistical and systematic uncertainties to the Monte Carlo prediction. The uncertainties are discussed in detail below. A table with the values of the uncertainties is included in Section 7.8.

All uncertainties are modeled as Gaussian distributions with one exception. The jet-energy scale is modeled as a Gaussian distribution, where the two halves, corresponding to upwards and downwards fluctuations, may have different widths to take into account the asymmetric nature of the jet-energy scale uncertainty. The distribution may thus be discontinuous at zero. The pile-up contribution is also modeled as a Gaussian distribution, allowing for upwards as well as downwards fluctuations. One might expect pile-up contributions to only increase the number of events after cuts, because pile-up always adds energy to the event, but there are e. g. cases in which objects which lead to an event veto profit from additional energy due to overlaid pile-up events, and pass a selection which they would not have without pile-up. Indeed, although increases in event numbers dominate, negative changes in the event numbers are found, too, when including pile-up.

The fluctuations covered by systematic uncertainties can either be independent between different background types and the signal, such as the purely stochastic fluctuations due to limited Monte Carlo statistics, or they can be correlated such as the uncertainty on the luminosity, which will lead to an up- or downscaling of all backgrounds and the signal Monte Carlo alike. In general, there are more possible levels of correlations between systematic uncertainties, between different physics processes, between different regions of phase space, or, on a larger scale, between different analyses or experiments. It is important to take cor-

relations between uncertainties into account, because fluctuations due to uncertainties will lead to larger variations in the outcome if the uncertainties are correlated. (Uncorrelated uncertainties are added in quadrature, whereas absolutely correlated uncertainties need to be added linearly.) The last column in Table 7.9 shows which systematic uncertainties are taken to be correlated across the signal and background Monte Carlo samples. The uncertainty on the signal normalization affects only the signal Monte Carlo sample, so there is no possible correlation.

7.7.1 Statistical Uncertainty

The statistical uncertainty is determined by the number of events available within the Monte Carlo samples. These numbers are given in Tables A.1 and A.2, together with the integrated luminosity the complete Monte Carlo dataset corresponds to. As for almost all samples (all but QCD) this luminosity is comfortably above the integrated luminosity of the data, the statistical uncertainty is not expected to be a dominant contribution, at least unless very harsh cuts are applied as in some of the signal regions. It can be directly evaluated from Monte Carlo, assuming Poisson distributions for the bin contents or the number of events at a certain step of the cutflow, respectively. If N_i is the number of events from Monte Carlo sample i surviving all cuts up to a certain stage, the uncertainty σ_i on the contribution from this sample is taken to be the square root of the number of events $\sqrt{N_i}$, which is just the standard deviation of a Poisson distribution with expectation value N_i . The total uncertainty is the weighted sum over all Monte Carlo samples. Denoting the weight of sample i with w_i , the total uncertainty σ_N is therefore given by

$$\sigma_N = \sqrt{\sum_i w_i^2 \sigma_i^2} = \sqrt{\sum_i w_i^2 N_i}. \quad (7.3)$$

This follows from the standard rule for the propagation of (uncorrelated) errors, which assumes Gaussian distributions for the individual uncertainties. This assumption is justified if N_i is sufficiently large.

7.7.2 Uncertainty on the Jet-Energy Scale

The Jet/Etmiss combined performance group provides an official tool called JESUNCERTAINTYPROVIDER, that can be used to assess the uncertainties on the Jet-Energy Scale [159, 164]. For data reconstructed with ATHENA release 15, as is used here, the values on which the uncertainties from this tool are based have been computed from Monte Carlo samples and in-situ measurements for the calorimeter response uncertainty. The updated JES uncertainty for release 16 would additionally include an individual treatment of isolated jets and close-by jets and take differences in the quark and gluon response into account [167].

In this analysis, version `JetUncertainties-00-02-00` of the JESUNCERTAINTYPROVIDER tool is used. The impact of the JES uncertainty is propagated to the event counts in the signal regions by rerunning the Monte Carlo twice, once with the JES scaled down by the given uncertainty and once with JES scaled up. Implementation-wise, the JESUNCERTAINTYPROVIDER tool is called once for each jet with the η and p_T of the jet and the number of vertices reconstructed for this event. The number of vertices is used to account for the energy contribution to calorimeter jets from multiple proton-proton interactions in pile-up events which is not included in the current JES calibration. Instead, it is taken to give a

separate contribution to the systematic uncertainty on the jet-energy scale. As the uncertainties have only been determined for a p_T range from 15 GeV to 1 TeV and for a η range within ± 4.5 , the tool only accepts these ranges for the values of the input parameters. For jets which fall outside these ranges, the closest values within the allowed ranges are therefore used to query the uncertainty. The tool returns an uncertainty Δ for the given parameters, and the p_T and energy of the jet are multiplied with a factor $1 \pm \Delta$ for the computation of the upwards and downwards fluctuation due to the JES uncertainty, respectively. The \cancel{E}_T is also updated to take the change in jet p_T into account.

7.7.3 Uncertainty on the Jet-Energy Resolution

To compute the uncertainty on the jet-energy resolution, another tool called JETENERGYRESOLUTIONPROVIDER is provided, which returns two numbers, the relative jet transverse momentum resolution σ_{p_T} and its uncertainty σ_r , both as function of jet p_T and rapidity y . The recipe to convert these numbers into a systematic uncertainty on the event count and reproduce the slightly worse resolution in data (see Section 4.4.2) is to calculate a smearing factor σ_s as

$$\sigma_s = \sqrt{(\sigma_{p_T} + \sigma_r)^2 - (\sigma_{p_T})^2}. \quad (7.4)$$

The energy and p_T of the jet is then smeared with a Gaussian of width σ_s centered around zero,

$$p_T \mapsto p_T \cdot (1 + \text{Gauss}(0, \sigma_s)), \quad (7.5)$$

and equivalently for the energy. Here, σ_{p_T} is assumed to correspond to the resolution in Monte Carlo σ_{MC} , and $(\sigma_{p_T} + \sigma_r)$ to the resolution in data σ_{data} . Equation (7.4) follows from the fact that in the smearing the standard deviations add quadratically, $\sigma_{\text{data}}^2 = \sigma_{\text{MC}}^2 + \sigma_s^2$.

The Monte Carlo is rerun with the smeared p_T and an updated \cancel{E}_T taking the change in the jet p_T into account, and the difference in the event count is used as the absolute systematic uncertainty. As the values provided by the JETENERGYRESOLUTIONPROVIDER have only been determined and validated for a p_T range from 20 GeV to 500 GeV and for a $|\eta|$ range within 2.8, the tool only accepts these ranges for the input parameters. For jets which fall outside these ranges, the closest values within the allowed ranges are therefore used to query the tool. Note that even though only jets that lie within the allowed parameter range of the tools will pass the cuts imposed in the event selection (see Section 7.3), in both cases, JES and Jet-Energy Resolution (JER), jets outside this region need to be considered, too, because of the propagation of the change in jet p_T to \cancel{E}_T . To do so, limiting the parameters as described above is necessary. The JER uncertainty is a correlated uncertainty, although it is computed event-wise with a smearing based on a random distribution and therefore not as obviously correlated as the JES uncertainty. But the change in event numbers induced by a deviation of the true JER from its assumed value within the uncertainty will affect all Monte Carlo samples in the same way as a deviation of the true JES from its assumed value does affect all samples in a correlated manner.

7.7.4 Uncertainty on the Integrated Luminosity

The computation of the integrated luminosity has been explained in Section 7.2. The luminosity uncertainty is taken from [165], which gives the official value of $\pm 3.4\%$ for the

Signal region	Uncertainty [%]	
	Z + jets	W + jets
A 2 jets + $m_{\text{eff}} > 500$ GeV	40	21
B 2 jets + $m_{T2} > 300$ GeV	68	71
C 3 jets + $m_{\text{eff}} > 500$ GeV	44	23
D 3 jets + $m_{\text{eff}} > 1000$ GeV	73	55

Table 7.10: Relative uncertainty in percent on the W + jets and Z + jets background in the four signal regions as given in [146].

average relative uncertainty on the central value of the updated luminosity measurement in 2010 (cf. Section 4.4.7). This uncertainty has a direct impact on the normalization, so that no rerunning of the Monte Carlo is necessary to estimate the variation in event numbers caused by this uncertainty. It affects signal and background in the same way, and is therefore taken to be a correlated uncertainty for signal and background.

7.7.5 Uncertainty from Pile-Up

For the top quark and vector boson samples, the uncertainty due to in-time pile-up effects is estimated by rerunning the selection on Monte Carlo samples including two additional overlaid pile-up interactions per event. From the difference in event numbers with respect to the nominal results without pile-up a relative uncertainty is computed. For the QCD background, this is not possible because no pile-up version of the PYTHIA QCD samples is available. There are ALPGEN QCD samples which are available both with and without pile-up, but these samples are rather small in terms of event numbers or corresponding sample luminosity, so that in three out of the four signal regions no events are left. A conservative estimate of the uncertainty from pile-up of 100 % has therefore been assigned to the PYTHIA QCD background.

7.7.6 Theoretical Uncertainties on the Background Normalization

Theoretical uncertainties include uncertainties on the distributions of kinematic variables within the sample and uncertainties on the predicted cross section of a Monte Carlo sample as a whole, which then give an uncertainty on the normalization of the events in this sample. For the analysis at hand, details about the uncertainties on the cross sections of the signal samples are presented in below. The uncertainties on the normalization of the Standard Model backgrounds can be expected to be under better control due to existing measurements and are assumed to be covered by the uncertainties taken from the official SUSY group studies summarized in Section 7.5. As described in Section 7.5.2, further studies within the ATLAS Supersymmetry group on the uncertainties of the Monte Carlo for the W + jets and Z + jets background have been done. As the uncertainties from the control region measurements are found to give a large contribution to the total uncertainty for these samples, these results are used here in addition to the uncertainties described above. The values in Table 7.10 are taken from [146], and state the relative uncertainty in percent on the W + jets and Z + jets background in the four signal regions from acceptance and control

region statistics (for $Z + \text{jets}$) and from lepton efficiency and control region statistics (for $W + \text{jets}$).

7.7.7 Theoretical Uncertainties on the Signal Normalization

The cross sections for the pair-production of SUSY particles in the scenarios specified by the grid points have been computed by the ATLAS SUSY group at NLO order using PROSPINO 2.1 [26, 198]. In addition to the cross sections themselves, theoretical uncertainties on these coming from several sources are provided [174]:

- `relUncPDF`: The relative uncertainty from the Parton Distribution Function (PDF) that is used. In the computation of the cross sections, the PDF set CTEQ6.6 [45, 184] from the CTEQ group has been used, which has 22 free parameters and is extrapolated from measurements and therefore entailed with its own systematic uncertainties that need to be propagated. The central fit PDF gives the nominal cross section and the uncertainty on this value is estimated by recomputing the cross sections for each of the 44 eigenvector PDF sets and combining the results as $1.645\Delta_{\text{PDF}} = \frac{1}{2}\sqrt{\sum_i(\sigma_+^i - \sigma_-^i)^2}$, where the sum runs over the 22 upwards and downwards variations which give σ_{\pm} . (This is the Hessian formalism for the uncertainty analysis [95], where the parton parameter space is spanned by a set of orthonormal eigenvectors. The factor 1.645 converts the 90% C. L. into a 68% C. L.) To obtain the relative uncertainty, Δ_{PDF} is divided by the central value of the cross section.
- `relScaleUncHalfQ`, `relScaleUnc2Q`: At NLO, the combined renormalization and factorization scale Q must be fixed by an external (hadronic) scale. PROSPINO uses the average mass of the produced massive particles or the transverse mass of the detected final-state particle as value of Q [26]. The relative uncertainty from the choice of this scale is determined by setting the scale once to twice and once to half the value used for the computation of the central value of the cross section, and recomputing the cross section. The differences of the resulting cross section with respect to the central value give the uncertainty.
- `relProspinoError`: The relative internal inaccuracy of the calculation, as it is returned by PROSPINO. This gives only a small contribution to the total uncertainty (cf. Table 7.12).

For those points of the 2010 MSSM grid which overlap with the 2011 MSSM grid, also the relative uncertainty $\Delta(\alpha_s)$ on the cross section from the uncertainty in the strong coupling constant α_s is available:

- `relUncAlphaS`: This uncertainty, arising from the uncertainty of the strong coupling constant, is defined as half the absolute difference between the cross sections obtained using the two extreme α_s variations AS_{-2} and AS_{+2} from CTEQ6.6AS [46].

For the 2010 mSUGRA grid with $\tan\beta = 3$, the uncertainty arising from α_s has not been computed centrally. This uncertainty therefore needs to be estimated from the $\tan\beta = 10$ grid, for which the distribution of the α_s uncertainties is shown in Figure 7.9. As the $\tan\beta = 10$ grid extends further out into the m_0 - $m_{1/2}$ plane than the $\tan\beta = 3$ grid does, and the α_s uncertainties increase with m_0 and $m_{1/2}$, only the subset of grid points from the




Figure 7.9: Distribution of the relative uncertainty from the variation of α_s for the subset of points from the $\tan\beta = 10$ grid that fall into the m_0 - $m_{1/2}$ range covered by the $\tan\beta = 3$ grid.

$\tan\beta = 10$ grid which fall into the range covered by the $\tan\beta = 3$ grid ($m_0 < 1160$ GeV, 90 GeV $< m_{1/2} < 430$ GeV) has been included in Figure 7.9. For each grid point, the α_s uncertainty is an average over all defined event classes (as defined in Section 7.6.3), weighted with the corresponding cross section, so that each grid point corresponds to one entry in the histogram in Figure 7.9.

To check how the α_s uncertainty depends on $\tan\beta$ in general, for a random sample of five grid points from the mSUGRA grid for $\tan\beta = 3$, the α_s uncertainty has been calculated here for the two dominant subprocesses using PROSPINO. Table 7.11 shows the results for the respective dominant subprocess. It is found that the α_s uncertainty $\Delta(\alpha_s)$ for $\tan\beta = 3$ is larger than for $\tan\beta = 10$, but only slightly so. Only for the point in this sample with very low m_0 and $m_{1/2}$, the deviation is larger, the uncertainty for $\tan\beta = 3$ being more than twice as big as for $\tan\beta = 10$, but the value of the relative uncertainty itself for this point is very small. A conservative value of 3% is therefore assumed for the α_s uncertainty in the following.

With respect to the final input to the limit setting, a total uncertainty on the normalization of each of the signal Monte Carlo samples is computed as the quadratic sum of the individual uncertainties described above. For the scale uncertainty (`relScaleUncHalfQ` and `relScaleUnc2Q`), the larger of the two values obtained when doubling or halving the scale is used. The uncertainty on the signal normalization is computed as a total uncertainty for the whole sample, where the uncertainties are weighted with the cross sections of the individual final states. Table 7.12 gives an overview of the size of the individual uncertainties in percent, by listing the minimum, maximum and average value of their distributions over all points of the two signal grids as well as the standard deviation. The scale and the PDF uncertainties clearly dominate among the systematic uncertainties that are considered. The uncertainties from the computation in PROSPINO itself and the uncertainty from the variation of the strong coupling constant are much smaller. As described above for the mSUGRA grid, a fixed value is assumed for `relUncAlphaS` for all grid points.

dominant final state	$\tan \beta = 3$ grid				$\tan \beta = 10$ grid			
	m_0	$m_{1/2}$	σ	$\Delta(\alpha_s)$	m_0	$m_{1/2}$	σ	$\Delta(\alpha_s)$
squark – squark	80	100	260	0.0013	100	90	381	0.0005
squark – gluino	200	190	8.66	0.0122	180	180	12.1	0.0112
squark – squark	180	400	0.102	0.0225	180	390	0.120	0.0225
squark – gluino	1000	190	0.387	0.0304	980	180	0.499	0.0288
squark – gluino	1000	340	0.0326	0.0479	980	330	0.0409	0.0460

Table 7.11: Relative uncertainty $\Delta(\alpha_s)$ on the cross section σ arising from the uncertainty in the strong coupling constant α_s for five grid points of the $\tan \beta = 3$ mSUGRA grid, as calculated with PROSPINO. The uncertainties are compared to those of grid points with $\tan \beta = 10$ and similar values of m_0 and $m_{1/2}$. All cross sections and uncertainties are given for the dominant subprocess stated in the first column. Masses are given in GeV, cross sections in pb.

Signal grid	Uncertainty	Min [%]	Max [%]	Mean [%]	Std. dev. [%]
mSUGRA	relUncPDF	3.0	24.3	10.0	5.4
	relScale	6.0	15.5	8.6	2.6
	relProspinoError	0.05	0.15	0.07	0.02
	relUncAlphaS			3.0	
MSSM	relUncPDF	1.5	49.9	8.1	7.1
	relScale	6.1	26.0	11.7	3.9
	relProspinoError	0.06	0.30	0.11	0.04
	relUncAlphaS	0.4	10.2	2.2	1.7

Table 7.12: Statistical properties of the distributions of the theoretical uncertainties on the cross sections in the two signal grids. Listed are the minimum (Min), maximum (Max) and average value (Mean) of the distributions of the uncertainties from all grid points, together with the standard deviation (Std. dev.).

7.8 Analysis Results in Numbers

This section summarizes the event counts in data and in Monte Carlo, following from the selection according to the cutflow presented in Section 7.3, as well as the statistical and systematic uncertainties on these numbersⁱ. From these, the inputs are distilled which are finally fed into the limit-setting calculation.

Table 7.13 lists the event counts from the cutflow on Monte Carlo for the different types of background samples, sorted from the left to the right by the event counts in the signal regions. The number of expected events n_{exp} is shown together with the absolute statistical uncertainty Δ_{stat} and the cut efficiency ϵ_{cut} on the respective samples with respect to the preceding cut^j. The rows which correspond to the four signal regions are highlighted. The table shows which of the backgrounds gives the largest contribution and allows to compare the efficiency of the cuts with respect to the suppression of the backgrounds. It demonstrates that a strong suppression of the otherwise dominating QCD background is achieved by cutting on \cancel{E}_T . The fraction of events with pairwise produced top quarks reaching the signal regions is the largest among the background samples which are considered. It is only small in terms of absolute event numbers due to the small production cross section, but when the LHC will run at the nominal center-of-mass energy, the background from top quark events will be of the same order as $Z + \text{jets}$ and $W + \text{jets}$.

Table 7.14 lists the systematic uncertainties in the four signal regions, broken up by the different background types and sources of uncertainty, which have been introduced above (cf. Table 7.9). The luminosity uncertainty has a fixed value for all background types and the signal. The uncertainties arising from pile-up, the jet-energy scale and resolution are computed in the way described in Section 7.7. The uncertainty on the background normalization is denoted as SUSY in this table because the values from the official Supersymmetry analysis are used (cf. Table 7.10). Not listed is the theory uncertainty on the signal normalization because this table contains only the numbers for the background Monte Carlo.

In Table 7.15 the final collection of numbers that make up the result of the analysis is given. It contains the number of events counted in data n_{data} in the signal regions, according to the selection due to the cutflow from Table 7.2, and the expected event counts n_{exp} in the signal regions from Monte Carlo from Table 7.13, together with the total uncertainties from all the different sources. The total uncertainties are derived from the sample uncertainties given in Table 7.14. Depending on the type of uncertainty and its correlation among different backgrounds as given in Table 7.9, the total uncertainty is the arithmetic or quadratic sum of the respective absolute uncertainties from all background types. For convenience, the numbers are again given as absolute and relative uncertainties. For signal regions A and C, which have the highest event counts, the dominating uncertainty comes from the uncertainty of the jet-energy scale. In signal regions B and D with much lower event counts, the uncer-

ⁱ Note that the distinction between statistical and systematic errors is not always sharp. In general, statistical errors are understood to be those that are affected by the dataset size, whereas systematic errors are independent of the dataset. In this sense, there is a statistical uncertainty due to the finite size of the Monte Carlo samples, which cannot be fixed by some auxiliary information, but it can be reduced by increasing the sample size, and there is a systematic uncertainty due to theory, e. g. when comparing different Monte Carlo generators, the value of which will not go to zero for infinite sample sizes. On the other hand, though there are uncertainties which look at least intuitively like a systematic uncertainty, for example the lepton energy scale, which is the largest uncertainty in the W mass measurement at the Tevatron [116, 117]. But this systematic uncertainty gets smaller the more luminosity is collected.

^j Note that not always the previous step is the preceding step in the cutflow, cf. Table 7.2.

cutflow step		Z+jets			W+jets			$t\bar{t}$			PYTHIA dijets			Single top		
		n_{exp}	Δ_{stat}	ϵ_{cut}	n_{exp}	Δ_{stat}	ϵ_{cut}	n_{exp}	Δ_{stat}	ϵ_{cut}	n_{exp}	n_{stat}	ϵ_{cut}	n_{exp}	Δ_{stat}	ϵ_{cut}
1	Raw events	137867.6	141.5		$1.05 \cdot 10^6$	450.7		5370.4	7.8		$4.52 \cdot 10^{11}$	$3.60 \cdot 10^8$		1250.4	6.8	
2	GRL	137867.6	141.5	1.00	$1.05 \cdot 10^6$	450.7	1.00	5370.4	7.8	1.00	$4.52 \cdot 10^{11}$	$3.60 \cdot 10^8$	1.00	1250.4	6.8	1.00
3	Main trigger	3418.9	25.3	0.02	11025.2	36.8	0.01	1541.3	3.6	0.29	5439501.5	32584.3	0.00	231.4	2.9	0.19
4	Jet cleaning	3401.6	25.3	0.99	10955.4	36.7	0.99	1530.6	3.6	0.99	5427600.0	32550.9	1.00	229.9	2.9	0.99
5	Primary vertex	3399.9	25.3	1.00	10951.2	36.7	1.00	1530.5	3.6	1.00	5427490.5	32550.8	1.00	229.7	2.9	1.00
6	Crack electrons	3353.8	25.2	0.99	10795.6	36.5	0.99	1509.7	3.6	0.99	5415772.5	32471.9	1.00	226.5	2.9	0.99
7	Veto on leptons	2306.3	23.3	0.69	4454.6	23.0	0.41	814.0	3.1	0.54	5353897.0	32326.6	0.99	104.5	2.0	0.46
8	Leading jet	1369.3	17.7	0.59	2353.9	16.1	0.53	573.5	2.7	0.70	2715003.3	14496.2	0.51	66.0	1.6	0.63
9	Monojet	552.8	13.3	0.40	541.5	9.1	0.23	3.1	0.1	0.01	72049.7	3463.0	0.03	3.4	0.4	0.05
10	Charge fraction	547.7	13.3	0.99	534.4	9.0	0.99	3.1	0.1	1.00	71096.6	3454.5	0.99	3.4	0.4	1.00
11	Second jet	707.2	10.7	0.52	1594.7	12.4	0.68	556.6	2.7	0.97	2544498.3	13033.8	0.94	59.3	1.5	0.90
12	Charge fraction	699.6	10.6	0.99	1560.1	12.2	0.98	551.5	2.7	0.99	2515436.5	12856.8	0.99	58.6	1.5	0.99
13	$\cancel{E}_T > 100 \text{ GeV}$	339.8	7.7	0.49	516.9	7.0	0.33	102.5	0.7	0.19	8165.4	212.8	0.00	14.6	0.7	0.25
14	$\Delta\phi(\text{jet}, \cancel{E}_T) > 0.4$	286.8	7.2	0.84	377.6	5.9	0.73	75.1	0.5	0.73	295.0	70.4	0.04	10.7	0.6	0.74
15	$\cancel{E}_T/m_{\text{eff}} > 0.3$	244.8	6.7	0.85	308.3	5.4	0.82	59.0	0.5	0.79	126.1	68.2	0.43	8.6	0.6	0.80
16	$m_{\text{eff}} > 500 \text{ GeV}$	50.3	3.0	0.21	47.9	2.1	0.16	8.5	0.2	0.14	7.2	3.8	0.06	1.0	0.2	0.12
17	$m_{T2} > 300 \text{ GeV}$	4.0	0.8	0.01	4.2	0.6	0.01	0.8	0.1	0.01	0.3	0.1	0.00	0.1	0.1	0.01
18	Third jet	229.5	5.4	0.17	552.6	6.6	0.23	494.8	2.6	0.86	1001039.9	6694.7	0.37	37.6	1.2	0.57
19	$\cancel{E}_T > 100 \text{ GeV}$	98.3	3.7	0.43	172.2	3.7	0.31	82.7	0.6	0.17	4125.6	170.2	0.00	7.0	0.5	0.19
20	$\Delta\phi(\text{jet}, \cancel{E}_T) > 0.4$	71.2	3.2	0.72	111.1	3.0	0.65	57.1	0.5	0.69	144.1	15.3	0.03	4.3	0.4	0.61
21	$\cancel{E}_T/m_{\text{eff}} > 0.25$	57.3	2.9	0.81	82.1	2.6	0.74	43.4	0.4	0.76	18.1	6.6	0.13	3.3	0.3	0.78
22	$m_{\text{eff}} > 500 \text{ GeV}$	26.4	2.0	0.46	33.5	1.6	0.41	15.3	0.2	0.35	7.2	3.8	0.40	1.1	0.2	0.32
23	$m_{\text{eff}} > 1 \text{ TeV}$	0.8	0.3	0.01	1.1	0.3	0.01	0.2	0.0	0.00	0.1	0.1	0.00	0.1	0.1	0.03
24	Fourth jet	52.0	2.2	0.04	136.2	3.2	0.06	349.8	2.3	0.61	232505.0	2536.3	0.09	19.7	0.8	0.30

Table 7.13: Event counts as they follow from the analysis cutflow, showing the expected event numbers n_{exp} from Monte Carlo, normalized to an integrated luminosity of $\mathcal{L}_{\text{int}}^{2010} = 33.4 \text{ pb}^{-1}$. Δ_{stat} is the absolute statistical uncertainty on this number, and ϵ_{cut} is the efficiency of the cut on this sample with respect to the preceding cut (cf. Table 7.2). The four rows defining the signal regions are highlighted.

tainty on the background normalization becomes the dominant source of uncertainty. In signal region D with the smallest event count, also the limited Monte Carlo statistics gives an important contribution. Another important uncertainty here is due to pile-up.

The table allows to compare the expected event counts in the signal regions from Standard Model backgrounds modeled in the Monte Carlo samples to the event counts that were actually measured in data. In signal regions A, C and D, the Standard Model background expectation from Monte Carlo exceeds the number of events measured in data. In signal region B, the number of events measured in data is slightly above the Standard Model background expectation from Monte Carlo, but also well consistent within uncertainties. This can be seen in Table 7.16, which is a summary of Table 7.15, summing over all sources of uncertainties.

Due to this finding, showing no significant excess of the number of events beyond the Standard Model background expectation, in the rest of this section the results will be interpreted in terms of exclusion limits. The two signal grids presented in Section 7.6 will be used to obtain the expected number of signal events for different realizations of parameters within the simplified supergravity and mSUGRA model.

7.9 Kinematic Distributions

A number of important distributions, on which the cuts in the analysis cutflow are based, are collected in Figures 7.10 through 7.18. All of them show the distribution of the respective variable in data and in Monte Carlo using the same scheme: The upper part of each plot shows several distributions in terms of event counts, the data being represented by the black dots, for which the error bars give statistical uncertainties^k. The background expectations from Monte Carlo are represented by the shaded histograms, stacking the QCD (PYTHIA di-jets), W + jets, Z + jets, top pair and single top background expectations on top of each other in different colors. They are normalized to the integrated luminosity of the 2010 dataset (cf. Equation (7.1)) and sorted, so that the largest contribution is on top. Note that the relative contributions from the different Standard Model backgrounds vary, depending on which cuts have already been applied, and the order of the stacking changes accordingly. QCD is the dominating background in the first plots, but is replaced in later steps of the cutflow by the other backgrounds. To give an impression of the expected additional event counts that would arise from a potential Supersymmetry signal, the red dashed line on top of the background expectations gives the Monte Carlo expectation for the supersymmetric benchmark scenario SU4, a scenario with comparably low masses for the supersymmetric particles and a large cross section. The vertical line with the arrow indicates the cut which is done on the variable shown in the histogram, the tip of the arrow pointing into the direction of the events that are being kept. In the lower part of each plot, observed data and Monte Carlo expectation are compared bin-by-bin to provide an overview of where the Monte Carlo under- or overestimates the number of events observed in data at one glance. The black markers in the lower part give the ratio of the number of data events and the expected number of events from the Monte Carlo expectation, here without including a potential signal, i. e. from Standard Model background processes only. The error bars above and below the markers show the uncertainties on the ratio, which stem from the statistical uncertainty on the number of

^k Note that this is by convention different from the numbers given in the tables, where statistical uncertainties are given explicitly only for the event counts expected from Monte Carlo.

Background	Uncertainty	Signal region								
		A		B		C		D		
W + jets	<i>Event count</i>	47.90		4.23		33.52		1.06		
	Statistics	2.06	(4.3 %)	0.60	(14.1 %)	1.63	(4.9 %)	0.28	(26.8 %)	
	Luminosity	1.63	(3.4 %)	0.14	(3.4 %)	1.14	(3.4 %)	0.04	(3.4 %)	
	JER	0.76	(1.6 %)	0.08	(1.9 %)	0.03	(0.1 %)	0.00	(0.0 %)	
	JES	up	9.47	(19.8 %)	1.20	(28.3 %)	6.97	(20.8 %)	0.15	(13.9 %)
		down	7.92	(16.5 %)	0.70	(16.6 %)	6.16	(18.4 %)	0.30	(28.0 %)
	Pile-up	2.41	(5.0 %)	0.42	(9.8 %)	1.72	(5.1 %)	0.24	(23.0 %)	
	SUSY	19.16	(40.0 %)	2.88	(68.0 %)	14.75	(44.0 %)	0.77	(73.0 %)	
Z + jets	<i>Event count</i>	50.29		3.97		26.41		0.81		
	Statistics	3.02	(6.0 %)	0.81	(20.5 %)	1.95	(7.4 %)	0.34	(42.3 %)	
	Luminosity	1.71	(3.4 %)	0.13	(3.4 %)	0.90	(3.4 %)	0.03	(3.4 %)	
	JER	0.44	(0.9 %)	0.08	(1.9 %)	0.64	(2.4 %)	0.00	(0.0 %)	
	JES	up	11.27	(22.4 %)	1.54	(38.8 %)	6.84	(25.9 %)	0.31	(38.7 %)
		down	9.22	(18.3 %)	0.73	(18.5 %)	4.88	(18.5 %)	0.00	(0.0 %)
	Pile-up	1.51	(3.0 %)	0.33	(8.3 %)	2.49	(9.4 %)	0.29	(35.6 %)	
	SUSY	10.56	(21.0 %)	2.82	(71.0 %)	6.07	(23.0 %)	0.45	(55.0 %)	
$t\bar{t}$	<i>Event count</i>	8.50		0.79		15.27		0.18		
	Statistics	0.18	(2.1 %)	0.06	(7.0 %)	0.24	(1.6 %)	0.03	(14.6 %)	
	Luminosity	0.29	(3.4 %)	0.03	(3.4 %)	0.52	(3.4 %)	0.01	(3.4 %)	
	JER	0.07	(0.8 %)	0.05	(5.9 %)	0.45	(3.0 %)	0.00	(0.0 %)	
	JES	up	1.78	(20.9 %)	0.22	(28.4 %)	3.83	(25.1 %)	0.07	(38.3 %)
		down	1.63	(19.2 %)	0.18	(22.5 %)	2.94	(19.2 %)	0.05	(25.5 %)
	Pile-up	0.09	(1.0 %)	0.02	(2.4 %)	0.21	(1.4 %)	0.07	(39.9 %)	
	SUSY	10.56	(21.0 %)	2.82	(71.0 %)	6.07	(23.0 %)	0.45	(55.0 %)	
Single top	<i>Event count</i>	1.02		0.12		1.07		0.11		
	Statistics	0.19	(18.7 %)	0.07	(54.3 %)	0.19	(18.1 %)	0.06	(55.9 %)	
	Luminosity	0.03	(3.4 %)	0.00	(3.4 %)	0.04	(3.4 %)	0.00	(3.4 %)	
	JER	0.05	(4.8 %)	0.04	(30.3 %)	0.00	(0.3 %)	0.00	(0.0 %)	
	JES	up	0.29	(28.1 %)	0.04	(30.3 %)	0.21	(19.5 %)	0.04	(32.3 %)
		down	0.20	(19.2 %)	0.04	(31.9 %)	0.28	(25.8 %)	0.07	(64.5 %)
	Pile-up	0.05	(4.9 %)	0.00	(1.0 %)	0.02	(1.6 %)	0.02	(20.9 %)	
	SUSY	10.56	(21.0 %)	2.82	(71.0 %)	6.07	(23.0 %)	0.45	(55.0 %)	
PYTHIA dijets	<i>Event count</i>	7.18		0.28		7.19		0.09		
	Statistics	3.81	(53.1 %)	0.13	(44.2 %)	3.81	(53.0 %)	0.07	(83.0 %)	
	Luminosity	0.24	(3.4 %)	0.01	(3.4 %)	0.24	(3.4 %)	0.00	(3.4 %)	
	JER	0.14	(2.0 %)	0.00	(0.4 %)	0.36	(5.0 %)	0.00	(0.0 %)	
	JES	up	2.83	(39.4 %)	0.09	(31.5 %)	0.07	(1.0 %)	0.07	(82.8 %)
		down	2.83	(39.4 %)	0.01	(4.2 %)	2.76	(38.3 %)	0.00	(0.0 %)
	Pile-up	7.18	(100.0 %)	0.28	(100.0 %)	7.19	(100.0 %)	0.09	(100.0 %)	
	SUSY	10.56	(21.0 %)	2.82	(71.0 %)	6.07	(23.0 %)	0.45	(55.0 %)	

Table 7.14: Systematic uncertainties as absolute event counts and relative numbers (in brackets) in the four signal regions, broken up by background type and source of uncertainty. The rows labelled “Event count” do not state an uncertainty but the number of Monte Carlo events from this sample reaching the respective signal region. “SUSY” stands for uncertainties adopted from the official analysis (see Section 7.5.2). Rounding makes small uncertainties appear as 0.00, which should be read as < 0.01 .

