

JOHANNES GUTENBERG-UNIVERSITÄT MAINZ

Search for Dijet Resonances with the Level-1 Topological Processor at ATLAS

DISSERTATION

Zur Erlangung des Grades
DOKTOR DER NATURWISSENSCHAFTEN

vorgelegt von

Johannes Frederic Damp,
geboren in Kirchheimbolanden

am Fachbereich Physik, Mathematik und Informatik
der Johannes Gutenberg-Universität
in Mainz



Mainz, 4. September 2020

Johannes Frederic Damp

Search for Dijet Resonances with the Level-1 Topological Processor at ATLAS

Gutachter:

Erster Gutachter: Entfernt aus datenschutzrechtlichen Gründen.

Zweiter Gutachter: Entfernt aus datenschutzrechtlichen Gründen.

Datum der mündlichen Prüfung: 20. Januar 2021

Abstract

Due to the limited bandwidth capacity of data acquisition systems, any analysis performed at a collider experiment has to rely on an efficient trigger system, which selects only the potentially interesting events. To cope with increasing data rates caused by higher instantaneous luminosities of up to $2.1 \cdot 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ in Run-2 of the LHC, trigger thresholds would have had to be increased, leading to both losses in signal efficiency and statistical limitations for a wide range of analyses. To avoid these losses, the FPGA-based Level-1 Topological Processor (L1Topo) has been introduced to ATLAS in Run-2. It extends the capabilities of the hardware-based first-level trigger by the ability to perform trigger decisions based on topological algorithms, for example angular selections or mass cuts. Using L1Topo, it was possible to maintain the trigger thresholds for Run-2 and to avoid losses in signal efficiencies while keeping the data rates low, which is of paramount importance for analyses targeting low- p_T signatures. This thesis reports on the firmware development and trigger performance of the topological processor during Run-2.

In addition, this thesis presents a novel approach to search for new physics in the dijet invariant mass spectrum. These resonance searches are performed by searching for a localized excess of events above a smoothly falling background of dijet events. Due to the limited available bandwidth of the data acquisition system, searches for dijet resonances suffer from statistical limitations in the mass range below 500 GeV. Using L1Topo, these statistical limitations can be fully avoided. By performing the analysis and creating invariant mass histograms directly on the first trigger level, the bandwidth limitations can be circumvented and every single collision event can be analyzed. This makes it possible to perform an analysis of the whole invariant mass spectrum using an unprecedented level of statistical precision. This thesis presents the development of the new approach and its implementation into the trigger and data acquisition system of ATLAS. Using this approach, data has been collected in 2018 with $\sqrt{s} = 13 \text{ TeV}$ in proton-proton collisions corresponding to a recorded offline luminosity of 60.6 fb^{-1} . A statistical analysis is performed to search for deviations from a data-driven background estimate. No evidence for new resonances is observed, the data is therefore used to set exclusion limits at a credibility level of 95 % on the product of the production cross section, the acceptance and the branching ratio, both for model-agnostic Gaussian resonance shapes and an axial-vector Z' dark matter mediator.

Kurzzusammenfassung

Aufgrund der begrenzten verfügbaren Bandbreite der Datenerfassungssysteme ist jedes Beschleunigerexperiment auf ein effizientes Triggersystem angewiesen, um nur die potenziell interessanten Ereignisse auszuwählen. Um mit den steigenden Datenraten, die durch höhere instantane Luminositäten von bis zu $2,1 \cdot 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ in Run-2 des LHCs verursacht werden, fertig zu werden, hätten die Triggerschwellen erhöht werden müssen, was zu Verlusten in der Signaleffizienz sowie zu statistischen Einschränkungen für eine Vielzahl von Analysen geführt hätte. Um diese Verluste zu vermeiden, wurde der Level-1 topologische Prozessor (L1Topo) für Run-2 in ATLAS eingeführt. Er erweitert die Fähigkeiten des hardwarebasierten Level-1-Triggers und ermöglicht es Triggerentscheidungen mit topologischen Algorithmen durchzuführen, zum Beispiel Winkel- oder Massenschnitte. Dies ermöglichte es, Verluste in Signaleffizienzen zu vermeiden und die Triggerschwellen in Run-2 niedrig zu halten, was von größter Wichtigkeit für zahlreiche Analysen bei ATLAS ist. In dieser Arbeit wird die Firmware-Entwicklung und die Leistungsfähigkeit der Trigger des topologischen Prozessors während Run-2 beschrieben.

Darüber hinaus stellt diese Arbeit einen neuartigen Ansatz zur Suche nach neuer Physik im Spektrum der invarianten Dijetmassen vor. Analysen von Dijetmassenspektren werden durch die Suche nach einem lokalisierten Überschuss von Ereignissen über einem glatt fallenden Dijethintergrund durchgeführt. Aufgrund der begrenzten verfügbaren Bandbreite des Datenerfassungssystems haben Dijetresonanzsuchen jedoch statistische Einschränkungen im Massenbereich unter 500 GeV. Mit Hilfe des L1Topo-Moduls können diese statistischen Einschränkungen jedoch umgangen werden. Mittels der Durchführung der Analyse und der Erstellung invarianter Massenhistogramme direkt auf der ersten Triggerstufe können die Bandbreitenbeschränkungen vermieden werden, und jedes einzelne Kollisionsereignis kann für die Analyse verwendet werden. Dies ermöglicht eine Analyse des gesamten invarianten Massenspektrums unter Verwendung einer noch nie dagewesenen statistischen Präzision. Mit diesem Ansatz wurden im Jahr 2018 Daten mit $\sqrt{s} = 13 \text{ TeV}$ in Proton-Proton-Kollisionen gesammelt, die einer aufgezeichneten Offlineluminosität von $60,6 \text{ fb}^{-1}$ entsprechen. Diese Arbeit stellt die Entwicklung des neuen Ansatzes und seine Implementierung in das Trigger- und Datenerfassungssystem von ATLAS vor. Es wird eine statistische Analyse durchgeführt, um nach Abweichungen von einer datengetriebenen Hintergrundschätzung zu suchen. In dieser Analyse werden keine Hinweise auf neue Resonanzen beobachtet. Daher werden die Daten dafür verwendet, um Ausschlussgrenzen mit einem Glaubwürdigkeitsniveau von 95 % auf das Produkt aus Wirkungsquerschnitt, Akzeptanz und Verzweigungsverhältnis zu setzen. Dies geschieht sowohl für modellunabhängige Gauß'sche Resonanzen, als auch für ein Modell der dunklen Materie, welches ein Axialvektor-Mediatorteilchen Z' vorhersagt.

Contributions of the Author

In a large collaboration like ATLAS, no result is the work of a single individual. Only by combining the efforts of thousands of scientists it is possible to not only record data from proton-proton collisions, but also to perform precision measurements and searches for new particles using this data. The central part of this thesis is focused on my own work, however this work was only possible due to contributions from numerous collaborators within ATLAS. An inclusion of their contributions is necessary in order to illustrate and understand the contents of this thesis, therefore my own contributions are stated explicitly in the following. Joining the L1Topo group in 2017, I made significant contributions to the algorithmic firmware and trigger algorithm commissioning for most of Run-2. My work made it possible to run L1Topo triggers in these years, which are of paramount importance for many analyses. In addition, I am contributing firmware and software to the upgrade of the topological processor for Run-3 of the LHC. In parallel, I focused on the development on a novel method to search for dijet resonances at low invariant masses. The development of the histogramming method was done as a joint work between Sebastian Artz [1] and me, both contributing in equal terms to firmware and software development. The final statistical analysis was performed by myself. My contributions to each chapter are explicitly listed in the following:

- Ch. 8: I contributed to the L1Topo firmware by implementing all algorithms for the trigger menus for the years 2017 and 2018 and by building the production firmware for data taking. In addition, I contributed to the commissioning and validation of those L1Topo triggers.
- Ch. 9: I was involved in the development of the general concept and the design of the various implemented histograms. I contributed to the histogramming firmware development by implementing the histogram category assignment as well as the calculation of the topological quantities. In addition, I was the main author of the standalone software package, which was later ported to the official L1Calo software. I was involved in the tests in the test rig as well as in the hot tower test runs with ATLAS.
- Ch. 10: I wrote most of the code responsible to apply data quality criteria and to select histograms, including the luminosity calculation, the implementation of filters and the framework for the graphical plotting of histograms. In addition, I derived the subtraction procedure.
- Ch. 11: The derivation of the mass histograms, the derivation of the binning and the combination of MC samples was performed by myself.
- Ch. 12: All studies on the properties of level-1 jets were done by myself, and all code to perform these studies was written by myself.
- Ch. 13: All shown simulation studies were performed by myself. For this, I implemented a bitwise software version of the histogramming running on top of the official simulation. In addition, I generated and implemented all aspects of the fast simulation-based approach.
- Ch. 14: I performed the background-only fit and limit setting. For this purpose, I adapted the
- Ch. 15 software framework developed by ATLAS for dijet resonance searches to the new data.

Contents

Abstract	i
Contributions of the Author	iii
List of Figures	x
List of Tables	xiv
Acronyms	xv
Glossary	xv
1. Introduction	1
1. Theoretical Foundations	3
2. Particles and Fundamental Interactions	5
2.1. The Standard Model of Particle Physics	5
2.1.1. Overview of the Standard Model	5
2.1.2. Electroweak Interaction	6
2.1.3. Spontaneous Symmetry Breaking	8
2.1.4. Quantum Chromodynamics	9
2.2. Beyond the Standard Model	11
2.2.1. Limitations of the Standard Model	11
2.2.2. Dark Matter	12
2.2.3. Theoretical Models for Dark Matter	14
2.2.4. Z' Dark Matter Mediator	15
3. Dijet Production in Proton-Proton Collisions	19
3.1. Proton-Proton Collisions	19
3.1.1. Introduction and Terminology	19
3.1.2. Structure of the Proton	20
3.1.3. The Factorization Theorem	22
3.1.4. Formation of Jets	23
3.2. Properties of Dijet Events	23
3.2.1. Kinematics of Two-Parton Scattering	24
3.2.2. Cross Section of Dijet Events	25
3.3. Monte Carlo Event Generation	26
3.3.1. Madgraph	28
3.3.2. Pythia	28

II. The Experiment	29
4. Technological Foundations	31
4.1. Introduction	31
4.2. Setup of an FPGA	32
4.2.1. Configurable Logic Blocks	32
4.2.2. DSP Slices	33
4.2.3. Block RAM	34
4.2.4. Multi Gigabit Transceivers	34
4.3. The IPBus Protocol	35
4.4. Firmware Development	36
4.5. Xilinx Vivado	37
4.5.1. Synthesis and Implementation	37
4.5.2. FPGA Design Characteristics: Resource Utilization and Timing	39
4.5.3. IPCores	39
4.5.4. Waveform Simulation	40
5. CERN and the LHC	41
5.1. The CERN Accelerator Complex	41
5.2. The Experiments at the LHC	42
5.3. LHC Performance	43
5.4. LHC Bunch Structure and Filling Schemes	45
6. The ATLAS Experiment	47
6.1. General Detector Overview	47
6.1.1. Coordinate System and Kinematic Variables	47
6.1.2. Particle Measurements with ATLAS	49
6.2. The Inner Detector	49
6.3. The Calorimeter System	50
6.3.1. The Electromagnetic Calorimeter	51
6.3.2. The Hadronic Calorimeter	52
6.4. The Muon System	54
6.5. Luminosity Detectors and Luminosity Measurement	55
6.6. The Trigger and Data Acquisition System	56
6.6.1. The Level-1 Trigger	57
6.6.2. The High Level Trigger	61
6.7. Detector Simulation	61
7. Jet Reconstruction in ATLAS	63
7.1. Jet Reconstruction in the Level-1 Trigger	63
7.2. Offline Jet Reconstruction	67
7.2.1. Tracks and Vertices	67
7.2.2. Jets	68
7.2.3. Jet Calibration	69
7.2.4. Jet Cleaning	73