		Signal region								
		A		B		C		D		
	n_{data}	87		11		66		2		
	n_{exp}	114.9		9.4		83.5		2.2		
Uncertainty	Statistics	5.3	(4.6 %)	1.0	(10.9 %)	4.6	(5.5 %)	0.5	(20.3 %)	
	Luminosity	3.9	(3.4 %)	0.3	(3.4 %)	2.8	(3.4 %)	0.1	(3.4 %)	
	JER	1.5	(1.3 %)	0.0	(0.1 %)	1.4	(1.7 %)	0.0	(0.0 %)	
	JES	up	25.6	(22.3 %)	3.1	(32.9 %)	17.9	(21.5 %)	0.6	(28.4 %)
		down	21.8	(19.0 %)	1.7	(17.7 %)	17.0	(20.4 %)	0.4	(18.5 %)
	Pile-up	11.2	(9.8 %)	1.1	(11.2 %)	11.6	(13.9 %)	0.7	(31.8 %)	
	SUSY	21.9	(19.0 %)	4.0	(42.9 %)	16.0	(19.1 %)	0.9	(39.6 %)	

Table 7.15: Summary of the numbers that constitute the result of the analysis. Given in the table are the event counts on data (n_{data}) and the expected counts on Monte Carlo (n_{exp}), together with the uncertainties on the Monte Carlo expectations in terms of absolute event counts and (in brackets) as percentage values. Rounding to one decimal place makes small uncertainties appear as 0.0, which should be read as < 0.1 .

		Event counts	
Signal region		n_{data}	$n_{\text{exp}} \pm \Delta_{\text{stat}} \pm \Delta_{\text{syst}}$
A	2 jets + $m_{\text{eff}} > 500$ GeV	87	$114.9 \pm 5.3^{+35.8}_{-33.1}$
B	2 jets + $m_{\text{T2}} > 300$ GeV	11	$9.4 \pm 1.0^{+5.2}_{-4.5}$
C	3 jets + $m_{\text{eff}} > 500$ GeV	66	$83.5 \pm 4.6^{+26.9}_{-26.2}$
D	3 jets + $m_{\text{eff}} > 1000$ GeV	2	$2.2 \pm 0.5^{+1.3}_{-1.2}$

Table 7.16: Summary of the analysis result in Table 7.15, giving the event counts observed in data (n_{data}) and expected from Monte Carlo (n_{exp}), together with the statistical (Δ_{stat}) and total systematic uncertainties (Δ_{syst}). The systematic uncertainties are asymmetric due to the asymmetric contribution from the uncertainty on the jet-energy scale.

data events. The shaded regions above and below the markers depict the total uncertainties on the Standard Model background expectation, summing all sources of uncertainties that have been discussed, including the statistical uncertainty from the limited Monte Carlo sample size. The distributions shown in the plots are the distributions before the cut indicated by the arrow is applied. The cuts in Figures 7.10 through 7.14 and Figures 7.15 through 7.18 correspond to the steps 13 through 17 and 19 through 22 in Table 7.2, respectively.

Specifically, the plots show the following distributions and cuts: Figure 7.10 shows the distribution of \cancel{E}_T before the cut at 100 GeV in the 2-jet channel, which reduces the QCD background significantly. Figure 7.11 shows the distribution of the angles $\Delta\phi(\text{jet}, \cancel{E}_T)$ between the three leading jets with $p_T > 40$ GeV and the direction of \cancel{E}_T in the transverse plane in the 2-jet channel, before the cut at 0.4 to get rid of events in which a hard jet is aligned with \cancel{E}_T , and the \cancel{E}_T potentially is due to a mismeasurement of the energy of this jet. Figure 7.12 shows the distribution of the ratio of the missing transverse energy and the effective mass in the 2-jet channel, before the cut at 0.3. Figure 7.13 shows the distribution of the effective mass in the 2-jet channel, before the cut at 500 GeV, which defines signal region A. Figure 7.14 shows the distribution of the transverse mass m_{T2} , the generalization of the transverse mass to pair decays, before the cut at 300 GeV in the 2-jet channel, which defines signal region B. Figure 7.15 shows the distribution of \cancel{E}_T , again before a cut at 100 GeV, but this time for events which have three jets with $p_T > 120$ GeV for the leading and $p_T > 40$ GeV for the second and third jet. Figure 7.16 shows the distribution of the angles $\Delta\phi(\text{jet}, \cancel{E}_T)$, before the cut at 0.4, this time in the three-jet channel. Figure 7.17 shows the distribution of the ratio of the missing transverse energy and the effective mass in the 3-jet channel, before the cut at 0.25. Figure 7.18 shows the distribution of the effective mass in the 3-jet channel before the cut at 500 GeV, which defines signal region C. The cut which defines signal region D in the 3-jet channel is applied to the same distribution but at 1 TeV, so that only two events are left in data after this cut, as can be seen in the same plot.

Note that the cuts which are indicated in the plots of $\Delta\phi(\text{jet}, \cancel{E}_T)$ are slightly misleading because an event passes this cut only if all of the three leading jets above 40 GeV fulfill $\Delta\phi(\text{jet}, \cancel{E}_T) > 0.4$, which cannot be represented in the one-dimensional plot. This explains why in Figure 7.11 it looks like QCD is still the dominating background after the cut on $\Delta\phi(\text{jet}, \cancel{E}_T)$, but according to Figure 7.12 it no longer is. This is also confirmed by Table 7.13.

7.10 Overview of Statistical Methods

Before continuing the evaluation of the numbers which constitute the outcome of the analysis in this thesis, in this section, the statistical methods will be presented that will be used to quantify the agreement between the data which has been taken and the predictions from the Standard Model, with or without the inclusion of Supersymmetry giving rise to physics processes beyond the Standard Model prediction. Some cross-checks will also be presented, which compare the performance of the two methods employed for the limit setting.

The important second step after having conducted an experiment is the interpretation of the results that were obtained. Especially in high-energy physics, where data-mining makes up a large fraction of the work, statistics plays a vital role. Depending on the nature and maturity of the experiment, potential methods for evaluating the outcome range from straightforward approaches like counting experiments with rectangular cuts to sophisticated techniques like neural networks, boosted decision trees or support vector machines,




Figure 7.10: Distribution of \cancel{E}_T before the cut at 100 GeV in the 2-jet channel. The upper plot shows the distribution observed in data as dots, with error bars indicating the statistical error, and the Monte Carlo expectations for the Standard Model backgrounds as stacked shaded histograms, with the expected additional counts for an exemplary supersymmetric signal shown as the dashed red line on top. The arrow depicts the cut that is applied to the variable on the horizontal axis. The lower part shows the ratio of data and Monte Carlo, with error bars indicating the uncertainty on the ratio arising from the statistical uncertainty on the observed counts in data. The shaded regions represent the systematic uncertainties.




Figure 7.11: Distribution of the angles $\Delta\phi(\text{jet}, \cancel{E}_T)$ between the three leading jets with $p_T > 40$ GeV and the direction of \cancel{E}_T in the transverse plane in the 2-jet channel, before the cut at 0.4. (Note that the cut at 0.4 in this histogram is slightly shifted with respect to the lower edge of the respective bin, but the discrepancy is negligible.)




Figure 7.12: Distribution of the ratio of the missing transverse energy and the effective mass in the 2-jet channel before the cut at 0.3.




Figure 7.13: Distribution of the effective mass in the 2-jet channel before the cut at 500 GeV, which defines signal region A.




Figure 7.14: Distribution of the transverse mass m_{T2} before the cut at 300 GeV in the 2-jet channel, which defines signal region B.




Figure 7.15: Distribution of E_T before the cut at 100 GeV in the 3-jet channel.

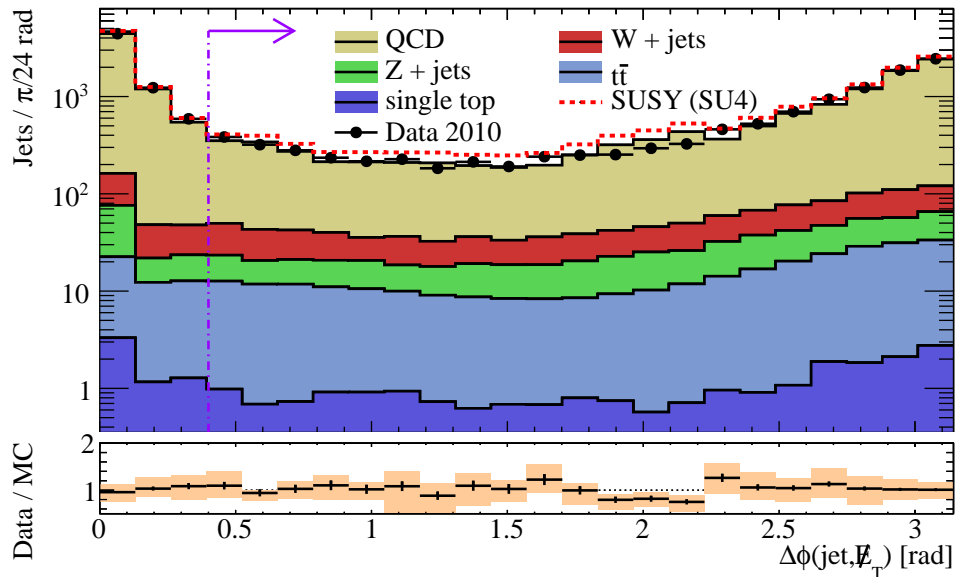


Figure 7.16: Distribution of the angles $\Delta\phi(\text{jet}, \cancel{E}_T)$ between the three leading jets with $p_T > 40 \text{ GeV}$ and the direction of \cancel{E}_T in the transverse plane in the 3-jet channel before the cut at 0.4.




Figure 7.17: Distribution of the ratio of the missing transverse energy and the effective mass in the 3-jet channel before the cut at 0.25.




Figure 7.18: Distribution of the effective mass in the 3-jet channel before the cut at 500 GeV, which defines signal region C. The cut which defines signal region D in the 3-jet channel is applied to the same distribution but at 1 TeV, reducing the number of observed events in data to only two.

which aim at extracting the maximum possible information from the measurements or are needed if no simple discriminating variable exists or the signal is much smaller than the background. In general, the better the understanding of the detector and its environment, the more advanced methods can be employed. For the first ATLAS data which has been collected so far, the focus is on techniques that are simple enough to still be reasonably intuitive. On the one hand, the understanding of the detector naturally has not yet reached a state that can be called complete, and most of the advanced techniques carry the risk of cloaking the causal relations of input and output due to their complexity. On the other hand, most of the advanced techniques are tailored towards a very specific analysis model, whereas at the moment more general approaches should be used, in order not to miss something new or bias the analysis towards a specific direction. As soon as the instantaneous luminosity levels out and stable detector operation is fully established, so that collecting data becomes day-to-day business, and when the first hints emerge where to look for what, the analysis will certainly be refined further.

Details about the limit-setting procedure based on the profile likelihood ratio, which has been employed by the ATLAS Supersymmetry group for the first analyses of 2010 and 2011 data, can be found in [147]. To implement this procedure, the ROOSTATS framework is used. A brief description of the main methods available in the ROOSTATS framework is given in [82]. One of the standard references for statistical methods in high-energy physics is [2]. Other references are cited in the discussion below.

7.10.1 Nomenclature

The following section will introduce a number of terms relevant to the discussion of the methods which are used to compute the limits.

One of the most important distributions for the following is the Poisson distribution. Its probability mass function is given by

$$\text{Pois}(k|\lambda) = \exp(-\lambda) \frac{\lambda^k}{k!}, \quad (7.6)$$

with $k \in \mathbf{N}$. It has one parameter $\lambda \in (0, \infty)$ and can be generalized to a continuous support by replacing the normalization $k!$ with $\Gamma(k+1)$. The Poisson distribution is an approximation of the Binomial distribution given in Equation (5.6) for rare events, i. e. for $n \rightarrow \infty$ and $np \rightarrow \lambda$. For large values of λ , due to the central limit theorem and up to a continuity correction, the Poisson distribution in turn can be approximated by a Gaussian distribution. Both expectation value and variance of the Poisson distribution are equal to λ , the standard deviation is thus $\sqrt{\lambda}$. The importance of this distribution stems from the fact that the number of entries in the bins of a histogram, or outcomes of counting experiments in general, are well modeled by a Poisson distribution for sufficiently large statistics. It also describes, for example, the distribution of radioactive decays per time of a source with constant activity.

Statistical Inference

The analysis of search results using statistical inference is done by hypothesis testing. In hypothesis testing, to show the existence of some effect, the opposite effect is hypothesized in the null hypothesis \mathcal{H}_0 , and it is then to be shown that \mathcal{H}_0 is not compatible with the experimental outcome. Thus, the discovery of new physics or the exclusion of such processes within certain bounds is done by rejecting the opposite hypothesis at a given level of certainty. One can either use the hypothesis of the existence of new physics processes in addition to the known Standard Model processes, quantify the effect of such new processes¹ and compute the probability that the data agree with this hypothesis. If the data agree well already with the Standard Model predictions, this probability will be low and one can reject the hypothesis of new physics at a high level of certainty, which will lead to an exclusion. Or, one can test the hypothesis that there is no physics beyond the known (background) processes described within the Standard Model. If the data do not agree with the background-only hypothesis, it will be excluded, effectively leading to the discovery of a new signal. This does not say whether the model for the new physics is correct, it only rejects the hypothesis that there is only background, but it leaves open the question, which of probably many possible explanations for the observed excess of events is correct.

Thus, for limit setting, the null hypothesis \mathcal{H}_0 states that both signal and background exist. This is called the signal-plus-background hypothesis \mathcal{H}_{S+B} in the following. The alternate hypothesis is that the hypothesized signal is absent or too small to be seen and that the data can be explained by previously established theories. This will be the background-only hypothesis \mathcal{H}_B .

Two types of errors are distinguished in hypothesis testing. A type I error occurs when a true null hypothesis is rejected on basis of the data. Taking the null hypothesis to be

¹ This implies assuming a certain signal strength. The signal strength is the ratio of the cross section used in the hypothesis over the expected cross section from the underlying model (Standard Model or new physics).

the signal-plus-background hypothesis, this leads to a false exclusion of a true signal (false exclusion rate). Type II errors mean failing to reject a false null hypothesis, corresponding then to a false discovery of an absent signal (false discovery rate). A test is called significant at a given level α , e. g. significant at the 5% level, if the probability for type I errors is less than or equal to α . The probability to not make a type II error, i. e. to correctly reject a false hypothesis, is called the power of a test. Taking again the null hypothesis to be \mathcal{H}_{S+B} , this is the exclusion potential. For making discoveries, the null hypothesis \mathcal{H}_0 would be \mathcal{H}_B , and the power of the test would be the discovery potential.

The level of agreement between the observed data and the hypothesis is quantified by the p -value. The p -value is the probability to obtain data which is at least as incompatible with the hypothesis to be tested as is the data at hand, under the assumption of the validity of the hypothesis. The hypothesis is reported to be excluded if the p -value is below a given threshold, i. e. if it is too unlikely to obtain data which is so incompatible with the hypothesis.

7.10.2 Frequentist and Bayesian Statistics

In the discussion of experimental results, limit setting and confidence levels, two different points of view concerning the interpretation of probability can be adopted, which are called Frequentist and Bayesian. The difference can be summarized like this [2]:

According to the Frequentist definition, the probability is the long-run expected frequency of occurrence, $p(A) = n_A/N$, where n_A is the number of times the event A has been found in N observations. In the Frequentist view, the population mean is fixed, but unknown, and can only be estimated from the data. A confidence interval, centered at the sample mean, is inferred from the distribution of the sample mean, using the Neyman or confidence belt construction [87], such that with a given coverage probability p_c , e. g. 95%, the true, unknown mean is in this interval, and if the experiment were repeated many times, and more confidence intervals constructed in the same way, a fraction p_c of these would contain the true mean. This should not be understood to mean that the true mean has a probability of $1 - p_c$ to lie outside the constructed interval, but that, if the true mean lies outside the constructed interval, the probability to obtain the measured result is $1 - p_c$. Saying that the true value lies within the interval with probability p_c is thus not a statement about the true value but about the interval limits. One problem of the Frequentist approach is that it relies on the repeatability of the experiment. Another is that it cannot incorporate limits on the measured parameter that are physically impossible to violate, such as a mass having to be non-negative. This is only possible in the Bayesian approach.

The Bayesian definition of probability is a subjective probability. Experimental evidence can and will modify an initial or *a priori* degree of belief. To do so, Bayes' theorem is invoked, which requires specifying the *a priori* distribution of the parameter values that are to be measured. Using Bayes' theorem, a credible interval is constructed, which then is reported to contain the true mean with a given probability. (Note the subtle difference with respect to the Frequentist view, where no concrete statement about the mean is being made.) If a proper prior is chosen, the outcome will be driven by the experimental observation, and the influence of using different priors needs to be checked, but will not lead to an unacceptable bias in the constructed limits. It should be noted that in general priors are not invariant under a change of parameter, although there are exceptions, using e. g. Jeffreys' rule [87]. For both the Frequentist and the Bayesian approach, it is clear that the larger the interval (and the less the information content of or equivalently the less restrictive the statement),

the higher the confidence at which this statement can be made.

The likelihood function is an important object in the context of statistical inference. For a joint probability density function (PDF) $f(\vec{x}|\vec{\theta})$ which encodes the probability of finding the measurement \vec{x} given the parameters of the model $\vec{\theta}$ for the data, the likelihood function L is constructed by evaluating the probability density function f with the observed data \vec{x}_{obs} and regarding it as a function of the parameters, i. e. $L = L(\vec{\theta}) = f(\vec{\theta}|\vec{x}_{\text{obs}})$. The likelihood is not a probability density function for the parameters. From a frequentist point-of-view this is not defined. In Bayesian statistics, a posterior probability density function for the parameters can be obtained from L by multiplication with a prior probability density function $\pi(\vec{\theta})$, using Bayes' theorem:

$$p(\vec{\theta}|\vec{x}) = \frac{f(\vec{x}|\vec{\theta})\pi(\vec{\theta})}{\int f(\vec{x}|\vec{\theta}')\pi(\vec{\theta}')d\vec{\theta}'} \quad (7.7)$$

In case of one parameter θ , this allows to compute a credible interval $[\theta_{\text{low}}, \theta_{\text{high}}]$ with a given probability of containing the true value of the parameter by solving

$$1 - \alpha = \int_{\theta_{\text{low}}}^{\theta_{\text{high}}} p(\theta|\vec{x}) d\theta \quad (7.8)$$

for θ_{low} and θ_{high} , where the threshold α is called the size of a test. For $\theta_{\text{low}} \rightarrow -\infty$ this gives an upper limit only. Note that in general the solution of Equation (7.8) is not unique. Often the smallest such interval is used.

The Bayesian approach allows to account for systematic uncertainties by including prior distributions on additional parameters in the likelihood function and marginalizing (integrating) over these. The additional parameters are called nuisance parameters. Limits on the measured parameter can also be included in the Bayesian limits by setting the prior probability to zero in the physically excluded parameter ranges. For example, for hypothesized signals that can only contribute positively to the expected event count, the prior for the signal strength could be chosen flat and non-negative. Within this framework, also control measurements, giving e. g. an estimate of the background, can readily be taken into account.

7.10.3 Methods for Setting Exclusion Limits

Two different methods are used in this thesis to set exclusion limits. The first is the so-called CL_s method, introduced by the LEP experiments for the Higgs search [65, 99]. This method is currently (2011) recommended by the ATLAS Statistics Forum. The second method is the Profile Likelihood Ratio method which is used by the ATLAS SUSY group in the first papers on searches for Supersymmetry in ATLAS data [147]. Both methods can be used to obtain exclusion limits at a given confidence level, i. e. a upper (or lower) limit on some quantity such that all values above (or below) this limiting value are excluded at the given confidence level. In the following, the mathematics behind the two methods is described. The technical details of the actual implementation can be found in Section A.3.

Unlike in experiments looking, for example, for the neutrino oscillations already mentioned in the introduction, where the signal hypothesis may predict an increased as well as a decreased event rate compared to the null hypothesis, it will be assumed here that new physics processes can only increase the number of events observed in the signal region.

Downward fluctuations are thus interpreted to not indicate the presence of a signal, and the signal strength is bound by zero from below.

CL_s Method

The CL_s method is a method for the calculation of confidence intervals based on two alternative hypotheses, the signal-plus-background hypothesis (\mathcal{H}_{S+B}), stating the existence of both signal and background, and the background-only hypothesis (\mathcal{H}_B), stating the existence of only background. It uses Poisson statistics and allows for an easy combination of results of independent searches. The CL_s method and an iterative approximate method to compute combined exclusion confidence levels based on CL_s have been described by Junk [65]. Read gives a generalized form, including probability density functions for the discriminating variables [99] and a discussion of the differences between CL_s and frequentist confidence intervals or Bayesian credible intervals in the context of Higgs searches at the Large Electron-Positron Collider [99, 100].

Binning the results of the analyses to be combined in their discriminant variables allows to treat each bin as a statistically independent search channel of a counting experiment, which can then be combined with the CL_s method. The likelihood for finding d_i events in channel i , assuming the existence of both signal and background, \mathcal{H}_{S+B} , is given by a product of Poisson probabilities,

$$L(d|s+b) = \prod_{i=1}^n \frac{\exp(-(s_i + b_i)) (s_i + b_i)^{d_i}}{d_i!}, \quad (7.9)$$

and in the same way, assuming \mathcal{H}_B instead,

$$L(d|b) = \prod_{i=1}^n \frac{\exp(-b_i) b_i^{d_i}}{d_i!}, \quad (7.10)$$

where n is the number of channels or bins.

A test statistic X is now defined, which discriminates signal-like outcomes from background-like ones using a likelihood ratio. In searches for new physics, an appropriate likelihood ratio is given by the ratio of the probability densities for the signal-plus-background hypothesis and the background-only hypothesis [65],

$$X = \frac{L(d|s+b)}{L(d|b)}. \quad (7.11)$$

X increases monotonically with the number of observed events and induces an ordering on the space of possible outcomes. It ranks the possible experimental outcomes from the least to the most signal-like or, equivalently, from the most to least background-like. This incorporates at this stage an implicit restriction to positive signals $s_i \geq 0$ as mentioned above. Its distribution can then be used to compute the probability of obtaining a certain outcome. Note that for vanishing backgrounds this test statistic cannot be used because it diverges for $\sum_i b_i = 0$. Using Equations (7.9) and (7.10), Equation (7.11) goes over into

$$X = \exp\left(\sum_{i=1}^n s_i\right) \prod_{i=1}^n \left(1 + \frac{s_i}{b_i}\right)^{d_i}. \quad (7.12)$$

Without destroying the ordering property, this can be written in logarithmic form,

$$Q := -2 \ln X = -2 \sum_{i=1}^n s_i - 2 \sum_{i=1}^n d_i \ln \left(1 + \frac{s_i}{b_i} \right), \quad (7.13)$$

which will also be used later in the implementation as it is numerically more stable. The additional factor -2 may be motivated from Wilks' theorem (see below), but here is purely conventional as Wilks' theorem is not used in this context. A more general form which takes into account the probability density functions of the discriminating variables is given in [99].

The probability to obtain a value of the test statistic which is less or equal to the value actually observed assuming the existence of signal and background \mathcal{H}_{S+B} is

$$P_{s+b}(X \leq X_{\text{obs}}) = \sum_{X(\{d'_i\}) \leq X_{\text{obs}}} \prod_{i=1}^n \frac{\exp(-(s_i + b_i)) (s_i + b_i)^{d'_i}}{d'_i!}. \quad (7.14)$$

The product is a weighting term, given by the likelihood in Equation (7.9), and $X_{\text{obs}} = X(\{d'_i\})$ are the observed numbers of events. In the same way, P_b is defined to give the probability that the test statistic is less or equal to the value that is observed in the data, assuming the existence of only background:

$$P_b(X \leq X_{\text{obs}}) = \sum_{X(\{d'_i\}) \leq X_{\text{obs}}} \prod_{i=1}^n \frac{\exp(-b_i) (b_i)^{d'_i}}{d'_i!}. \quad (7.15)$$

Note that X is written as a function of $\{d'_i\}$ only because the signal and background expectations, $\{s_i\}$ and $\{b_i\}$, which also go into the computation of X as parameters, stay the same in both cases, when computing CL_{s+b} and when computing CL_b . In the continuous limit, this is sometimes written as

$$P_{s+b}(X \leq X_{\text{obs}}) = \int_{-\infty}^{X_{\text{obs}}} \frac{dP_{s+b}}{dX} dX, \quad (7.16)$$

with dP_{s+b}/dX being the probability density of X under \mathcal{H}_{S+B} , and in the same way for $P_b(X \leq X_{\text{obs}})$, *mutatis mutandis*.

The confidence value for the background-only hypothesis, denoted CL_b , is

$$\text{CL}_b = P_b(X \leq X_{\text{obs}}). \quad (7.17)$$

Note that calling this the confidence in the background hypothesis (as is done in the first publications on CL_s) can be misleading because high values of CL_b close to one show lack of agreement of the data with the background-only hypothesis, and will be used to reject the background-only hypothesis and to quote a potential discovery: If $1 - \text{CL}_b < \alpha$, one would exclude the background-only hypothesis at a confidence level of $1 - \alpha$, interpreting $1 - \text{CL}_b$ as the probability that the background fluctuates to give a distribution of observed events at least as signal-like as the one observed in data. Note that the background in the search result is accounted for in two ways. It goes into the test statistic, and it is used in the computation of the confidences.

CL_{s+b} is the confidence level for the possibility of simultaneous presence of new particle production and background,

$$\text{CL}_{s+b} = P_{s+b}(X \leq X_{\text{obs}}). \quad (7.18)$$

Small values of CL_{s+b} disfavor the signal-plus-background hypothesis \mathcal{H}_{S+B} , so to exclude \mathcal{H}_{S+B} at a given confidence level $1 - \alpha$, the condition $CL_{s+b} < \alpha$ needs to be fulfilled.

The criticism about CL_{s+b} is that when e. g. using $\alpha = 0.05$, the exclusion potential goes down only to (and stays at) 5%, even for expected signal rates which are very small compared to the expected background [99]. This means that if the data has a downward fluctuation so that the number of observed events is smaller than the number of expected background events, any signal may be excluded at a high confidence level. This is correct from the Frequentist point of view because the result of the interpretation is known to be wrong in 5% of the cases, and Frequentist intervals are statements about the data, not about the signal. But this also means that by increasing the background expectation, an experiment can be improved *a priori* so that it would turn out to be better to have an inferior experiment with a higher expected background level, and even worse, the exclusion result can be improved *a posteriori* by increasing the background expectation [100].

The solution to avoid these unintuitive properties is to require that experiments without sensitivity to a particular model should not be able to exclude this model. One therefore defines the “Modified Frequentist confidence level” CL_s as

$$CL_s := CL_{s+b} / CL_b, \quad (7.19)$$

which can be thought of as an approximate confidence in the signal-only hypothesis [100], and in terms of p -values can be formulated as

$$CL_s = \frac{p\text{-value of signal-plus-background hypothesis } \mathcal{H}_{S+B}}{1 - p\text{-value of background-only hypothesis } \mathcal{H}_B}. \quad (7.20)$$

Just like with CL_{s+b} , the signal-plus-background hypothesis \mathcal{H}_{S+B} is excluded at a confidence level of $1 - \alpha$ if $CL_s < \alpha$. As CL_b is smaller than or equal to one, CL_s is always greater than or equal to CL_{s+b} , i. e. the models excluded by using CL_s are a subset of those excluded using CL_{s+b} so that the upper limit using CL_s is higher and therefore weaker and more conservative. Being a ratio of probabilities, CL_s may become larger than one and is not a proper probability itself.

An important figure of merit of the average expected performance of an experiment are the expected CL_s and CL_{s+b} values under assumption of the background-only hypothesis^m. If the expected CL_s value (or CL_{s+b} , depending on the choice of which criterion to use) lies above the threshold for the exclusion of a model, then the experiment is not expected to be able to exclude this model even in case the signal is truly absent. It still may be able to exclude this model due to underfluctuations in the number of observed events, if the sensitivity is not too small such that the renormalization in the CL_s method prohibits the exclusion. In general, it is therefore interesting to compare the expected and observed exclusion limits to judge whether the experiment potentially should be able to extend its currently observed exclusion limits. In order to get an impression of the typical variations in the exclusion potential, in addition to the expected CL_s values the change in the expected CL_s when the background fluctuates up or down by one standard deviation is of interest. These fluctuations are often indicated as one-sigma bands above and below the expected CL_s limits, e. g. when plotting limits on production cross sections as function of the particle mass (usually the ratio to the theoretically expected cross section is shown).

^m They are computed as the fraction of the distribution of the test statistic which lies above the median of the test statistic, assuming the validity of \mathcal{H}_B . The expected CL_b therefore is 0.5, the quantile value of the median.




Figure 7.19: Example distributions of the test statistic $Q = -2 \ln(X)$ under the two hypotheses \mathcal{H}_{S+B} (dashed line) and \mathcal{H}_B (solid line), with $d = 60$, $b = 50$ and $s = 30$ (left) and $d = 40$, $b = 50$ and $s = 2$ (right). The vertical line marks the observed value of the test statistic.

Figure 7.19 shows examples of the distribution of the test statistic $Q = -2 \ln(X)$ for the two hypotheses \mathcal{H}_{S+B} and \mathcal{H}_B , where X is calculated according to Equation (7.12), and its cumulative distribution is governed by Equations (7.14) and (7.15), respectively. In this type of plot, the distribution of Q for \mathcal{H}_{S+B} is always shifted left with respect to \mathcal{H}_B , as can be seen from Equation (7.13) because d_i on average is larger when evaluating the distribution for \mathcal{H}_{S+B} . The histograms have been filled performing $2 \cdot 10^5$ toy experiments, varying the observed number of events randomly under the two hypotheses and computing each time Q . As only one channel is used here for simplicity, strong binning effects are visible due to the discreteness of the event numbers that go into the test statistic. In the left plot, $d = 60$ observed, $b = 50$ expected background and $s = 30$ expected signal events have been used, resulting in a good separation of the distributions for the two alternative hypotheses. Both CL_s and CL_{s+b} are around 0.01 (CL_b is 0.9), so having chosen $\alpha = 0.05$, the signal hypothesis is excluded at a 95% confidence level (C. L.). In general, for large significances the curves for \mathcal{H}_{S+B} and \mathcal{H}_B are well separated so that in case of the absence of signal, CL_b is close to one and the exclusion potential of CL_s converges to that of CL_{s+b} . For low sensitivity, the denominator in Equation (7.19) makes CL_s larger. An example of this is shown in the right plot in Figure 7.19, where the signal expectation $s = 2$ is very small compared to the background expectation $b = 50$ and the number of observed events is taken to be $d = 40$. For this example, the signal would still be excluded at 95% C. L. using CL_{s+b} ($CL_{s+b} = 0.037$), but not using CL_s ($CL_s = 0.56$), because CL_b is also very small ($CL_b = 0.066$).

So far, no uncertainties on the signal and background expectations have been taken into account. Usually no uncertainties on the confidence level are quoted, but instead the confidence limits are modified to allow for the experimental uncertainties, meaning in the case of CL_s that the higher the uncertainties, the broader the two distributions of the test statistic, the longer the tails, the larger the overlap and the weaker the exclusion limits will be. This can be implemented by doing toy experiments like they were used to fill the histograms in Figure 7.19 and smearing the expected number of signal and background events with the corresponding uncertainties which will yield broader distributions for the test statistic. Correlations between uncertainties can be taken into account by fluctuating these numbers in a correlated manner in each pseudoexperiment. In the end, these toy experiments are an




Figure 7.20: CL_s from an analytic calculation for the simple case of only one search channel and without taking any uncertainties into account. The horizontal and the vertical axis are the number of background events b and signal events s , respectively. The lines are the 95% C. L. contours for different values of the observed number of events d .

approximation to averaging over all possible outcomes for signal and background given by the probability distributions of their systematic uncertainties.