III. The Level-1 Topological Processor at ATLAS	75
8. The Level-1 Topological Processor	77
8.1. Hardware Overview	77
8.2. Firmware Overview	78
8.2.1. The Real-Time Data Path	78
8.2.2. The Control Path and Infrastructure	79
8.2.3. The Readout Path	81
8.3. The Algorithmic Firmware	81
8.3.1. Data Formats	83
8.3.2. Sort and Select Algorithms	84
8.3.3. Topological Decision Algorithms	87
8.3.4. The Invariant Mass Algorithm	88
8.4. The Trigger Menu and Firmware Implementation	90
8.5. Commissioning, Simulation and Validation	90
8.6. Trigger Performance	92
IV. The Analysis	95
9. L1Topo Histogramming	97
9.1. Motivation for a Triggerless Analysis	97
9.2. The Histogramming Concept	99
9.2.1. Kinematic Selection and Histograms	101
9.3. Firmware Implementation	103
9.3.1. Histograms	104
9.3.2. Resource Utilization	105
9.4. Software Implementation	106
9.4.1. Integration to the L1Calo Online Software	107
9.4.2. Readout Interval	107
9.5. Histogramming Tests	108
9.5.1. Standalone Tests	109
9.5.2. Tests in the Test Rig	109
9.5.3. Tests in ATLAS	109
10. Towards a Histogramming Data Analysis	113
10.1. Data Quality Checks	113
10.1.1. Good Run Lists in ATLAS	113
10.1.2. The Coma Database	114
10.1.3. Luminosity Block Association and Luminosity Calculation	115
10.1.4. Histogram Filters and Histogram Sums	117
10.2. Recorded Data	117
10.2.1. Beam Splashes	117
10.2.2. Cosmic Data	118
10.2.3. Physics Data	121
10.2.4. Physics Data: The Invariant Mass Categories	123
10.2.5. Physics Data: Luminosity and Pile-Up Profile	125

10.3. Subtraction Procedure	126
11. The Invariant Mass Spectrum	129
11.1. The Spectrum in Data	129
11.2. Monte Carlo Samples	130
11.2.1. Signal Samples	130
11.2.2. Background Samples	130
11.3. Analysis Binning	131
12. Level-1 Jet Properties	135
12.1. Spatial Resolution	135
12.2. Jet Energy Scale and Resolution	136
12.2.1. Jet Energy Calibration	137
12.2.2. Effect of the Jet Energy Calibration	141
12.2.3. Jet Energy Scale Uncertainty	142
12.3. Jet Energy Resolution	143
12.4. Reconstruction Efficiency	144
12.5. Effect of the Jet Cleaning on L1 Jets	146
12.6. Pile-up in Jets	147
13. Monte Carlo Approach	149
13.1. Full Simulation	149
13.1.1. Full Simulation Data Comparison	150
13.2. Sliding Window Algorithm Simulation	152
13.2.1. Sample Generation	153
13.2.2. Jet Elements and Noise Cuts	153
13.2.3. The Sliding Window Algorithm	154
13.2.4. Energy Scale Correction	155
13.2.5. Simulation of Pileup	158
13.2.6. Study of Response Unsmoothness	160
13.2.7. The Invariant Mass Categories Revisited	161
14. The Background-Only Fit	163
14.1. Statistics	163
14.1.1. The Likelihood	163
14.1.2. The Background Estimate	164
14.1.3. The Sliding Window Fit (SWiFit)	164
14.1.4. Test Statistics	165
14.1.5. Uncertainties on the Background Estimate	167
14.1.6. Potential Discovery and Systematic Uncertainties	168
14.2. Performing the Background Estimate	168
14.3. Invariant Mass in Bins of $\langle\mu\rangle$	169
14.4. Setup of the Background-only Fit	173
14.4.1. Selection of the Fit Range	174
14.4.2. Selection of the Window Width	176
14.4.3. Summary of the Fit Strategy	177
14.5. Results: Performing the Fit to Data	179

15. Setting Exclusion Limits	181
15.1. Bayesian Limits Framework	181
15.1.1. Theory of Bayesian Limit Setting	181
15.1.2. Choice of the Priors	182
15.1.3. Treatment of Systematics	183
15.1.4. Implementation of the Limit Setting	185
15.2. Results: Setting Limits on Data	185
15.2.1. Generic Limits	186
15.2.2. Model-dependent Limits	187
15.3. Results and Comparison to other Searches	188
16. Conclusion and Future Plans	191
16.1. Summary and Conclusion	191
16.2. Possible Improvements and Outlook for Run-3	192
A. Experiment and Object Reconstruction	195
A.1. FPGA Resources	195
A.2. Object Reconstruction	195
A.3. L1Topo Firmware: Data Types	197
A.4. Firmware Algorithms	198
A.5. Jet Cleaning	200
B. L1Topo Histogramming	201
B.1. Firmware Implementation	201
B.2. Standalone Software	204
B.3. Additional Histogramming Tests	205
C. Precision of the Invariant Mass Algorithm	207
D. Statistical Analysis	209
D.1. Fit Uncertainty Calculation	209
D.2. Estimated Level-1 E_T	210
D.3. Definition of the Bins	211
D.4. Reweighting Procedure	212
D.5. Background Estimate	213
D.6. Markov Chain Monte Carlo	213
Bibliography	215
Acknowledgements	229
Curriculum Vitae	231

List of Figures

2.1.	Gauge group of the Standard Model	6
2.2.	Particle content of the Standard Model	7
2.3.	Overview of different measurements of $\alpha_S(Q^2)$	10
2.4.	Rotation curves of NGC 6503	13
2.5.	Distribution of matter in the bullet cluster	14
2.6.	Temperature anisotropy measurement of the CMB by the Planck satellite	15
2.7.	Feynman diagrams for processes involving dark matter	16
2.8.	Combined 95 % constraints on M_R and m_χ	17
3.1.	Illustration of all involved stages in a proton-proton collision	20
3.2.	Depiction of the DGLAP equations	21
3.2.	Depiction of the NNPDF3.1 NNLO PDF set	22
3.3.	Graphical representation of the factorization theorem	23
3.4.	Schematic illustration of a two-to-two scattering process	24
3.5.	Diagrams contributing to production of dijet final states	25
3.6.	Leading order matrix elements for two-to-two massless parton scattering	26
3.7.	ATLAS measurements of the dijet cross section as a function of m_{jj} and y^*	27
4.1.	Simplified FPGA architecture	32
4.2.	Six input LUT in Xilinx FPGAs	33
4.3.	Schematic of a LUT connected to I/O and storage elements	33
4.4.	Simplified architecture of a DSP	34
4.5.	Data flows for a true dual-port BRAM	35
4.6.	Median write/read latency and throughput for IPBus requests	36
4.7.	Comparison of the elaborated design and synthesized design for a 2 bit adder	38
4.8.	Example of a Vivado simulation with Waveform window	40
5.1.	Overview of the LHC accelerator	42
5.2.	Mean number of interactions per bunch crossing in the ATLAS detector	44
5.3.	Instantaneous and integrated luminosity at the LHC	45
5.4.	Bunch group set 1991 from run 336548	46
6.1.	Sectional view of the ATLAS detector	48
6.2.	ATLAS inner detector	50
6.3.	ATLAS calorimeter system	51
6.4.	ATLAS electromagnetic calorimeter barrel module	52
6.5.	Development of a hadronic shower	53
6.6.	Module of the Tile calorimeter	54
6.7.	ATLAS muon system	55
6.8.	Standard Model total production cross section measurements	57
6.9.	ATLAS TDAQ overview	58
6.10.	Trigger tower granularity	59

6.11.	Calorimeter coverage of the Jet/Energy Modules	59
6.12.	L1 topological quantities	60
7.1.	Calorimeter pulse shapes at ATLAS	64
7.2.	Mean shift of the signal pedestal at the PPMs of ATLAS	65
7.3.	Noise thresholds of the JEP energy LUT	66
7.4.	Trigger tower combination into jet elements	67
7.5.	Local maximum finding of the L1 trigger jet algorithm	68
7.6.	L1 jet trigger algorithm window sizes	68
7.7.	Overview of the ATLAS jet calibration scheme	69
7.8.	Average jet energy response of uncalibrated offline jets	71
7.9.	Combined jet energy scale uncertainty at ATLAS	72
8.1.	Run-2 L1Topo system inside the ATCA crate	78
8.2.	Picture of an unplugged L1Topo module	78
8.3.	Schematic view of the L1Topo RTDP firmware	79
8.4.	Depiction of the IPBus Bridge	80
8.5.	Structure of the algorithmic firmware of L1Topo	82
8.6.	Jet TOB data format from the CMX	83
8.7.	Sorting matrix for 6 input channels	85
8.8.	Sorting matrix for 6 input channels with presorting	85
8.9.	Illustration of the L1Topo select block	86
8.10.	Scheme of the L1Topo jet sorting algorithm	86
8.11.	Schematic design created by Vivado of the invariant mass algorithm	89
8.12.	L1Topo firmware resource utilization	90
8.13.	Photograph of the L1Calo test rig	91
8.14.	Photograph of L1Topo in the test rig	91
8.15.	Mismatch rates of L1Topo hardware and simulation decisions	92
8.16.	L1Topo $H \rightarrow \tau\tau$ trigger rates and efficiency	93
8.17.	Rate reduction and signal efficiency of L1Topo dimuon triggers	94
8.18.	L1Topo SimpleCone trigger efficiencies	94
9.1.	Evolution of prescales of the L1_J25 trigger	97
9.2.	Invariant mass spectrum of the offline analysis	98
9.3.	Collider limits on the Z' mediator coupling as a function of mass	98
9.4.	Summary plot of ATLAS limits for the dark matter mediator	99
9.5.	L1Topo histogramming concept	100
9.6.	Sketch of the RTDP and histogramming firmware	104
9.7.	Resource utilization of the histogramming firmware	106
9.8.	Arrangement of lumiblocks for run 357962	108
9.9.	Exemplary overview of histograms in the Online Histogramming browser	110
9.10.	Histogramming test results for single-jet events	111
9.11.	E_T^{miss} spectrum recorded in run 336506	112
10.1.	Luminosity calculation scheme	116
10.2.	Exemplary plots created using histogram filters and COMA information	118
10.3.	$\eta - \phi$ histograms recorded in events with beam splashes	119
10.4.	L1Topo histograms from cosmic runs	120

10.5.	Luminosity and pile-up as a function of the histogram number for a single run	121
10.6.	L1Topo histograms from physics runs	122
10.7.	Invariant mass for central jets with $ y^* < 0.3$	123
10.8.	Data collected in a hot tower run to debug the invariant mass categories	124
10.9.	Comparison of COMA and offline luminosity	125
10.10.	Comparison of COMA and offline pile-up profile	125
10.11.	Example of the subtraction procedure for an E_T histogram	127
10.12.	Comparison of the Tile calibration effect with a possible signal	127
11.1.	Total invariant mass spectrum of 2018 histograms for the analysis	129
11.2.	Correlation of cross section and coupling to quarks for Z' samples	130
11.3.	Truth jet p_T distribution for full simulation QCD samples	132
11.4.	Total invariant mass spectrum for QCD full simulation samples	132
11.5.	Exemplary fit to determine σ/μ for $m_{Z'} = 350$ GeV	133
11.6.	Fit of the relative mass resolution	133
11.7.	Total invariant mass spectrum with the derived binning	134
12.1.	Spatial distances between level-1 jets and matched reconstructed offline jets	136
12.2.	Energy correlation between level-1 jets and matched offline jets	137
12.3.	Measured energy response of level-1 jets	138
12.4.	Energy response for L1 jets in two exemplary bins of E_T^{L1} and $ \eta^{L1} $	139
12.5.	Exemplary fits of the response as a function of the estimated L1 jet E_T	140
12.6.	L1 jet calibration factors as a function of E_T and $ \eta^{L1} $	141
12.7.	Mass scale of dijet pairs at L1	142
12.8.	Effect of a jet calibration on the mass spectrum at level-1	143
12.9.	Comparison of calibrated and uncalibrated spectrum of minimum bias dijet events	144
12.10.	Comparison of various calibrated spectra of minimum bias dijet events	144
12.11.	Relative energy resolution of level-1 jets	145
12.12.	Reconstruction efficiency of level-1 jets	146
12.13.	Jet cleaning efficiency for calibrated anti- k_T $R = 0.4$ jets	147
12.14.	Jet cleaning efficiency for level-1 jets	147
12.15.	Energy ratio $E_T^{4 \times 4} / E_T^{8 \times 8}$ in bins of $\langle \mu \rangle$	148
12.16.	Mean number of jets in bins of $\langle \mu \rangle$	148
12.17.	Response in bins of the ratio $E_T^{4 \times 4} / E_T^{8 \times 8}$	148
13.1.	Comparison of L1Topo and full simulation spectrum	150
13.2.	Energy spectrum using reconstructed jets	151
13.3.	Comparison of level-1 mass histograms from L1Topo, from recorded data and simulation	151
13.4.	Comparison of kinematic quantities between QCD full simulation and histogramming	152
13.5.	Madgraph+Pythia spectrum in different slices	154
13.6.	Example of a sliding window algorithm event	155
13.7.	Comparison of the energy and mass spectrum between data and SLW simulation	156
13.8.	Fit of the quadratic resolution difference	157
13.9.	Energy response distribution at the different correction stages	157
13.10.	Invariant mass spectrum from Delphes with energy scale correction	158