The CL_s method only yields a binary decision whether or not a given scenario (in terms of signal and background expectations) is excluded at a given certainty level. From this, upper limits on the signal strength can be found by starting in the excluded region and reducing the number of expected signal events gradually, which will give a monotonically increasing CL_s value. This is continued until CL_s is no longer smaller than the given threshold to find the upper limit on the number of signal events at a given certainty level. Note that the confidence intervals constructed in this way do neither have the same interpretation as Frequentist, nor as Bayesian intervals [99].

For the very simple case of only one search channel ($n = 1$) and no uncertainties, CL_s can be calculated analytically without effort. The p -values from Equation (7.20) can then be computed directly from the cumulative distribution function of the Poisson distribution. In the numerator of Equation (7.20), the CL_{s+b} values are symmetric under exchange of s and b , and so is CL_s , as long as CL_b in the denominator is close to one. For small s though, the symmetry of CL_s is broken. This may seem counterintuitive because CL_b only depends on b and not on s , but due to the symmetry of CL_{s+b} , small s correspond to large b , for which the p -value of the background-only hypothesis increases and CL_b thus goes down. This can be seen in Figure 7.20, which shows exclusion limits obtained from the analytic calculation in the plane spanned by the number of expected signal events s and expected background events b . Five different hypothetical numbers of observed events d are compared, giving five contours in the plane where the computed value of CL_s is 0.05, corresponding to an exclusion at a confidence level of 95%. In the plot in this figure, everything to the upper right of the respective contour lines is excluded at the given confidence level.

Profile Likelihood Ratio Method

The second method which will be used later to evaluate and interpret the outcome of the analysis of ATLAS data is the Profile Likelihood Ratio (PLR) method. It is described e. g. in [42]. Like CL_s , which uses a likelihood ratio in the test statistic X , the PLR method is also based on a likelihood ratio. But unlike CL_s , the central step in the PLR method involves fitting the data using a model that is not prescribed by the method itself. This model has two classes of parameters. The first class are the parameters of interest, which are parameters that the evaluation of the measurements is aiming at, for example particle masses or cross sections. The parameters in the second class are called nuisance parameters, for which values are determined by the fit to the data, too, but these values are not considered relevant and only introduced to allow the model to give a better description of the data by including additional degrees of freedom. Systematic uncertainties like signal efficiencies or the jet-energy scale are typical examples of parameters in this second class.

The model that is used in the PLR method to describe the data is encoded in terms of a likelihood function that assigns a likelihood value to every possible set of parameters and measurements:

$$L(\vec{\mu}, \vec{v}) = \prod_{i \in \text{events}} P(\vec{x}_i | \vec{\mu}, \vec{v}). \quad (7.21)$$

$P(\vec{x} | \vec{\mu}, \vec{v})$ is the probability to obtain a measurement \vec{x} , given the values $\vec{\mu}$ for the parameters of interest and \vec{v} for the nuisance parameters, which characterize the shape of the individual PDFs that together constitute P . This is the most general form, where the product runs over all events. In a binned analysis, it would run over all measurement channels or bins of the relevant histograms instead. The typical form of P is a product of Poisson probabilities for the observation of d_j events in channel j , with an expected number of s_j signal and b_j background events, multiplied by additional PDFs describing subsidiary measurements and the probabilities to observe a set of given values for the nuisance parameters. Restricting to the signal strength parameter as the only parameter of interest μ^n , the binned form of the likelihood can be written as

$$L(\mu, \vec{v}) = \prod_{j=1}^n \text{Pois}(d_j | \mu f_s(s_j, \vec{v}) + f_b(b_j, \vec{v})) \cdot \prod_k p(v_k), \quad (7.22)$$

where n is the number of channels as before, and k runs over the nuisance parameters. The signal strength μ is a simple scaling factor for the number of expected signal events. $\mu = 0$ corresponds to the background-only hypothesis, $\mu = 1$ is the nominal signal hypothesis. Since the nuisance parameters can have an influence on the number of expected events, the Poisson probabilities have not been written in a way that they depend directly on s and b , but additional wrapper functions f_s and f_b have been inserted, which incorporate this dependency and will be written out below. The second product contains the probabilities p to observe given values of the nuisance parameters, which here have been assumed to factor. These probability distributions are prior distributions in the sense of Bayesian statistics and are often taken to be uniform or Gaussian.

Using the likelihood function L , the profile likelihood ratio λ is defined as:

$$\lambda(\mu) = \frac{L(\mu, \hat{\vec{v}})}{L(\hat{\mu}, \hat{\vec{v}})}. \quad (7.23)$$

ⁿ The following can also be generalized to more than one parameter of interest [42].

In the denominator, $\hat{\mu}$ and $\hat{\vec{v}}$ are the maximum-likelihood estimators for μ and \vec{v} , which are obtained by maximizing the likelihood L over all possible values of μ and \vec{v} . The numerator is called the profile likelihood function. $\hat{\vec{v}}$ in the numerator is the conditional maximum-likelihood estimator, which maximizes L for a given value of μ . (The computation of the maximum is often called a maximum-likelihood fit.) Note that the likelihood ratio λ in the end only depends on the parameter of interest μ . As the likelihood is a product of probabilities, from the above definition it is clear that $0 \leq \lambda(\mu) \leq 1$. $\lambda(\mu)$ approaches one if the data agrees well with the hypothesized value of μ . In many analyses, the number of signal events is assumed to be non-negative, and therefore negative values of μ are excluded because they would correspond to negative signal contributions. As probabilities usually involve products of small numbers, it is convenient to use the logarithm of the likelihood ratio

$$t(\mu) = -2 \ln \lambda(\mu) \quad (7.24)$$

as test statistic, which, as in the case of the CL_s method, by convention includes a factor -2 . t attains its minimum value of zero at $\mu = \hat{\mu}$, and larger values of t correspond to an increasing incompatibility between the data and the hypothesized value of the signal strength μ .

Knowing the probability distribution $f(t(\mu)|\mu)$ of the Log-Likelihood Ratio (LLR) $t(\mu)$ for a fixed value μ ,

$$p_\mu = \int_{t_{\text{obs}}}^{\infty} f(t(\mu)|\mu) dt(\mu) \quad (7.25)$$

gives the probability to obtain a value for $t(\mu)$ equally or less compatible with the data than the value t_{obs} which has been observed under the assumption of the signal strength μ . If the p -value for a given μ , p_μ , is found to lie below a critical value α , then this value of the signal strength μ is said to be excluded at a confidence level $1 - \alpha$. From this, an upper limit on μ as the largest value for μ so that $p_\mu \leq \alpha$ can be derived.

It is therefore important to know $f(t(\mu)|\mu)$. A number of approximations for the distributions of different variants of the LLR and analytic expressions for the resulting upper limits are presented in [42]. These approximations are based on results by Wald [112] and Wilks [114], stating that under certain optimality conditions and for sufficiently large statistics, $t(\mu) = -2 \ln \lambda(\mu)$ follows a χ^2 distribution with $h - m$ degrees of freedom, h being the total number of parameters and m being the number of nuisance parameters, i. e. here $h - m = 1$. This approximation is also used in the implementation of the PLR method that is described below.

7.10.4 Comparison of the CL_s and PLR Methods

Two different methods to compute limits have been presented above. As the two methods are optimized with regard to slightly different aspects, no exact agreement is expected, but still the limits that are obtained should be consistent. This will be studied and discussed in this section. First, a summary of important differences between CL_s and PLR is given.

- The main purpose of CL_s is dealing with underfluctuations in case of small sensitivity, and to avoid exclusion of signals if the exclusion is not justified due to a lack of sensi-




Figure 7.21: Comparison of exclusion limits from the CL_s and the PLR method for different number of observed events (given in the legend) as function of the number of background events b and signal events s .

tivity. In PLR, no such mechanism is implemented⁹ and this will later become apparent through the fact that the limits obtained from CL_s are slightly more conservative than those from the PLR method.

- The procedure to compute limits using the CL_s method only gives an upper limit on the number of signal events, which is fully acceptable when setting exclusion limits. For sufficiently high signal-to-background ratios, PLR may also yield a lower limit on the signal, depending on the number of observed events. This is visible in Figure 7.21, which is discussed below.
- With respect to the implementation, the PDF for the nuisance parameters in the ROOSTATS framework can be modeled to any level of complexity using custom PDFs. This allows to use an asymmetric distribution e. g. to better take into account the nature of the asymmetric JES uncertainties. In the CL_s implementation employed here (TLIMIT), this is not possible because all uncertainties are modeled as Gaussian distributions. Of course, this is not a general shortcome of the CL_s method, which could as well be implemented to use other distributions for the uncertainties when doing the pseudo-experiments.

In Figure 7.21, the exclusion limits set by the CL_s and PLR method are compared. On the horizontal axis the number of background events b and on the vertical axis the number of signal events s is varied between 0 and 80. The contours shown in the plot correspond to five different numbers of observed events. They mark where the CL_s or p -value from PLR, respectively, are equal to 0.05. Everything to the top right of each line is excluded at

⁹ In principle, the ratio of two hypotheses could be used with PLR in a similar manner as done in CL_s , too, instead of using the p -value of the hypothesis \mathcal{H}_{S+B} alone.

a confidence level of 95%. The contour lines are symmetric under $s \leftrightarrow b$, except for very small signal numbers. The plot shows that the limits from the two different methods give the expected agreement, the exclusion from CL_s being more conservative for low expected signal strengths than the exclusion from PLR. It also shows, as mentioned above, that PLR, due to the parabolic shape of the LLR, may give an additional second contour at lower $s + b$ values. In these cases, the outmost line is the exclusion limit of interest here.

Note that for the CL_s method all uncertainties, and for the PLR method all but one, namely the JES uncertainty, are modeled using symmetric Gaussian distributions. This would in principle allow to combine all uncertainties into one, adding them up in the appropriate way given by their correlations. The main advantage of keeping the uncertainties as distinct numbers is an easier handling of many channels with different combinations of correlations between uncertainties, which does not apply in the following, where only one channel is used.

7.11 Interpretation of Analysis Results

In Section 7.10, the methods have been presented which will be employed in the following to perform the final step of the search for Supersymmetry in ATLAS 2010 data in the zero-lepton channel. CL_s and PLR will now be applied to the results of the cutflow and the uncertainties summarized in Section 7.8. As no statistically significant excess over the Standard Model background expectations is observed in the data from 2010, the event counts are interpreted by setting exclusion limits on processes that go beyond the Standard Model. Here, the observable that comprises the search result simply is the number of events that satisfy all cuts that define one of the four signal regions. The limits will first be set in a model-independent way, and then using the Monte Carlo grids presented in Section 7.6 to model the supersymmetric signal events.

7.11.1 Model-Agnostic Limits

Without choosing a specific set of parameters for Supersymmetry or even without specifying the type of Beyond Standard Model (BSM) physics, a model-independent limit on the cross section of BSM physics processes can be derived^E. Of course, the value of this limit, being the result of a particular analysis, depends on the cuts which are applied to suppress Standard Model backgrounds, and these are optimized for, and therefore prefer, particular signal event topologies. Thus, the limit is not fully model-independent, but at least it is model-agnostic.

Although in this approach the number of signal events reaching the signal regions defined by the cutflow is left free, still an estimate of the relative uncertainties on the signal is needed. Two different values are used and compared in this section, 20% and 40%, which approximately reflect the spread of the quadratic sum of all signal uncertainties at typical points within the mSUGRA signal grid. These total uncertainties are used in an uncorrelated way with respect to the background uncertainties, and therefore have conservatively been rounded up. On top of this uncorrelated uncertainty estimate, the uncertainty on the

^E This means apart from the assumption that they lead to an increase in event counts, not to a destructive interference with known physics processes.

Method	Δ_{tot}	Signal Region							
		A		B		C		D	
		n_{signal}	σ [pb]	n_{signal}	σ [pb]	n_{signal}	σ [pb]	n_{signal}	σ [pb]
PLR	0.00	53.4	1.6	14.1	0.4	39.9	1.2	4.6	0.1
	0.20	56.8	1.7	15.4	0.5	42.6	1.3	4.9	0.1
	0.40	74.8	2.2	24.3	0.7	56.6	1.7	6.7	0.2
CL_s	0.00	58.2	1.7	12.8	0.4	44.0	1.3	5.7	0.2
	0.20	62.7	1.9	14.2	0.4	47.5	1.4	6.1	0.2
	0.40	85.8	2.6	20.3	0.6	64.5	1.9	9.0	0.3

Table 7.17: Model-agnostic upper limits on the number of signal events n_{signal} and the effective cross section σ in the four signal regions, from PLR and CL_s , comparing three different values of the total systematic uncertainty Δ_{tot} .

luminosity is included when computing the limits, and in addition to the two values given above, the results assuming no uncertainty but the one from the luminosity are also shown.

The PLR method directly returns an upper limit on the number of signal events n_{signal} , fixing the signal strength to $\mu = 1$ and leaving n_{signal} free as the parameter of interest, then using the approximation for the distribution of the test statistic as described above. To find an upper limit on n_{signal} following the CL_s prescription, the set of possible values of n_{signal} needs to be scanned, which is done here using a simple bisection method with a fixed number of 10 steps. This is considerably faster than a scan of equidistant points. As the CL_s implementation is based on Monte Carlo, the resulting CL_s values will exhibit random fluctuations which may cause the bisectioning to exclude the wrong half of the interval, but this will only happen when the interval sizes are already of the size of the Monte Carlo fluctuations. The sequence of intervals has been checked with respect to this effect, and the effect has been found to be negligible, as can be seen in Figure 7.22. This plot shows, for a total uncertainty of 40% plus the luminosity uncertainty, the development of CL_s as it is computed when running the bisection method for the four different signal regions. Each dot corresponds to the CL_s value computed for the midpoint of the current interval. The horizontal line marks $\text{CL}_s = 0.05$, and the vertical line gives the final upper limit as the mid of the interval obtained from the last iteration. The fluctuations from the Monte Carlo evaluation are visible in the plot, but in the relevant region the dots do not fluctuate by so much that a significant impact on the limits is expected. From the plot, a systematic uncertainty arising from the bisection method seems to be no larger than ± 1 event. (The slope of the dependence of CL_s on n_{signal} for the other two total uncertainties, 0% and 20%, is larger and the uncertainty on the limit from the method therefore is smaller.)

The upper limits on the number of signal events n_{signal} from the two methods are summarized in Table 7.17 for the three different total uncertainties Δ_{tot} , which do not include the correlated luminosity uncertainty. The table also lists the corresponding upper limit on the cross section calculated via $\sigma = n_{\text{signal}}/\mathcal{L}_{\text{int}}$, where \mathcal{L}_{int} is the total integrated luminosity from Section 7.2. Note that these limits naturally cannot take into account the signal efficiency because no particular form of the signal or topology of signal events is assumed.




Figure 7.22: Evolution of CL_s in the determination of the upper limit on the number of signal events using a bisection method. Each point corresponds to the CL_s value computed for the midpoint of the current interval. The colors correspond to the four different signal regions. The horizontal line marks $CL_s = 0.05$, the vertical line the final upper limit. In this plot, a total uncertainty of 40 % (plus the luminosity uncertainty) is assumed.

The limit is therefore understood to be on an effective cross section including signal efficiency and other factors as appropriate. The confidence level is 95 % for both PLR and CL_s . The higher the systematic uncertainties, the larger, i. e. weaker, are the limits on the number of signal events and the cross section. Note that CL_s always gives slightly weaker limits than PLR except for signal region B. The harshest event selection, corresponding to signal region D, gives the best limits among the four signal regions, excluding BSM physics processes with effective cross sections above 0.1 to 0.3 pb. The dependence of these limits on the total uncertainty Δ_{tot} is strong, as can be seen in the table. No additional uncertainty on these limits has been evaluated.

7.11.2 Exclusion of Specific SUSY Scenarios

In this section, the results summarized in Section 7.8 are used as inputs to set limits and thereby exclude certain ranges of values for the parameters of specific models of supersymmetric physics processes. As the parameter space even in minimal Supersymmetry models with additional constraints has too high a number of dimensions, only scans of low-dimensional subspaces of the full parameter space are possible. This is done here using the signal grids from Section 7.6, which implement two constrained models. For the mSUGRA model, the grid points span the m_0 - $m_{1/2}$ plane, the MSSM grid consists of points in the gluino mass-squark mass plane.

Selecting the Signal Region

As the four signal regions are not independent, it needs to be decided how to deal with their correlations. The approach chosen here is to select the best signal region for each grid point

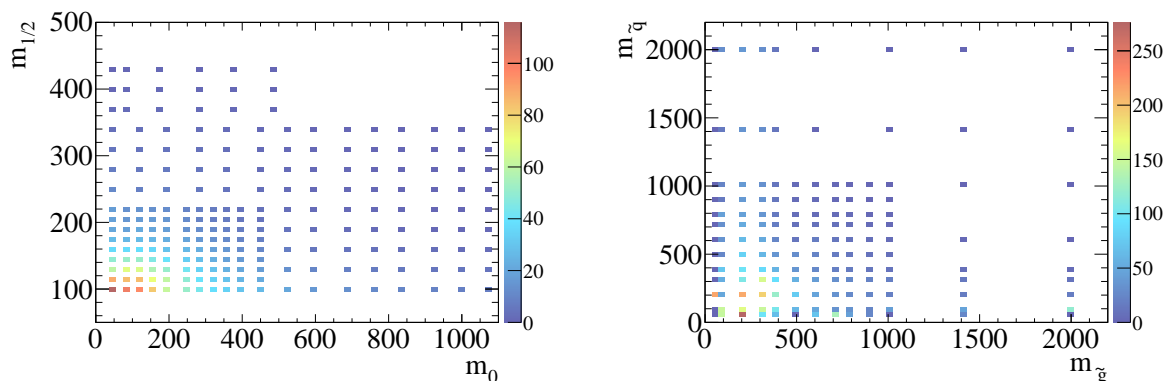


Figure 7.23: Best s/\sqrt{b} ratio (i. e. largest among the four signal regions) for the Monte Carlo samples in the signal grids. Left: mSUGRA grid, right: MSSM grid.

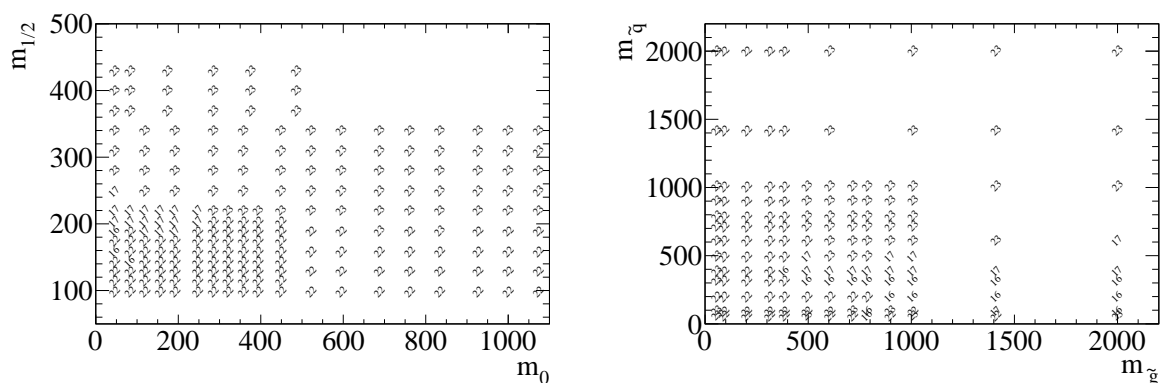


Figure 7.24: Signal region with the best s/\sqrt{b} ratio, denoted by the number of its last defining step in the cutflow. Left: mSUGRA grid, right: MSSM grid.

individually. This must be done without looking at the data, i. e. from the expected exclusion potential alone, to avoid a bias similar to the look-elsewhere effect. Thus, for each of the four signal regions, the expected and observed limits are computed independently, and the one with the best expected limit is adopted as the nominal result, as has been done e. g. in [94]. Several different measures can be used to determine the best exclusion potential: the largest expected ratio of signal over background events, the smallest expected CL_s value and the smallest expected p -value from the PLR method.

Figure 7.23 shows, for each of the Monte Carlo samples comprising the mSUGRA grid (left plot) and MSSM grid (right plot), the best, i. e. the highest, of the four s/\sqrt{b} values obtained for each of the four signal regions at every point of the grid. The plots in Figure 7.24 show which of the four signal region provides this s/\sqrt{b} ratio. In these two plots, the number of the last step in the cutflow from Table 7.2, defining the respective signal region is printed, so that the numbers 16, 17, 22 and 23 denote signal regions A, B, C and D, respectively.

From the right plot in Figure 7.23 for the MSSM grid, one can see that the ratio s/\sqrt{b} is largest in a vertical and a horizontal strip with gluino or squark masses around 200 GeV. For higher masses, as well as for the grid points where one of the masses is very low, s/\sqrt{b} becomes smaller. As the background is the same for all grid points, this means that the

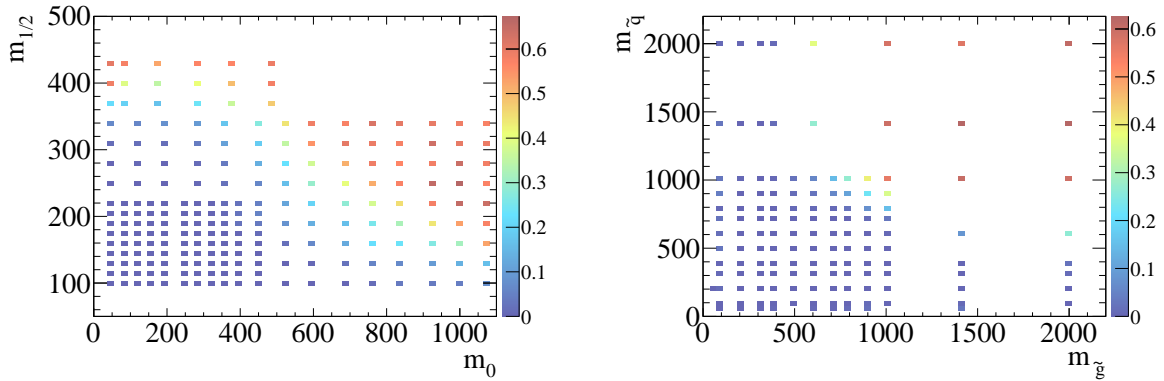


Figure 7.25: Best CL_s value (i. e. smallest amongst all four signal regions) expected under the assumption of background only for the Monte Carlo samples in the signal grids. Left: mSUGRA grid, right: MSSM grid.

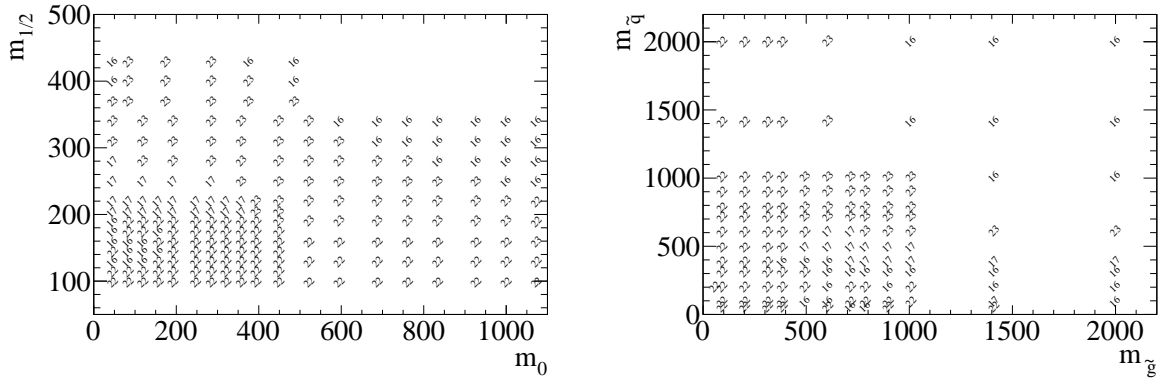


Figure 7.26: Signal region with the best CL_s value expected under the assumption of background only, denoted by the number of its last defining step in the cutflow. Left: mSUGRA grid, right: MSSM grid.

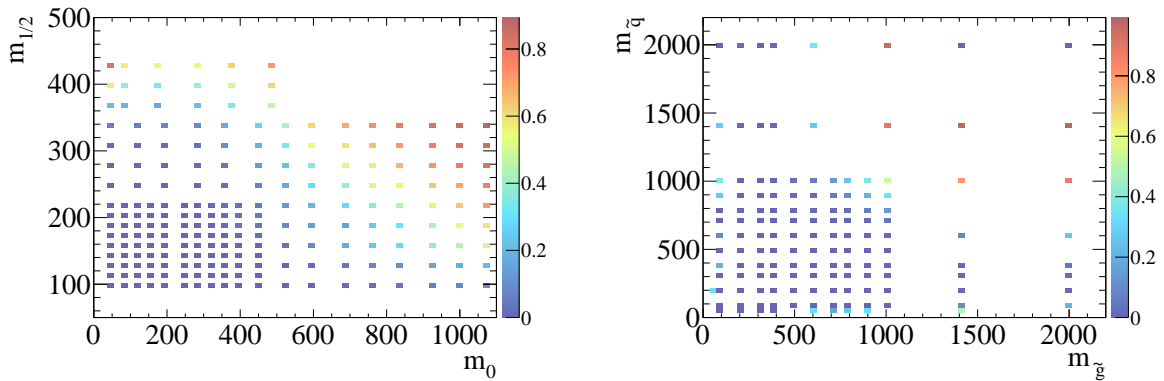


Figure 7.27: Best PLR value (i. e. smallest amongst all four signal regions) expected under the assumption of background only for the Monte Carlo samples in the signal grids. Left: mSUGRA grid, right: MSSM grid.

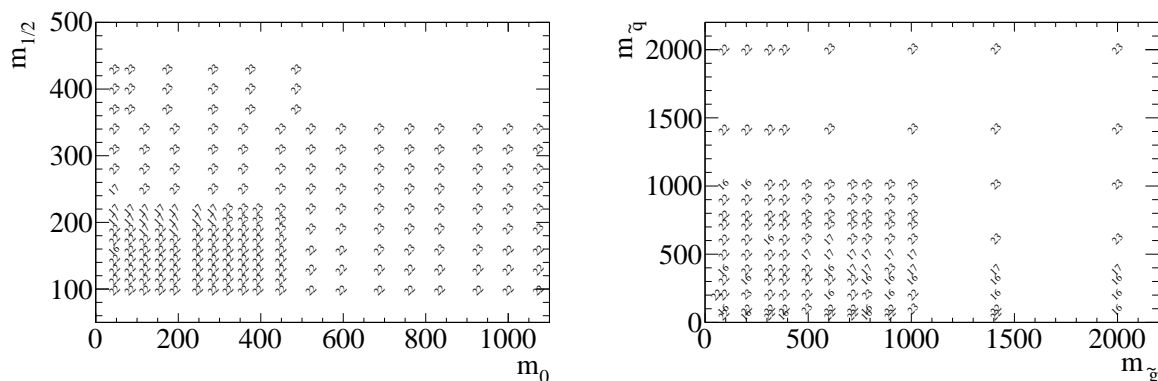


Figure 7.28: Signal region with the best PLR value expected under the assumption of background only, denoted by the number of its last defining step in the cutflow. Left: mSUGRA grid, right: MSSM grid.

signal yield in these regions is lower. Taking into account that the cross section falls off monotonically with increasing masses (cf. Figure 7.7), this in particular means that the much larger cross section for low masses cannot compensate for the loss in the signal efficiency for low squark and gluino masses, due to the softer spectrum of the decay products, so that the product of larger cross section and smaller signal efficiency still yields a smaller number of events expected in the signal regions. For the mSUGRA grid in the left plot, the behavior is different and no such peak at intermediate masses can be seen. Here, the ratio s/\sqrt{b} decreases in general radially, with the highest values in the region where both m_0 and $m_{1/2}$ are small.

In Figures 7.25 and 7.26, the same type of plots is shown, this time choosing the signal region according to the best CL_s value, where the best CL_s value is the smallest CL_s value expected under the assumption of the background-only hypothesis and computed from the expected number of signal and background events without considering the observed data. Again, Figure 7.25 shows the CL_s values obtained for the mSUGRA (left) and MSSM grid (right), and Figure 7.26 shows which signal region yields this CL_s value and will thus be used in the limit-setting procedure to determine the CL_s value for this particular grid point. Finally, in Figures 7.27 and 7.28 the same plots are shown, this time choosing the best signal region in terms of the smallest expected PLR value under the background-only hypothesis.

The main difference between a selection based on s/\sqrt{b} , CL_s or PLR is that CL_s and PLR take the uncertainties into account, whereas s/\sqrt{b} is purely based on the expected number of signal and background events in the signal regions. For most grid points, the signal region chosen according to the three measures is the same, but there are differences, with signal regions A and B chosen only rarely by s/\sqrt{b} , and more often by CL_s in both grids, and by PLR in the MSSM grid. Note that for some of the MSSM grid points, all of the signal regions have a expected number of signal events of zero. For these points $s/\sqrt{b} = 0$, and neither the CL_s nor the PLR value can be computed. They are therefore left out in the CL_s and PLR plots.




Figure 7.29: 95% C. L. exclusion limits in the mSUGRA plane spanned by m_0 and $m_{1/2}$. The red line marks the observed exclusion limit obtained from the CL_s method in this analysis, the solid blue line gives the expected exclusion limit and the dashed and dotted blue lines mark the $\pm 1\sigma$ bands from two different methods as explained in the text. Shaded areas indicate exclusion results from other experiments.




Figure 7.30: 95% C. L. exclusion limits in the mSUGRA plane spanned by m_0 and $m_{1/2}$. The brown line marks the observed exclusion limit obtained from the PLR method in this analysis, the solid green line gives the expected exclusion limit and the dashed green lines mark the $\pm 1\sigma$ band. Shaded areas indicate exclusion results from other experiments.




Figure 7.31: 95% C. L. exclusion limits in the mSUGRA plane spanned by m_0 and $m_{1/2}$, comparing in particular the results obtained in this analysis from the CL_s method with those from PLR. The line styles and color coding are the same as in Figures 7.29 and 7.30.

Exclusion Plots for the mSUGRA Grid

Figure 7.29 shows the 95% C. L. exclusion limits in the mSUGRA plane in terms of the expected and observed exclusion limits, as they are obtained from the CL_s method in this analysis. The expected limits mark the range of parameter space which the analysis is expected to be able to exclude at the given confidence level, the observed limits mark the range that is excluded based on the observed data. Differences between these two quantities may either arise from statistical fluctuations in the observed data, from sources of systematic bias which have not been taken into account properly, or, in case the observed limit stays below the expected limits, from signal processes. Figure 7.30 shows the limits which have been computed using the PLR method, and Figure 7.31 shows a comparison of the expected and observed limits from CL_s and PLR in one plot. In Figure 7.31, it can be seen that the expected limits for both CL_s and PLR are the same, but the observed limit from CL_s is more conservative than the one from PLR as expected.

To give a feeling for the size of the impact of statistical fluctuations, the $\pm 1\sigma$ bands (colored dashed lines) demonstrate the variation of the expected limits arising from up- and downwards fluctuations of the event counts by one standard deviation. These variations are computed in two different ways. For the CL_s limit from TLIMIT, they can be computed using the built-in methods from TConfidenceLevel (cf. Section A.3.1), which varies the expected background up and down by one standard deviation according to the distribution of the test statistic for the assumption of only background. The second approach used here to compute the bands is to vary the background expectation up and down manually by a value that corresponds to a fluctuation of the background by one standard deviation according to the total uncertainties on the background, and to re-run the limit calculation. This approach is possible for both CL_s and PLR. The two alternatives are different conceptually and therefore not *a priori* expected to give the same results. They are compared in Figure 7.29 because for PLR only the latter approach can be used. As can be seen in this plot, for the downwards fluctuations of the background, allowing to exclude a larger range, the two methods give similar results, for the upwards fluctuations they do not.

Results from searches for Supersymmetry have been published by past and present experiments. These publications set limits on many different quantities, such as masses of supersymmetric particles and cross sections, and use various different supersymmetric models and parameter sets. Care therefore needs to be taken to gather a set of results which is consistent in the choice of model and the values for its parameters. For comparison, results from other collider experiments are shown as shaded areas in the exclusion plots in Figure 7.29 through Figure 7.31. The macros which draw these regions are taken from the official code of the ATLAS Supersymmetry group. In the following, it is described where these results can be found in literature and what are the respective model settings. Note that parts of the plane are also excluded by theory constraints. A small region with $m_0 \lesssim 60$ GeV and $m_{1/2} \lesssim 80$ GeV is excluded because there is no electroweak symmetry breaking, but this region is smaller than the region excluded by DØ [90].