13.11. Mean energy density of pile-up events as used in Delphes	159
13.12. Comparison of distributions with and without pile-up	159
13.13. Response calibration jumps and their effect on the mass spectrum	161
13.14. Polynomial fit of the induced peak height and unsmoothness	161
13.15. Invariant mass categories in the SLW simulation	162
14.1. Sketch of the functionality of the SWiFt procedure	165
14.2. Application of the SWiFt approach to a single histogram	169
14.3. Invariant mass spectra in all of 2018 in $\langle\mu\rangle$ bins	170
14.4. Invariant mass spectra in $\langle\mu\rangle$ bins, reweighted to the same luminosity	170
14.5. Invariant mass spectra for a Z' signal in $\langle\mu\rangle$ bins	171
14.6. Z' mass spectra in $\langle\mu\rangle$ bins, reweighted to the same luminosity	171
14.7. Simplified depiction of different combination of jets from pile-up events	171
14.8. Examples of fits of the $\langle\mu\rangle$ dependency	173
14.9. Extrapolated histogram compared to the original data spectrum	174
14.10. Result of a turn-on fit to the low-mass region	175
14.11. χ^2 p -values as a function of the uppermost bin	175
14.12. SWiFt fit with window half-widths of 5 bins and 14 bins	176
14.13. χ^2 p -value as a function of the window half-width	177
14.14. BumpHunter p -value as a function of the window half-width	177
14.15. SWiFt fit for different window widths with injected Gaussian signals	178
14.16. Fit result of the $\langle\mu\rangle$ -extrapolated spectrum	179
14.17. Distribution of the χ^2 test statistic	180
14.18. Distribution of the BumpHunter test statistic	180
14.19. Relative uncertainties of the background estimate	180
15.1. 95 % upper limits on $\sigma \times A \times \text{BR}$ for Gaussian signals	186
15.2. Comparison of 95 % upper limits for various signal widths	187
15.3. Marginalized posterior distribution for the signal scaling parameter	188
15.4. Comparison of prior and posterior for the function choice uncertainty and the luminosity uncertainty	189
15.5. Upper limits on $\sigma \times A \times \text{BR}$ for Z' samples	190
15.6. Comparison of L1Topo and ATLAS limits	190
16.1. Fit result of the invariant mass spectrum in the signal region in simulation	192
16.2. Limits summary plot with potential limits from the signal region	193
16.3. Jets reconstructed on jFEX	194
A.1. Depiction of the level-1 electron/photon/tau algorithms	196
A.2. Muon trigger algorithms	196
A.3. Depiction of the SimpleCone algorithm	199
A.4. Depiction of an energy sum tree	199
B.1. Schematic depiction of the histogramming top-level module	201
B.2. BRAM configurations for the histogramming firmware	204
B.3. Simulation of the bin mapping functionality	205
B.4. Simulation of the histogramming with IPBus	206

B.5.	Measurement of the latency of histogram read requests	206
C.1.	L1Topo LUT rounding errors	207
C.2.	Box plot of the invariant mass errors due to rounding	208
D.1.	Response distribution in a different binning	210
D.2.	Fits of the response as a function of the mean reconstructed jet energy	210
D.3.	Fit result using the alternate fitting function	213
D.4.	Depiction of the Metropolis algorithm used in MCMC	214

List of Tables

5.1.	Overview of LHC parameters	43
8.1.	Resource utilization of the JetSort algorithm	87
8.2.	Examples for topological algorithms performed on L1Topo	87
9.1.	Table of the histogram categories in y^* and $ \eta_{1,2} $	102
9.2.	Overview of the implemented invariant mass histogram bins and bin widths	102
9.3.	Table of implemented control histograms	103
9.4.	Resources utilized by the histogramming firmware	106
10.1.	Excerpt from the COMA table of all ATLAS runs	114
10.2.	Excerpt from the COMA table on lumiblock start times, end times and durations	115
10.3.	Excerpt from the COMA table on lumiblock-level luminosity and $\langle\mu\rangle$	115
11.1.	Overview of QCD full simulation samples	131
11.2.	Final binning derived for the analysis	134
13.1.	Slice definitions for samples generated with Madgraph	153
A.1.	Resources of the XC7VX690TFFG1927--3 FPGA	195
A.2.	Jet η lookup table	197
B.1.	List of IPBus registers to configure/check various aspects of the histogramming	203
D.1.	Definition of the $\langle\mu\rangle$ bins used for the pile-up extrapolation procedure	211

Acronyms

ADC analog-to-digital converter 64	LHC Large Hadron Collider 41
DAQ Data Acquisition 56	MC Monte Carlo 26
ECAL Electromagnetic Calorimeter 49	PDF Parton Distribution Function 21
FIFO First In, First Out 81	PMT photomultiplier tube 53
FSR Final State Radiation 19	QCD Quantum Chromodynamics 9
GUI Graphical User Interface 37	RAM Random-Access Memory 33
HCAL Hadronic Calorimeter 49	RTDP Real-Time Data Path 77
HDL Hardware Description Language 31	UDP User Datagram Protocol 35
I/O Input/Output 31	VHDL Very High Speed Integrated Circuit Hardware Description Language 36
IP Internet Protocol 35	WIMP Weakly Interacting Massive Particle 14
ISR Initial State Radiation 19	

Glossary

ALU (Arithmetic Logic Unit)

Digital electronic circuit which performs arithmetic and bitwise operations on binary integers. [33](#)

ASIC (Application-Specific Integrated Circuit)

Integrated circuit customized for a specific use case instead of general purpose usage. [31](#)

ATCA (Advanced Telecommunications Computing Architecture)

Open standard for communication equipment, defining the setup of racks as well as boards. [35](#)

BAT (Bayesian Analysis Toolkit)

Software package for data analysis based on Bayes theorem. [185](#)

BC (Bunch Crossing)

Crossing of LHC proton bunches in the ATLAS detector. [43](#)

BCID (Bunch Crossing Identifier)

Unique identifier between 0 and 3563 for every bunch crossing at ATLAS. [45](#)

BCM (Beam Conditions Monitor)

ATLAS forward detector for monitoring beam conditions and luminosity measurements. [55](#)

BH (BumpHunter)

Algorithm to perform a hypothesis test for the possible presence of a signal excess. [166](#)

BRAM (Block Memory)

Logic resource on an FPGA, used to store large amounts of data. [34](#)

CLB (Configurable Logic Block)

Basic logic resource on an FPGA, containing LUTs and Flip-Flops. [32](#)

CMB (Cosmic Microwave Background)

Background radiation originating from the propagation of photons in the early universe. [13](#)

CMX (Common Merger eXtended)

Module of the first level trigger, which combines and sorts objects from the JEP and CP. [59](#)

COMA (Conditions Metadata for ATLAS)

Relational database storing metadata information down to the lumiblock level for ATLAS runs. [114](#)

CP (Cluster Processor)

Module of the first level trigger which identifies electron/ γ and tau objects. [58](#)

CPM (Cluster Processor Module)

Single module of the CP covering a small part of the calorimeter. [58](#)

CRC (Cyclic Redundancy Check)

Method to calculate check sums for raw data, used to detect transmission errors. [35](#)

CSC (Cathode Strip Chamber)

Multi-wire proportional chambers used in the muon detector. [54](#)

CTP (Central Trigger Processor)

Module of the first level trigger, which forms the final trigger decision. [57](#)

DSP (Digital Signal Processor)

Logic resource on an FPGA, combination of multiplier and adder. [33](#)

EM scale (electromagnetic scale)

Energy scale at calorimeters which yields correct results for electromagnetic, but not for hadronic objects. [52](#)

FPGA (Field-Programmable Gate Array)

Reprogrammable integrated circuit, used widely in the ATLAS level-1 trigger. [31](#)

GRL (Good Run List)

List of lumiblocks in ATLAS runs with good recording conditions fit for analyses. [113](#)

HEC (Hadronic End Cap)

Part of the hadronic calorimeter of ATLAS, covering a region of $1.375 \leq |\eta| < 3.2$. [73](#)

HLT (High Level Trigger)

Second level of the ATLAS trigger based on software algorithms. [56](#), [58](#)

IBL (insertable *B*-layer)

Innermost layer of the ATLAS pixel detector. [50](#)

IPBus

Packet-based control protocol for communication with FPGAs. [31](#)

IP Core (Intellectual Property Core)

Pre-implemented entity for firmware design shipped with Vivado. [39](#)

IS (Information Service)

System to distribute status and monitoring information between ATLAS software applications. [57](#)

JEM (Jet/Energy Module)

Module of the JEP. [59](#)

JEP (Jet/Energy-Sum Processor)

Module of the first level trigger which builds jets and energy sum objects. [58](#)

JES (Jet Energy Scale)

Energy scale of jets, containing additional calibrations compared to the EM scale. [69](#)

jFEX (Jet Feature Extractor)

Module of the upgraded level-1 trigger for Run-3 to identify jets. [193](#)

L1A (L1 Accept)

Logical OR of all trigger items, indicates whether an event is sent from level-1 to the HLT. [60](#)

L1Calo (Level-1 Calorimeter Trigger)

Part of the level-1 trigger which reconstructs calorimeter objects. [57](#)

L1Muon (Level-1 Muon Trigger)

Part of the level-1 trigger which reconstructs muons. [57](#)

L1Topo (Level-1 Topological Processor)

Part of the level-1 trigger applying topological algorithms to combinations of objects. [57](#)

LAr (Liquid Argon)

Part of the ATLAS calorimeter, used in the electromagnetic calorimeter and the hadronic end-cap and forward calorimeter. [51](#)

lumiblock (luminosity block)

Data taking interval of about 1 min, corresponding to the shortest time for which the integrated luminosity can be estimated reliably in ATLAS. [61](#)

LUCID (Luminosity measurement using Cerenkov Integrating Detectors)

Detector at ATLAS used for luminosity measurements. [55](#)

LUT (lookup table)

FPGA resource to implement combinational logic. Also used to denote arrays that replace runtime computation with simpler array indexing operations. [32](#)

MCMC (Markov Chain Monte Carlo)

Method to construct and sample from arbitrary complex posterior distributions. [185](#)

MDA (Monitoring Data Archive)

System to retrieve all available monitoring histograms at the end of a run and store them permanently. [57](#)

MDT (Monitored Drift Tube)

High-pressure proportional drift tubes, part of the ATLAS muon system. [54](#)

MGT (Multi Gigabit Transceiver)

Serializer/Deserializer operating at serial bit rates above 1 Gb/s. [34](#)

MUCTPI (Muon CTP Interface)

Part of the level-1 trigger processing muon candidates from different detector regions. [59](#)

MUCTPIToTopo

Extension of the MUCTPI which provides muon objects to L1Topo. [60](#)

OHS (Online Histogramming Service)

Application which allows sharing of monitoring histograms during ATLAS data taking. [57](#)

PPM (PreProcessor Module)

Part of the ATLAS level-1 trigger which processes calorimeter signals for the L1Calo algorithms. [58](#)

PS (Proton Synchrotron)

Pre-accelerator of the LHC at CERN. [41](#)

ROD (Readout Driver)

Part of the data acquisition system to obtain and format data from detector front-end buffers. [56](#)

RoI (Region of Interest)

Interesting regions identified by the level-1 trigger to seed HLT algorithms. [58](#)

ROS (Readout System)

Part of the data acquisition system of ATLAS which stores data for the HLT. [57](#)

RPC (Resistive Plate Chamber)

Muon trigger chambers in the barrel region, based on the detection of gas ionization. [55](#)

RTL (Register Transfer Level)

Design abstraction to model a digital circuit as the flow of signals between registers. [37](#)

SCT (Semiconductor Tracker)

Silicon microstrip detector and part of the inner detector. [49](#)

SLW (Sliding Window)

Algorithm in the level-1 trigger to identify local maxima of energy depositions. [66](#)

SPS (Super Proton Synchrotron)

Pre-accelerator of the LHC at CERN, following on the PS. [41](#)

SWiFit (Sliding Window Fit)

Modification of the global fit procedure, where multiple background estimates are performed to windows smaller than the total mass range. [164](#)

TDAQ (Trigger and Data Acquisition)

System of the ATLAS detector, responsible for selection and storage of collision events. [56](#)

TGC (Thin Gap Chamber)

Multi-wire proportional chambers, used as the muon trigger chamber in the end-cap region. [55](#)

TLA (Trigger Level Analysis)

Analysis with partially reconstructed events using only HLT jet information. [98](#)

TOB (Trigger Object)

Combination of E_T , η and ϕ identified in the level-1 trigger. [59](#)

TRT (Transition Radiation Tracker)

Part of the inner detector, used for tracking and electron identification. [50](#)

TTC (Timing, Trigger and Control)

System for the distribution of synchronous timing information at the LHC. [61](#)

Introduction

The Standard Model of particle physics is one of the most successful scientific theories. It describes the fundamental particles and their interactions, except for gravitation, using the framework of a quantum field theory. It has been developed and formulated starting in the 1970s and has since been confirmed with great precision by a large variety of experimental data, compare for example references [2, 3]. The last missing puzzle piece, the Higgs boson, has been discovered in 2012 using the Large Hadron Collider (LHC) at CERN [4, 5]. Despite its enormous success, there are numerous shortcomings of the Standard Model. Various parameters, for example the masses of the fundamental particles, are not predicted by the Standard Model and have to be inferred from experiments instead. In addition, several experimentally observed physics phenomena cannot be explained in the scope of the Standard Model. For example, the Standard Model lacks an explanation for the existence of dark matter or for the matter-antimatter asymmetry in the universe. A variety of theories have been developed as extensions of the Standard Model to address these open questions. Some of these theories predict the existence of new particles, which could be observed at the LHC. However, the production rates of these particles can be very low, and they can be produced in previously uncovered phase space regions. In order to explore these, the LHC beam energy has been increased for Run-2 of the LHC. In addition, the luminosity has been increased to gain the necessary statistical precision to discover rare phenomena. An increase in luminosity comes with an increase in pile-up, which means the occurrence of multiple proton-proton interactions during the same beam crossing. This introduces new challenges for the Trigger and Data Acquisition (TDAQ) system of ATLAS, which is responsible for the selection and storage of the collision data. In order to utilize the available bandwidth as good as possible, the collision rate of 40 MHz is reduced to the order of 1 kHz, by selecting events with interesting signatures. The ATLAS trigger system reduces the data rate in two steps, where the first level is based on fast, custom-built electronics hardware. A fundamental piece of the first-level trigger is the Level-1 Topological Processor (L1Topo), which has been introduced to the trigger system in Run-2 to cope with the increase in trigger rates [6]. L1Topo is provided with information from the ATLAS calorimeters and the muon detectors, and performs topological trigger selections on combinations of its inputs. The usual solution to cope with rising luminosities is to raise the trigger selection thresholds, which can result in decreased signal efficiencies and statistical limitations for analyses with low- p_T signatures. L1Topo provides the unique capability of performing topological selections, for example angular selections or mass cuts, already on the first trigger level, to record data at high luminosities without losing signal efficiency or raising trigger thresholds.

A particularly challenging signature are events with two jets in the final state, so-called *dijet events*. While these events can arise from a wide range of processes via the strong interaction, they can also be the result of possible new physics. Dijet final states can be the product of a new particle

which links the particles of the Standard Model to the dark matter sector, and can possibly decay into a pair of quarks. By recording the invariant mass of each dijet event, such a new particle could be observed as a resonance in the invariant mass spectrum. These so-called *dijet resonance searches* can be performed either targeting a specific model or completely model-independent, and can therefore use purely phenomenological signal parametrizations. However, due to the enormous Standard Model dijet production rate, only a very small fraction of all dijet events can be saved by the TDAQ system, especially at low invariant masses. This limitation results in a loss of sensitivity for masses in the $\mathcal{O}(100 \text{ GeV})$ range, which is of particular interest for dark matter mediator searches. Several approaches have been developed in ATLAS to circumvent these statistical limitations. A Trigger Level Analysis (TLA), which uses only partially recorded events limited to jet information [7], and an analysis of dijet events with additional initial state photons [8], are used. While the TLA can reach masses down to 450 GeV, it is still statistically limited in the lower range as it runs in the second trigger level. The dijet+photon analysis can reach lower masses, though it is statistically limited by the low cross section for the initial state photon.

In this thesis, a novel approach which utilizes L1Topo to circumvent the statistical limitations is presented. As the topological processor runs in the first level of the trigger system, it processes data from every single collision that occurs in the ATLAS detector. By using the information provided from the calorimeters, invariant dijet masses can be directly calculated on L1Topo. In order to circumvent the data rate constraints, these invariant masses are histogrammed over a certain time interval, and read out afterwards for an offline analysis. Using this method, it is possible to perform a dijet analysis already online on the first trigger level, allowing to use the full amount of available data and explore otherwise inaccessible regions of phase space.

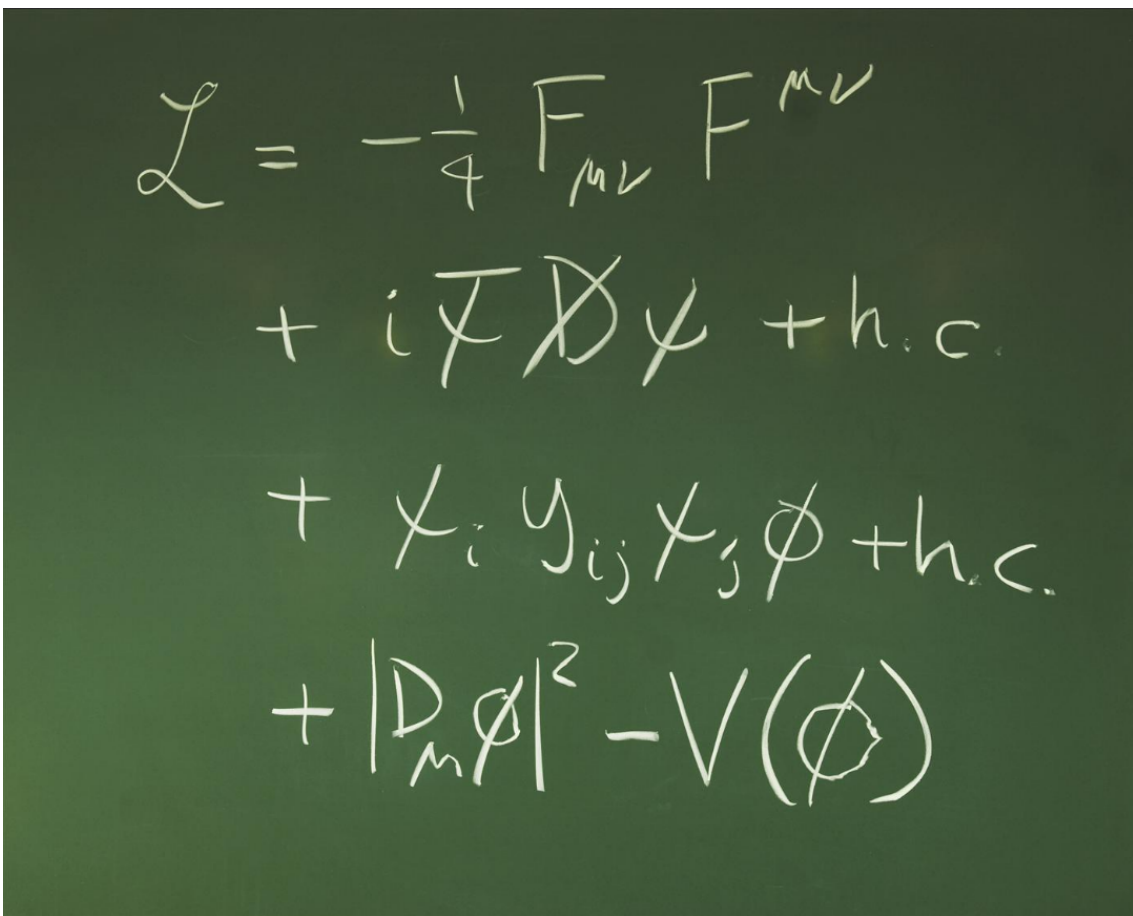
This thesis is structured in four parts. The first part introduces the Standard Model of particle physics and its possible extensions in Chapter 2, followed by a review of dijet events created in proton-proton collisions in Chapter 3. Afterwards, the second part introduces the required experimental foundations. Starting with a summary of the most important technological information on FPGAs in Chapter 4, the LHC and the ATLAS experiment are introduced subsequently in Chapter 5 and Chapter 6, with a particular focus on the calorimeters and the trigger system. The part concludes with a discussion of jet reconstruction with the ATLAS detector in Chapter 7. Afterwards, Part III focuses on the L1Topo and gives an introduction to the hardware and firmware design used in Run-2 of the LHC in Chapter 8. Part IV, the last part of the thesis, is dedicated to the histogram-based dijet analysis using L1Topo. It starts with a description of the histogramming concept, its implementation in firmware and software and the performed tests in Chapter 9. Afterwards, the fundamental concepts of the analysis are introduced in Chapter 10, presenting first results of the recorded data. Based on these results, Chapter 11 derives the invariant mass spectrum for the statistical analysis. Chapter 12 then studies the properties of level-1 jets, for example their energy scale and spatial resolution. The last chapters discuss strategies for a background estimate and the statistical analysis of the recorded data, where Chapter 13 presents an approach based on simulated events and Chapter 14 presents the statistical procedure used to search for signs of new physics. Chapter 15 concludes this by setting exclusion limits based on the observed data. Finally, a summary of this thesis and an outlook to possible improvements for Run-3 is given in Chapter 16.

Part I.

Theoretical Foundations

If people do not believe that mathematics is simple, it is only because they do not realize how complicated life is.

— John von Neumann


$$\begin{aligned}\mathcal{L} = & -\frac{1}{4} F_{\mu\nu} F^{\mu\nu} \\ & + i\bar{\psi} \not{D} \psi + h.c. \\ & + \bar{\psi}_i Y_{ij} \psi_j \phi + h.c. \\ & + |D_\mu \phi|^2 - V(\phi)\end{aligned}$$

Lagrangian density of the Standard Model of particle physics written on a blackboard, from reference [9]. This form of the Lagrangian is famous for being on the CERN coffee mug [10]. The first line describes the force carriers of the interactions, the second line describes the fundamental fermions and their interactions. The third and fourth line belong to the Higgs mechanism, the third line describing the generation of masses for fermions, the fourth line the Higgs particle and the generation of masses for bosons.

Particles and Fundamental Interactions

In Section 2.1, the first part of this chapter, the Standard Model of elementary particle physics and its interactions are introduced. The discussion starts with Section 2.1.1, where the underlying principles and the particle content of the Standard Model are described. Afterwards, the theoretical formulation of the interactions is explained, starting with a summary of the electroweak sector and its symmetry breaking in Sections 2.1.2 and 2.1.3. A specific focus is placed on a discussion of the strong interaction in Section 2.1.4, as this forms the foundation for the understanding of dijet production in proton-proton collisions. The material presented in these sections is largely based on references [11–14]. Following this, the limitations and unanswered questions of the Standard Model are discussed in Section 2.2, the second part of this chapter. A particular focus is placed on a discussion of experimental evidence for dark matter and possible theoretical models. The information presented on these topics is based on references [15–18].

2.1. The Standard Model of Particle Physics

2.1.1. Overview of the Standard Model

From a mathematical viewpoint, the Standard Model is described by a gauge quantum field theory, where excitations of quantum fields are interpreted as fundamental particles. Fermions are described by spin-1/2 fields ψ , gauge bosons are represented by spin-1 vector fields A_μ and the Higgs boson by a spin-0 field H . The Standard Model comprises all known fundamental particles and their strong, electromagnetic and weak interactions. These interactions are understood as the exchange of spin-1 gauge bosons. Gravitational interactions are not part of the Standard Model, as currently no successful approach to formulate a renormalizable theory of quantum gravity is known.