For the mSUGRA grid, five different colors mark five different regions excluded by earlier experiments. The legend of the plot lists the name of the corresponding experiments and in brackets the channel of the search from which these limits were obtained. The first two exclusion regions are set by the LEP experiments from a search for signatures of two scalar leptons or two charginos, respectively. The following paragraph explains where the limits included in the official exclusion plots come from. The LEP slepton exclusion region shown

in Figure 7.29 matches the exclusion limit which is cited in [90] as the result of the LEP2 slepton searches for $\tan\beta = 3$, $A_0 = 0$ and $\text{sign}\mu < 0$ with reference to [125], and in [92] as the result of the LEP slepton searches for $\tan\beta = 3$, $A_0 = 0$ and $\text{sign}\mu > 0$ with reference to [124] and [126]. Similar limits are shown as result of a slepton search by the ALEPH collaboration, interpreted in a minimal supergravity setting with $\tan\beta = 5$ and 10 , $A_0 = 0$ and both $\text{sign}\mu < 0$ and $\text{sign}\mu > 0$ in [88], the exclusion limits from the slepton search on $m_{1/2}$ for $\text{sign}\mu < 0$ being slightly weaker than for $\text{sign}\mu > 0$. In the same plot in [88], an exclusion region from chargino searches is presented, which extends over the whole m_0 range in the m_0 - $m_{1/2}$ plane, again for $\tan\beta = 5$ and 10 , $A_0 = 0$ and both signs of μ , and again the limits are weaker in $m_{1/2}$ for $\text{sign}\mu < 0$. For the chargino search, the border of the exclusion region differs by about 30 GeV. An update on these exclusion limits is given in [89, 120] for $\tan\beta = 10$ and 30 , $A_0 = 0$ and both signs of μ , showing small improvements for the chargino limits at $\tan\beta = 10$. [93] cites an exclusion of the chargino masses below 103.5 GeV as LEP direct limit, with reference to [124], and shows this in relation with limits on $m_{1/2}$ up to 170 GeV for $\tan\beta = 3$, $A_0 = 0$ and $\text{sign}\mu > 0$ and small m_0 . [90] also includes such limits under the label LEP2 chargino limits, with reference to [125], for $\tan\beta = 3$, $A_0 = 0$ and $\text{sign}\mu < 0$ but with significantly smaller exclusion reach only up to 130 GeV in $m_{1/2}$. [92] includes LEP chargino limits up to 170 GeV in $m_{1/2}$ for $\tan\beta = 3$, $A_0 = 0$ and $\text{sign}\mu > 0$, with reference to [124] and [126]. In conclusion, the chargino limits cited in the exclusion plots from the LEP experiments are appropriate for $\text{sign}\mu > 0$, but if $\text{sign}\mu < 0$ as in the scenario assumed in the mSUGRA grid used for this analysis, the limits from LEP for $m_{1/2}$ seem to be weaker than indicated by the official contours.

The other three exclusion regions included in the plots are from Tevatron experiments. The exclusion regions from the $D\bar{D}$ experiment are from a search of 2.3 fb^{-1} of data for production of a chargino $\tilde{\chi}_1^\pm$ and the second-lightest neutralino $\tilde{\chi}_2^0$, with final states containing three charged leptons and missing transverse energy, interpreted in the mSUGRA framework with $\tan\beta = 3$, $A_0 = 0$ and $\text{sign}\mu > 0$ [92], and from a search of 2.1 fb^{-1} of data for squarks and gluinos in final states with jets and missing transverse energy, interpreted under the same assumptions but with $\text{sign}\mu < 0$ [90]. The exclusion region from CDF is the result of a search for gluinos and squarks in 2.0 fb^{-1} of Tevatron data, assuming minimal supergravity with $\tan\beta = 5$, $A_0 = 0$ and negative sign of μ [121], which has been published as [94] but without the plot of the m_0 - $m_{1/2}$ plane.

The thick dashed-dotted black line is the result of the official zero-lepton analysis of 2010 data by the ATLAS Supersymmetry group [105]. The slim dashed black line is the corresponding result from CMS on the same amount of data from 2010 [39]. The dashed grey lines labelled with particle type (\tilde{g} or \tilde{q}) and mass in GeV indicate the gluino and squark masses for $\tan\beta = 10$. As the mSUGRA and MSSM models have completely different mass spectra, the limits on the squark and gluino masses read off from this plot need not correspond to the limits shown below.

Exclusion Plots for the MSSM Grid

Figure 7.32 shows the 95% C. L. exclusion limits in the MSSM plane, again in terms of the expected and observed exclusion limits, as they are obtained from the CL_s method in this analysis. The $\pm 1\sigma$ bands are shown as colored dashed lines. Figures 7.33 and 7.34 show the limits which have been computed using the PLR method and a comparison of the expected and observed limits from CL_s and PLR for the MSSM grid in one plot.




Figure 7.32: 95% C. L. exclusion limits in the MSSM plane spanned by $m_{\tilde{g}}$ and $m_{\tilde{q}}$. The solid red line marks the observed exclusion limits obtained from the CL_s method in this analysis, the solid blue line gives the expected exclusion limit and the dashed and dotted blue lines mark the $\pm 1\sigma$ bands from two different methods as explained in the text. Shaded areas indicate exclusion results from other experiments.




Figure 7.33: 95 % C. L. exclusion limits in the MSSM plane spanned by $m_{\tilde{g}}$ and $m_{\tilde{q}}$. The brown line marks the observed exclusion limits obtained from the PLR method in this analysis, the solid green line gives the expected exclusion limit and the dashed green lines mark the $\pm 1\sigma$ band. Shaded areas mark exclusion results from other experiments.




Figure 7.34: 95 % C. L. exclusion limits in the MSSM plane spanned by $m_{\tilde{g}}$ and $m_{\tilde{q}}$, comparing in particular the results obtained in this analysis from the CL_s method with those from PLR. The line styles and color coding are the same as in Figures 7.32 and 7.33.

Four exclusion limits from earlier experiments have been included in the plots. At LEP, the L3 experiment has set a lower bound on squark masses of about 100 GeV from 450 pb⁻¹ of data in the MSSM framework. It extends over the gluino mass range up to 550 GeV, where the squark mass becomes smaller than the neutralino mass in their setting [91]. The other contours are from Run I and Run II of the Tevatron experiments. The contour from Run I is governed by the CDF result, interpreting 84 pb⁻¹ of data in a general MSSM framework with $\tan\beta = 3$ [35], looking for final states including three or more jets and missing transverse energy from the two LSPs. The Run II results from DØ are based on 2.1 fb⁻¹ of data, setting limits on the squark and gluino masses in an mSUGRA interpretation with $\tan\beta = 3$, $A_0 = 0$ and $\mu < 0$ [90]. The results from CDF use 2 fb⁻¹ and $\tan\beta = 5$, $A_0 = 0$ and $\mu < 0$ [94]. The lower part of the boundary of the Run II exclusion limits close to the diagonal is due to the fact that there is no mSUGRA solution for that part of parameter space in their setting.

The limit setting in the MSSM grid using the PLR method suffers from low Monte Carlo statistics in the region of low signal efficiency and large cross sections. This enhances fluctuations in the p -values, which in the end lead to problems in the contour finding. When using only the signal regions which contain at least 10 raw signal Monte Carlo events, corresponding to roughly one per mille of the available Monte Carlo statistics, also PLR gives stable results, thus this approach is used here. Figures 7.33 and 7.34 cannot exclude the region below $m_{\tilde{q}} < 100$ GeV or $m_{\tilde{g}} < 200$ GeV, but all of this region is excluded by the CL_s method, for which the effect was not observed, as can be seen in Figure 7.32.

7.12 Summary and Conclusion

The Supersymmetry analysis presented in this chapter has evaluated data taken with the ATLAS detector in 2010, corresponding to an integrated luminosity of $\mathcal{L}_{\text{int}}^{2010} = 33.4$ pb⁻¹ of proton-proton collisions at $\sqrt{s} = 7$ TeV. It validates the official analysis, which has been done in the ATLAS Supersymmetry group, and is optimized for supersymmetric final states with no leptons and high-energetic jets and large missing transverse energy. The event counts in four signal regions are used to set limits on mass parameters in two different supersymmetric settings, mSUGRA and MSSM. Limits are derived with two different methods, CL_s and PLR, which give consistent results. The limits are expressed as exclusion contours at a 95% confidence level in Figures 7.29 through 7.34, where for comparison results from other experiments and the official analyses of 2010 data from ATLAS and CMS are included.

The comparison to the exclusion contour from the official ATLAS analysis in these plots shows that the limits set in the analysis presented here are weaker. The excluded mass ranges are smaller by roughly 10 to 15% for the mSUGRA grid and 20% for the MSSM grid. This is expected and mainly due to the inclusion of pile-up effects, which have not been considered in the official analysis. The pile-up effects are taken into account here in a conservative approach in terms of an additional uncertainty on all backgrounds for which corresponding Monte Carlo samples are available. A reweighting of the 2010 Monte Carlo according to the pile-up distribution in data is not possible because no suitable Monte Carlo samples are available. Furthermore, an updated estimate of the luminosity has been used here. This has two consequences: The uncertainty on the luminosity is smaller, thus improving the limit. As the uncertainty on the luminosity is small compared to most of the other uncertainties, this has only a negligible impact though. On the other hand, the new luminosity estimate gives a 4% smaller integrated luminosity, so that the estimation of the number of

background events from Monte Carlo is smaller, whereas the number of observed events is the same as in the official analysis. This also leads to a weaker limit.

For the continuation of the analysis with 2011 data, a number of changes are due. Besides using updated object definitions and calibrations, the most important changes are made necessary by the increase in the instantaneous luminosity, which has two major consequences. Firstly, higher rates will make triggers with too low thresholds unaffordable so that these will become prescaled. In order not to lose sensitivity due to the prescales, triggers with higher thresholds will have to be used instead. As the analysis strategy currently assumes the trigger to be fully efficient, the offline cuts need to follow the online requirements, which means that harsher offline cuts need to be applied, depending on the onset of the plateau for the new triggers. Section 5.6 discusses the updates in the trigger strategy for the zero-lepton analysis for 2011 and the efficiency of the triggers which have been employed in the analyses of 2011 data. Reweighting of Monte Carlo events to account for the trigger efficiency may become important for the zero-lepton analysis when increasing the statistics by lowering the offline cuts into the turn-on region of the primary trigger becomes attractive, or, put the other way round, when the trigger thresholds become so high, that the offline selection can no longer follow without inducing a significant loss of sensitivity in the search. Note that a steep turn-on curve will be of disadvantage here because the systematic uncertainties due to the trigger efficiency become large close to the turn-on region. Secondly, the increase in luminosity also entails an increase in the level of in-time pile-up, as discussed e. g. in the context of the \cancel{E}_T model in Section 6.5. Taking into account these pile-up effects in an appropriate way will then be necessary, in order not to deteriorate the analysis performance. For these higher levels of pile-up, a reweighting of the Monte Carlo events according to the distribution of the number of concurrent proton-proton interactions in the recorded events needs to be done, and in fact Monte Carlo samples with a broad spectrum of additional pile-up vertices have been produced and are in use for the on-going analyses of 2011 data in ATLAS. For the 2010 analysis, in particular for the MSSM grid, the spacing of the grid points was relatively large for intermediate and high values of the squark and gluino masses. This will also be improved in the 2011 Monte Carlo which has a finer granularity.

In conclusion, it is interesting to see that the relatively small amount of data collected in 2010 in terms of integrated luminosity already allows to set limits which are competitive with the existing limits from earlier experiments and often even go far beyond these, although these are based on a many times larger integrated luminosity but at a lower center-of-mass energy. Considering the fact that in 2011 already more than a hundred times more data has been collected by ATLAS, this demonstrates the impact that the physics results from the LHC will have.

A Appendix

A.1 Trigger Emulation

Writing code that is able to reproduce the decision of the online trigger system has a number of benefits and use cases outlined below. This kind of simulation of the trigger behavior will be called trigger emulation in the following, to make clear the distinction from the official trigger simulation that it used to simulate the trigger decision in Monte Carlo. The main difference between the trigger emulation which is used in this thesis for different purposes and the trigger simulation is that the trigger emulation only emulates the hypothesis algorithms, but it does not emulate the feature-extraction algorithms. This would be far more complicated, because it would require the reconstruction of physics objects from low-level detector data. Instead, the emulation uses the trigger objects reconstructed online by the feature-extraction algorithms of the trigger system. This brings some limitations because it is not possible to emulate trigger decisions in events for which the feature-extraction algorithms, which construct the required trigger objects, have not been run online so that these objects are not available. The official trigger simulation used in Monte Carlo covers the full trigger system, and thus also the feature-extraction algorithms. It does not only make use of the same code base which is also run online in the HLT, but for Level 1 also needs to simulate the trigger algorithms, which for performance reasons are implemented in hardware. This means that there may be more or less subtle differences between the performance of the actual trigger and its counterpart in Monte Carlo.

Although it only emulates a part of the trigger system, the trigger emulation can still be of use in a number of cases. First of all, it helps to understand in detail, how and based on what quantities the decision of the trigger is made. In particular, a validation of the kind which will be presented below would show whether the trigger emulation correctly captures all relevant features of the trigger hypothesis algorithms. The main use with respect to studies of the trigger performance is that the trigger emulation allows to study trigger definitions which do not exist in the online trigger menu, within the limitations outlined above.

What is always possible within these limitations is to emulate triggers with higher thresholds or additional cuts on existing trigger objects. The main pitfall which has to be taken care of is the principle of early rejection, which means that if none of the chains within a certain slice, e. g. the jet slice, has issued an accept for a given event at one of the trigger levels below the highest (Level 1 and Level 2 of the ATLAS trigger system), the feature extraction at the subsequent levels will not be executed. This means that the corresponding trigger objects will not be reconstructed. The trigger emulation will then always reject this event, although, if the feature-extraction had been run, they might have created objects which would have fulfilled the trigger selection. In this case, the decision of the trigger emulation gives the wrong result for this event. This will in particular cause problems for studies which select event samples based on an orthogonal trigger, because it is not guaranteed that the trigger objects needed for the trigger performance study are available at all trigger levels, but only those belonging to the orthogonal trigger which has accepted all of the events. This means

that for this type of trigger, the feature-extraction must have been executed at all trigger levels. On the other hand, by selecting a sample of events with the same type of trigger that is under study, but using a lower threshold in the selection, it can be ensured that the needed objects are available on all levels. A special case are the \cancel{E}_T triggers, for which the feature-extraction algorithms are run for all events at all trigger levels such that the emulation of these triggers is possible in all events that are recorded.

A.1.1 Implementation

In this thesis, the computation of trigger efficiencies is in general based on the emulation of triggers. Therefore, a validation of the emulation is advisable. Its results are presented in this subsection. First, a description is given of which triggers are emulated and how this is done, followed then by plots of the validation results, which show full agreement between the “official” trigger decision stored in the event record and the emulated trigger decision.

The triggers which are studied in this thesis, and thus have been implemented in the emulation code, are \cancel{E}_T triggers, jet triggers, multi-jet triggers, and combined jet + \cancel{E}_T triggers. The emulation will be briefly sketched in the following. The \cancel{E}_T triggers, being based on an event-global quantity, are the easiest type of trigger to emulate. The measurements of \cancel{E}_T at Level 1, Level 2 and Event Filter can be compared directly against the respective thresholds, where it has to be taken into account whether the muon contributions to \cancel{E}_T at Level 2 and Event Filter are included or not. In its current implementation at the time of writing, the \cancel{E}_T computation in the HLT does not make use of the muon correction (cf. Section 4.3.2). Cuts can be applied at all levels for the full chain \cancel{E}_T triggers, or only at Event Filter for the \cancel{E}_T part of the jet-seeded combined jet + \cancel{E}_T triggers. The latter type is referred to as EF_EFonly_xe25_noMu (for a threshold of 25 GeV) in this thesis (cf. Section 7.3). An event will pass Level 1 of the \cancel{E}_T trigger chains, regardless of the configured thresholds, if the flag signaling an overflow in the computation of \cancel{E}_x or \cancel{E}_y at Level 1 is raised, and also pass Level 2 if an overflow of \cancel{E}_T or $\sum E_T$ at Level 1 occurred, because Level 2 does not recompute \cancel{E}_T , but uses the Level 1 result (cf. Section 4.3.2).

The emulation of jet triggers is more sophisticated (cf. Section 4.3.3), because there can be an arbitrary number of jet objects at each trigger level, and the trigger objects constructed by the feature-extraction algorithms have to be matched between the different trigger levels^a. This matching associates Level 1 RoIs, Level 2 proto-jets and jets reconstructed by the Event Filter, yielding a jet object which spans all three levels. Only RoIs within $|\eta| < 3.2$ are considered. A jet trigger only accepts an event if the same jet object exceeds the thresholds at all required trigger levels. If one trigger jet exceeds the trigger threshold at Level 1, but not at Level 2, and another trigger jet vice versa, this is not sufficient to issue a trigger signal for this event. Different types of matching have been implemented. The matching based on RoI words makes use of a unique identifier assigned to every jet RoIs at Level 1, which is copied to jet objects seeded by this RoI and allows an unambiguous matching. Problems may be encountered if the RoI words are missing in the data-format being used, in which case a matching based on ΔR can be employed instead, or if several jet algorithms are run in parallel in the HLT, so that the same jet appears in several variants in the trigger jet collection. To resolve this, additional flags are necessary which provide information about which jet has been constructed by which algorithm. In the end, there has never been a need for matching

^a For the same reason, the computation of jet trigger efficiencies is technically more involved than the computation of \cancel{E}_T trigger efficiencies, cf. Section 5.3.

of jet trigger objects between Level 2 and Event Filter because in 2010 jet trigger objects reconstructed at Event Filter level does not affect the trigger decision, and in 2011, after the introduction of the full-scan algorithm at Event Filter level, simply no matching is done between Level 2 and Event Filter.

Based on the existing implementation of jet and \cancel{E}_T triggers, the emulation of multi-jet and jet + \cancel{E}_T triggers is straightforward. The jet + \cancel{E}_T triggers are a simple logical AND of the corresponding jet and \cancel{E}_T triggers, although the thresholds may differ from the single object trigger chains, and there is the special type of jet-seeded jet + \cancel{E}_T triggers in 2010, which does not use cuts on \cancel{E}_T at Level 1 and Level 2. The multi-jet triggers which are vital for the multi-jet Supersymmetry searches, but have not been employed for the zero-lepton analysis presented in this thesis, extend on the jet triggers by requiring more than one matched jet at each level satisfying the respective threshold ladder. In addition to the above, on data the trigger emulation always checks that the event carries a valid bunch group code, which roughly speaking means that proton-proton collisions are expected in this bunch crossing and corresponds to the flags BGRP0 and BGRP1 being set at Level 1, a requirement which should be included in the definition of any trigger on non-empty bunch crossings.

A.1.2 Validation

The validation of the trigger emulation is done on both Monte Carlo and data. If the emulation is correct, the main difference between Monte Carlo and real data comes from prescales. The fact that in Monte Carlo all prescales are set to unity at all trigger levels makes a full validation possible, where no false positives (type I errors) and no false negatives (type II errors) should occur. A falsely positive decision means that the trigger emulation yields a positive decision for an event where the actual trigger decision read from the event record is negative. In data, this may either be due to an error in the emulation or, for prescaled triggers, due to prescales. Although the raw trigger decision before applying the prescales is also available in the event record, this cannot be used for validating the full trigger chain for triggers which are prescaled at Level 1 or Level 2, because if the prescales lead to a veto of an otherwise positive trigger decision at one of the lower levels, the subsequent levels are not executed and thus no trigger decision is available for these levels. This ambiguity makes the validation on data less informative. A falsely negative decision means that the trigger emulation yields a negative decision for an event where the actual trigger decision read from the event record is positive. This would always indicate a problem in the emulation.

The validation results of the trigger emulation on Monte Carlo (mc10) are shown in Figure A.1. On the horizontal axis, all tested trigger chains are listed with their official names. They include jet triggers with the full-scan algorithm at Event Filter introduced in 2011 (prefix EF_j) and \cancel{E}_T triggers (prefix EF_xe). Some results for Level 1 and Level 2 are also shown. The vertical axis has four bins, which give the four possible combinations of true and false positive and negative results of the emulated trigger decision. The event sample used for this plot has MC ID 107712 and contains $Z \rightarrow \nu\nu$ events plus additional jets. No discrepancy between the emulation and the actual trigger decision can be seen, all bins with false decisions are empty. (No number is printed for bins with zero entries.)

The plot in Figure A.2 is structured analogously, but this time showing the results of the validation of the trigger emulation on data from 2010. A number of important differences with respect to Figure A.1 can be seen: No Event Filter jet chains are listed in this plot because the Event Filter has not been used to reject events in the jet slice in 2010. It was therefore suf-




Figure A.1: Validation of the trigger emulation in Monte Carlo from 2011 (mc10), using an event sample with $Z \rightarrow \nu\nu$ plus additional jets. The four bins on the vertical axis give the number of false negatives and positives and the number of events for which the emulation and stored trigger decision are both positive or both negative, i. e. the emulation result is correct.

ficient to simulate the jet chains up to Level 2. Several jet + \cancel{E}_T triggers are included in this plot, whereas no jet + \cancel{E}_T triggers are listed in the previous plot for 2011. The validation for the updated jet + \cancel{E}_T triggers has not been implemented for 2011, because they are a simple logical AND of \cancel{E}_T and jet triggers and considered unproblematic. A striking difference are the many false positives in this plot, which are, of course, coming from the prescales as explained above. The false positives only occur for the lowest thresholds within the groups of different types of triggers, because these are the ones which are prescaled first^b. The same check has also been done using the whole data from period 2010 H, showing agreement between emulation and the actual trigger decision, with the usual exceptions of false positives due to prescales, for a large number (about $6 \cdot 10^8$ in total) of tested trigger decisions.

As was said at the beginning of this section, the same code that has been used to produce the validation plots is used for the computation of trigger efficiencies which are presented in this thesis. The emulation of the trigger decision in the computation of the efficiencies has the advantage that the efficiencies can be computed independently of changes in the trigger menu which was run online. Moreover, efficiencies for newly introduced trigger chains can be computed on data which was taken before the new chains were defined to maximize the available statistics, bearing in mind the above mentioned limitations of the emulation. Finally, note that the trigger emulation must never be used for sample triggers, because events that have been rejected online cannot be recovered by emulating the sample trigger, at least not without the risk of introducing a bias from the collection of triggers which constitute the trigger menu: If a trigger used as a sample trigger were emulated, the composition of the resulting event sample would depend on the triggers which have been running in parallel, rather than the sample trigger itself, because only those events have been recorded and can be recovered which have fired another trigger, and therefore they represent only a particular (biased) subsample of all possible events.

A.2 Production of Trigger Rate Plots from COOL

All data needed to produce plots of the trigger rates can be retrieved from the COOL database. In this database, the data is organized in folders and is usually indexed by run and luminosity block. It can be accessed by specifying a range of Interval Of Validity (IOV) timestamps. In general, an IOV timestamp is a 63-bit unsigned integer used as unique identifier, obtained by summing up the run number multiplied by 2^{32} and the event number. Here however, for most folders the IOV timestamp consists of the run and luminosity-block number. Each folder may contain several channels, with each channel having the same structure i. e. the same fields and type of data.

The folder TRIGGER/LVL1/Menu, for example, has four fields: the IOV timestamp of the run, the channel number, the item name and a version tag. Each channel corresponds to one of the 256 possible triggers (trigger items) of the L1 Trigger. This folder is used to map the channel identifier to the name of the L1 trigger item, which is needed for the look-ups in other folders. The rates at Level 1 are computed from the ratio of two trigger counters which

^b Prescales are usually applied at the lowest possible trigger level. It has been verified that, for a given trigger chain which has prescales applied at Level 1 or 2 only, the false positives vanish if that respective trigger from this chain is used as sample trigger, which is the trigger at the highest level of this chain with prescales other than one. Another possibility to avoid false positives would be to not count events in which the execution of a trigger chain was stopped due to prescales, but the trigger decision before prescales is not included in the data format used as input for the validation.

are filled by the CTP. The denominator is the number of LHC turns, $n_{\text{LHC turns}}$, within the given luminosity block from the folder TRIGGER/LVL1/CTPCORELBDATA. This is the maximum number of times a Level-1 trigger could have fired. The numerator is the L1 counter n_{L1}^A for the respective trigger A , which counts how often this L1 trigger actually fires. Three different counters are available for each trigger, counting the trigger signals before and after application of the prescales (but before any dead time or busy veto or the trigger-mask is applied) and the final number of L1 accepts. The L1 rate ν_{L1}^A for a trigger A is thus given by

$$\nu_{\text{L1}}^A = \frac{n_{\text{L1}}^A}{n_{\text{LHC turns}}} \cdot \nu_{\text{orbit}}, \quad (\text{A.1})$$

where $\nu_{\text{orbit}} = 11.245$ kHz is the revolution frequency of the LHC, i. e. the number of times each bunch travels around the LHC ring in one second (cf. Table 3.1). The impact of the reduced center-of-mass energy on this value is calculated in Section A.4.2.

The prescales for Level 1 are taken from the TRIGGER/LVL1/Prescales folder, which holds the prescale values per L1 trigger item. The raw rate of a L1 trigger before prescaling can either be computed by using the corresponding counter before prescale, or by multiplying the counter with L1 accepts after prescale by the prescale value. Using the L1 accept counter will include the dead time in the rate, which will result in a lower value, but for normal data taking this difference has been found to only be of the order of one to two percent.

The HLT menu is needed to obtain the association between the different chains in the computation of the aggregate prescale for L2 and EF chains and for the mapping between trigger name and an integer used as index which in case of the HLT is called the chain counter. It is stored in the folder TRIGGER/HLT/Menu, from which the chain name, the chain counter and the name of the lower chain are the most important information. The lower chain is, for L2 chains and EF chains, the name of the L1 item or L2 chain, which seeds the L2 chain or EF chain, respectively. The prescales for the HLT chains themselves are stored in the folder TRIGGER/HLT/Prescales, which is indexed by the chain counter. The HLT rates are stored in the folders TRIGGER/LUMI/LVL2COUNTERS and EFCOUNTERS. As they are stored as Binary Large Objects (BLOBs) they need to be decoded with the TrigStatsFromCool module from ATHENA, which directly returns the rates and counts.

The instantaneous luminosity per run and luminosity block is stored in two separate folders, one of which contains luminosity information from online sources (TRIGGER/LUMI/LBLESTONL), the other from offline sources (LBLESTOFL). Each folder has several channels, corresponding to different sources or algorithms for the luminosity estimate. For the trigger rate plots always the value from channel 0 is used, which is the default instantaneous luminosity in units of 10^{30} Hz/cm², taken from the currently preferred luminosity algorithm at any given time.

The average number of interactions per bunch crossing $\langle \mu \rangle$ can also be obtained from these two folders. It is computed from the instantaneous luminosity $\mathcal{L}_{\text{inst}}$ according to

$$\langle \mu \rangle = \frac{\mathcal{L}_{\text{inst}} \cdot \sigma_{\text{MB}}}{\nu_{\text{orbit}} \cdot n_{\text{coll}}}, \quad (\text{A.2})$$

where σ_{MB} is the inclusive minimum-bias cross section which is assumed to be $\sigma_{\text{MB}} = 71.5$ mb for the LHC at the current center-of-mass energy of 7 TeV. n_{coll} is the number of colliding bunches in the respective run. In this computation, the value of $\langle \mu \rangle$ is averaged over all bunch crossings, i. e. it assumes that the size and spread of the bunches and thus the luminosity is the same for all bunches.

It should be mentioned that the rates in the COOL database are known to be incorrect in some rare cases, due to errors which may occur when collecting the input for writing the rates to COOL. The trigger rates measured online are not only stored in the COOL database, but also in another archive. Plots of the trigger rates on a daily basis can be accessed using an online tool called WEB TRIGGER RATE PRESENTER [204]. The rates obtained from COOL have been cross-checked against these plots in several random checks and have been found to agree. In any case, in the trigger rate plots created from COOL not single runs are used but whole periods, and thus the impact from these rare errors is not expected to render the results unusable. The trigger rates are available from the COOL database starting with data-taking period 2010 G.

A.3 Implementation of the CL_s and PLR Methods

This section describes technical details of the implementation of the CL_s and the PLR methods, which are used to evaluate the results of the analysis in Section 7. The mathematical background of the methods is presented in Section 7.10.3. For all computations version 5.30/00 of ROOT is used.

A.3.1 CL_s Method

The computation of exclusion limits with the CL_s method makes use of the TLIMIT class, which is part of the ROOT framework [181]. Internally, a Monte Carlo (MC) method is used to perform pseudoexperiments, fluctuating the signal and background expectations within the systematic errors and computing the probability of having that particular outcome.

The prototype of the relevant member function of the TLIMIT class is `TConfidenceLevel* ComputeLimit(TLimitDataSource* data[, options])`. The most important option is the number of MC steps or pseudoexperiments. The TLimitDataSource object holds, for each channel, the hypothesized number of signal and background events as well as the number of events observed in data, plus the relative systematic uncertainties, which are all modeled by Gaussian distributions. Correlations between systematic uncertainties can be specified by assigning the same name to a set of uncertainties, which then in each MC step share a common random number drawn from a standard normal distribution, i. e. with a mean of zero and a variance of one, multiplied by the value of the relative uncertainty. The TConfidenceLevel object stores the results of the MC pseudoexperiments and allows to retrieve several derived quantities, most importantly CL_s , but also CL_{s+b} and the expected values of CL_s and CL_{s+b} under the background-only hypothesis.

The methods of TConfidenceLevel which return the observed CL_s and CL_{s+b} values provide an option called `use_sMC` that is not documented. This option is deactivated (i. e. set to `false`) by default. Nonetheless, in the way the TLimit class is used here, it has been found that the results are only stable if this option is set to `true`. This is also in agreement with comparisons to a private implementation of the CL_s method, which does not account for uncertainties. It is therefore assumed that this option needs to be set to `true`. The methods of TConfidenceLevel which return the expected CL_s and CL_{s+b} values, such as `TConfidenceLevel::GetExpectedCLs_b`, which returns the expected CL_s under the background-only hypothesis, do not have the `use_sMC` option, but accept an optional argument called `sigma`. This option is also not documented, but is assumed to return the expected CL_s

and CL_{s+b} values when letting the expected background fluctuate up or down by a given number of standard deviations according to the distribution of the test statistic. This has also been validated by testing against a private implementation.

A.3.2 PLR Method

The implementation of the profile likelihood method is done using the `ProfileLikelihoodCalculator` class from the ROOSTATS framework [82]. It builds upon ROOFIT and ROOT, and itself relies on MINUIT [81] as minimizer. The `ProfileLikelihoodCalculator` is fed with a `RooDataSet` object holding the observed number of events and a `ModelConfig` object describing the data model in terms of the nuisance parameters, the parameter of interest and the PDF. Again, as the signal is only expected to increase the number of observed events, the signal strength μ as parameter in the likelihood fit is limited by zero from below.

The model PDF is basically a product of a Poisson distribution as shown in Equation (7.22) for the special case of only one search channel, $n = 1$, and several standard normal distributions for the nuisance parameters:

$$L = \text{Pois}(n_{\text{obs}} | \mu f_s(s, \vec{v}) + f_b(b, \vec{v})) \cdot \prod_{u \in \mathcal{N}} \text{Normal}(v_u). \quad (\text{A.3})$$

In this equation, \mathcal{N} is the set of statistical and systematic uncertainties listed below, which correspond to the uncertainties given in Table 7.9. All of these are modeled by a standard normal distribution.

$$\mathcal{N}_c = \{\text{luminosity, JER, JES, pile-up}\} \quad (\text{A.4})$$

$$\mathcal{N}_s = \{\text{statistical (signal), theoretical (signal)}\} \quad (\text{A.5})$$

$$\mathcal{N}_b = \{\text{statistical (background), SUSY (background)}\} \quad (\text{A.6})$$

$$\mathcal{N} = \mathcal{N}_c \cup \mathcal{N}_s \cup \mathcal{N}_b \quad (\text{A.7})$$

Again, ‘‘SUSY’’ stands for the additional uncertainties on the background normalization. The theoretical uncertainties for the signal concern the signal cross section. \mathcal{N}_c contains the correlated uncertainties, the other two sets the uncorrelated uncertainties for signal and background. The expected number of signal and background events is modified by the nuisance parameters according to

$$f_s(s, \vec{v}) = s \cdot \left(1 + \sum_{u \in \mathcal{N}_c} \sigma_u^s v_u + \sum_{u \in \mathcal{N}_s} \sigma_u^s v_u\right) \quad (\text{A.8})$$

$$f_b(b, \vec{v}) = b \cdot \left(1 + \sum_{u \in \mathcal{N}_c} \sigma_u^b v_u + \sum_{u \in \mathcal{N}_b} \sigma_u^b v_u\right) \quad (\text{A.9})$$

Here, σ_u^s and σ_u^b specify the relative uncertainty for signal and background, respectively. By using standard normal distributions, which are then multiplied by the relative uncertainties, it is possible to implement uncertainties that are correlated, but of different size for signal and background. In the above equations, the JES uncertainty is written as a symmetric uncertainty. The actual implementation allows for an asymmetric JES uncertainty, and the addend $\sigma_{\text{JES}}^s v_{\text{JES}}$ in Equations (A.8) needs to be replaced by

$$\sigma_{\text{JES}}^s v_{\text{JES}} \rightarrow \begin{cases} \sigma_{\text{JES, up}}^s v_{\text{JES}} & \text{for } v_{\text{JES}} \geq 0, \\ \sigma_{\text{JES, down}}^s v_{\text{JES}} & \text{otherwise,} \end{cases} \quad (\text{A.10})$$

and analogously for the background in Equation (A.9). This means that an upward fluctuation is assumed to be equally likely as a downward fluctuation, but the size of the uncertainty can be scaled independently.

The `ProfileLikelihoodCalculator::GetInterval` method returns a `LikelihoodInterval` object, from which, for a given confidence level, lower and upper limits on the parameter of interest μ can be queried, which are computed using the approximation of $-2 \ln \lambda(\mu)$ as a χ^2 distribution introduced above. Just opposite to the procedure described above for the CL_s method, where a confidence interval was to be derived from the binary decision whether a parameter value is excluded or not, here, a p -value for the signal-plus-background hypothesis is sought. To compute this p -value, the value of $-2 \ln \lambda(\mu)$ for the nominal signal strength $\mu = 1$ is converted into a p -value using the above approximation. If this value is smaller than α , the signal-plus-background hypothesis is reported to be excluded at a certainty level of $1 - \alpha$. The expected limits from the PLR method under the assumption of the background-only hypothesis \mathcal{H}_B are obtained by setting the number of observed events equal to the background expectation from Monte Carlo. The expected p -value for this hypothesis is then computed in the same way as described above.

A.4 Computation of LHC Parameters

A.4.1 Required Magnetic Field Strength

The following calculation shows that the magnetic field needed to keep the protons on a track with the radius of the LHC tunnel is proportional to the beam energy with a proportionality constant of approximately 1 Tesla/TeV. The calculation is very simple and can be done classically by requiring the Lorentz force \vec{F}_L and centripetal force \vec{F}_C to be equal:

$$\vec{F}_L = q\vec{v} \times \vec{B} \stackrel{!}{=} -m\vec{\omega} \times \vec{r} \times \vec{\omega} = \vec{F}_C. \quad (\text{A.11})$$

\vec{v} is the velocity, $\vec{\omega}$ the vectorial angular velocity, \vec{r} the position vector and q the charge of the protons in a homogeneous magnetic field \vec{B} . In the high-energy approximation with $E \approx |\vec{p}|c$ and assuming \vec{B} to be perpendicular to \vec{v} , it follows that

$$\Rightarrow B = \frac{E}{qrc} \quad (\text{A.12})$$

$$= \frac{E \cdot 1.6 \cdot 10^{-7} \text{ J/TeV}}{2.0 \cdot 10^{-7} \text{ m}^2/\text{s}} \quad (\text{A.13})$$

$$\Rightarrow B[\text{T}] = 0.8 E[\text{TeV}]. \quad (\text{A.14})$$

A.4.2 Impact of the Center-of-Mass Energy on the Revolution Frequency

From Table 3.1, it can be seen that the revolution frequency of the proton bunches around the LHC ring is 11.245 kHz at an energy E of 7 TeV per beam. At the time of writing, the energy is only 3.5 TeV per beam (cf. Section 3.3.2), only half the value of the design energy. As a consequence, the protons are slower and the revolution frequency is lower. The frequency of 11.245 kHz is an input to the computation of the trigger rates (cf. Section A.2), thus the question is what the impact of the lower proton energy on this computation is.

The relativistic factor of the protons at $\sqrt{s} = 14$ TeV is

$$\gamma_{14} = E/m_p = 7461, \quad (\text{A.15})$$

and the relative velocity of the protons with respect to the speed of light c is thus

$$\frac{v_{14}}{c} = \sqrt{1 - 1/\gamma_{14}^2} = 1 - 9 \cdot 10^{-9}. \quad (\text{A.16})$$

For $\sqrt{s} = 7$ TeV, one obtains

$$\gamma_7 = E/m_p = 3730, \quad (\text{A.17})$$

and the relative velocity of the protons with respect to the speed of light is

$$\frac{v_7}{c} = \sqrt{1 - 1/\gamma_7^2} = 1 - 3.6 \cdot 10^{-8}. \quad (\text{A.18})$$

The absolute smallness of the change in the velocity makes clear, already at this step, that reducing the center-of-mass energy to half its design value has no significant impact on the revolution frequency. Using 11.245 kHz is therefore correct within the precision at which the trigger rates are stored and evaluated. The exact value is 11.2455 kHz, given by $C_{\text{LHC}} c/v_7$, where C_{LHC} is the circumference of the LHC ring.

A.5 Analytic Forms of ΣE_T and \cancel{E}_T

In this section, the mathematical proofs for the analytic forms of the probability density functions for ΣE_T and \cancel{E}_T will be derived, which have been used in the model presented in Section 6.

A.5.1 Distribution of ΣE_T for Several Concurrent Interactions

The analytic form for ΣE_T is motivated by the fact that there are μ interactions taking place at the same time, so that their energy depositions are attributed to the same bunch crossing and therefore add up. Each of the interactions is assumed to be independent of the others and contribute an energy according to the same probability density function (PDF). This PDF is assumed to be an exponential distribution with rate parameter λ , defined as

$$f_{\text{exp}}(x|\lambda) := \begin{cases} \lambda \exp(-\lambda x) & \text{for } x \geq 0, \\ 0 & \text{for } x < 0. \end{cases} \quad (\text{A.19})$$

Theorem.

Let S_n be the sum of $n > 0$ independent random variables $X_i, i \in \{1, \dots, n\}$,

$$S_n := \sum_{i=1}^n X_i, \quad (\text{A.20})$$

each of which is distributed according to $f_{X_i}(x|\lambda) = f_{\text{exp}}(x|\lambda)$ with a common value for λ . Then, the PDF f_{S_n} of the random variable S_n is

$$f_{S_n}(x|\lambda) := \begin{cases} \frac{\lambda^n}{(n-1)!} x^{n-1} \exp(-\lambda x) & \text{for } x \geq 0, \\ 0 & \text{for } x < 0. \end{cases} \quad (\text{A.21})$$

Proof. (By Mathematical Induction.)

Initial Step: When $n = 1$, then $S_n = X_1$, and thus by definition

$$f_{S_1}(x|\lambda) = f_{X_1}(x|\lambda) = f_{\text{exp}}(x|\lambda). \quad (\text{A.22})$$

Equation (A.21) thus holds for $n = 1$.

Inductive Step: The inductive assumption is that Equation (A.21) holds for a given $n \geq 1$. It is to be shown that from this it follows that Equation (A.21) also holds for $n + 1$.

In general, the addition rule for two independent random variables U and V with probability density functions f_U and f_V is

$$\begin{aligned} f_W(x) &= f_{U+V}(x) \\ &= (f_U * f_V)(x) := \int_{-\infty}^{\infty} f_U(y) f_V(x-y) \, dy, \end{aligned} \quad (\text{A.23})$$

where f_W is the probability density function of the random variable $W = U + V$. Applying the addition rule (A.23) to

$$S_{n+1} = \left(\sum_{i=1}^n X_i \right) + X_{n+1}, \quad (\text{A.24})$$

it follows that

$$\begin{aligned} f_{S_{n+1}}(x|\lambda) &= \left(f_{S_n} * f_{X_{n+1}} \right) (x|\lambda) \\ &= \int_{-\infty}^{\infty} f_{S_n}(y|\lambda) * f_{\text{exp}}(x-y|\lambda) \, dy \\ &= \int_0^x \frac{\lambda^n}{(n-1)!} y^{n-1} \exp(-\lambda y) \lambda \exp(-\lambda(x-y)) \, dy \\ &= \frac{\lambda^{n+1}}{(n-1)!} \exp(-\lambda x) \int_0^x y^{n-1} \, dy \\ &= \frac{\lambda^{n+1}}{n!} x^n \exp(-\lambda x), \end{aligned} \quad (\text{A.25})$$

where the integration range in the third line is due to the probability density functions f_{S_n} and f_{exp} being zero for negative arguments. □

Under the model assumptions outlined above, the distribution of $\sum E_T$ in events with μ concurrent interactions is therefore given by the probability density function f_{S_μ} , i. e. in the notation of Section 6,

$$P\left(\sum E_T = x|\mu\right) \equiv f_{S_\mu}(x|\lambda). \quad (\text{A.26})$$

A.5.2 Distribution of \mathbb{E}_T from Resolution Effects

In the following, the analytic form for \mathbb{E}_T under the assumption of only fake \mathbb{E}_T from resolution effects will be derived. The assumption of the presence of only this type of fake \mathbb{E}_T translates into the distribution of the x - and y -component of the missing energy in the transverse plane, \mathbb{E}_x and \mathbb{E}_y , following a Gaussian distribution of same width $\sigma_x = \sigma_y = \sigma$. The distribution of \mathbb{E}_T then follows from the geometrical relation $\mathbb{E}_T = \sqrt{(\mathbb{E}_x)^2 + (\mathbb{E}_y)^2}$. As was done above, the probability distribution will be derived generically and then applied to the physical question.

Theorem.

Let X and Y be independent random variables following a Gaussian distribution with mean zero and width σ^2 ,

$$f_X(x|\sigma) = f_Y(x|\sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{x^2}{2\sigma^2}\right). \quad (\text{A.27})$$

The probability density function of the random variable $U = \sqrt{X^2 + Y^2}$ is then given by

$$f_U(x|\sigma) := \begin{cases} \frac{x}{\sigma^2} \exp\left(-\frac{x^2}{2\sigma^2}\right) & \text{for } x \geq 0, \\ 0 & \text{for } x < 0. \end{cases} \quad (\text{A.28})$$

Proof.

The general rule for transforming a bivariate probability density function $f_{X,Y}$ is

$$f_{U,V}(u,v) = \sum_{i=1}^k f_{X,Y}(h_{1i}(u,v), h_{2i}(u,v)) |J_i|, \quad (\text{A.29})$$

where h_{1i} and h_{2i} are the inverse functions of the variable transformation $U = g_1(X,Y)$, $V = g_2(X,Y)$ on the i -th set of a partition A_i of the support of $f_{X,Y}$. The partition is chosen such that the inverse of $g_{1,2}$ on each A_i exists, i. e. $g_1|_{A_i}$ and $g_2|_{A_i}$ are injective functions, mapping A_i onto the support B of $f_{U,V}$. J_i is the Jacobian determinant of the inverse transformation on A_i [4]. The transformation chosen here is

$$(X,Y) \mapsto (U,V) = (\sqrt{X^2 + Y^2}, Y). \quad (\text{A.30})$$

As (X,Y) and $(-X,Y)$ are mapped onto the same point by this transformation, the partition

$$\begin{aligned} A_1 &= \{(X,Y) \in \mathbb{R}^2 | X < 0\} \\ A_2 &= \{(X,Y) \in \mathbb{R}^2 | X \geq 0\} \end{aligned} \quad (\text{A.31})$$

of the support of $f_{X,Y}$ is used. Hence, on A_1 ,

$$h_{11}(u,v) = -\sqrt{u^2 - v^2} \quad (\text{A.32})$$

$$h_{21}(u,v) = v \quad (\text{A.33})$$

$$\Rightarrow J_1 = \begin{vmatrix} \frac{2u}{2\sqrt{u^2 - v^2}} & \frac{-2v}{2\sqrt{u^2 - v^2}} \\ 0 & 1 \end{vmatrix} = \frac{u}{\sqrt{u^2 - v^2}}. \quad (\text{A.34})$$

On A_2 , evidently $h_{12} = -h_{11}$ and $h_{22} = h_{21}$, thus $J_2 = -J_1$. Using Equation (A.27) and the independence of X and Y , so that $f_{X,Y} = f_X f_Y^\varepsilon$, it follows from Equation (A.29) that

$$f_{U,V}(u,v) = \frac{1}{2\pi\sigma^2} \frac{2u}{\sqrt{u^2 - v^2}} \exp\left(-\frac{u^2}{2\sigma^2}\right). \quad (\text{A.35})$$

The desired probability density function of $U = \sqrt{X^2 + Y^2}$ is obtained by marginalization,

$$\begin{aligned} f_U(u) &= \int_{\text{supp } V} f_{U,V}(u,v) \, dv \\ &= \frac{u}{\pi\sigma^2} \exp\left(-\frac{u^2}{2\sigma^2}\right) \int_{-u}^u \frac{1}{\sqrt{u^2 - v^2}} \, dv \\ &= \frac{u}{\sigma^2} \exp\left(-\frac{u^2}{2\sigma^2}\right), \end{aligned} \quad (\text{A.36})$$

where in the second line $v \leq u$ has been used, and f_U is defined on $u \geq 0$. □

Equation (A.28) can be applied directly to state the distribution of \cancel{E}_T in events with only fake \cancel{E}_T , for which, as stated above, the x and y component are assumed, from the central limit theorem, to be given by a Gaussian distribution because they are a sum over a large number of independent measurements of the energy deposited in the cells of the calorimeter. The distribution of \cancel{E}_T in the notation of Section 6 is

$$P(\cancel{E}_T = x|\sigma) \equiv f_U(x|\sigma). \quad (\text{A.37})$$

A.6 Supplementary Plots: \cancel{E}_T Model

In the description of the model for the contribution of detector resolution effects to fake \cancel{E}_T in Section 6, some of the plots are included in the main part of the thesis only for either 2010 or 2011. To confirm that the conclusions from the data from these two data-taking periods are consistent, here, the two corresponding plots for the respective other data-taking period are shown, highlighting similarities and differences where applicable. Moreover, additional plots provide more details.

Figure A.3 shows the distribution of the number of reconstructed primary vertices $n_{\text{pv, reco}}$ for four small event samples with a narrow distribution of $\langle\mu\rangle$. The spread of $\langle\mu\rangle$ is 0.1 here. This is the same plot as in Figure 6.7 but for 2011. The data points are again fitted with a Poisson with free rate parameter μ and the fit result $\langle\mu\rangle_{\text{fit}}$ is given in the legend, together with the ratio of $\langle\mu\rangle_{\text{fit}}$ and the nominal $\langle\mu\rangle$ taken from the COOL database expressed as a percentage value in brackets. The plot mainly differs in the values of $\langle\mu\rangle$ being higher in 2011 than in

^ε Here, it is exploited that the means of $f_X(x|\sigma)$ and $f_Y(y|\sigma)$, x_0 and y_0 , respectively, are zero by assumption, such that in the numerator of the exponential function in Equation (A.35):

$$(x - x_0)^2 + (y - y_0)^2 = x^2 + y^2 = u^2.$$

The reason to assume a common width parameter σ for both f_X and f_Y is obvious.




Figure A.3: Distribution of the number of reconstructed primary vertices $n_{\text{pv, reco}}$ in data from 2011 for four small event samples with a narrow distribution of $\langle \mu \rangle$. The data points are fitted with a Poisson distribution (solid lines) with free rate parameter μ and the fit result $\langle \mu \rangle_{\text{fit}}$ is given in the legend together with the ratio of $\langle \mu \rangle_{\text{fit}}$ over the nominal $\langle \mu \rangle$ expressed as a percentage value.

2010. Another interesting observation is the continuation of the trend that with higher $\langle \mu \rangle$ the ratio of $\langle \mu \rangle_{\text{fit}} / \langle \mu \rangle$ gets smaller, which is an indication for the vertex reconstruction performance becoming worse for busier events, where the probability of accidentally merging vertices becomes higher.

Figure A.4 shows the fit which is used to extract the slope of $\sum E_T$ as input parameter of the model for 2010 data. This plot is analogous to Figure 6.8, which is for 2011, with the difference that for 2010 the $\langle \mu \rangle$ in data in general is smaller, so that here events with $\langle \mu \rangle \approx 2.0$ are used. The data in the plot is taken with a minimum-bias trigger. The fit shown in the plot gives a slope of $\lambda_{\text{fit}} = 0.0405(9)$ for 2010, which is in very good agreement with the value of $\lambda_{\text{fit}} = 0.0410(4)$ obtained from 2011 data. While for 2011 the fit yields a value of $\langle \mu \rangle_{\text{fit}} = 4.96(5)$ for the average number of concurrent interactions, which is compatible with the nominal value $\langle \mu \rangle \approx 5$, for 2010 the fit converges to $\langle \mu \rangle_{\text{fit}} = 1.52(8)$, which is significantly lower than the nominal value $\langle \mu \rangle \approx 2.0$. This may indicate a better description of the data by the model for higher values of $\langle \mu \rangle$, but as well be due to the different trigger selections. The latter possibility cannot be ruled out because the number of events taken with random triggers in 2010 or minimum-bias triggers in 2011 is too small. Figure A.5 confirms that there is a general difference between 2010 and 2011. In this plot, the vertical axis is the value for $\langle \mu \rangle$ obtained from a fit of the $\sum E_T$ spectrum above 50 GeV in the same way as was done in Figure A.4. This is plotted as function of the nominal $\langle \mu \rangle$, which is taken from the COOL database and is computed from the instantaneous luminosity. Every point corresponds to a sample of events with a narrow distribution of $\langle \mu \rangle$ ($|\langle \mu \rangle - \langle \mu \rangle_0| < 0.1$, where $\langle \mu \rangle_0$ is the target value of the nominal $\langle \mu \rangle$ from COOL). The error bars show the size of the uncertainty from the fit of $\sum E_T$. The five points for $\langle \mu \rangle < 4$ are from 2010, the other four from 2011. It can be seen that the correlation of the fit result with the nominal value of $\langle \mu \rangle$ is better in 2011 than in 2010, although for $\langle \mu \rangle \approx 7$ the fit to 2011 data also gives too small a value.




Figure A.4: Comparison of fits to the Event Filter $\sum E_T$ spectrum in events from 2010 data with $\langle \mu \rangle \approx 2.0$. Shown are the distributions in subsamples selected by the number of reconstructed vertices $n_{pv, reco}$ between 1 (red) and 2 through 6 (gray) as well as the full sample (blue). The data points represent the distributions in data, the curves are fits of the same type as explained for Figure 6.8.




Figure A.5: Correlation of the value for $\langle \mu \rangle_{fit}$ obtained from a fit of the full $\sum E_T$ spectrum and the nominal $\langle \mu \rangle$ taken from the COOL database. The error bars indicate the uncertainties from the fit. The dashed line marks the diagonal.




Figure A.6: Rates of one high-threshold \cancel{E}_T trigger, EF_xe60_noMu, in periods 2011 G – K, as function of the average number of interactions per bunch crossing $\langle\mu\rangle$. The runs are grouped and colored by the number of colliding bunches which are given in brackets in the legend.

Periods 2011 I – K have been chosen in Figure 6.18 to display the rates of high threshold \cancel{E}_T triggers because in these periods the number of colliding bunches is approximately the same value (1317 or 1333) for all runs. In Figure A.6, the rates of one of the triggers, EF_xe60_noMu, are shown for a larger dataset with runs from periods 2011 G – K, again coloring groups of runs with different numbers of colliding bunches differently. It can be seen that there are again discontinuities in the rate whenever the number of colliding bunches changes, although the rates are plotted as function of $\langle\mu\rangle$, for which at the beginning of the discussion of the \cancel{E}_T model it has been shown that the low threshold triggers do not exhibit discontinuities (cf. Figure 6.2). In fact, the rates of the high threshold triggers exhibit discontinuities when plotted as function of $\langle\mu\rangle$ as well as when plotted as function of the instantaneous luminosity. This can be explained by the rate being composed of two contributions, one of which depends in a continuous manner on the instantaneous luminosity per bunch, i. e. the event activity expressed by the average number of interactions, and the other continuously on the instantaneous luminosity, independently of the number of colliding bunches. As discussed before, the first part would be attributed to fake \cancel{E}_T from noise as well as from mismeasurements of jets, the second part would be attributed to real \cancel{E}_T from invisible particles produced in the hard scattering process. For the low threshold triggers, the rates are completely dominated by fake \cancel{E}_T , so that the discontinuities due to the small fraction of real \cancel{E}_T events are not visible.

A.7 Supplementary Plots: Trigger Efficiencies in Monte Carlo

This section collects plots of trigger efficiencies that have not been included in the main part of the thesis. They are shown here for completeness.

Figure A.7 is the analog of Figure 5.14, but has as offline reference the \cancel{E}_T variant which is used by the ATLAS Supersymmetry group for the analysis of 2011 data. This variant is

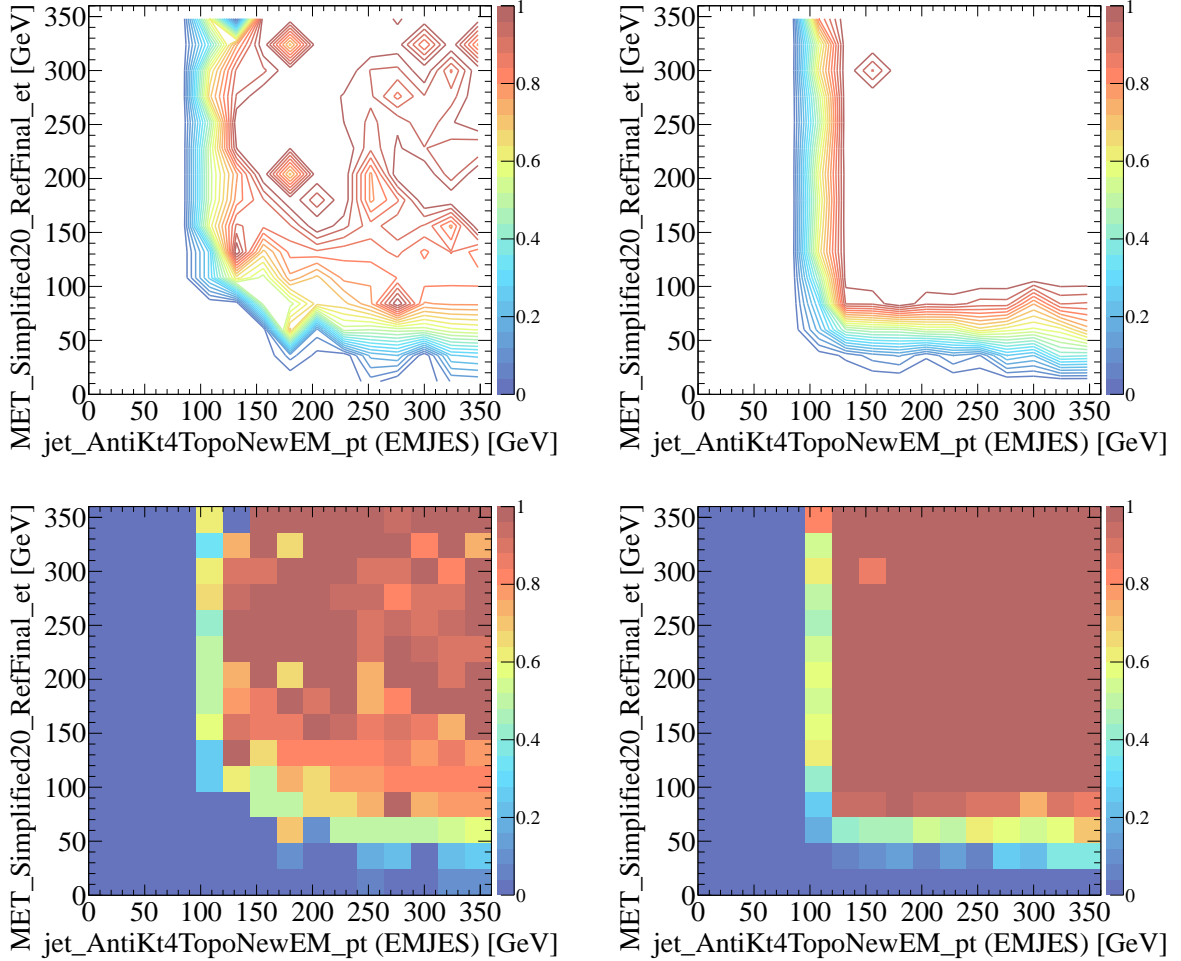


Figure A.7: Two-dimensional efficiency plots of the combined jet + \cancel{E}_T trigger EF_j75_a4tc_EFFS_xe45_loose_noMu in Monte Carlo events, for a QCD sample (left) and for a $Z \rightarrow \nu\nu + \text{jets}$ sample (right). The two upper plots show the efficiencies as contours, which highlights details in the regions where the efficiencies change, in particular in the turn-on regions. The two lower plots use a color-coding of the efficiencies, thereby making it easier to distinguish regions with low and high efficiencies. The SUSY definition of \cancel{E}_T is used as offline reference on the vertical axis (cf. Figure 5.14).




Figure A.8: Projections of the efficiency of the combined jet + \cancel{E}_T trigger EF_j75_a4tc_EFFS_xe45_loose_noMu in Monte Carlo events onto the \cancel{E}_T (left) and jet p_T (right) axis, after cuts on the respective orthogonal variable at 130 GeV, for the combined QCD sample and for the combined $Z \rightarrow \nu\nu + \text{jets}$ sample. The SUSY definition of \cancel{E}_T is used as offline reference here (cf. Figure 5.15).

called MET_Simplified20_RefFinal (cf. Section 4.4.3). In the same way, Figures A.8 and A.9 are the analogs of Figures 5.15 and 5.16 with the SUSY definition of missing transverse energy instead of the LocHadTopo variant.

For the $Z \rightarrow \nu\nu + \text{jets}$ sample, the turn-on curves for the two offline references look very similar. In particular, the projection onto \cancel{E}_T shows some small differences for the QCD samples. The turn-ons for J5 and J6 are sharper, but do not reach a plateau efficiency of 100%. The combination in Figure A.8 shows a lower efficiency of the trigger at the beginning of the plateau when plotted as function of the SUSY \cancel{E}_T definition.

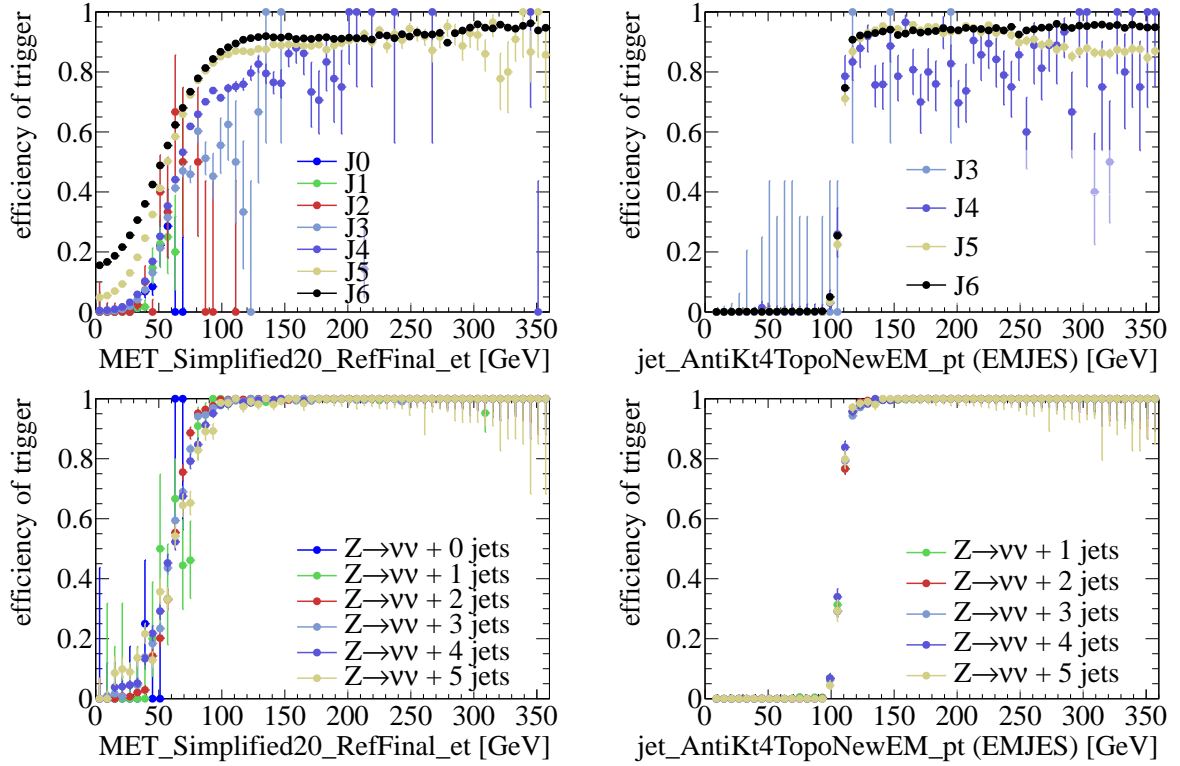


Figure A.9: Projections of the efficiency of the combined jet + \cancel{E}_T trigger EF_j75_a4tc_EFFS_xe45_loose_noMu in Monte Carlo events for each subsample of the QCD sample (top) and the $Z \rightarrow \nu\nu + \text{jets}$ sample (bottom) individually. Left: onto offline \cancel{E}_T , right: onto jet p_T , after cuts on the respective orthogonal variable at 130 GeV. The SUSY definition of \cancel{E}_T is used as offline reference here (cf. Figure 5.16).

A.8 Overview of Monte Carlo Datasets

Tables A.1 and A.2 list all Monte Carlo samples that were used in the analysis in Section 7 to model the backgrounds from Standard Model processes. The symbolic description of the processes follows the common convention that in case that no charge state, generation or flavor is given explicitly, all possibilities are understood to be included which are in agreement with the conservation laws from the Standard Model. In addition to the process name, the subprocess description and the ID of the Monte Carlo sample, the tables include the final cross section σ used in the analysis for the normalization of the Monte Carlo to the data, the cross section σ_{AMI} stored in the AMI database, and the MC generator efficiency ϵ_{gen} . The factor k_{eff} is the ratio of the final cross section after application of all normalization factors and the cross section from the AMI database, given by $k_{\text{eff}} = \sigma / (\sigma_{\text{AMI}} \cdot \epsilon_{\text{gen}})$. N_{sample} is the number of simulated events in the subsample. To get a feeling whether the number of events in Monte Carlo is sufficient, the integrated luminosity \mathcal{L}_{MC} that corresponds to the size of the samples is given, where $\mathcal{L}_{\text{MC}} = N_{\text{sample}} / \sigma$. Only for the dijet samples with extremely large cross sections, the sample luminosity \mathcal{L}_{MC} is smaller than the integrated luminosity $\mathcal{L}_{\text{int}}^{2010} = 33.4 \text{ pb}^{-1}$ (cf. Section 7.2) of the dataset to be evaluated. The dataset tags of the MC samples are also given (cf. Section A.10.4).

All top quark samples in Table A.2 are generated with MC@NLO. A characteristic feature of MC@NLO is the appearance of negative event weights arising from the subtraction method [75]. This is a consequence of the way the events are generated in the integration over phase space. For the MC@NLO samples, the unweighted total event count is given. Due to the negative-weight events, the value of \mathcal{L}_{MC} given in the Table A.2 is smaller than the one obtained from the unweighted total event count. No errors are available on the cross sections given in Tables A.1 and A.2. It is therefore difficult to estimate the number of significant digits, but they will probably be less than the number of digits included in the tables. Nevertheless, it is common to present these numbers cutting them at the same position, rather than introducing some arbitrary rounding, so this habit is adopted here. The number of events in the table is exact.