In a quantum field theory, the dynamics of the quantum fields are expressed in terms of a Lagrangian density \mathcal{L} . From \mathcal{L} , the fundamental interactions are derived using the principle of local gauge invariance, which requires invariance of the Lagrangian density under certain local gauge transformations described by the *gauge group* of the theory. For the Standard Model, the gauge group is given by $SU(3)_C \times SU(2)_L \times U(1)_Y$. The requirement of local gauge invariance leads to the introduction of spin-1 gauge boson fields, which are associated with the generators of the lie algebra of the respective symmetry group, as shown in Figure 2.1. The eight massless spin-1 particles G_μ^a arising from the $SU(3)_C$ are called *gluons*, the corresponding quantum number C is called color charge. All particles transforming with respect to this part of the gauge group are carrying color charge, and do therefore couple to gluons. The corresponding interaction is called strong interaction, with the gluons as force mediators. Three spin-1 bosons W_μ^i are associated

$$\begin{array}{ccccc}
 SU(3)_C & \times & SU(2)_L & \times & U(1)_Y \\
 \downarrow & & \downarrow & & \downarrow \\
 8G_\mu^a & & 3W_\mu^i & & 1B_\mu \\
 a = 1 \dots 8 & & i = 1, 2, 3 & &
 \end{array}$$

Figure 2.1.: The gauge group of the Standard Model is shown with the association of gauge boson fields to the generators of the algebra of the group. For each generator, one gauge boson is introduced. Adapted from [11].

with the group $SU(2)_L$, where the index L indicates that only left-handed particles are carrying the corresponding quantum number T , the *weak isospin*. The factor $U(1)_Y$ gives rise to an additional spin-1 field B_μ , with the corresponding charge Y , the weak hypercharge. The four bosons W_μ^i and B_μ are directly related to the well known physical bosons W^\pm , Z and γ , which mediate the weak and electromagnetic interactions, as described in more detail in Section 2.1.2. Particles gain their non-zero masses via the so-called *Higgs mechanism* [19–21], where the gauge symmetry of the Standard Model is spontaneously broken to a $SU(3)_C \times U(1)_Q$ symmetry, which is discussed in Section 2.1.3. The Higgs mechanism predicts the existence of a spin-0 boson, the Higgs boson. It has been discovered in 2012 by the ATLAS and the CMS experiments, respectively [4, 5]. While the bosons are force mediators, the fermions function as the building blocks of matter. The interactions of these fermions are given by their transformation properties with respect to the gauge group $SU(3)_C \times SU(2)_L \times U(1)_Y$. There are three generations of each fermion type, which are identical in all quantum numbers, but have different masses. *Leptons* are defined as those fermions which do not carry the $SU(3)_C$ color charge. In total, six leptons are known: the electron e , the muon μ , the tau lepton τ and the corresponding neutrinos ν_e , ν_μ and ν_τ . Quarks, on the other hand, are defined as those elementary fermions which participate in the strong interaction. There are six quarks in the standard model, the up quark u , the down quark d , the strange quark s , the charm quark c , the bottom quark b and the top quark t . For each fermion an antiparticle exists which has the same mass, but the opposite sign for additive quantum numbers like the electric charge. All fundamental particles, their quantum numbers and their interactions are summarized in Figure 2.2.

2.1.2. Electroweak Interaction

In the 1960s Sheldon Glashow, Steven Weinberg and Abdus Salam developed a unified theory of the electromagnetic and weak interaction [23–25]. It is based on a $SU(2)_L \times U(1)_Y$ gauge symmetry. The weak isospin T is introduced based on the experimental observation that charged weak interactions only couple to left-handed particles. Its third component T_3 is related to the weak hypercharge Y and the electric charge Q via

$$Q = T_3 + \frac{1}{2}Y. \quad (2.1)$$

By convention, the weak interaction is constructed in a way that the charged gauge bosons couple to the component T_3 of the weak isospin. Left-handed fermions, transforming under the two-dimensional representation of the gauge group $SU(2)_L$, are grouped in doublets with $T_3 = \pm 1/2$. Right-handed fermions transform with respect to the trivial one-dimensional representation,

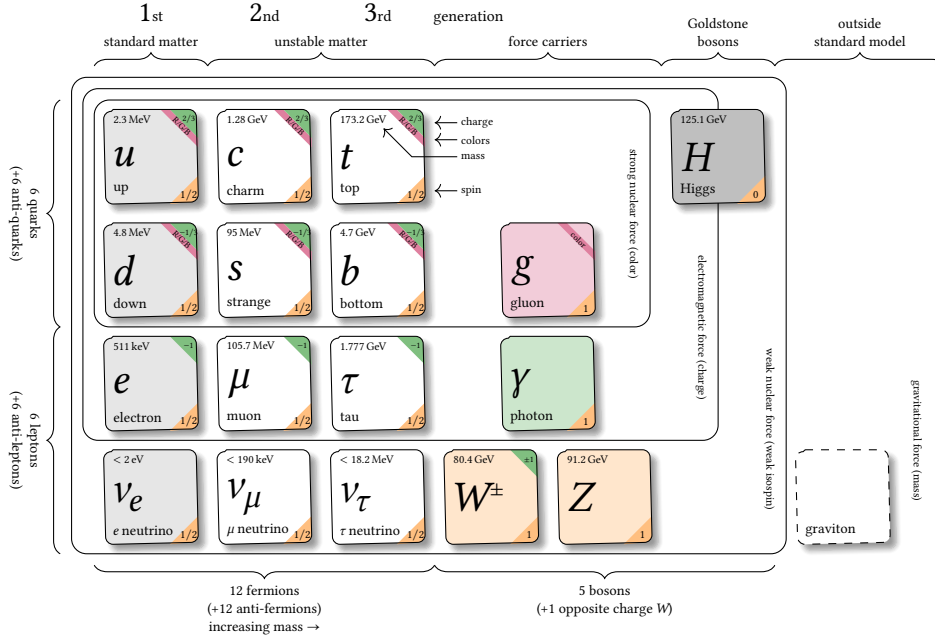


Figure 2.2.: Particle content of the Standard Model, taken from reference [22]. Properties like the charge, the spin or the mass are shown for each particle. The particles are grouped by their generation as well as by their interactions.

and are therefore singlets with $T = 0$, which do not participate in weak interactions mediated by charged bosons. The requirement of $SU(2)_L \times U(1)_Y$ gauge invariance leads to the following Lagrangian

$$\mathcal{L}_{EW} = \sum_k \bar{\Psi}_L^k i \gamma^\mu (\partial_\mu - i g \frac{\tau_j}{2} W_\mu^j - i g' \frac{Y}{2} B_\mu) \Psi_L^k + \sum_l \bar{\psi}_R^l i \gamma^\mu (\partial_\mu - i g' \frac{Y}{2} B_\mu) \psi_R^l - \frac{1}{4} W_{\mu\nu}^j W_j^{\mu\nu} - \frac{1}{4} B_{\mu\nu} B^{\mu\nu}, \quad (2.2)$$

where the sum over flavors k of left-handed field doublets Ψ_L^k has been distinguished from the right-handed fields ψ_R^l . The generators of the group $SU(2)_L$ are given as $\frac{\tau_j}{2}$, where τ_j are the Pauli matrices for $j = 1, 2, 3$. Two coupling constants are contained in the Lagrangian, g for $SU(2)_L$ and g' for $U(1)_Y$. The field tensors corresponding to the bosons B_μ and W_μ^j are given as

$$B_{\mu\nu} = \partial_\mu B_\nu - \partial_\nu B_\mu$$

$$W_{\mu\nu}^j = \partial_\mu W_\nu^j - \partial_\nu W_\mu^j - g \epsilon^{jkl} W_{\mu,k} W_{\nu,l}. \quad (2.3)$$

However, the bosons B_μ and W_μ^j are not experimentally observed, instead they mix to form the mass eigenstates W_μ^\pm , Z_μ and A_μ :

$$W_\mu^\pm = \frac{1}{\sqrt{2}} (W_\mu^1 \mp i W_\mu^2) \quad (2.4)$$

$$Z_\mu = W_\mu^3 \cos \theta_W - B_\mu \sin \theta_W \quad (2.5)$$

$$A_\mu = W_\mu^3 \sin \theta_W + B_\mu \cos \theta_W. \quad (2.6)$$

which are identified as the W^\pm and Z bosons and the photon. The *weak mixing angle* θ_W is directly related to the weak coupling constants g and g' as well as the elementary charge e :

$$e = g' \cos \theta_W = g \sin \theta_W. \quad (2.7)$$

Experimental observations [26–29] confirm that the bosons W^\pm and Z are massive, an introduction of mass terms into the Lagrangian 2.2 is therefore necessary. However, the presence of mass terms would violate its gauge invariance. An additional mechanism is therefore necessary to explain the non-zero boson masses, the Higgs mechanism. By spontaneously breaking the $SU(2)_L \times U(1)_Y$ symmetry, masses for the gauge bosons are created.

2.1.3. Spontaneous Symmetry Breaking

The symmetry breaking is achieved by introducing an isospin doublet ϕ ($T = \frac{1}{2}, Y = 1$) of two complex scalar fields ϕ^+ and ϕ^0 as

$$\phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix}, \quad (2.8)$$

which is called the *Higgs field*. The additional $SU(2)_L \times U(1)_Y$ -invariant Lagrangian for ϕ is given as

$$\mathcal{L}_{\text{Higgs}} = (\mathcal{D}^\mu \phi)^\dagger (\mathcal{D}_\mu \phi) + \underbrace{\mu^2 \phi^\dagger \phi + \lambda (\phi^\dagger \phi)^2}_{V(\phi^\dagger, \phi)}. \quad (2.9)$$

The covariant derivative \mathcal{D}_μ introduced in Equation (2.9) reads as follows

$$\mathcal{D}_\mu = \partial_\mu + ig \frac{\tau_j}{2} W_\mu^j + ig' \frac{Y}{2} B_\mu. \quad (2.10)$$

The potential $V(\phi^\dagger, \phi)$ defined in Equation (2.9) is constructed with parameter values $\mu^2 < 0$ and $\lambda > 0$. In this set of variables, V has a local maximum at the origin as well as a circle of degenerate minima with $\phi^\dagger \phi = -\mu^2/2\lambda \equiv v^2$ around it. If one of these minima is chosen as the ground state, the $SU(2)_L \times U(1)_Y$ symmetry of the system is broken to a $U(1)_Q$ symmetry. An expansion of ϕ around the vacuum expectation value of the neutral field yields

$$\phi(x) \approx \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + H(x) \end{pmatrix}. \quad (2.11)$$

The boson described by the field H is a physical neutral scalar, the Higgs boson with a mass of $\sqrt{2}|\mu|$. The remaining three of the four degrees of freedom of ϕ are absorbed to create mass terms for the massive electroweak gauge bosons

$$m_W = \frac{1}{2}vg \quad m_Z = \frac{1}{2}v\sqrt{g^2 + g'^2}, \quad (2.12)$$

while the photon remains massless. In total, four free parameters are needed to describe the electroweak sector, the $SU(2)_L \times U(1)_Y$ gauge couplings g and g' and two free parameters of the Higgs potential, for example μ and v .

The Higgs mechanism can also be used to generate fermion masses by introducing Yukawa couplings between the scalar Higgs field and the fermions. For all Dirac fermions, gauge invariant mass terms can be constructed by combining left- and right-handed fields as follows

$$\mathcal{L}_{\text{Yukawa}} = -g_f (\bar{\Psi}_L \phi \psi_R + \bar{\psi}_R \phi^\dagger \Psi_L) \text{ or } \mathcal{L}_{\text{Yukawa}} = g_f (\bar{\Psi}_L \phi_c \psi_R + \bar{\psi}_R \phi_c^\dagger \Psi_L) \quad (2.13)$$

where $\phi_c = -i\sigma_2 \phi^*$. However, the fermion masses are not predicted by the Standard Model, and their values have to be inferred from experiments instead.

2.1.4. Quantum Chromodynamics

The theory describing the strong interaction is called Quantum Chromodynamics (QCD). It is a gauge field theory with a non-abelian symmetry group $SU(3)_C$, which has eight generators t_a with a color index $a = 1, \dots, 8$. While quarks transform under the fundamental representation of $SU(3)_C$, where gluons are commonly represented via the Gell-Mann matrices $t_a = \lambda_a/2$, the gluons itself transform under the adjoint representation, which is defined by $[t_a, t_b] = i f_{abc} t_c$, with the structure constants f_{abc} of the $SU(3)_C$ gauge group. The gauge-invariant Lagrangian of QCD can be written as follows

$$\mathcal{L} = \sum_q \bar{q}_j (i \gamma_\mu \partial^\mu \delta_{jk} - g_s \gamma_\mu G_\mu^{jk} - \delta_{jk} m_f) q_k - \frac{1}{4} G_{\mu\nu}^a G_a^{\mu\nu}, \quad (2.14)$$

where the sum runs over all quark flavors q . The quark fields q carry color indices $j, k = 1, 2, 3$ and the a th gluon field tensor is given as

$$G_{\mu\nu}^a = \partial_\mu G_\nu^a - \partial_\nu G_\mu^a - g_s f_{abc} G_\mu^b G_\nu^c. \quad (2.15)$$

The last term of $G_{\mu\nu}^a$ describes the self-coupling of gluons, a consequence of the non-Abelian structure of $SU(3)_C$. The only free parameter of the QCD Lagrangian, the coupling g_s , is related to the strong coupling constant α_S via

$$\alpha_S = \frac{g_s^2}{4\pi} \quad (2.16)$$

In a graphical representation, the QCD Lagrangian can be depicted as follows [14]:

$$\mathcal{L} = \sum \left[\begin{array}{c} j \longrightarrow k \\ \delta_{jk} \end{array} + \begin{array}{c} j \longrightarrow \begin{array}{l} \nearrow k \\ \searrow a \end{array} \\ g_s t_{jk}^a \end{array} \right] + \left[\begin{array}{c} a \text{-----} b \\ \delta_{ab} \end{array} + \begin{array}{c} a \text{-----} b \\ \searrow c \\ g_s f_{abc} \end{array} + \begin{array}{c} a \text{-----} b \\ \searrow c \\ \nearrow d \\ g_s^2 f^{abe} f_{cde} \end{array} \right] \quad (2.17)$$

The fermionic part of the Lagrangian in the first brackets contains a term for the propagation of free quarks as well as a term describing the coupling between quarks and gluons, which is proportional to the coupling strength g_s . In the gluonic part in the second brackets three terms exist, the first one describing the propagation of a free gluon. The second and the third term are gluon self-coupling terms for three- and four-gluon coupling and contain the structure constants f_{abc} . The three gluon vertex is proportional to g_s , while the four gluon coupling is proportional to g_s^2 .