Process	Subprocess	MC set ID	σ [pb]	σ_{AMI} [pb]	ϵ_{gen}	k_{eff}	N_{sample}	\mathcal{L}_{MC}	Dataset tag
$Z \rightarrow e^+e^- + \text{jets}$	$N_p = 0$	107650	830.13	664.10	1.000	1.25	304 216	366	e529_s765_s767_r1302_r1306_p305
	$N_p = 1$	107651	166.24	132.99	1.000	1.25	63 440	382	e529_s765_s767_r1302_r1306_p305
	$N_p = 2$	107652	50.28	40.23	1.000	1.25	19 497	388	e529_s765_s767_r1302_r1306_p305
	$N_p = 3$	107653	13.92	11.14	1.000	1.25	5 499	395	e529_s765_s767_r1302_r1306_p305
	$N_p = 4$	107654	3.62	2.89	1.000	1.25	1 499	414	e529_s765_s767_r1302_r1306_p305
	$N_p \geq 5$	107655	0.94	0.75	1.000	1.25	500	532	e529_s765_s767_r1302_r1306_p305
$Z \rightarrow \mu^+\mu^- + \text{jets}$	$N_p = 0$	107660	830.13	663.79	1.000	1.25	303 947	366	e529_s765_s767_r1302_r1306_p305
	$N_p = 1$	107661	166.24	132.95	1.000	1.25	62 996	379	e529_s765_s767_r1302_r1306_p305
	$N_p = 2$	107662	50.28	40.38	1.000	1.25	18 993	378	e529_s765_s767_r1302_r1306_p305
	$N_p = 3$	107663	13.92	11.16	1.000	1.25	5 497	395	e529_s765_s767_r1302_r1306_p305
	$N_p = 4$	107664	3.62	2.90	1.000	1.25	1 499	414	e529_s765_s767_r1302_r1306_p305
	$N_p \geq 5$	107665	0.94	0.76	1.000	1.24	499	531	e529_s765_s767_r1302_r1306_p305
$Z \rightarrow \tau^+\tau^- + \text{jets}$	$N_p = 0$	107670	830.13	662.50	1.000	1.25	303 359	365	e529_s765_s767_r1302_r1306_p305
	$N_p = 1$	107671	166.24	133.94	1.000	1.24	63 481	382	e529_s765_s767_r1302_r1306_p305
	$N_p = 2$	107672	50.28	40.30	1.000	1.25	19 492	388	e529_s765_s767_r1302_r1306_p305
	$N_p = 3$	107673	13.92	11.03	1.000	1.26	5 497	395	e529_s765_s767_r1302_r1306_p305
	$N_p = 4$	107674	3.62	2.80	1.000	1.29	1 499	414	e529_s765_s767_r1302_r1306_p305
	$N_p \geq 5$	107675	0.94	0.78	1.000	1.20	499	531	e529_s765_s767_r1302_r1306_p305
$Z \rightarrow \nu\bar{\nu} + \text{jets}$	$N_p = 0$	107710	35.96	3538.80	0.008	1.32	2 999	83	e530_s765_s767_r1302_r1306_p305
	$N_p = 1$	107711	549.69	731.35	0.607	1.24	44 487	81	e530_s765_s767_r1302_r1306_p305
	$N_p = 2$	107712	247.42	222.50	0.878	1.27	39 491	160	e530_s765_s767_r1302_r1306_p305
	$N_p = 3$	107713	74.71	62.00	0.966	1.25	11 995	161	e530_s765_s767_r1302_r1306_p305
	$N_p = 4$	107714	18.58	15.87	0.990	1.18	7 993	430	e530_s765_s767_r1302_r1306_p305
	$N_p \geq 5$	107715	6.02	4.38	0.999	1.38	2 500	415	e530_s765_s767_r1302_r1306_p305

Process	Subprocess	MC set ID	σ [pb]	σ_{AMI} [pb]	ϵ_{gen}	k_{eff}	N_{sample}	\mathcal{L}_{MC}	Dataset tag
$W \rightarrow e\nu_e + \text{jets}$	$N_p = 0$	107680	8288.15	6870.50	1.000	1.21	1 381 931	167	e511_s765_s767_r1302_r1306_p305
	$N_p = 1$	107681	1550.14	1293.80	1.000	1.20	258 408	167	e511_s765_s767_r1302_r1306_p305
	$N_p = 2$	107682	452.09	376.60	1.000	1.20	188 896	418	e511_s765_s767_r1302_r1306_p305
	$N_p = 3$	107683	120.97	101.29	1.000	1.19	50 477	417	e511_s765_s767_r1302_r1306_p305
	$N_p = 4$	107684	30.33	25.25	1.000	1.20	12 991	428	e511_s765_s767_r1302_r1306_p305
	$N_p \geq 5$	107685	8.27	7.12	1.000	1.16	3 449	417	e511_s765_s767_r1302_r1306_p305
$W \rightarrow \mu\nu_\mu + \text{jets}$	$N_p = 0$	107690	8288.15	6871.10	1.000	1.21	1 386 038	167	e511_s765_s767_r1302_r1306_p305
	$N_p = 1$	107691	1550.14	1294.70	1.000	1.20	255 909	165	e511_s765_s767_r1302_r1306_p305
	$N_p = 2$	107692	452.09	376.08	1.000	1.20	187 860	416	e511_s765_s767_r1302_r1306_p305
	$N_p = 3$	107693	120.97	100.72	1.000	1.20	50 887	421	e511_s765_s767_r1302_r1306_p305
	$N_p = 4$	107694	30.33	25.99	1.000	1.17	12 991	428	e511_s765_s767_r1302_r1306_p305
	$N_p \geq 5$	107695	8.27	7.13	1.000	1.16	3 498	423	e511_s765_s767_r1302_r1306_p305
$W \rightarrow \tau\nu_\tau + \text{jets}$	$N_p = 0$	107700	8288.15	6873.30	1.000	1.21	1 365 491	165	e511_s765_s767_r1302_r1306_p305
	$N_p = 1$	107701	1550.14	1295.40	1.000	1.20	254 753	164	e511_s765_s767_r1302_r1306_p305
	$N_p = 2$	107702	452.09	375.07	1.000	1.21	188 446	417	e511_s765_s767_r1302_r1306_p305
	$N_p = 3$	107703	120.97	101.77	1.000	1.19	50 472	417	e511_s765_s767_r1302_r1306_p305
	$N_p = 4$	107704	30.33	25.76	1.000	1.18	12 996	428	e511_s765_s767_r1302_r1306_p305
	$N_p \geq 5$	107705	8.27	7.00	1.000	1.18	3 998	483	e511_s765_s767_r1302_r1306_p305
$W \rightarrow b\bar{b} + \text{jets}$	$N_p = 0$	106280	3.90	3.31	1.000	1.18	6 499	1 665	e524_s765_s767_r1302_r1306_p305
	$N_p = 1$	106281	3.17	2.68	1.000	1.19	5 500	1 734	e524_s765_s767_r1302_r1306_p305
	$N_p = 2$	106282	1.71	1.38	1.000	1.24	2 997	1 755	e524_s765_s767_r1302_r1306_p305
	$N_p \geq 3$	106283	0.73	0.66	1.000	1.11	1 500	2 049	e524_s765_s767_r1302_r1306_p305

Table A.1: Detailed list of the Monte Carlo samples for the $W + \text{jets}$ and $Z + \text{jets}$ background.

Process	Subprocess	MC set ID	σ [pb]	σ_{AMI} [pb]	ϵ_{gen}	k_{eff}	N_{sample}	\mathcal{L}_{MC}	Dataset tag
$t\bar{t}$	semi- / fully leptonic	105200	89.40	144.12	0.556	1.12	999 387	8 648	e510_s765_s767_r1302_r1306_p305
	fully hadronic	105204	71.39	144.28	0.444	1.11	149 899	1 626	e540_s765_s767_r1302_r1306_p305
Single top	t -channel $W \rightarrow e\nu_e$	108340	7.15	7.15	1.000	1.00	9 993	831	e508_s765_s767_r1302_r1306_p305
	t -channel $W \rightarrow \mu\nu_\mu$	108341	7.15	7.18	1.000	1.00	9 997	837	e508_s765_s767_r1302_r1306_p305
	t -channel $W \rightarrow \tau\nu_\tau$	108342	7.15	7.13	1.000	1.00	10 000	829	e508_s765_s767_r1302_r1306_p305
	s -channel $W \rightarrow e\nu_e$	108343	0.47	0.47	1.000	1.00	9 950	17 962	e534_s765_s767_r1302_r1306_p305
	s -channel $W \rightarrow \mu\nu_\mu$	108344	0.47	0.47	1.000	1.00	9 996	17 979	e534_s765_s767_r1302_r1306_p305
	s -channel $W \rightarrow \tau\nu_\tau$	108345	0.47	0.47	1.000	1.00	9 996	18 081	e534_s765_s767_r1302_r1306_p305
	Wt channel, all W decays	108346	14.58	14.58	1.000	1.00	14 995	911	e508_s765_s767_r1302_r1306_p305
Pythia QCD dijets	J0	105009	$1.262 \cdot 10^{10}$	$9.857 \cdot 10^9$	1.000	1.28	1 379 184	<1	e468_s766_s767_r1303_r1306_p305
	J1	105010	$8.679 \cdot 10^8$	$6.781 \cdot 10^8$	1.000	1.28	1 395 383	<1	e468_s766_s767_r1303_r1306_p305
	J2	105011	$5.247 \cdot 10^7$	$4.099 \cdot 10^7$	1.000	1.28	1 398 078	<1	e468_s766_s767_r1303_r1306_p305
	J3	105012	$2.808 \cdot 10^6$	$2.194 \cdot 10^6$	1.000	1.28	1 397 430	<1	e468_s766_s767_r1303_r1306_p305
	J4	105013	112 261	87 704	1.000	1.28	1 397 401	12	e468_s766_s767_r1303_r1306_p305
	J5	105014	3 008	2 350	1.000	1.28	1 391 612	463	e468_s766_s767_r1303_r1306_p305
	J6	105015	43.0	33.6	1.000	1.28	1 347 654	31 316	e468_s766_s767_r1303_r1306_p305
J7	105016	0.18	0.14	1.000	1.28	1 125 428	6 280 290	e468_s766_s767_r1303_r1306_p305	

Table A.2: Detailed list of the Monte Carlo samples for the top quark and QCD background. For the MC@NLO samples, i. e. all top quark samples in this table, the unweighted total event count is given.

	-2000006	-2000005	-2000004	-2000003	-2000002	-2000001	-1000037	-1000024	-1000006	-1000005	-1000004	-1000003	-1000002	-1000001	1000001	1000002	1000003	1000004	1000005	1000006	1000021	1000022	1000023	1000024	1000025	1000035	1000037	2000001	2000002	2000003	2000004	2000005	2000006	\tilde{t}_R	
-2000006																																			\tilde{b}_R
-2000005																																			\tilde{c}_R
-2000004		2.0	2.7	2.9	2.9						1.3	1.7	1.7	1.8	3.5	3.9	2.7	2.4			3.4	3.4							2.5	3.1	1.6	5.2	\tilde{s}_R		
-2000003		2.7	2.6	3.2	3.2						1.6	2.0	2.1	2.1	3.8	4.2	3.0	2.7			3.8	3.1							2.8	3.3	5.2	1.6	\tilde{u}_R		
-2000002		2.9	3.2	2.9	3.3						1.8	2.1	2.3	2.3	3.9	4.3	3.2	2.9			3.9	3.8		0.0					3.0	5.2	2.1	2.0	\tilde{d}_R		
-2000001		2.9	3.2	3.3	3.0						1.8	2.2	2.3	2.3	4.0	4.4	3.2	2.8			3.9	3.3							5.2	3.5	2.2	1.9	$\tilde{\chi}_2^-$		
-1000037																																		$\tilde{\chi}_1^-$	
-1000024																0.5																		\tilde{t}_1	
-1000006																																			\tilde{b}_1
-1000005																																			\tilde{c}_L
-1000004		1.3	1.6	1.8	1.8						2.0	2.7	2.8	2.9	2.5	2.9	2.8	5.2			3.4	2.2							3.5	3.9	2.7	2.4	\tilde{s}_L		
-1000003		1.7	2.0	2.1	2.2						2.7	2.6	3.2	3.2	2.8	3.2	5.2	3.2			3.8	2.5							3.8	4.2	3.0	2.7	\tilde{u}_L		
-1000002		1.7	2.1	2.3	2.3						2.8	3.2	2.9	3.3	3.0	5.2	2.1	1.8			3.9	2.6							3.9	4.3	3.1	2.8	\tilde{d}_L		
-1000001		1.8	2.1	2.3	2.3						2.9	3.2	3.3	2.9	5.2	3.4	2.1	1.7			3.9	2.7							4.0	4.4	3.2	2.9	\tilde{d}_L		
1000001		3.5	3.8	3.9	4.0						2.5	2.8	3.0	5.2	4.2	5.1	3.8	3.6			4.7	3.2		0.3					3.6	4.1	2.8	2.5	\tilde{d}_L		
1000002		3.9	4.2	4.3	4.4			0.5			2.9	3.2	5.2	3.4	5.1	5.0	4.2	3.9			5.1	3.6			0.3				4.0	4.5	3.2	2.9	u_L		
1000003		2.7	3.0	3.2	3.2						2.8	5.2	2.1	2.1	3.8	4.2	2.6	2.8			3.8	2.5							2.8	3.2	1.9	1.7	s_L		
1000004		2.4	2.7	2.9	2.8						5.2	3.2	1.8	1.7	3.6	3.9	2.8	2.0			3.4	2.1							2.5	2.9	1.6	1.4	c_L		
1000005																																			b_1
1000006																																			t_1
1000021		3.4	3.8	3.9	3.9						3.4	3.8	3.9	3.9	4.7	5.1	3.8	3.4			6.3	4.3							4.7	5.1	3.8	3.4	g		
1000022		3.4	3.1	3.8	3.3						2.2	2.5	2.6	2.7	3.2	3.6	2.5	2.1			4.3								3.8	4.8	3.1	3.4	χ_1^0		
1000023																																			χ_2^0
1000024															0.3																				χ_1^+
1000025					0.0											0.3																			χ_3^0
1000035																																			χ_4^0
1000037																																			χ_2^+
2000001		2.5	2.8	3.0	5.2						3.5	3.8	3.9	4.0	3.6	4.0	2.8	2.5			4.7	3.8							4.1	4.9	3.8	3.5	\tilde{d}_R		
2000002		3.1	3.3	5.2	3.5						3.9	4.2	4.3	4.4	4.1	4.5	3.2	2.9			5.1	4.8				0.7			4.9	5.0	4.2	3.9	u_R		
2000003		1.6	5.2	2.1	2.2						2.7	3.0	3.1	3.2	2.8	3.2	1.9	1.6			3.8	3.1							3.8	4.2	2.6	2.7	s_R		
2000004		5.2	1.6	2.0	1.9						2.4	2.7	2.8	2.9	2.5	2.9	1.7	1.4			3.4	3.4							3.5	3.9	2.7	2.0	c_R		
2000005																																			b_R
2000006																																			t_R

Table A.3: Overview of all pairs of supersymmetric particles found in the final states in the Monte Carlo samples comprising the MSSM grid. The numbers give the decadic logarithm of the number of events in which the two supersymmetric particles, given by the respective row and column, have been found. The tilde commonly used to mark supersymmetric partners of Standard Model particles has been suppressed here. All final states fall into one of the classes defined in Table 7.8.

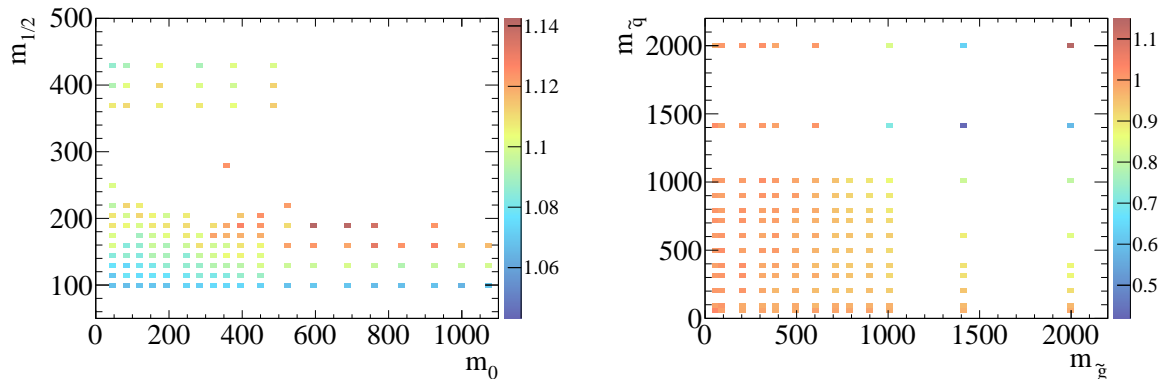


Figure A.10: Ratio of the leading-order cross section for all identified final states over the generator cross section. Left: mSUGRA grid, right: MSSM grid.

A.8.1 Overview of Final States in the Monte Carlo Signal Grids

Tables A.3 and A.4 give an overview in which the frequency is listed at which given pairs of supersymmetric particles are found in the final states from all the Monte Carlo samples comprising the two signal grids discussed in Section 7.6.3. The rows and columns are labelled both with the particle numbers from the Monte Carlo particle numbering scheme of the Particle Data Group [87] and the usual particle symbols. Note that the scalar quarks, apart from the scalar top and bottom quark, are written using indices indicating the chirality. In the particle numbering scheme from the Particle Data Group, the particle numbers ± 1000005 , ± 2000005 , ± 1000006 and ± 2000006 stand for the mixed states $\tilde{b}_{1,2}$ and $\tilde{t}_{1,2}$, respectively, the lighter mixed state having the smaller number.

From Table A.3, it can be seen that almost all final states fall into the classes with IDs 1 through 4 ($\tilde{q}\tilde{g}$, $\tilde{g}\tilde{g}$, $\tilde{q}\tilde{q}$ and $\tilde{q}\tilde{q}$), for which the cross sections are known. Thus, only a negligibly small number of events has to be discarded due to missing cross section information. This is different for the mSUGRA grid, for which the final state counts are shown in Table A.4. Here, for a larger number of final states cross sections have been computed, but still there are combinations which do not fall into any of the classes defined in Table 7.8. They are marked in the table with a negative sign. The impact of these states is discussed in Sections 7.6.3 and A.8.2.

A.8.2 Impact of Unidentified Events

The approach in Section 7.6 does not take the cross section into account, and it is conceivable that if the ignored subprocesses have a comparably large cross section, the fraction of neglected processes is much larger in terms of cross section than in terms of event counts. Therefore, as a second cross-check one can try to estimate the cross section of the neglected processes. The NLO cross section of the full Monte Carlo sample cannot be used for this, because if it were known, the problem would not have arisen in the beginning. But a computation of the cross section of the full sample is available at LO from the MC generator HERWIG++ via the AMI. Of course, the LO generator cross section also only includes the processes that have been included in the event generation.

Figure A.10 shows the ratio of the LO cross section for all identified final states over the




Figure A.11: Distribution of jet p_T in the subsample J3 containing PYTHIA QCD Monte Carlo. The solid histogram line gives the average of the first and second leading jet, the dashed histogram line gives the third leading jet (if present). The two vertical dashed lines mark the lower and upper bound of the \hat{p}_T interval that is covered by J3.

generator cross section. The LO cross section for all identified final states is computed using the k -factor between the HERWIG++ LO cross section and PROSPINO NLO cross section for the MSSM grid and the PROSPINO LO-NLO k -factor for the mSUGRA grid, for which the k -factor with respect to HERWIG++ is not provided in the cross section file.

The gap in the plot for the mSUGRA grid is due to the fact that for these samples the generator cross section for some reason has not been stored in the AMI database. This plot shows only values above 1, which may be due to the fact that here the PROSPINO LO-NLO k -factor are used to compute the LO cross section from the NLO PROSPINO cross section so that the LO cross sections of two different generators, HERWIG++ and PROSPINO, are compared. One could try to renormalize the ratio to restore a mean of one, but being meant only as a second cross-check of the magnitude of an effect that in any case cannot lead to wrong results, the spread of the values shall suffice to convey the basic message of the plot, i. e. that there are no outliers which would indicate that a large fraction of the cross section were missing. In the right plot, which is for the MSSM grid, the normalization seems to be more reasonable, and the plot shows in general the expected behavior. The ratio decreases from the left to the right and exhibits the same tendency as the corresponding plot in Figure 7.8, although the magnitude of the effect appears to be larger. In conclusion, the plots have to be interpreted with some care, and their informative value is unfortunately not too high, but at least they do not hint that the cross sections of the unidentified states comprise an surprisingly large fraction.

A.9 Jet Spectrum in Pythia QCD Samples

Figure A.11 shows the distribution of p_T for the leading jets in one of the subsamples of PYTHIA QCD Monte Carlo. The subsample chosen is J3 so that according to Table 7.7, it

includes events with $70 \text{ GeV} < \hat{p}_T < 140 \text{ GeV}$. What this means in practice can be seen from the plot, together with some additional interesting observations. The solid line in the plot is the average of the p_T of the first and second leading jet from all events of this sample. The p_T is reconstructed offline at EM+JES scale (the `AntiKt4TopoEMJES` collection is used, cf. Section 4.4.2) and thus smeared by detector effects, some of which will lead to an underestimation of the momentum, others to an overestimation. If parts of the jet go into an uninstrumented region or are otherwise out of acceptance of the detector the corresponding part of the energy will be lost and lead to a decrease of the visible momentum. The same tendency is caused by invisible particles such as neutrinos created within the jet that carry away momentum but cannot be detected and particles which are created at such an angle that they lie outside the radius of the cone which is used for the reconstruction. On the other hand, overlaps with other jets or particles which happen to fall into the same region of the calorimeter that is summed up to give the energy of the jet will add energy to the jet on top of the energy of the original parton, and in addition to this, there are statistical fluctuations from the energy measurements in the calorimeter. Another important contribution will come from initial- and final-state radiation that smears out the effects of the original kinematics selection already at parton level [97]. All of these effects will lead to a broadening of the jet resolution, but the overall normalization of the scale is corrected by the EM+JES factor, so that the averaged peak correctly appears at 70 GeV , marked by the left vertical line as the lower bound of the \hat{p}_T interval. The sorting of the reconstructed jets will lead to a systematic bias which pushes the peak for the leading jets to higher and the peak for the second leading jets to lower values, therefore the average of the two leading jets is used in Figure A.11. p_T values outside $70 \text{ GeV} < p_T < 140 \text{ GeV}$ are due to mismeasurements and the limited resolution. This can also be seen from the fact that within this interval, the slope is roughly constant, but for higher p_T , the distribution falls off much faster. The dashed line shows that the third leading jet is already much softer, peaking close to zero. There is a cut-off at low p_T values, because only jets reconstructed with a $p_T > 7 \text{ GeV}$ are written into the jet collections.

A.10 Computing

This section documents technical aspects of the software which has been employed in the context of this thesis, in particular the ROOT framework, which is widely used in high-energy physics, and ATHENA, the framework developed and used within the ATLAS collaboration.

A.10.1 ROOT

One of the most widely used general frameworks in high-energy physics is the ROOT framework, which was developed starting in 1994 in the context of the NA49 experiment at CERN. It is designed to be able to handle the large amounts of data [101, 181] which are produced in modern collider experiments. ROOT is written in C++ and relies on the CINT interpreter [5] to provide a command-line interface and script processor that accepts C++-style commands. In addition to this interactive environment, the core features of ROOT include a proprietary file format that allows to store and mine large amounts of data using vertical data storage techniques (see below), a graphical user-interface for the visualization of data mainly in terms of histograms and projections, and libraries that provide a wide

spectrum of methods from basic mathematical functions to advanced tools for statistical analyses. Coming from the physics community, its main area of application still lies therein.

The native ROOT file format is a compressed binary format which is machine-independent and optimized for quick data access in large files. ROOT files contain both the data and a description of the type of data. An expression that has been coined in this context is to *persistify* data. In object-oriented programming languages like C++, complex data is represented as objects which are instantiated from classes. These objects cannot only store data, but also provide methods to process and access the data and reflect complex relations between data and objects. The problem is that C++ does not have built-in means to save these objects from memory to file. Of course, it is possible to store primitive data types and to dump the binary representation of an object as it is stored in memory, but all references that are stored as pointers, and therefore depend on the absolute position of the object in memory, are then rendered useless. This binary representation is also highly dependent on the platform, the compiler and revision of the code and cannot be interpreted without access to the underlying C++ code in which the object is defined. Objects therefore normally only exist in a transient state, as long as the process that created them is running, but of course the objects created as results of the processing of physical data need to be saved to be able to distribute them and access them again later.

This is made possible by the reflection capabilities^d of ROOT and the underlying CINT interpreter, respectively. They allow to create *dictionaries* describing the data, from which C++ classes can automatically be produced which can be used to write transient objects to a file and read them back again. Backwards-compatibility is a feature that is very important in experiments which are supposed to collect data over decades, and is ensured by ROOT's support of schema evolution. Another feature of the ROOT file format is that it allows to organize data in a nested folder structure within the ROOT file itself.

The data that needs to be stored in collider experiments can neatly be split into independent *events*, each of which corresponds to the collection of all data attributed to particles that have been produced in the same bunch-crossing. In practice, the association of reconstructed^e objects to a certain bunch crossing is not trivial at all, due to the high frequency of collisions within the detector. Particles which are produced in the collision are usually highly relativistic, but even in the short time a particle with a velocity near the speed of light needs to cross the ATLAS detector volume, several more interactions will have taken place because of the targeted collision frequency of 40 MHz. This means that at any point in time, the remnants of collisions from several subsequent bunch-crossings will leave traces in the different detector systems concurrently. Every collision event is described by the reconstructed objects that have been associated with it and by event-wise quantities such as the missing transverse energy, status flags for detector components etc. The event record is the collection of all recorded data which belong to a single event.

Instead of writing objects into ROOT files directly, these are normally organized in TTrees. A TTree is a collection of branches, each of which can be thought of as a vector of objects of a given type. All entries in all of these vectors with the same index are understood to belong to the same event. The data of each event is therefore distributed over all vectors. Of course,

^d In computer science, reflection means that a program can access and modify its own structure, in order to be able, for example, to retrieve information about the type of an object or the relative position of data members at run-time. The reflection database of ROOT is Reflex which provides a type description layer on top of C++ [201].

^e The reconstruction algorithms are described in Section 4.4.

the number of reconstructed particles of a certain type may be different for every event, and therefore the entries of these vectors are not objects, but vectors of different sizes themselves. The advantage of storing data in this vertical structure is that branches can be switched off selectively and the processing event by event can be considerably sped up by this.

A.10.2 The Athena Framework

Every experiment that produces raw data in such amounts that they can only be handled with computers needs a software framework for the evaluation of the data. ROOT is a general framework that is independent of the experiment. In addition to and often on top of ROOT, every larger experiment in high-energy physics has its own proprietary framework that is adapted to its specific requirements.

In ATLAS, the main software framework is called ATHENA^f. ATHENA is an enhanced version of the Gaudi framework [52], originally developed by the LHCb experiment [127] with the envisioned long lifetime of the software and maintainability in mind. Being the main software of an active experiment, ATHENA is under constant development. Roughly once a year a new major revision is released. The releases relevant for this thesis are number 15 and 16 [182]. Most of ATHENA is coded in C++, and the configuration is done through Python scripts, which allow changes in configuration parameters without recompiling. This could also be done using text files, but scripts offer much more flexibility.

ATHENA is written as a general framework for the wide range of applications in high-energy physics. It is used to run detector simulations and create the Monte Carlo in the central production system, it is used in the online system to steer the detector, in particular in the data-taking system and HLT, it incorporates the event data model and the reconstruction algorithms used to evaluate the raw detector output, is used for the reprocessing of data on the basis of improved calibrations, and is also intended to be suitable for end-user analysis.

An overview of the structure and the major components of the ATHENA framework is given in Figure A.12, and used here to exemplify the purpose of the central parts an event-based analysis framework in general consists of. The application manager is the central algorithm that manages and coordinates the execution of all other components. The second important steering unit is the event loop manager, which coordinates the processing of all events present in the input data one after the other. All kinds of debug and log messages are fed into the message service, which, depending on its configuration, decides which messages to present to the user. Every event is passed through a number of algorithms, which retrieve their configuration from the job options service and access the data within the event through the event data service which is called `StoreGate` in ATHENA. Additional event-independent information about the detector (its geometry, alignment or calibrations) and static data such as particle properties are provided by the detector data service, particle property and similar services. Although ATHENA is not framework used for the analysis presented in this analysis, but instead the much lighter framework SFRAME [199] is used, many of these concepts can also be found there, too, and can be generalized to most event-based analyses.

Special methods are implemented for classes which are defined in the ATHENA framework, allowing to persistify objects that were instantiated from these class through the Persistence Service (cf. Figure A.12). These persistifiers can alleviate the need to use exactly

^f Not to be confused with one of the experiments at the Antiproton Decelerator (AD) at CERN bearing the same name, which successfully produced and detected cold antihydrogen atoms [20].

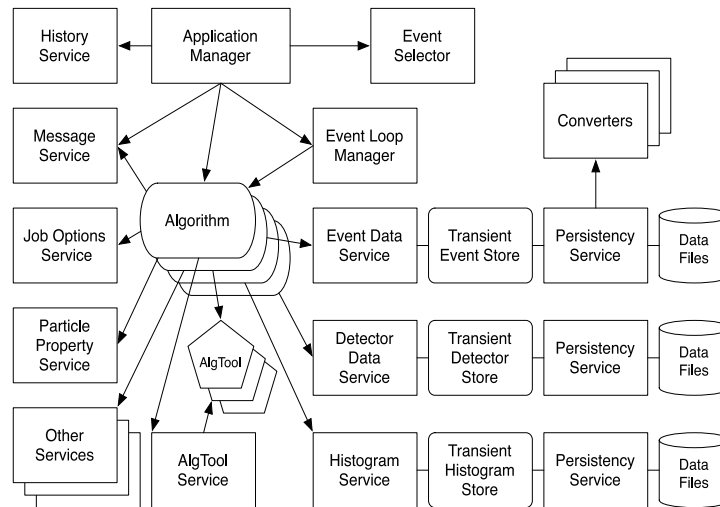


Figure A.12: Overview of the major components of the ATHENA framework and their relationships in terms of navigability and usage.

the same code revision the data was stored with, when reading the data back from a file. The downside of this is that most of the data formats used in ATLAS cannot be read using ROOT only, i. e. without the ATHENA framework. An overview of the data formats in ATLAS is given in the following.

A.10.3 Data formats

Data taken with the ATLAS detector is made available for analysis in several formats, which mainly differ in volume, the level of detail and the level of processing of the detector output. Three main classes of formats can be distinguished. The raw output from the detector is stored in raw data (RAW) and Raw Data Objects (RDO) files, before any reconstruction of physics objects is done. These files contains low-level output from the detectors which for most analyses is unnecessarily intricate to use. From this output, physics objects are created in the reconstruction step of data processing (cf. Section 4.4). These objects are stored in the Event Summary Data (ESD) and AOD data format, which is able to hold the description of physics objects like electrons or tracks of charged particles and the relation between these objects, such that e. g. for a muon, the tracks associated to it in the process of reconstruction, and in turn for the track, the position of hits in different subdetectors can be stored and accessed in the analysis. The AODs are a format derived from the ESDs and contain a subset of the information. Due to the flexibility of these formats that requires special converters for persistifying the information, they cannot be accessed by ROOT directly, but only in combination with the ATHENA libraries. Many analyses therefore use the simple ROOT-readable ntuple (NTUP) data format, which obviates the need to use ATHENA because these files, also called (flat) *n-tuples*, contain only simple vectors and objects that can be processed by ROOT directly. This comes at the price of reduced flexibility and often an again increased data volume, because instead of being able to make links between objects, part of the information has to be stored redundantly in different places. Many versions of ntuples, differing

by their variable content, are automatically produced in the central production, from the higher data formats such as AOD and ESD. One of these ntuple variants are the SUSY ntuples, NTUP_SUSY, which are adapted to the requirements of the ATLAS Supersymmetry group.

Note that this distinction is done from the point-of-view of analysis, i. e. from the objects contained within the files rather than the underlying physical file format, which is the ByteStream format for RAW, a format which reflects the format in which the data are delivered from the detector read-out systems, and the POOL ROOT format for RDO, AOD and ESD files, which are object-oriented representations, where RDO is an object-oriented representation of the ByteStream data, i. e. not in terms of physics objects.

The data formats described above, besides technical aspects in terms of accessibility, are different compromises between data size and level of detail. Even when using the NTUP format, which is supposed to be the smallest representation with regard to event size, reducing the amount of data that needs to be processed in subsequent analysis steps may be desired or required. For the reduction of the data volume, three different approaches are available:

- Events that fail very basic cuts like the trigger requirement or (for data) the Good Runs List (GRL) selection can be discarded from the dataset. This is called *skimming*. The efficiency of this reduction depends strongly on the cut and, on Monte Carlo, on the sample.
- Objects that do not satisfy the criteria to be identified as a certain physics object can be removed from object container. This is called *thinning*. For example, jets that do not have a certain minimum p_T will not be stored in the jet object collections, so this kind of thinning is done right after the object reconstruction. From a technical point of view, thinning is more difficult to implement in flat ntuples because the entries connected to an object need to be removed from all branches.
- Variables, i. e. object or event properties, that are not needed in the analysis can be removed (*slimming* of the dataset). This is easily possible in flat ntuples thanks to the structure of this data format, where data is not stored in event objects, but every variable is a vector and entries at the same offset in all vectors are assumed to belong to the same event. It is thus possible to remove variables by deleting the corresponding vector without having to loop over all events.