Using the terms from Equation (2.17), Feynman diagrams can be written down for arbitrary QCD processes. Predictions for these processes can then be made by summing up all relevant diagrams in a perturbation series and using Feynman rules to calculate the resulting cross section. In order to yield finite, meaningful results, a quantum field theory needs to be *renormalizable*. In renormalization, the divergences arising from loop diagrams in the perturbation series are absorbed into a redefinition of the coupling strength, which then depends on the scale of momentum transfer Q^2 at which the interaction occurs and an unphysical renormalization scale μ_R^2 . Physics results should not depend on a particular choice of this unphysical scale [30]. This requirement leads to a first order differential equation for the β function, which describes the scale dependence of the coupling. A solution is found using a perturbative expansion in powers of α_S :

$$\beta(\alpha_S) = Q^2 \frac{\partial \alpha_S}{\partial Q^2} = -(\beta_0 \alpha_S^2 + \beta_1 \alpha_S^3 + \mathcal{O}(\alpha_S^4)). \quad (2.18)$$

To determine the perturbative coefficients β_0, β_1, \dots the calculation of all diagrams with one or more loops in the gluon propagator is required. Using only 1-loop corrections, Equation (2.18) can be solved analytically and the running of the strong coupling can be written as

$$\alpha_S(Q^2) = \frac{\alpha_S(\mu_R^2)}{1 + \beta_0 \alpha_S(\mu_R^2) \ln \frac{Q^2}{\mu_R^2}}, \text{ with } \beta_0 = \frac{11n_c - 2n_f}{12\pi}, \quad (2.19)$$

where $n_c = 3$ describes the number of colors and $n_f \leq 6$ the number of quark flavors participating in interactions at the scale Q^2 . The resulting value for β_0 is larger than zero, leading to α_S decreasing asymptotically with increasing Q^2 . Figure 2.3 shows the values of $\alpha_S(Q^2)$ based on QCD calculations in comparison to experimental results.




Figure 2.3.: Overview of different measurements of α_S as a function of the energy scale Q , from reference [31]. The degree of perturbation theory used for the α_S extraction for each experiment is indicated in parentheses.

At small distances corresponding to energy scales of $\mathcal{O}((100 \text{ GeV})^2)$ α_S is small, for example $\alpha_S(M_Z^2) \approx 0.12$, and quarks and gluons can be treated as quasi-free particles, and therefore perturbation theory can be applied. This phenomenon is referred to as *asymptotic freedom*. However, even at these scales higher-order corrections cannot be neglected. Due to this it is often necessary to calculate QCD processes beyond leading order, resulting in expensive calculations as a consequence of the high number of diagrams involved. On the other side of the spectrum $Q^2 \approx \mathcal{O}(1 \text{ GeV}^2)$, α_S gets larger with increasing distances, and perturbative methods can no longer be applied, as diagrams from all orders show significant contributions. In this region, gluon self-interaction leads to strongly bound states of quarks and gluons. When trying to separate two bound quarks, attractive interactions via the exchange of virtual gluons between the quarks create a gluon color field located in a narrow tube. At relatively large distances, the energy density

in the tube is constant and the total energy therefore proportional to the distance between the quarks. Phenomenologically, the energy content in the tube is described by a potential of the form

$$V(r) \approx \kappa r, \text{ where } \kappa \sim 1 \text{ GeV/fm} \quad (2.20)$$

At some point of spatial separation, it is energetically favorable to form color-neutral states by creating additional $q\bar{q}$ pairs from the QCD vacuum compared to further increasing the energy in the gluon field. The resulting property of the strong interaction is called *color confinement* and describes the phenomenon that colored objects do not appear as free particles, instead they are always observed in color-neutral bound states like mesons or baryons.

2.2. Beyond the Standard Model

2.2.1. Limitations of the Standard Model

Even though the Standard Model agrees with a large variety of experimental results, there are several shortcomings of the theory, some of them in form of experimentally observed phenomena which cannot be explained in the scope of the Standard Model. Experimental deficits include the matter-antimatter asymmetry, which describes the observation that the universe is almost exclusively made up of matter, even though the Big Bang should have created equal amounts of matter and antimatter [32]. In order to explain the observed difference, a mechanism which treats particles and antiparticles in different ways is required. A mechanism of the Standard Model providing this difference is CP violation: physics processes, where all particles are exchanged with their antiparticles, are not fully equivalent to the spatially mirrored process. However, the CP violation in the weak interaction of the Standard Model can only cause a small part of the observed asymmetry, another source for CP violation is therefore required [32].

Another experimental shortcoming is related to observations of neutrino oscillations, for example in [33, 34], which showed that neutrinos can oscillate between different flavors. The theoretical description of this mechanism requires non-zero mass differences between neutrinos, leading to the requirement of small, but non-vanishing neutrino masses. In the Standard Model, fermion masses are generated using Yukawa couplings, compare (2.13). However, no right-handed neutrinos exist in the Standard Model to create neutrino masses via Yukawa interactions. A modification of the Standard Model is therefore required to introduce a mechanism explaining the non-zero masses for the neutrinos.

In addition, the Standard Model does not provide a viable candidate for dark matter, even though compelling experimental evidence exists for its existence. This is discussed in more detail in Section 2.2.2. Furthermore, the nature of dark energy, which is responsible for the accelerating expansion of the universe, is currently unknown [31].

Besides these experimental observations, there are a couple of theoretical shortcomings of the Standard Model, which do not affect its ability to predict experimental results, but rather indicate a lack of deeper understanding of the underlying principles. Numerous parameters, for example the masses of the fundamental particles, are not predicted by the Standard Model and have to be inferred from experiments instead. The origin of these parameter values remain unknown.

Another theoretical problem is related to the strong CP problem. It describes the observation that the CP symmetry is preserved in QCD, even though it would be possible to introduce a CP -violating term into the Lagrangian of QCD without affecting its gauge invariance [35]. This new term depends on a new parameter θ , whose value is a priori unknown and not constrained.

Experimental results on the electric dipole moment of the neutron, for example in reference [36], show that θ is extremely small and close to zero. However, this particular value introduces the necessity of fine-tuning and a violation of naturalness. Currently, no explanation for the observed value of θ is known.

Moreover, the theoretical formulation of the interactions is a point which theories beyond the Standard Model try to address. The successful unification of the electromagnetic and weak interactions, as discussed in Section 2.1.2, gave rise to the idea that all three forces of the Standard Model could be manifestations of the same unifying force. This so-called Grand Unified Theory is believed to be relevant at scales of $\sim 10^{16}$ GeV, where the running couplings of all three interactions are of equal size. In formulations of this theory, the symmetry group of the Standard Model is seen as a part of a higher dimensional group [37], e.g. $SU(5)$, which breaks down to the known $SU(3)_C \times SU(2)_L \times U(1)_Y$ group below a certain energy scale. Grand unified theories often predict a decay of the proton, however experiments found a lower limit on its lifetime of $\approx 10^{34}$ y, depending on the exact decay mode [31].

The inability to include the gravitational interaction with the Standard Model is another one of its shortcomings. While gravitation is much weaker than the other interactions at all energy scales studied up to today, it is believed to become comparable in strength when the Planck scale of $\sim 10^{19}$ GeV is reached, where the Standard Model breaks down. An inclusion of the theory of General Relativity with the Standard Model to a so-called Theory of Everything has been unsuccessful so far.

2.2.2. Dark Matter

Compelling evidence for physics beyond the Standard Model is provided by the existence of dark matter in the universe [38]. Direct hints on dark matter come from observations of *rotation curves*, which show the velocity dependence of stars and gas on their distance to the center of the galaxy. This velocity can be written as follows using Newton's gravitational law

$$v(r) = \sqrt{\frac{GM(r)}{r}} \quad (2.21)$$

where $M(r)$ describes the total mass enclosed in a radius of r around the center of the galaxy. Most of the visible mass is distributed close to the galactic center, implying $M(r) \approx \text{constant}$ and a velocity dependence of $v \sim 1/\sqrt{r}$ in the outer regions of the galaxy where only small amounts of visible matter are present. However, the observations are different, $v(r)$ is detected as approximately constant, compare the rotation curves shown in Figure 2.4. The discrepancies can be explained by introducing a large non-luminous halo component to the matter content of the galaxy, called dark matter.

While it is also possible to explain the observed rotation curves by a modification of the gravitational forces involved, there are other phenomena which are difficult to explain using theories of modified gravity, for example the distribution of matter in the so-called *bullet cluster*. It consists of two colliding galaxies, both containing visible mass in form of stars and hot gas. An image of the bullet cluster is shown in Figure 2.5. While the visible mass has been decelerated in the collision due to the electromagnetic interaction, the gravitational centers of the galaxies were not slowed down. These observations could neither be explained with the visible matter nor with modified gravitational laws [41], and are therefore seen as a direct evidence for an additional dark matter content in the galaxy clusters, which is not allowed to participate in the electromagnetic interaction.




Figure 2.4.: Rotation curves of the galaxy NGC 6503, from reference [39]. The measured values of the rotation velocity are shown together with the three-component fit (solid line). The individual components, namely gas, disk and dark matter halo are shown as dotted, dashed and dashed-dotted lines, respectively. While the disk contribution shows the expected $1/\sqrt{r}$ dependence for large r , the contribution from the dark matter halo has to be included to the gas and disk contributions to explain the observed values.

While these observations present evidence for the existence of dark matter, no conclusions can be drawn on its abundance in the universe. Instead, it can be inferred from observations of the Cosmic Microwave Background (CMB). The CMB is a background radiation originating from the propagation of photons in the early universe [38]. Immediately after the Big Bang, the universe was a hot plasma opaque to photons, as these were continuously scattered off free electrons. After the recombination epoch, where electrons combined with protons to form hydrogen, the number of free electrons dropped and the universe became transparent for the photons. As a result, they decoupled and started to travel freely through the universe, forming the cosmic background radiation. The CMB is known to be isotropic at the 10^{-5} level, and follows the form of the spectrum of black body radiation with a temperature of 2.726 K. Small anisotropies in the CMB are the result of small density fluctuations in the plasma of the early universe [42]. The accuracy of cosmological models describing the development of the early universe can therefore be checked by comparing their predicted anisotropies to the observed anisotropies, since the predictions are very sensitive to modifications of the tested models and their parameters. Precise measurements of the CMB have been performed using the Planck satellite [43]. Figure 2.6 shows the sky map of the measurements, with temperature fluctuation in the order of a few hundred μK . The analysis of these observations yield support for the Big Bang hypothesis, in particular for the so-called ΛCDM model of cold dark matter with the cosmological constant Λ . This model contains the relative abundance of dark matter as a parameter, allowing to infer it from a fit to the observed CMB anisotropies. The resulting dark matter abundance is given as $\Omega_c \approx 0.26$, the remaining energy content of the universe is visible, baryonic matter ($\Omega_b \approx 0.05$) and dark energy



Figure 2.5.: The distribution of the matter in the galaxy cluster 1E 0657–56, commonly referred to as the bullet cluster, is shown. X-ray image data (red) is superimposed with matter distributions calculated from gravitational lensing (blue) and images of visible light (galaxies). From reference [40].

($\Omega_\Lambda = 0.69$).

2.2.3. Theoretical Models for Dark Matter

As discussed in the previous section, compelling evidence for dark matter exists. However, the question on what exactly dark matter is and how it relates to the known particles of the Standard Model remains open. Regardless of their exact nature, possible candidate particles must obey several constraints from the experimental observations: they must be stable on cosmological time scales, they must interact only very weakly with electromagnetic radiation, and they must provide the correct abundance Ω_c . A large variety of candidates and theoretical models trying to address these questions exist [12], in this thesis the focus is placed on the so-called Weakly Interacting Massive Particles (WIMPs). WIMPs are particles with masses in a range of $\mathcal{O}(10 \text{ GeV})$ to $\mathcal{O}(1 \text{ TeV})$. They are a candidate for cold dark matter, and provide the correct abundance Ω_c as a result of a *freeze-out* procedure, when weak scale interactions of dark matter particles χ with themselves are assumed. This is described briefly in the following. In the early universe, high temperatures lead to the formation and annihilation of WIMPs and ordinary matter in a thermal equilibrium. When the universe expanded, the temperature dropped and at some point the mean thermal energy was too low for the dark matter formation process to continue, while annihilation was still possible, leading to a decrease in the number of WIMPs. After further expansion of the universe, the expansion rate became faster than the annihilation rate, leading to a ceasing of the annihilation process and effectively freezing the number of WIMPs in the universe. This process is called the *freeze out* of the WIMPs. The measured abundance of dark matter does then correspond to the abundance of WIMPs after the freeze out.

A fundamental component of this scenario is the existence of non-vanishing couplings between Standard Model particles and dark matter [17]. The required interactions can possibly be mediated by a new kind of mediator, linking Standard Model and dark matter particles. These interactions motivate a variety of experimental searches, either directly for the WIMP or for the mediator.




Figure 2.6.: The temperature anisotropy measurement of the CMB is shown, as measured by the Planck satellite in 2015 [44].

A simplified axial-mediator dark matter model, which is used as a benchmark model in a wide range of experimental searches, is explained in the following section.

2.2.4. Z' Dark Matter Mediator

In this thesis a simple extension of the Standard Model is considered, where an additional $U(1)$ gauge symmetry is introduced. A dark matter candidate particle is assumed to have charges only under this new group [16]. Under the assumption that some Standard Model particles do also carry this charge, an additional gauge boson mediating interactions between Standard Model particles and dark matter is introduced. In the case of a dark matter particle χ with mass m_χ , and a spin-1 mediator Z' with axial-vector couplings to Standard Model and dark matter particles, the corresponding Lagrangian is given as

$$\mathcal{L}_{Z'} = g_q \sum_{q=u,d,s,c,b,t} Z'_\mu \bar{q} \gamma^\mu \gamma^5 q + g_\chi Z'_\mu \bar{\chi} \gamma^\mu \gamma^5 \chi, \quad (2.22)$$

where the coupling g_q to quarks is assumed to be universal and g_χ is the coupling to dark matter particles. In total, the model contains four parameters, the masses m_χ and $M_{Z'}$, and the coupling strengths g_χ and g_q . Since couplings to leptons have been neglected, this model is also-called a *leptophobic* Z' model [17]. The main advantage of this benchmark model is that its parameter space $(g_\chi, g_q, m_\chi, M_{Z'})$ can be explored complimentary by using results from direct detection experiments, dark matter abundance predictions, monojet searches and dijet searches. While dijet searches are in principle signature driven, i.e. sensitive to a large range of models rather than fine-tuned to some specific model, the simplified dark matter mediator model is used to interpret the results in the scope of a dark matter search and to compare it with the results of other searches and predictions from cosmology [15]. To understand how the model allows for a wide range of complementary searches, the Feynman diagrams arising from Equation (2.22) can be examined. Figure 2.7 shows four different channels which are used for the various searches. Constraints to dark matter in the four dimensional parameter space $(g_\chi, g_q, m_\chi, M_R)^1$ arising

¹In some sources, the axial-vector mediator is denoted as R instead of Z' .




Figure 2.7.: Four possible Feynman diagrams arising from $\mathcal{L}_{Z'}$ are shown. The process of dark matter annihilation is shown in (a), which sets the relic abundance. In (b), a dark matter particle scatters off a Standard Model particle, as used in direct detection experiments, (c) shows the monojet signature with additional gluon radiation and (d) the resonance signature through Z' quark couplings.

from the various search methods are shown in Figure 2.8. Some regions can be excluded from theoretical considerations: areas with mediator masses much higher than dark matter masses are excluded by dark matter particle overproduction, as these areas would yield an excessive relic abundance. These areas are drawn in red in Figure 2.8. When making predictions with a simplified model, one has to consider that such a model is often only a low-energy effective description of a more complex theory. To ensure that this low-energy assumption still holds and that perturbation theory still applies, the dark matter mass is not allowed to be much higher than the mediator mass [15]. The resulting perturbativity boundary is shown in gray. Note that this boundary does not actually put hard constraints on the underlying physics, but rather constrains the parameter space the simplified model is sensitive in. The direct detection experiments are especially sensitive at low mediator masses, where a large density of dark matter particles is predicted. The monojet channel is most sensitive in a region with small values of m_χ and m_R . However, its sensitivity drops at the line $m_\chi = m_R/2$, as the decay channel $Z' \rightarrow \chi\bar{\chi}$ shown in Figure 2.7(c) is kinematically disfavored for values of $m_\chi > m_R/2$. Additionally, high values of m_R require a high-momentum gluon to recoil against the mediator, making this region difficult to access with a monojet signature. In contrast to this, the dijet searches cover a wider range of m_χ values, being limited at high values of m_R due to the requirement of mediator production and at low masses due to trigger rate limitations.

In addition, the sensitivity of the various channels varies with the absolute and relative sizes of the couplings. While the sensitivity of the dijet channel suffers hugely when the value of g_q is decreased, the monojet search is less effected. Especially in the case $m_R > m_\chi/2$ the mediator decay to dark matter is favored over the decay to quarks, limiting the dijet sensitivity. However, the dijet search is still the search most sensitive to mediator masses and dark matter masses in the WIMP range of $\mathcal{O}(100 \text{ GeV})$. In particular, for some choices of the couplings, for example in Figure 2.8(d), large areas of parameter space not excluded exist, especially at lower mediator masses and low quark couplings.




Figure 2.8.: Combined 95 % constraints in the m_χ - M_R plane from different experimental approaches. R denotes an axial vector mediator particle like the Z' presented in this section. Experimental constraints are set from direct detection (orange, dotted) by the LUX experiment [45], from monojet signatures (green, dashed) measured by the CMS experiment [46] and from dijet searches (blue, dash-dotted) by UA2 [47], ATLAS [48, 49], CDF [50] and CMS [51]. In addition, constraints from the cosmological relic density are shown in red and the perturbative boundary is shown in gray. For the left (right) column, the value of $g = (g_\chi g_q)^{1/2}$ is fixed to 1 (0.25). The upper (lower) row shows coupling ratios of g_χ/g_q of 1 (9). Taken from reference [15].

Dijet Production in Proton-Proton Collisions

In the first part of this chapter, an introduction to proton-proton collisions is given. Starting with Section 3.1.1, the basic terminology used in this context is established. Phenomenological methods which are applied to describe the internal structure of protons are introduced in Section 3.1.2. Afterwards, Section 3.1.3 introduces the factorization theorem, which describes how to calculate cross sections for processes in proton-proton collisions. Section 3.1.4 explains the formation of jets from quarks or gluons, as these can arise in the final states of proton-proton collisions. The second part of the chapter is dedicated to a brief introduction to the theory of dijet events, starting with a description of some commonly used kinematic variables in Section 3.2.1. Section 3.2.2 explains the possible processes resulting in dijet final states and discusses methods to distinguish between QCD background and possible events from signal processes with the same signature, for example the Z' process introduced in Section 2.2.4. The first part uses material from references [30, 52, 53], while the second part is primarily based on the material in references [14, 54]. This chapter concludes with a discussion of the simulation and event generation of events in proton-proton collisions in Section 3.3.

3.1. Proton-Proton Collisions

3.1.1. Introduction and Terminology

Protons are composite objects with a complex internal structure which cannot be described via perturbative methods of QCD. The description of a proton-proton collision is therefore more complex than a collision between point-like elementary particles. An illustration of a proton-proton collision is given in Figure 3.1. In this figure, the incoming protons p_1 and p_2 are indicated by three black lines and a gray ellipse each. A primary interaction can occur between two partons, the constituents of the proton. This primary interaction is called *hard scattering process*. The hard scattering process between two partons, in this case two gluons, is shown in black in the center of the image. Additional radiation of particles can happen either for the incoming particles, i.e. in the initial state, or for the outgoing particles in the final state. The former is called Initial State Radiation (ISR) while the latter is called Final State Radiation (FSR). Repeated emission of QCD radiation results in the development of a so-called *parton shower*. The colored partons of the parton shower will eventually form color-neutral hadrons in a process called *hadronization*. Some of these hadrons are unstable, and will therefore decay further into stable particles, for example lighter hadrons or photons. As the QCD radiation is preferably emitted at small angles,




Figure 3.1.: Illustration of all involved stages in a proton-proton collision, adapted from [55]. A $gg \rightarrow gg$ process is shown.

the resulting hadrons form a collimated bundle, a so-called *jet*. In addition, secondary interactions can occur between the remnants of the protons, as shown in orange in the lower part of Figure 3.1. These additional interactions are part of the so-called underlying event, the partons created in these interactions are usually softer than those created in the hard scattering process. Nevertheless, formation of parton showers and hadronization also occurs in the underlying event, resulting in a contamination of the hard scattering process.

3.1.2. Structure of the Proton

A proton is a bound state of three valence quarks, namely two up quarks and one down quark, which are bound via the exchange of gluons. These gluons can split into $q\bar{q}$ pairs, so-called sea quarks. In addition, radiation of further gluons from other gluons or valence quarks is possible. All these components of a proton, valence quarks, sea quarks and gluons, are referred to as partons. The processes occurring inside a proton often take place at non-perturbative low energy scales. Phenomenological models are therefore used to describe the behavior of partons inside of protons, for example the *parton model* introduced in the following. In a good approximation, protons at hadron colliders can be described in the infinite momentum frame, where the partons have a momentum collinear to the proton's momentum. A parton is then described as carrying a longitudinal momentum fraction x of the total proton momentum. As the hard scattering does not occur between the composite protons but rather between two partons a and b , the energy

available for the collision is not given by the full center-of-mass energy \sqrt{s} , but by the partonic center-of-mass energy $\sqrt{\hat{s}} = \sqrt{x_a x_b s}$, where x_a and x_b are the momentum fractions that the partons carry. It is therefore necessary to describe the distribution of momentum fractions carried by specific partons, this is parametrized by the *Parton Distribution Functions (PDFs)* $f_a(x_a)$. As they are non-perturbative by construction, the PDFs are not predicted by QCD and have to be determined from experimental data instead. Besides having a dependence on the momentum fraction x , the PDFs also depend on the value of momentum transfer Q^2 at which the interaction occurs. This can be understood as a higher Q^2 value corresponds to a higher spatial resolution, which means that additional partons can be resolved, leading to a scale dependence of the PDFs. While no prediction of the PDFs at a certain scale Q_0^2 can be made, their scaling behavior with respect to Q^2 is described by the *DGLAP* equations [52]

$$Q^2 \frac{df_a(x_a, Q^2)}{dQ^2} = \sum_{b \in \{q, g\}} \int_{x_a}^1 \frac{dz}{z} \frac{\alpha_S}{2\pi} \hat{P}_{ba}(z) f_b\left(\frac{x_a}{z}, Q^2\right), \quad (3.1)$$

where $z = x_a/x_b$ describes the ratio of the momentum fractions of the partons a and b . The functions $\hat{P}_{ba}(z)$ introduced in Equation (3.1) are the splitting functions, which describe the probability for a parton b to undergo a splitting resulting in a parton a . The splitting functions can be calculated perturbatively. In total, four different splitting functions appear in the DGLAP equations, \hat{P}_{qq} , \hat{P}_{qg} , \hat{P}_{gq} and \hat{P}_{gg} . They describe the probability that a quark q with momentum fraction x_q can result either from a quark with the same flavor and momentum fraction $x_{q'} > x_q$ which radiated a gluon, or from a gluon with momentum fraction $x_g > x_q$ that split into a $q\bar{q}$ pair. Analogously, a gluon can result from a quark with higher momentum fraction $x_q > x_g$ radiating a gluon or from a gluon with $x_{g'} > x_g$ radiating a gluon. Equation (3.1) can be understood in a pictorial representation, which is shown in Equation 3.2:

Equation 3.2.: A depiction of the DGLAP equations is shown, describing the evolution of PDFs with Q^2 . The striped blob represents the incoming hadron, in which a parton b is found with a probability $f_b(x/z, Q^2)$ and undergoes a splitting $\hat{P}_{ba}(z)$. The splitting functions are represented by the according leading order parton branching. While the first line corresponds to $a = q$, the second line corresponds to $a = g$. Adapted from [52].