A.10.4 Nomenclature of Dataset Names in ATLAS

The general format of the name of an ATLAS dataset is [135]

$$\textit{Project}.[\textit{OtherFields}].\textit{DataType}.\textit{Version}[/]$$

If the dataset name ends with a slash, then the dataset is a container holding several datasets. The project part of the dataset name gives a general idea of what is contained in the dataset. The following project names are of special importance in the context of this thesis:

- *mcvv*: Datasets which contain Monte Carlo events begin with *mcvv*, where *vv* gives the release version of the Monte Carlo, which roughly corresponds to the two last digits of the year in which the Monte Carlo production has been in started.

- *datayy*: This is the project name of datasets containing data taken with the ATLAS detector, where *yy* specifies the year in the which the data has been taken.
- *user*: Datasets which are not centrally produced, but uploaded or created by users, start with *user* to make them easily distinguishable from official productions.

The full name of a Monte Carlo dataset will have the form

mcNN_subProject.datasetNumber.physicsShort.prodStep.dataType.Version

The project part of the name typically ends with a specification of the center-of-mass energy used in the simulation, which can be 900GeV, 7TeV, 10TeV, or 2TeV for 2.36 TeV. *datasetNumber* is a 6 digit decimal number used as unique identifier of Monte Carlo datasets. It is associated with a particular type of physics process, which is also described in a slightly easier to remember textual form in *physicsShort*. The *datasetNumber* uniquely identifies the class of events simulated in this sample across different release versions of the Monte Carlo. *prodStep* gives the production step used to create the data, which is also reflected in the *Version* field. *dataType* specifies the type of data, e. g. RAW, ESD, AOD, or NTUP, and is explained below, as is the *Version* field.

The full name of a dataset with real data will have the form

dataNN_subProject.runNumber.streamName.prodStep.dataType.Version

The fields are again mostly self-explanatory: *runNumber* is the number of the run the data belongs to and *streamName* is the name of the data stream. *subProject*, *prodStep*, *dataType* and *Version* are the same as for Monte Carlo datasets.

The *dataType* field is important to select the data format (see above) best suited for the analysis. The *Version* field encodes the configuration of each of the processing steps that were passed by this dataset. It consists of several entries separated by underscores, each consisting of a letter and a number the meaning of which can be resolved via the AMI. For Monte Carlo datasets, the *Version* field typically contains a combination of e, s and r, describing the event generation, detector simulation and reconstruction, respectively. For datasets containing real data, the first two steps do not apply, of course.

A.10.5 Databases in ATLAS

Apart from storing the detector output, from which the events and physics objects are reconstructed, running a large machine like the ATLAS detector requires a lot of additional bookkeeping. Part of the information stored in these databases is metadata, i. e. second-level information about other data stored elsewhere. For example, this can be information about the distribution of the datasets over the various grid sites. Another part is primary information like the status of the detector, summaries of the measurements or the configuration of the trigger system.

The Atlas Metadata Interface (AMI) is a catalogue, not only of the datasets with Monte Carlo simulated and real events. It also stores the software version tags that have been used in the simulation of Monte Carlo and processing of data. The ATLAS trigger database (TRIGGERDB) is a relational database that holds the information that is needed to configure the trigger system in online data taking [110]. This information needs to be deployed to the Central Trigger Processor (CTP) and all worker nodes at Level 2 and Event Filter, and needs

to be changeable during data taking in realtime. It is also needed in the simulation of Monte Carlo events, and the information is replicated to the COOL database (see below), from which it is picked up in the reprocessing of the data and stored also in the event records.

The ATLAS Conditions Database, called COOL for short after the name of the underlying application programming interface, records information about the state of the detector at the time when events are collected [127], such as the alignment of a tracking device, temperatures measured by a sensor or the trigger configuration. Rather than being an event-based record of detector data, the data in COOL database describes the state of the detector during given intervals of time called Interval Of Validity. The COOL database is versioned, which means that for some objects in the database there may exist several different versions for the same IOV. This may be the case, for example, when calibration or alignment parameters are updated, or an updated luminosity computation becomes available, which supersedes the old calibration valid for the same time. COOL is a common project of both ATLAS and the LHCb experiment at the LHC. Metadata about the detector configuration and conditions are also stored in the Conditions Metadata (COMA) database [183], which offers a compact overview of subsets of important information from other databases.

Image Credits

The following list gives external sources of figures which have been used in this thesis:

- Figure 2.1: Plot taken from [69]. The plots have been rearranged within the figure. © 1993 by The American Physical Society.
- Figure 3.2: Available from <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResults>, version of 16.12.2011. Original title: "ATLAS instantaneous luminosity". Original description: "ATLAS instantaneous luminosity profiles as measured online for representative LHC fills with 7 TeV centre-of-mass energy in 2010. The gray shaded curves give the delivered luminosity, the green shaded curves show the delivered luminosity during stable beam conditions allowing ATLAS to turn on their tracking devices, and the yellow shaded curves give the recorded luminosity with the entire detector available. [...] The luminosity values shown have been calibrated with van-der-Meer beam-separation scan data. The error on the luminosity is estimated to be 11 %, dominated by the uncertainty in the beam intensities."
- Figure 3.3: Downloaded from https://lar-elec-automatic-validation.web.cern.ch/lar-elec-automatic-validation/cgi-bin/PulseShape.sh?sub_sub_partition=EMBA&ft=0&sl=7&ch=124&campaign=185&gain=M on 28.09.2011.
- Figure 3.4: Downloaded from CERN Document Server, record LHC-PHO-1993-005, <http://cdsweb.cern.ch/record/841542> on 10.06.2011. Original title: "Layout of the LEP tunnel including future LHC infrastructures. L'ensemble du tunnel LEP avec les futures infrastructures LHC." © 1993 CERN.
- Figure 3.5: Available from <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResults>, version of 16.12.2011. Original title: "Total Integrated Luminosity in 2011". Original description: "Cumulative luminosity versus day delivered to (green), and recorded by ATLAS (yellow) during stable beams and for pp collisions at 7 TeV centre-of-mass energy in 2011. The delivered luminosity accounts for the luminosity delivered from the start of stable beams until the LHC requests ATLAS to turn the sensitive detector off to allow a beam dump or beam studies. Given is the luminosity as determined from counting rates measured by the luminosity detectors. These detectors have been calibrated with the use of the van-der-Meer beam-separation method, where the two beams are scanned against each other in the horizontal and vertical planes to measure their overlap function."
- Figure 3.6: Plot taken from [128], page 34. This is a modified version of the plot in [33], page 95.
- Figure 3.7: Downloaded from the CERN Document Server, record CERN-GE-0803012, <http://cdsweb.cern.ch/record/1095924> on 02.11.2011. Original title: "Computer generated image of the whole ATLAS detector". ATLAS Experiment © 2008 CERN.

- Figures 3.8: Downloaded from the CERN Document Server, record CERN-GE-0803014, <http://cdsweb.cern.ch/record/1095926> on 02.11.2011. Original title: "Computer generated image of the ATLAS inner detector". ATLAS Experiment © 2008 CERN.
- Figure 3.9: Downloaded from the CERN Document Server, record CERN-GE-0803015, <http://cdsweb.cern.ch/record/1095927> on 04.11.2011. Original title: "Computer Generated image of the ATLAS calorimeter". ATLAS Experiment © 2008 CERN.
- Figure 3.10: Downloaded from the CERN Document Server, record CERN-GE-0803017, <http://cdsweb.cern.ch/record/1095929> on 04.11.2011. Original title: "Computer generated image of the ATLAS Muons subsystem". ATLAS Experiment © 2008 CERN.
- Figure 3.11: Plot taken from [43], page 163.
- Figure 4.1: Figure taken from [140], page 2.
- Figure A.12: Figure taken from [127], page 28.

Bibliography

- [1] I. J. R. Aitchison. *Supersymmetry in Particle Physics: An Elementary Introduction*. Cambridge University Press, Cambridge, 2007.
- [2] R. J. Barlow. *Statistics: A Guide to the Use of Statistical Methods in the Physical Sciences*. The Manchester Physics Series. John Wiley & Sons, Chichester, 1989. Reprinted 1999.
- [3] P. Binétruy. *Supersymmetry: Theory, Experiment, and Cosmology*. Oxford Graduate Texts. Oxford University Press, Oxford, 2006.
- [4] G. Casella and R. L. Berger. *Statistical Inference*. Duxbury/Thomson Learning, Pacific Grove (US-CA), second edition, 2002.
- [5] M. Goto. *C++ Interpreter – CINT*. CQ publishing, 1997. In Japanese, cited as given on <http://root.cern.ch/drupal/content/publication-list>.
- [6] R. K. Ellis, W. J. Stirling, and B. R. Webber. *QCD and Collider Physics*, volume 8 of *Cambridge Monographs on Particle Physics, Nuclear Physics and Cosmology*. Cambridge University Press, Cambridge, 1996.
- [7] W. Greiner and J. Rafelski. *Spezielle Relativitätstheorie*. Number 3a in Theoretische Physik. Verlag Harri Deutsch, Frankfurt am Main, third, revised edition, 1992.
- [8] D. Griffiths. *Introduction to Elementary Particles*. Wiley-VCH, Weinheim, second, revised edition, 2008.
- [9] F. Halzen and A. D. Martin. *Quarks & Leptons: An Introductory Course in Modern Particle Physics*. John Wiley & Sons, 1984.
- [10] K. Kleinknecht. *Detektoren für Teilchenstrahlung*. B. G. Teubner Verlag, Wiesbaden, fourth, revised edition, 2005.
- [11] P. Langacker. *The Standard Model and Beyond*. Series in High Energy Physics, Cosmology, and Gravitation. CRC Press, Boca Raton (US-FL), 2010.
- [12] B. R. Martin and G. Shaw. *Particle Physics*. The Manchester Physics Series. John Wiley & Sons, Chichester, second edition, 1997.
- [13] M. Maggiore. *A Modern Introduction to Quantum Field Theory*. Number 12 in Oxford Master Series in Physics. Oxford University Press, Oxford, 2005.
- [14] A. Messiah. *Quantenmechanik, Band 2*. Verlag Walter de Gruyter, Berlin, second edition, 1985.
- [15] M. E. Peskin and D. V. Schroeder. *An Introduction To Quantum Field Theory*. Addison-Wesley, Reading (US-MA), 1995.

- [16] R. Stock (editor). *Encyclopedia of Applied High Energy and Particle Physics*. Wiley-VCH, Weinheim, 2009.
- [17] S. S. AbdusSalam, B. C. Allanach, H. K. Dreiner, J. Ellis, U. Ellwanger, et al. “Benchmark Models, Planes, Lines and Points for Future SUSY Searches at the LHC”. arXiv.org e-print archive, arXiv:hep-ph/1109.3859, 2011.
- [18] M. L. Mangano, F. Piccinini, A. D. Polosa, M. Moretti, and R. Pittau. “ALPGEN, a generator for hard multiparton processes in hadronic collisions”. *Journal of High Energy Physics*, 07(2003)001, 2003. doi:10.1088/1126-6708/2003/07/001.
- [19] C. Amsler et al. “Review of Particle Physics”. *Physics Letters B*, 667(1–5): 1, 2008. doi:10.1016/j.physletb.2008.07.018.
- [20] The ATHENA Collaboration. “Production and detection of cold antihydrogen atoms”. *Nature*, 419: 456–459, 2002. doi:10.1038/nature01096.
- [21] The ATLAS Collaboration. “The ATLAS Experiment at the CERN Large Hadron Collider”. *Journal of Instrumentation*, 3(08): S08003, 2008. doi:10.1088/1748-0221/3/08/S08003.
- [22] H. Baer, E.-K. Park, and X. Tata. “Collider, direct and indirect detection of supersymmetric dark matter”. *New Journal of Physics*, 11(10): 105024, 2009. doi:10.1088/1367-2630/11/10/105024.
- [23] J. Bagnasco, M. Dine, and S. Thomas. “Detecting technibaryon dark matter”. *Physics Letters B*, 320(1–2): 99–104, 1994. doi:10.1016/0370-2693(94)90830-3.
- [24] A. Barr, C. Lester, and P. Stephens. “A variable for measuring masses at hadron colliders when missing energy is expected; m_{T2} : The Truth behind the glamour”. *Journal of Physics G: Nuclear and Particle Physics*, 29: 2343–2363, 2003. doi:10.1088/0954-3899/29/10/304.
- [25] A. Barr and C. Gwenlan. “The race for supersymmetry: Using m_{T2} for discovery”. *Physical Review D*, 80(7): 074007, 2009. doi:10.1103/PhysRevD.80.074007.
- [26] W. Beenakker, R. Höpker, and M. Spira. “PROSPINO: A Program for the Production of Supersymmetric Particles in Next-to-leading Order QCD”. arXiv.org e-print archive, arXiv:hep-ph/9611232, 1996.
- [27] W. Beenakker, R. Höpker, M. Spira, and P. M. Zerwas. “Squark and gluino production at hadron colliders”. *Nuclear Physics B*, 492(1–2): 51–103, 1997. doi:10.1016/S0550-3213(97)80027-2.
- [28] W. Beenakker, M. Krämer, T. Plehn, M. Spira, and P. M. Zerwas. “Stop production at hadron colliders”. *Nuclear Physics B*, 515(1–2): 3–14, 1998. doi:10.1016/S0550-3213(98)00014-5.
- [29] The Muon ($g - 2$) Collaboration. “Final report of the E821 muon anomalous magnetic moment measurement at BNL”. *Physical Review D*, 73(7): 072003, 2006. doi:10.1103/PhysRevD.73.072003.

-
- [30] The UA2 Collaboration. "Measurement of jet production properties at the CERN $\bar{p}p$ Collider". *Physics Letters B*, 144(3–4): 283–290, 1984. doi:10.1016/0370-2693(84)91822-7.
- [31] R. Bhargava and C. Khatri. "The distribution of product of independent beta random variables with application to multivariate analysis". *Annals of the Institute of Statistical Mathematics*, 33(1): 287–296, 1981. doi:10.1007/BF02480942.
- [32] M. Cacciari, G. P. Salam, and G. Soyez. "The anti- k_t jet clustering algorithm". *Journal of High Energy Physics*, 04(2008)063, 2008. doi:10.1088/1126-6708/2008/04/063.
- [33] J. M. Campbell, J. W. Huston, and W. J. Stirling. "Hard interactions of quarks and gluons: a primer for LHC physics". *Reports on Progress in Physics*, 70(1): 89–193, 2007. doi:10.1088/0034-4885/70/1/R02.
- [34] D. Casadei. "Estimating the selection efficiency". arXiv.org e-print archive, arXiv:physics.data-an/0908.0130, 2010.
- [35] The CDF Collaboration. "Search for Gluinos and Scalar Quarks in $p\bar{p}$ Collisions at $\sqrt{s} = 1.8$ TeV Using the Missing Energy plus Multijets Signature". *Physical Review Letters*, 88(4): 041801, 2002. doi:10.1103/PhysRevLett.88.041801.
- [36] A. H. Chamseddine, R. L. Arnowitt, and P. Nath. "Locally Supersymmetric Grand Unification". *Physical Review Letters*, 49: 970–974, 1982. doi:10.1103/PhysRevLett.49.970.
- [37] H.-C. Cheng, K. T. Matchev, and M. Schmaltz. "Bosonic supersymmetry? Getting fooled at the CERN LHC". *Physical Review D*, 66(5): 056006, 2002. doi:10.1103/PhysRevD.66.056006.
- [38] The CMS Collaboration. "The CMS high level trigger". *The European Physical Journal C*, 46: 605–667, 2006. doi:10.1140/epjc/s2006-02495-8.
- [39] Raffaele Tito D’Agnolo for the CMS Collaboration. "Supersymmetry Searches in multijet events with CMS". *Journal of Physics: Conference Series*, 323(1): 012009, 2011. doi:10.1088/1742-6596/323/1/012009.
- [40] B. Combridge, J. Kripfganz, and J. Ranft. "Hadron production at large transverse momentum and QCD". *Physics Letters B*, 70(2): 234–238, 1977. doi:10.1016/0370-2693(77)90528-7.
- [41] The ATLAS Collaboration. "Studies of the performance of the ATLAS detector using cosmic-ray muons". *The European Physical Journal C*, 71(3): 1–36, 2011. doi:10.1140/epjc/s10052-011-1593-6.
- [42] G. Cowan, K. Cranmer, E. Gross, and O. Vitells. "Asymptotic formulae for likelihood-based tests of new physics". *The European Physical Journal C*, 71(2): 1554, 2011. doi:10.1140/epjc/s10052-011-1554-0.
- [43] The ATLAS Collaboration. "Expected Performance of the ATLAS Experiment — Detector, Trigger and Physics". arXiv.org e-print archive, arXiv:hep-ex/0901.0512, 2009. Report number: CERN-OPEN-2008-020.

- [44] J. Pumplin, D. R. Stump, J. Huston, H.-L. Lai, P. Nadolsky, et al. “New Generation of Parton Distributions with Uncertainties from Global QCD Analysis”. *Journal of High Energy Physics*, 07(2002)12, 2002. doi:10.1088/1126-6708/2002/07/012.
- [45] P. M. Nadolsky, H.-L. Lai, Q.-H. Cao, J. Huston, J. Pumplin, et al. “Implications of CTEQ global analysis for collider observables”. *Physical Review D*, 78(1): 013004, 2008. doi:10.1103/PhysRevD.78.013004.
- [46] H.-L. Lai, J. Huston, Z. Li, P. Nadolsky, J. Pumplin, et al. “Uncertainty induced by QCD coupling in the CTEQ global analysis of parton distributions”. *Physical Review D*, 82(5): 054021, 2010. doi:10.1103/PhysRevD.82.054021.
- [47] G. Degrossi, S. Heinemeyer, W. Hollik, P. Slavich, and G. Weiglein. “Towards high-precision predictions for the MSSM Higgs sector”. *The European Physical Journal C*, 28(1): 133–143, 2003. doi:10.1140/epjc/s2003-01152-2.
- [48] The DØ Collaboration. “High- p_T jets in $\bar{p}p$ collisions at $\sqrt{s} = 630$ and 1800 GeV”. *Physical Review D*, 64(3): 032003, 2001. doi:10.1103/PhysRevD.64.032003.
- [49] S. Dimopoulos and D. Sutter. “The supersymmetric flavor problem”. *Nuclear Physics B*, 452(3): 496–512, 1995. doi:10.1016/0550-3213(95)00421-N.
- [50] The ATLAS Collaboration. “Electron performance measurements with the ATLAS detector using the 2010 LHC proton-proton collision data”. arXiv.org e-print archive, arXiv:hep-ex/1110.3174, 2011. Submitted to European Physical Journal C.
- [51] S. Klous. “Event streaming in the online system: Real-time organization of ATLAS data”. In *Real Time Conference (RT), 2010 17th IEEE-NPSS*, pages 1–8, 2010.
- [52] G. Barrand, I. Belyaev, P. Binko, M. Cattaneo, R. Chytracsek, et al. “GAUDI — A software architecture and framework for building HEP data processing applications”. *Computer Physics Communications*, 140(1–2): 45–55, 2001. doi:10.1016/S0010-4655(01)00254-5.
- [53] F. Gianotti. “Physics during the first two years of the LHC”. *New Journal of Physics*, 9(9): 332, 2007. doi:10.1088/1367-2630/9/9/332.
- [54] S. L. Glashow, J. Iliopoulos, and L. Maiani. “Weak Interactions with Lepton-Hadron Symmetry”. *Physical Review D*, 2(7): 1285–1292, 1970. doi:10.1103/PhysRevD.2.1285.
- [55] R. Haag, J. T. Łopuszański, and M. Sohnius. “All possible generators of supersymmetries of the S -matrix”. *Nuclear Physics B*, 88(2): 257–274, 1975. doi:10.1016/0550-3213(75)90279-5.
- [56] L. Hall, J. Lykken, and S. Weinberg. “Supergravity as the messenger of supersymmetry breaking”. *Physical Review D*, 27(10): 2359–2378, 1983. doi:10.1103/PhysRevD.27.2359.
- [57] G. Corcella, I. G. Knowles, G. Marchesini, S. Moretti, K. Odagiri, et al. “HERWIG 6: an event generator for hadron emission reactions with interfering gluons (including supersymmetric processes)”. *Journal of High Energy Physics*, 01(2001)010, 2001. doi:10.1088/1126-6708/2001/01/010.

-
- [58] G. Corcella, I. Knowles, G. Marchesini, S. Moretti, K. Odagiri, et al. “HERWIG 6.5 Release Note”. arXiv.org e-print archive, arXiv:hep-ph/0210213, 2002.
- [59] M. Bähr, S. Gieseke, M. Gigg, D. Grellscheid, K. Hamilton, et al. “Herwig++ physics and manual”. *The European Physical Journal C*, 58(4): 639–707, 2008. doi:10.1140/epjc/s10052-008-0798-9.
- [60] M. Bähr, S. Gieseke, M. Gigg, D. Grellscheid, K. Hamilton, et al. “Herwig++ 2.3 Release Note”. arXiv.org e-print archive, arXiv:hep-ph/0812.0529, 2008.
- [61] P. W. Higgs. “Broken Symmetries and the Masses of Gauge Bosons”. *Physical Review Letters*, 13: 508–509, 1964. doi:10.1103/PhysRevLett.13.508.
- [62] N. Berger, T. Bold, T. Eifert, G. Fischer, S. George, et al. “The High-Level-Trigger Steering of the ATLAS Experiment”. *IEEE Transactions on Nuclear Science*, 55(1): 165–171, 2008. doi:10.1109/TNS.2007.913483.
- [63] H. Baer, F. E. Paige, S. D. Protopopescu, and X. Tata. “ISAJET 7.69: A Monte Carlo Event Generator for pp , $\bar{p}p$, and e^+e^- Reactions”. arXiv.org e-print archive, arXiv:hep-ph/0312045, 2003.
- [64] J. M. Butterworth, J. R. Forshaw, and M. H. Seymour. “Multiparton interactions in photoproduction at HERA”. *Zeitschrift für Physik C*, 72(4): 637–646, 1996. doi:10.1007/s002880050286.
- [65] T. Junk. “Confidence level computation for combining searches with small statistics”. *Nuclear Instruments and Methods in Physics Research Section A*, 434(2–3): 435–443, 1999. doi:10.1016/S0168-9002(99)00498-2.
- [66] N. Kidonakis. “Single top quark production cross section at hadron colliders”. In *Proceedings of the XVIIIth International Workshop on Deep-Inelastic Scattering and Related Subjects. April 19–23, 2010. Convitto della Calza, Firenze, Italy. Published online at <http://pos.sissa.it/cgi-bin/reader/conf.cgi?confid=106>, id. 196*, 2010.
- [67] J. E. Kim and H. P. Nilles. “The μ -problem and the strong CP-problem”. *Physics Letters B*, 138(1–3): 150–154, 1984. doi:10.1016/0370-2693(84)91890-2.
- [68] P. Langacker. “Structure of the Standard Model”. In P. Langacker (editor), *Precision tests of the standard electroweak model*, volume 14 of *Advanced series on directions in high energy physics*, pages 15–36. World Scientific, Singapore, 1995.
- [69] P. Langacker and N. Polonsky. “Uncertainties in coupling constant unification”. *Physical Review D*, 47(9): 4028–4045, 1993. doi:10.1103/PhysRevD.47.4028.
- [70] The LEP Working Group for Higgs Boson Searches. “Search for the Standard Model Higgs boson at LEP”. *Physics Letters B*, 565: 61–75, 2003. doi:10.1016/S0370-2693(03)00614-2.
- [71] C. G. Lester and D. J. Summers. “Measuring masses of semi-invisibly decaying particles pair produced at hadron colliders”. *Physics Letters B*, 463(1): 99–103, 1999. doi:10.1016/S0370-2693(99)00945-4.

- [72] L. Evans and P. B. (editors). “LHC Machine”. *Journal of Instrumentation*, 3(08): S08001, 2008. doi:10.1088/1748-0221/3/08/S08001.
- [73] O. Adriani, L. Bonechi, M. Bongi, G. Castellini, R. D’Alessandro, et al. “Measurement of zero degree single photon energy spectra for $\sqrt{s} = 7$ TeV proton-proton collisions at LHC”. *Physics Letters B*, 703(2): 128–134, 2011. doi:10.1016/j.physletb.2011.07.077.
- [74] The ATLAS Collaboration. “Luminosity determination in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector at the LHC”. *The European Physical Journal C*, 71(4): 1–37, 2011. doi:10.1140/epjc/s10052-011-1630-5.
- [75] S. Frixione and B. R. Webber. “Matching NLO QCD computations and parton shower simulations”. *Journal of High Energy Physics*, 06(2002)029, 2002. doi:10.1088/1126-6708/2002/06/029.
- [76] S. Frixione, P. Nason, and B. R. Webber. “Matching NLO QCD and parton showers in heavy flavour production”. *Journal of High Energy Physics*, 08(2003)007, 2003. doi:10.1088/1126-6708/2003/08/007.
- [77] S. Frixione, E. Laenen, P. Motylinski, and B. R. Webber. “Single-top production in MC@NLO”. *Journal of High Energy Physics*, 03(2006)092, 2006. doi:10.1088/1126-6708/2006/03/092.
- [78] S. Frixione, E. Laenen, P. Motylinski, C. White, and B. R. Webber. “Single-top hadroproduction in association with a W boson”. *Journal of High Energy Physics*, 07(2008)029, 2008. doi:10.1088/1126-6708/2008/07/029.
- [79] C. Macesanu. “The Phenomenology of Universal Extra Dimensions at Hadron Colliders”. *International Journal of Modern Physics A*, 21(11): 2259–2296, 2006. doi:10.1142/S0217751X06030886.
- [80] The ATLAS Collaboration. “Performance of Missing Transverse Momentum Reconstruction in Proton-Proton Collisions at 7 TeV with ATLAS”. arXiv.org e-print archive, arXiv:hep-ex/1108.5602, 2011. Submitted to European Physical Journal C.
- [81] F. James and M. Roos. “Minuit: A System for Function Minimization and Analysis of the Parameter Errors and Correlations”. *Computer Physics Communications*, 10(6): 343–367, 1975. doi:10.1016/0010-4655(75)90039-9.
- [82] L. Moneta, K. Cranmer, G. Schott, and W. Verkerke. “The RooStats project”. In *Proceedings of the 13th International Workshop on Advanced Computing and Analysis Techniques in Physics Research*, Jaipur, India, 2010. Published online at <http://pos.sissa.it/cgi-bin/reader/conf.cgi?confid=93>.
- [83] A. Sherstnev and R. S. Thorne. “Parton distributions for LO generators”. *The European Physical Journal C*, 55(4): 553–575, 2008. doi:10.1140/epjc/s10052-008-0610-x.
- [84] A. D. Martin, W. J. Stirling, R. S. Thorne, and G. Watt. “Parton distributions for the LHC”. *The European Physical Journal C*, 63(2): 189–285, 2009. doi:10.1140/epjc/s10052-009-1072-5.

-
- [85] E. Noether. "Invariante Variationsprobleme". *Kgl. Ges. d. Wiss. Nachrichten, Math.-phys. Klasse, Göttingen*, 2: 235–257, 1918.
- [86] K. A. Olive. "Dark Matter in SuperGUT Unification Models". *Journal of Physics: Conference Series*, 315(1): 012021, 2011. doi:10.1088/1742-6596/315/1/012021.
- [87] Nakamura, K. and others (Particle Data Group). "Review of Particle Physics". *Journal of Physics G: Nuclear and Particle Physics*, 37: 075021, 2010. doi:10.1088/0954-3899/37/7A/075021.
- [88] The ALEPH Collaboration. "Search for supersymmetric particles in e^+e^- collisions at \sqrt{s} up to 202 GeV and mass limit for the lightest neutralino". *Physics Letters B*, 499(1–2): 67–84, 2001. doi:10.1016/S0370-2693(01)00015-6.
- [89] The ALEPH Collaboration. "Absolute mass lower limit for the lightest neutralino of the MSSM from e^+e^- data at \sqrt{s} up to 209 GeV". *Physics Letters B*, 583(3–4): 247–263, 2004. doi:10.1016/j.physletb.2003.12.066.
- [90] The DØ Collaboration. "Search for squarks and gluinos in events with jets and missing transverse energy using 2.1 fb^{-1} of $p\bar{p}$ collision data at $\sqrt{s} = 1.96 \text{ TeV}$ ". *Physics Letters B*, 660(5): 449–457, 2008. doi:10.1016/j.physletb.2008.01.042.
- [91] The L3 Collaboration. "Search for scalar leptons and scalar quarks at LEP". *Physics Letters B*, 580(1–2): 37–49, 2004. doi:10.1016/j.physletb.2003.10.010.
- [92] The DØ Collaboration. "Search for associated production of charginos and neutralinos in the trilepton final state using 2.3 fb^{-1} of data". *Physics Letters B*, 680(1): 34–43, 2009. doi:10.1016/j.physletb.2009.08.011.
- [93] The CDF Collaboration. "Search for Supersymmetry in $p\bar{p}$ Collisions at $\sqrt{s} = 1.96 \text{ TeV}$ Using the Trilepton Signature for Chargino-Neutralino Production". *Physical Review Letters*, 101(25): 251801, 2008. doi:10.1103/PhysRevLett.101.251801.
- [94] The CDF Collaboration. "Inclusive Search for Squark and Gluino Production in $p\bar{p}$ Collisions at $\sqrt{s} = 1.96 \text{ TeV}$ ". *Physical Review Letters*, 102(12): 121801, 2009. CDF Note 9229. doi:10.1103/PhysRevLett.102.121801.
- [95] J. Pumplin, D. Stump, R. Brock, D. Casey, J. Huston, et al. "Uncertainties of predictions from parton distribution functions. II. The Hessian method". *Physical Review D*, 65(1): 014013, 2001. doi:10.1103/PhysRevD.65.014013.
- [96] G. Piacquadio, K. Prokofiev, and A. Wildauer. "Primary vertex reconstruction in the ATLAS experiment at LHC". *Journal of Physics: Conference Series*, 119(3): 032033, 2008. doi:10.1088/1742-6596/119/3/032033.
- [97] T. Sjöstrand, S. Mrenna, and P. Skands. "PYTHIA 6.4 physics and manual". *Journal of High Energy Physics*, 05(2006)026, 2006. doi:10.1088/1126-6708/2006/05/026.
- [98] M. Quiros. "Introduction to extra dimensions". arXiv.org e-print archive, arXiv: hep-ph/0606153, 2006.

- [99] A. L. Read. “Modified Frequentist Analysis of Search Results (The CL_s Method)”. In F. James, Y. Perrin, and L. Lyons (editors), *Proceedings of the Workshop on Confidence Limits*, volume CERN-2000-005, pages 81–99, CERN, Geneva, Switzerland, 2000.
- [100] A. L. Read. “Presentation of search results: the CL_s technique”. *Journal of Physics G: Nuclear and Particle Physics*, 28(10): 2693–2704, 2002. doi:10.1088/0954-3899/28/10/313.
- [101] R. Brun and F. Rademakers. “ROOT — an object oriented data analysis framework”. *Nuclear Instruments and Methods in Physics Research Section A*, 389(1–2): 81–86, 1997. doi:10.1016/S0168-9002(97)00048-X.
- [102] V. C. Rubin and W. K. Ford, Jr. “Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions”. *Astrophysical Journal*, 159: 379–403, 1970. doi:10.1086/150317.
- [103] V. C. Rubin, W. K. Ford, Jr., and N. Thonnard. “Rotational properties of 21 Sc galaxies with a large range of luminosities and radii, from NGC 4605 ($R = 4$ kpc) to UGC 2885 ($R = 122$ kpc)”. *Astrophysical Journal*, 238: 471–487, 1980. doi:10.1086/158003.
- [104] The SNO Collaboration. “Electron energy spectra, fluxes, and day-night asymmetries of ^8B solar neutrinos from measurements with NaCl dissolved in the heavy-water detector at the Sudbury Neutrino Observatory”. *PRC*, 72(5): 055502, 2005. doi:10.1103/PhysRevC.72.055502.
- [105] The ATLAS Collaboration. “Search for squarks and gluinos using final states with jets and missing transverse momentum with the ATLAS detector in $\sqrt{s} = 7$ TeV proton-proton collisions”. *Physics Letters B*, 701(2): 186–203, 2011. doi:10.1016/j.physletb.2011.05.061.
- [106] S. P. Martin. “A Supersymmetry Primer”. arXiv.org e-print archive, arXiv: hep-ph/9709356, 1997.
- [107] K. Lane. “Two Lectures on Technicolor”. arXiv.org e-print archive, arXiv: hep-ph/0202255, 2002.
- [108] G. ‘t Hooft. “Renormalizable Lagrangians for massive Yang-Mills fields”. *Nuclear Physics B*, 35(1): 167–188, 1971. doi:10.1016/0550-3213(71)90139-8.
- [109] The DØ Collaboration. “Direct measurement of the mass difference between top and antitop quarks”. *Physical Review D*, 84(5): 052005, 2011. doi:10.1103/PhysRevD.84.052005.
- [110] A. dos Anjos, P. Bell, D. Berge, J. Haller, S. Head, et al. “The Configuration System of the ATLAS Trigger”. *IEEE Transactions on Nuclear Science*, 55(1): 392–398, 2008. doi:10.1109/TNS.2007.913934.
- [111] The ATLAS Collaboration. “Performance of the ATLAS Trigger System in 2010”. arXiv.org e-print archive, arXiv:hep-ex/1110.1530, 2011. To appear in *European Physical Journal C*.

-
- [112] A. Wald. “Tests of statistical hypotheses concerning several parameters when the number of observations is large”. *Transactions of the American Mathematical Society*, 54: 426–482, 1943. doi:10.1090/S0002-9947-1943-0012401-3.
- [113] J. Wess and B. Zumino. “A Lagrangian Model Invariant Under Supergauge Transformations”. *Physics Letters B*, 49(1): 52, 1974. doi:10.1016/0370-2693(74)90578-4.
- [114] S. S. Wilks. “The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses”. *The Annals of Mathematical Statistics*, 9(1): pp. 60–62, 1938. Available online: <http://www.jstor.org/stable/2957648>.
- [115] The DØ Collaboration. “Measurements of inclusive $W +$ jets production rates as a function of jet transverse momentum in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV”. *Physics Letters B*, 705(3): 200–207, 2011. doi:10.1016/j.physletb.2011.10.011.
- [116] The Tevatron Electroweak Working Group. “Updated Combination of CDF and DØ Results for the Mass of the W Boson”. arXiv.org e-print archive, arXiv:hep-ex/0908.1374, 2009.
- [117] A. V. Kotwal and J. Stark. “Measurement of the W Boson Mass at the Tevatron”. *Annual Review of Nuclear and Particle Science*, 58: 147–175, 2008. doi:10.1146/annurev.nucl.58.110707.171227.
- [118] F. Zwicky. “Die Rotverschiebung von extragalaktischen Nebeln”. *Helvetica Physica Acta*, 6: 110–127, 1933.
- [119] F. Zwicky. “On the Masses of Nebulae and of Clusters of Nebulae”. *Astrophysical Journal*, 86: 217–246, 1937. doi:10.1086/143864.
- [120] The ALEPH Collaboration. “Update on Searches for Charginos and Neutralinos in e^+e^- Collisions at \sqrt{s} up to 208 GeV and Mass Limit for the Lightest Neutralino”. Report No. ALEPH 2001-080 / CONF 2001-050, CERN, 2002. http://lepsusy.web.cern.ch/lepsusy/www/lspmsugra_summer02/02-06.2/mSUGRA_208.html.
- [121] The CDF Collaboration. “Search for Gluino and Squark Production in Multijets Plus Missing E_T Final States”. CDF Note 9229, Fermilab, 2008. http://www-cdf.fnal.gov/physics/exotic/r2a/20080214.squark_gluino/squark_gluino.html.
- [122] The CMS Collaboration. “Measuring Electron Efficiencies at CMS with Early Data”. CMS Physics Analysis Summaries CMS-PAS-EGM-07-001, CERN, 2008.
- [123] The CMS Collaboration. “Combination of CMS searches for a Standard Model Higgs boson”. CMS Physics Analysis Summaries CMS-PAS-HIG-11-032, CERN, 2011.
- [124] The LEP2 SUSY Working Group. “Combined LEP Chargino Results, up to 208 GeV for large m_0 ”. Note LEPSUSYWG/01-03.1, CERN, 2001. http://lepsusy.web.cern.ch/lepsusy/www/inos_moriond01/charginos_pub.html.
- [125] The LEP SUSY Working Group. “Interpretation of the results in Minimal SUGRA”. Note LEPSUSYWG/02-06.2, CERN, 2002. http://lepsusy.web.cern.ch/lepsusy/www/lspmsugra_summer02/02-06.2/mSUGRA_208.html.

- [126] The LEP2 SUSY Working Group. “Combined LEP Selectron/Smuon/Stau Results, 183–208 GeV”. Note LEPSUSYWG/04-01.1, CERN, 2004. http://lepsusy.web.cern.ch/lepsusy/www/sleptons_summer04/slep_final.html.
- [127] The ATLAS Collaboration. “ATLAS Computing: Technical Design Report”. ATLAS-TDR-017, CERN-LHCC-2005-022, CERN, Geneva, 2005.
- [128] The ATLAS Collaboration. “ATLAS High-Level Trigger, Data Acquisition and Controls: Technical Design Report”. ATLAS-TDR-016, CERN-LHCC-2003-022, CERN, Geneva, 2003.
- [129] The ATLAS Collaboration. “ATLAS detector and physics performance: Technical Design Report, Volumes I and II”. ATLAS-TDR-014, CERN-LHCC-99-014; ATLAS-TDR-015, CERN-LHCC-99-015, CERN, Geneva, 1999.
- [130] The ATLAS Collaboration. “ATLAS Level-1 Trigger: Technical Design Report”. ATLAS-TDR-012, CERN-LHCC-98-014, CERN, Geneva, 1998.
- [131] O. S. Brüning, P. Collier, P. Lebrun, S. Myers, R. Ostojic, et al. (editors). *LHC Design Report*, volume 1: the LHC Main Ring. CERN, Geneva, 2004. CERN-2004-003-V-1.
- [132] M. Benedikt, P. Collier, V. Mertens, J. Poole, and K. Schindl (editors). *LHC Design Report*, volume 3: the LHC Injector Chain. CERN, Geneva, 2004. CERN-2004-003-V-3.
- [133] C. Hensel and A. Mann. “Efficiency Estimates for the Combined jet + E_T^{miss} Trigger of the ATLAS Experiment Calculated from pp Collisions with $\sqrt{s} = 7 \text{ TeV}$ ”. ATLAS Note ATL-COM-DAQ-2011-027, CERN, Geneva, 2011.
- [134] M. Begel, M. Campanelli, R. Chislett, G. Kasieczka, J. Machado Miguens, et al. “Jet Trigger Performance with 2011 Data”. ATLAS Note ATL-COM-DAQ-2011-063, CERN, Geneva, 2011.
- [135] S. Albrand, D. Barberis, F. Gianotti, C. Guyot, R. Hawkings, et al. “ATLAS Dataset Nomenclature”. ATLAS Note ATL-GEN-INT-2007-001, ATL-COM-GEN-2007-003, CERN, Geneva, 2007.
- [136] M. Baak, C. Guyot, M. Hauschild, R. Hawkings, B. Heinemann, et al. “Data Quality Status Flags and Good Run Lists for Physics Analysis in ATLAS”. ATLAS Note ATL-COM-GEN-2009-015, CERN, Geneva, 2009.
- [137] D. Casadei, G. Lewis, A. Mann, A. Mincer, J. Morel, et al. “Study of pile-up effect on missing transverse energy trigger”. ATLAS Note ATL-COM-DAQ-2010-218, CERN, Geneva, 2010.
- [138] G. Romeo, A. Schwartzman, R. Piegai, T. Carli, and R. Teuscher. “Jet Energy Resolution from In-situ Techniques with the ATLAS Detector Using Proton-Proton Collisions at a Center of Mass Energy $\sqrt{s} = 7 \text{ TeV}$ ”. ATLAS Note ATL-COM-PHYS-2011-240, CERN, Geneva, 2011.
- [139] D. Casadei, I. Aracena, S. Banerjee, P.-H. Beauchemin, K. Cranmer, et al. “The ATLAS missing E_T trigger for the initial LHC operation”. ATLAS Note ATL-DAQ-INT-2010-004, CERN, Geneva, 2010.

-
- [140] S. Klous. "Event streaming in the online system". ATLAS Note ATL-DAQ-PROC-2010-017, CERN, Geneva, 2010.
- [141] D. Casadei, I. Aracena, S. Banerjee, P.-H. Beauchemin, S. Calvet, et al. "The implementation of the ATLAS missing E_T triggers for the initial LHC operation". ATLAS Note ATL-DAQ-PUB-2011-001, CERN, Geneva, 2011.
- [142] T. Berillari. "Status of the ATLAS liquid argon calorimeter and its performance". ATLAS Note ATL-LARG-PROC-2011-010, CERN, Geneva, 2011.
- [143] W. Lampl, S. Laplace, D. Lelas, P. Loch, H. Ma, et al. "Calorimeter Clustering Algorithms: Description and Performance". ATLAS Note ATL-LARG-PUB-2008-002, ATL-COM-LARG-2008-003, CERN, Geneva, 2008.
- [144] M. Aharrouche, C. Anastopoulos, V. Dao, F. Derue, F. Dudziak, et al. "Expected electron performance in the ATLAS experiment". ATLAS Note ATL-PHYS-INT-2010-126, CERN, Geneva, 2010.
- [145] S. Allwood-Spires, M. Barisonzi, H. Beauchemin, R. Bruneliere, J. Buchanan, et al. "Monte Carlo samples used for top physics". Technical Report ATL-PHYS-INT-2010-132, CERN, Geneva, 2010.
- [146] F. Ahles, I. Aracena, S. Asai, S. Ask, M. Baak, et al. "Search for squarks and gluinos using final states with jets and missing transverse momentum with the ATLAS detector in $\sqrt{s} = 7$ TeV proton-proton collisions". ATLAS Note ATL-PHYS-INT-2011-009, CERN, Geneva, 2011.
- [147] R. Bruneliere, M. Baak, J. Lundberg, M. C. Rammensee, and T. J. Khoo. "Setting exclusion limits in ATLAS supersymmetry searches with a likelihood ratio based method". Technical Report ATL-PHYS-INT-2011-032, CERN, Geneva, 2011.
- [148] S. Asai, S. Ask, M. Baak, A. J. Barr, R. Bruneliere, et al. "Search for squarks and gluinos using final states with jets and missing transverse momentum with the ATLAS experiment in $\sqrt{s} = 7$ TeV proton-proton collisions: supporting documentation". ATLAS Note ATL-PHYS-INT-2011-055, CERN, Geneva, 2011.
- [149] S. Asai, S. Ask, M. Baak, A. J. Barr, R. Bruneliere, et al. "Further search for squarks and gluinos using final states with jets and missing transverse momentum with the ATLAS experiment in $\sqrt{s} = 7$ TeV proton-proton collisions: supporting documentation". ATLAS Note ATL-PHYS-INT-2011-085, CERN, Geneva, 2011.
- [150] B. Resende. "Muon identification algorithms in ATLAS". ATLAS Note ATL-PHYS-PROC-2009-113, CERN, Geneva, 2009.
- [151] The ATLAS Collaboration. "ATLAS Monte Carlo tunes for MC09". ATLAS Note ATL-PHYS-PUB-2010-002, CERN, Geneva, 2010.
- [152] The ATLAS Collaboration. "Combination of Higgs Boson Searches with up to 4.9 fb^{-1} of pp Collisions Data Taken at $\sqrt{s} = 7$ TeV with the ATLAS Experiment at the LHC". ATLAS Note ATLAS-COM-CONF-2011-195, CERN, Geneva, 2011.

- [153] The ATLAS Collaboration. “Electron and photon reconstruction and identification in ATLAS: expected performance at high energy and results at 900 GeV”. ATLAS Note ATLAS-CONF-2010-005, CERN, Geneva, 2010.
- [154] The ATLAS Collaboration. “Characterization of Interaction-Point Beam Parameters Using the pp Event-Vertex Distribution Reconstructed in the ATLAS Detector at the LHC”. ATLAS Note ATLAS-CONF-2010-027, CERN, Geneva, 2010.
- [155] The ATLAS Collaboration. “Data-Quality Requirements and Event Cleaning for Jets and Missing Transverse Energy Reconstruction with the ATLAS Detector in Proton-Proton Collisions at a Center-of-Mass Energy of $\sqrt{s} = 7$ TeV”. ATLAS Note ATLAS-CONF-2010-038, CERN, Geneva, 2010.
- [156] The ATLAS Collaboration. “Performance of the missing transverse energy reconstruction in proton-proton collisions at center-of-mass energy of $\sqrt{s} = 7$ TeV with the ATLAS detector”. ATLAS Note ATLAS-CONF-2010-039, CERN, Geneva, 2010.
- [157] The ATLAS Collaboration. “Properties of Jets and Inputs to Jet Reconstruction and Calibration with the ATLAS Detector Using Proton-Proton Collisions at $\sqrt{s} = 7$ TeV”. ATLAS Note ATLAS-CONF-2010-053, CERN, Geneva, 2010.
- [158] The ATLAS Collaboration. “Jet energy resolution and selection efficiency relative to track jets from in-situ techniques with the ATLAS Detector Using Proton-Proton Collisions at a Center of Mass Energy $\sqrt{s} = 7$ TeV”. ATLAS Note ATLAS-CONF-2010-054, CERN, Geneva, 2010.
- [159] The ATLAS Collaboration. “Jet energy scale and its systematic uncertainty for jets produced in proton-proton collisions at $\sqrt{s} = 7$ TeV and measured with the ATLAS detector”. ATLAS Note ATLAS-CONF-2010-056, CERN, Geneva, 2010.
- [160] The ATLAS Collaboration. “Performance of the Missing Transverse Energy Reconstruction and Calibration in Proton-Proton Collisions at a Center-of-Mass Energy of 7 TeV with the ATLAS Detector”. ATLAS Note ATLAS-CONF-2010-057, CERN, Geneva, 2010.
- [161] The ATLAS Collaboration. “Luminosity Determination Using the ATLAS Detector”. Technical Report ATLAS-CONF-2010-060, CERN, Geneva, 2010.
- [162] The ATLAS Collaboration. “Muon Reconstruction Performance”. ATLAS Note ATLAS-CONF-2010-064, CERN, Geneva, 2010.
- [163] The ATLAS Collaboration. “Performance of the ATLAS Jet Trigger in the Early $\sqrt{s} = 7$ TeV Data”. ATLAS Note ATLAS-CONF-2010-094, CERN, Geneva, 2010.
- [164] The ATLAS Collaboration. “Update on the jet energy scale systematic uncertainty for jets produced in proton-proton collisions at $\sqrt{s} = 7$ TeV measured with the ATLAS detector”. ATLAS Note ATLAS-CONF-2011-007, CERN, Geneva, 2011.
- [165] The ATLAS Collaboration. “Updated Luminosity Determination in pp Collisions at $\sqrt{s} = 7$ TeV using the ATLAS Detector”. ATLAS Note ATLAS-CONF-2011-011, CERN, Geneva, 2011.

-
- [166] The ATLAS Collaboration. “A measurement of the ATLAS muon reconstruction and trigger efficiency using J/ψ decays”. ATLAS Note ATLAS-CONF-2011-021, CERN, Geneva, 2011.
- [167] The ATLAS Collaboration. “Jet energy scale and its systematic uncertainty in proton-proton collisions at $\sqrt{s} = 7$ TeV in ATLAS 2010 data”. ATLAS Note ATLAS-CONF-2011-032, CERN, Geneva, 2011.
- [168] The ATLAS Collaboration. “Muon Momentum Resolution in First Pass Reconstruction of pp Collision Data Recorded by ATLAS in 2010”. ATLAS Note ATLAS-CONF-2011-046, CERN, Geneva, 2011.
- [169] The ATLAS Collaboration. “Muon reconstruction efficiency in reprocessed 2010 LHC proton-proton collision data recorded with the ATLAS detector”. ATLAS Note ATLAS-CONF-2011-063, CERN, Geneva, 2011.
- [170] The ATLAS Collaboration. “Luminosity Determination in pp Collisions at $\sqrt{s} = 7$ TeV using the ATLAS Detector in 2011”. Technical Report ATLAS-CONF-2011-116, CERN, Geneva, 2011.
- [171] The ATLAS Collaboration. “Search for neutral MSSM Higgs bosons decaying to $\tau^+\tau^-$ pairs in proton-proton collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector”. ATLAS Note ATLAS-CONF-2011-132, CERN, Geneva, 2011.
- [172] C. Hensel, F. Kohn, M. Hamer, and A. Quadt. “Muon trigger efficiencies with a combined tag & probe and bootstrap method”. ATLAS Note ATL-COM-PHYS-2010-916, CERN, Geneva, 2010.
- [173] S. Mättig. “The Online Luminosity Calculator of ATLAS”. ATLAS Note ATL-DAQ-PROC-2011-009, CERN, Geneva, 2011.
- [174] R. H. L. van der Leeuw, O. Igonkina, and P. de Jong. “Systematic Uncertainties in mSUGRA NLO cross-section calculations at $\sqrt{s} = 7$ TeV”. ATLAS Note ATL-PHYS-INT-2010-086, CERN, Geneva, 2010.
- [175] N. Makovec. “Tight and loose bad jet rejection for release 15 data-set”. Presentation given in the Jet/Etmiss Phone Conference, available online, <https://indico.cern.ch/getFile.py/access?contribId=2&resId=1&materialId=slides&confId=102852>, accessed 01.09.2011.
- [176] M. Nessi. “News from Chamonix Workshop”. Presentation given in the ATLAS Weekly Meeting, available online, <https://indico.cern.ch/getFile.py/access?contribId=1&resId=1&materialId=slides&confId=72798>, last accessed 07.11.2011.
- [177] T. Plehn, private communication.
- [178] M. Hamer, private communication.
- [179] I. Aracena et al. “The ATLAS missing E_T trigger performance with initial LHC runs at $\sqrt{s} = 7$ TeV”. Technical Report ATL-COM-DAQ-2010-191, CERN, Geneva, 2010.
- [180] C. T. Hill and E. H. Simmons. “Strong dynamics and electroweak symmetry breaking”. *Physics Reports*, 381(4–6): 235–402, 2003. doi:10.1016/S0370-1573(03)00140-6.

Bibliography

Web Citations

Internet resources are a vital part of contemporary research. However, web citations are naturally volatile and subject to change. They have thus been avoided whenever possible. For each reference listed here, the date of the information retrieval is specified. Some of the resources are password protected.

- [181] I. Antcheva and O. Couet (editors). "ROOT User's Guide (5.26.00)". Available online, <http://root.cern.ch/drupal/content/users-guide/>, accessed 22.07.2011.
- [182] S. Goldfarb. "Atlas Offline Software Release Status". Website, <http://atlas-computing.web.cern.ch/atlas-computing/projects/releases/status/>, accessed 16.05.2011.
- [183] E. Gallas. ATLAS TWiki: "ConditionsMetadata", <https://twiki.cern.ch/twiki/bin/viewauth/Atlas/ConditionsMetadata>, last accessed 02.11.2011 (restricted access).
- [184] C.-P. Yuan. "CTEQ6 Parton Distribution Home Page". Website, <http://hep.pa.msu.edu/cteq/public/cteq6.html>, last accessed 10.08.2011.
- [185] E. Torrence. "ATLAS Data Quality Summary". Website, <https://atlas.web.cern.ch/Atlas/GROUPS/DATAPREPARATION/DQSummary/>, accessed 25.02.2011 (restricted access).
- [186] ILC Global Design Effort. "ILC — Facts and figures". Website, <http://www.linearcollider.org/about/What-is-the-ILC/Facts-and-figures>, last accessed 05.11.2011.
- [187] ATLAS Luminosity Working Group. "ATLAS Luminosity Calculator". Online tool, <https://atlas-lumicalc.cern.ch/>, last accessed 11.08.2011 (restricted access).
- [188] H. Baer, F. E. Paige, S. D. Protopopescu, and X. Tata. "ISAJET 7.82: A Monte Carlo Event Generator for pp , $p\bar{p}$, and e^+e^- Interactions". Website, <http://www.nhn.ou.edu/~isajet/>, last accessed 11.08.2011.
- [189] T. Carli. ATLAS TWiki: "HowToCleanJets", <https://twiki.cern.ch/twiki/bin/view/AtlasProtected/HowToCleanJets>, last accessed 01.09.2011 (restricted access).
- [190] T. Carli. ATLAS TWiki: "JetEtmissDataAnalysisRecommendationSummer2010", <https://twiki.cern.ch/twiki/bin/viewauth/AtlasProtected/JetEtmissDataAnalysisRecommendationSummer2010>, last accessed 20.10.2011 (restricted access).
- [191] CERN Press Office. "CERN releases analysis of LHC incident". Available online, <http://press.web.cern.ch/press/pressreleases/releases2008/PR14.08E.html>, last accessed 07.11.2011.

- [192] M. Donszelmann. “LHC Shutdown Schedule”. Website, <https://atlas.web.cern.ch/Atlas/Upgrade/Schedule.php>, last accessed 05.11.2011 (restricted access).
- [193] M. Klein. “The LHeC: Deep Inelastic Electron-Nucleon Scattering at the LHC”. Website, <http://www.ep.ph.bham.ac.uk/exp/LHeC/>, last accessed 05.11.2011.
- [194] T. J. Khoo. ATLAS TWiki: “MSSMSquarkGluinoGrid”, <https://twiki.cern.ch/twiki/bin/view/AtlasProtected/MSSMSquarkGluinoGrid>, last accessed 24.11.2011 (restricted access).
- [195] C. Lester. “The Oxbridge MT2 / Stransverse Mass Library”. Available online, <http://www.hep.phy.cam.ac.uk/~lester/mt2/index.html>, last accessed 01.11.2011.
- [196] MCP Convenors. ATLAS TWiki: “ATLAS Muon Combined Performance — Guidelines for Analysis in Release 15”, <https://twiki.cern.ch/twiki/bin/view/AtlasProtected/MCPAnalysisGuidelinesRel15>, last accessed 01.09.2011 (restricted access).
- [197] NIST/SEMATECH. “e-Handbook of Statistical Methods”. Available online, <http://www.itl.nist.gov/div898/handbook/eda/section3/eda366h.htm>, last accessed 14.10.2011.
- [198] T. Plehn. “Prospino webpage”. Available online, <http://www.thphys.uni-heidelberg.de/~plehn/index.php?show=prospino>, last accessed 09.08.2011.
- [199] A. Krasznahorkay. “SourceForge.net: sframe”. Website, <http://sourceforge.net/apps/mediawiki/sframe/>, last accessed 10.11.2010.
- [200] E. Torrence. “ATLAS Data Summary — 2011 Records”. Website, <https://atlas.web.cern.ch/Atlas/GROUPS/DATAPREPARATION/DataSummary/2011/records.html>, last accessed 18.11.2011 (restricted access).
- [201] The ROOT Team. “Reflex”. Website, <http://root.cern.ch/drupal/content/reflex>, last accessed 18.08.2011.
- [202] ATLAS TWiki: “SUSYCrosssections”, <https://twiki.cern.ch/twiki/bin/view/AtlasProtected/SUSYCrosssections>, last accessed 17.08.2011 (restricted access).
- [203] ATLAS TWiki: “TopMC2009”, <https://twiki.cern.ch/twiki/bin/view/AtlasProtected/TopMC2009>, last accessed 19.08.2011 (restricted access).
- [204] ATLAS Trigger Group. “Web Trigger Rates Presenter Archive”. Available online, https://atlas-trigconf.cern.ch/wtrp_archive/index.html, last accessed 21.09.2011 (restricted access).

Acronyms

ALFA	Absolute Luminosity For ATLAS	55
AMI	Atlas Metadata Interface	194
AOD	Analysis Object Data	75
BLOB	Binary Large Object	A-7
BSM	Beyond Standard Model	234
CDM	Cold Dark Matter	16
CERN	Conseil Européen pour la Recherche Nucléaire	39
CLIC	Compact Linear Collider	45
CMS	Compact Muon Solenoid	46
COMA	Conditions Metadata	A-35
COOL	ATLAS Conditions Database	79
CPU	Central Processing Unit	61
CSC	Cathode Strip Chamber	53
CTP	Central Trigger Processor	55
EF	Event Filter	59
EM	electromagnetic	63
ESD	Event Summary Data	A-32
FEX	Feature Extraction	61
GMSB	Gauge-Mediated Supersymmetry Breaking	28
GRL	Good Runs List	77
GUT	Grand Unified Theory	20
HLT	High-Level Trigger	61
ILC	International Linear Collider	45
IOV	Interval Of Validity	A-6
JER	Jet-Energy Resolution	207
JES	Jet-Energy Scale	71
L1	Level 1	59
L2	Level 2	59
LAr	Liquid Argon Gas	50
LB	Luminosity Block	77

Acronyms

LEP	Large Electron-Positron Collider	26
LHC	Large Hadron Collider	29
LLR	Log-Likelihood Ratio	232
LO	Leading Order	81
LSP	Lightest Supersymmetric Particle	25
LUCID	Luminosity measurement Using a Čerenkov Integrating Detector	55
MBTS	Minimum-Bias Trigger Scintillators	54
MC	Monte Carlo	79
MDT	Monitored Drift Tube	53
MSSM	Minimal Supersymmetric Standard Model	20
mSUGRA	Minimal Supergravity	27
NLO	Next-to-Leading Order	81
NLSP	Next-Lightest Supersymmetric Particle	28
NMSSM	Next-to-Minimal Supersymmetric Standard Model	28
NNLO	next-to-next-to-leading order	193
NTUP	ROOT-readable ntuple	A-32
PDF	Parton Distribution Function	192
PLR	Profile Likelihood Ratio	231
QCD	Quantum Chromodynamics	6
QED	Quantum Electrodynamics	6
QFT	Quantum Field Theory	12
RDO	Raw Data Objects	A-32
RGE	Renormalization Group Equation	28
RMS	Root Mean Square	63
RoI	Region of Interest	61
ROB	Read-Out Buffer	61
RPC	Resistive Plate Chamber	54
RPV	R-Parity Violating	28
SCT	Silicon Microstrip Tracker	48
SUSY	Supersymmetry	17
TDAQ	Trigger and Data-Acquisition System	83
TDR	Technical Design Report	165
TGC	Thin Gap Chamber	54
TRT	Transition Radiation Tracker	48
UED	Universal Extra Dimensions	18
WIMP	Weakly Interacting Massive Particle	17

XML	Extensible Markup Language.....	78
ZDC	Zero-Degree Calorimeter.....	55

Further explanations of acronyms can be found in the appendix of [127, 128].

Index

- Accelerator chain, 41
- ALFA, 55
- ALICE, 46
- Anti- k_t algorithm, 72
- Asymptotic freedom, 13
- ATHENA framework, A-31
- ATLAS, 47

- Background normalization, 192
- Barn, 29
- Bayesian statistics, 224
- β -distribution, 97
- Bootstrapping, 99
- Bootstrapping condition, 99
- Bunch crossing, 38
- Bunch train, 42
- ByteStream format, A-33

- Calorimeter, 32
- Central Trigger Processor, 59
- Chiral supermultiplet, 22
- CL_s method, 226
- CMS, 46
- Colliding bunch, 42
- Collinear safe, 72
- Combined trigger, 65
- Compact Linear Collider, 45
- Control region, 179
- Coordinate system, 47
- Cutflow, 182

- Dark energy, 16
- Dark matter, 16
- Data formats, A-32
- Data stream, 61
- Data-taking period, 77
- Dataset nomenclature, A-33
- Dead material, 35, 71
- Dead time, 59
- ΔR , 47

- Dirac equation, 5
- Discovery potential, 224
- Dummy vertex, 154

- Early-rejection principle, 61
- Effective mass, 76
- Egamma stream, 62
- Electromagnetic calorimeter, 50
- EM+JES calibration, 71
- Empty bunch, 42
- Enhanced bias runs, 84
- Event, A-30
- Event Builder, 61
- Event Filter, 61
- Event generator, 79
- Event reconstruction, 66
- Event record, A-30
- Exclusion potential, 224
- Extended technicolor, 17
- Extra dimensions, 18

- Fake \cancel{E}_T , 35
- Fake rate, 67
- Feature-extraction algorithm, 61
- Filled bunch, 42
- Filter efficiency, 80
- Fine-tuning problem, 16
- Forced accept, 59
- Forward calorimeter, 51
- Frequentist statistics, 224
- Full-chain jet + \cancel{E}_T trigger, 116

- Gauge supermultiplet, 22
- Goldstone boson, 8
- Good Runs List, 77
- Gravity supermultiplet, 22

- Hadronic calorimeter, 51
- Hadronic jet, 33
- Heavy quark, 13

- Hidden sector, 26
- Hierarchy problem, 16
- Higgs mechanism, 7
- High Level Trigger, 61
- Hypothesis algorithm, 61

- In-time pile-up, 38
- Incident of 2008, 42
- Infrared safe, 72
- Inner Detector, 48
- Instantaneous luminosity, 36
- Integrated luminosity, 36
- International Linear Collider, 45
- Invariant mass, 29
- Invisible particle, 34

- Jet cleaning, 184
- Jet-energy resolution, 72, 207
- Jet-energy scale, 71
- Jet-energy scale uncertainty, 206
- Jet-seeded jet + \cancel{E}_T trigger, 116
- JetTauEtmiss stream, 62

- Kaluza-Klein theory, 18
- Klein-Gordon equation, 4

- Lagrangian density, 4
- Lagrangian formalism, 3
- Large Hadron Collider, 40
- LAr hole, 52
- Lepton veto, 188
- Level 1, 59
- Level 2, 61
- LHC time bucket, 41
- LHCB, 46
- LHCF, 46
- Light quark, 13
- Likelihood function, 225
- Local gauge invariance, 5
- Logarithmic reweighting, 102
- LUCID, 55
- Luminosity, 36
- Luminosity block, 77
- Luminosity measurement, 78

- MET_EMJES_RefFinal_CellOutEM, 75
- MET_LocHadTopo, 75
- MET_Simplified20_RefFinal, 75

- MET_Topo, 74
- MinBias stream, 62
- Minimal supergravity model, 27
- Minimal Supersymmetric Standard Model, 22
- Minimum-bias event, 82
- Minimum-bias trigger, 65, 98
- Minimum-bias trigger scintillators, 54
- Missing transverse energy, 35
- Missing transverse momentum, 35
- MOEDAL, 47
- Monte Carlo counting, 97
- Monte Carlo method, 79
- Monte Carlo tune, 192
- MSSM grid, 201
- mSUGRA grid, 200
- Muon spectrometer, 36
- Muons stream, 62

- N-tuple, A-32
- Natural units, 29
- Noether's Theorem, 5
- Nuisance parameter, 231
- Null hypothesis, 223

- Object identification, 185
- One-sigma band, 228
- Orthogonal trigger, 98
- Out-of-time pile-up, 38
- Overlap removal, 187

- p -value, 224
- Parameter of interest, 231
- Particle Content
 - MSSM, 22
 - Standard Model, 9
- Parton distribution function, 13
- Pass-through factor, 58
- Pile-up, 37
- Pile-up uncertainty, 208
- Plateau efficiency, 93
- Plateau region, 93
- PLR method, 231
- Poincaré group, 5
- POOL, A-33
- Prescale factor, 58
- Primary trigger, 59
- Proca equation, 5

-
- Profile likelihood ratio, 231
Pseudorapidity, 47
Purity, 67
- QCD normalization, 195
Quark flavor, 9
- R -parity, 25
 R -parity violation, 28
Random trigger, 65, 98
Rapidity, 47
Real \cancel{E}_T , 35
Reconstruction efficiency, 67
ROOT, A-29
Running (of coupling constants), 11
- Secondary trigger, 59
Signal efficiency, 88
Signal region, 178, 189
Signal strength, 231
Signal yield, 89
Signal-to-background ratio, 190
Singlino, 28
Solenoid magnet, 50
Solid state detector, 31
Spontaneous symmetry-breaking, 7
Staco muons, 69
Statistical uncertainty, 206
Stransverse mass, 76
Strong interaction, 12
SU4 benchmark point, 88
Sudakov form factor, 81
Super Large Hadron Collider, 45
Super-Higgs mechanism, 27
Superfield, 22
Supergravity, 27
Superpotential, 24
Superspace, 22
Supersymmetry, 19
Supersymmetry breaking, 25
Systematic uncertainties, 205
- Technicolor, 17
Theoretical uncertainty, 208
Time bucket, 41
Toroid magnet, 53
TOTEM, 47
Tracking device, 30
- Transparent mode, 58
Transverse energy, 29
Transverse mass, 76
Transverse momentum, 29
Trigger chain, 61
Trigger menu, 59
Trigger nomenclature, 66
Trigger requirements, 57
Trigger steering, 61
True \cancel{E}_T , 35
Truth information, 80
TTree, A-30
Turn-on curve, 93
Turn-on region, 93
Type I error, 223
Type II error, 224
- Underlying event, 82
Unique rate, 59
Universal extra dimensions, 18
- Vector supermultiplet, 22
Vertex merging, 157
Very Large Hadron Collider, 45
- Weak interaction, 12
Weinberg-Salam model, 8
- XS trigger, 173
- ZDC, 55