To fully determine the PDFs, their dependence on x has to be determined from data. A wide range of experimental results are used for this purpose, for example results from deep inelastic scattering by the H1 and ZEUS collaborations at the HERA accelerator [56], fixed target experiments [57], or from hadron collider experiments [58, 59]. The PDFs are then modeled using a generic functional parametrization or a neural network at a low scale $Q_0^2 = \mathcal{O}(1 \text{ GeV}^2)$. Using these parametrizations an experimental observable, for example a cross section, is calculated and

compared to experimental data to determine the best fit parameters. The resulting PDFs are then evolved up in Q^2 using the DGLAP equations. The PDFs as determined from experiment are shown in Figure 3.2 for the NNPDF PDF set [60]. At low values of x , the gluon PDF is dominant, while the valence quark PDFs dominate at higher values of x . The general behavior is similar for both Q^2 scenarios, however for higher Q^2 values the contribution at low x get larger, as more partons can be resolved in this case.




Figure 3.2.: The NNPDF3.1 NNLO (next-to-next-to-leading order) PDFs, for momentum transfers of 10 GeV^2 (left) and 10^4 GeV^2 (right). Taken from reference [60].

3.1.3. The Factorization Theorem

In order to completely describe proton-proton interactions at given energies, both the perturbatively accessible hard-scattering process of the partons and the non-perturbative regime of partons in the proton need to be considered. The central concept is to separate the hard process, described by a cross section $\hat{\sigma}$ in a perturbation series, from the non-perturbative part, described by PDFs $f(x, Q^2)$. This factorization procedure is depicted in Figure 3.3. In leading order, the short-distance cross section $\hat{\sigma}$ corresponds to the $2 \rightarrow 2$ parton cross section. Higher order terms of the perturbative corrections contain contributions from effects which can occur long before the hard scattering itself, for example the radiation and splitting of additional partons as described in Equation 3.2. These contributions are usually factored out and absorbed into the description of the incoming hadrons, i. e. the PDFs. The remaining contributions involve only processes at high Q^2 -values. In this way, the partonic cross section is completely independent on the incoming hadrons. For this procedure, it is necessary to introduce the so-called *factorization scale* μ_F^2 , which can be interpreted as the scale which separates the soft low- Q^2 from the hard high- Q^2 physics. Below μ_F , collinear and infrared divergences of QCD are absorbed into the PDFs. Mathematically, the factorization is described by the *factorization theorem*:

$$\sigma_{AB} = \sum_{a,b} \int_{x_a=0}^1 \int_{x_b=0}^1 dx_a dx_b f_{a/A}(x_a, \mu_F^2) f_{b/B}(x_b, \mu_F^2) \hat{\sigma}_{ab}. \quad (3.3)$$

It describes the factorization of the scattering cross section σ_{AB} of two partons a, b contained in protons A, B into the PDFs $f_{a/A}$, $f_{b/B}$ and the cross section of the hard scattering process $\hat{\sigma}_{ab}$. The




Figure 3.3.: Graphical representation of the factorization of a process in a proton-proton collision. Two partons a, b with momentum fractions x_a and x_b participate in the hard scattering process. The probabilities to find these partons in the incoming protons are given by the PDFs $f_{a/A}(x_a, \mu_F^2)$ and $f_{b/B}(x_b, \mu_F^2)$. These partons are then participating in the hard scattering process, described by the cross section $\hat{\sigma}_{ab}$. Adapted from reference [53].

partonic cross section $\hat{\sigma}_{ab}$ can be calculated as a perturbation series in the running coupling α_S :

$$\hat{\sigma}_{ab} = [\hat{\sigma}_0 + \alpha_S(\mu_R^2)\hat{\sigma}_1 + \dots]_{ab} \quad (3.4)$$

If all orders of perturbation theory were considered, the results would be independent on the unphysical choice of μ_F and μ_R . However, in reality the choice of these parameters has to be considered as an uncertainty on the theoretical results.

3.1.4. Formation of Jets

In proton-proton collisions, final states containing multiple high-energetic partons are created frequently. However, these final state partons are not experimentally observable as free particles due to confinement. Their experimental signatures are instead special objects, so-called *jets*. In general, a jet is a very collimated shower of color neutral particles containing mostly hadrons, but also photons and leptons. Jets evolve from high-energetic partons which can split to produce additional partons, similar to the processes described by the DGLAP equations 3.2. The resulting partons are typically emitted nearly collinear to the original parton, leading to a high degree of collimation. Each splitting reduces the parton energies and leads to increasing distances between the partons. When the energies of the partons reach the energy regime of $\mathcal{O}(1 \text{ GeV})$, confinement results in the hadronization of the partons which bind together to form color-neutral hadrons. To infer the properties of the original partons, the resulting jets have to be measured and reconstructed in an experiment. The exact definition of how a jet can be reconstructed is not unique, but depends on the choice of a jet algorithm, which clusters the energy of the shower particles. Information on jet algorithms is given in Chapter 7.

3.2. Properties of Dijet Events

In the following, events with two partons in the final state are reviewed. An example for such a process is the decay of the dark matter mediator into two quarks as shown in Figure 2.7(d). Such events can result in final states with two jets, which are then called dijet events. The most dominant contribution to events with dijet final states are $1+2 \rightarrow 3+4$ parton scattering processes in QCD, for example as shown in Figure 3.1. In the following, the kinematic properties of these processes are reviewed in the case of the high energy limit, where the parton masses can be neglected compared to their momenta.

3.2.1. Kinematics of Two-Parton Scattering

For a $1 + 2 \rightarrow 3 + 4$ parton scattering process, several kinematic relations hold regardless of the exact details of the interaction and the involved particles. A suitable choice of reference system for this process is the center-of-mass frame, the scattering process in this frame is illustrated in Figure 3.4.




Figure 3.4.: Schematic illustration of a two-to-two scattering process in the center-of-mass frame. The four momenta p_i of the partons are given as well as the scattering angle θ^* .

In this frame, the Lorentz invariant Mandelstam variables, which are given as $\hat{s} = (p_1 + p_2)^2$, $\hat{t} = (p_1 - p_3)^2$ and $\hat{u} = (p_2 - p_3)^2$, can be expressed as a function of \hat{s} and the scattering angle θ^* only:

$$\hat{t} = -\frac{1}{2}\hat{s}(1 - \cos\theta^*), \quad \hat{u} = -\frac{1}{2}\hat{s}(1 + \cos\theta^*). \quad (3.5)$$

The scattering angle θ^* is directly related to the rapidities¹ of the outgoing partons $\pm y^*$ in the center-of-mass frame [54]:

$$\cos\theta^* = \tanh y^* \quad (3.6)$$

The rapidity in the center-of-mass frame y^* can be expressed in dependence on the rapidities $y_{3,4}$ of the outgoing partons in the laboratory frame

$$y^* = \frac{|y_3 - y_4|}{2}. \quad (3.7)$$

With Equation (3.7), the scattering angle in the center-of-mass frame can be directly related to the rapidities of the final-state partons in the laboratory system:

$$\cos\theta^* = \tanh\left(\frac{|y_3 - y_4|}{2}\right) = \tanh y^* \quad (3.8)$$

As will be discussed in the following, the angular dependence of the dominant contributions to the cross section of the $1 + 2 \rightarrow 3 + 4$ parton scattering process are completely described by the angle θ^* . Equation (3.8) relates this quantity to measurable quantities in the laboratory frame, namely the final state rapidities y_3 and y_4 , and therefore allows analyses to utilize this variable. In the following, the cross section for $1 + 2 \rightarrow 3 + 4$ parton processes is derived as well as its angular dependence.

¹The rapidity y of a particle with energy E and momentum p is defined as $y = \frac{1}{2} \log \frac{E+p_z}{E-p_z}$.

3.2.2. Cross Section of Dijet Events

According to the factorization theorem Equation (3.3), the cross section for the $1 + 2 \rightarrow 3 + 4$ scattering process can be written as the product of a PDF part and a hard scattering part. Using this, the differential cross section for $2 \rightarrow 2$ parton scattering can be expressed as follows [30]

$$\frac{\partial^2 \sigma}{\partial m_{jj}^2 \partial \cos \theta^*} = \sum_{a,b} \int_{x_a=0}^1 \int_{x_b=0}^1 dx_a dx_b f_{a/A}(x_a, \mu_F^2) f_{b/B}(x_b, \mu_F^2) \delta(x_a x_b s - m_{jj}^2) \frac{d\hat{\sigma}_{ab}}{d \cos \theta^*} \quad (3.9)$$

where the delta distribution δ ensures four-momentum conservation and

$$\frac{d\hat{\sigma}_{ab}}{d \cos \theta^*} = \overline{\sum} \frac{1}{32\pi m_{jj}^2} |\mathcal{M}(ab \rightarrow kl)|^2 \frac{1}{1 + \delta_{kl}}. \quad (3.10)$$

$\overline{\sum} |\mathcal{M}|^2$ corresponds to the spin averaged sum of matrix elements which are calculated from all possible diagrams, and the prefactor $1/(1 + \delta_{kl})$ accounts for double counting in case of identical final state partons. To arrive at an expression for the cross section, all possible processes up to a certain order in perturbation theory have to be considered. Many Feynman diagrams can lead to two parton final states, examples of first order diagrams are shown in Figure 3.5. The matrix elements of the possible processes differ due to the different vertex and color factors for the involved interactions. All possible leading order matrix elements can be expressed in dependence on the three Mandelstam variables only [54]. Using Equation (3.5), they can be expressed as a function of the center-of-mass scattering angle $\cos \theta^*$. Figure 3.6 shows the resulting cross section values for the different contributions as a function of θ^* . The dominant processes are $gg \rightarrow gg$, $gq \rightarrow gq$ and $q\bar{q} \rightarrow q\bar{q}$. For all processes besides the $q_1\bar{q}_1 \rightarrow q_2\bar{q}_2$ subprocess, which is lacking a t -channel diagram, a characteristic t -channel pole can be observed. The overall angular dependence driven by the t -channel diagrams is approximately given as

$$\frac{d\sigma}{d \cos \theta^*} \sim \frac{1}{\hat{t}^2} \sim \frac{1}{(1 - \cos \theta^*)^2}, \quad (3.11)$$

which enhances the matrix elements at small scattering angles or at high values of $|\cos \theta^*|$. The differential dijet cross section $\partial^2 \sigma / \partial m_{jj} \partial y^*$ has been measured by the ATLAS experiment [61]. Results of these measurements are shown in Figure 3.7. The measured cross sections agree with the QCD predictions and range from 10^3 pb/GeV for invariant masses of 300 GeV down




Figure 3.5.: A selection of leading order diagrams contributing to production of dijet final states is shown. Diagrams for the s -channel, t -channel and u -channel are shown.




Figure 3.6.: Leading order matrix elements for two-to-two massless parton scattering. The values for the Matrix elements are taken from reference [54].

to 10^{-5} pb/GeV for invariant masses of 6 TeV for jets with $y^* < 0.5$. It therefore dominates significantly over non-QCD processes with the same signature. The cross section is monotonically and smoothly falling over many orders of magnitude. In presence of an additional contribution, for example a Z' mediator resonance, the differential cross section is enhanced for invariant masses around $m_{Z'}$, thus this particle could be observed as an excess in the invariant mass distribution. In order to enhance such an excess, the angular structure of dijet events as described by Equation (3.11) can be utilized. While QCD $2 \rightarrow 2$ parton scattering processes are dominated by t -channel poles, signal processes of interest are often given by s -channel production and decay of dijet resonances, as shown in Figure 2.7(d). Due to the absence of a t -channel diagram, the angular dependence of these signal processes is in general much flatter than for the QCD background. A strategy commonly used by dijet searches is therefore to enhance the signal to background ratio by performing kinematic selections on the center-of-mass scattering angle $|\cos \theta^*|$ to reduce the QCD t -channel contributions. For an actual event selection in data, $\cos \theta^*$ is translated into the rapidity y^* by utilizing Equation (3.8). Optimal cut values on y^* are found to be at ≈ 0.6 , which optimizes the sensitivity S/\sqrt{B} for a wide variety of signal models, where S describes the number of signal events and B the number of background events [54]. A more stringent cut of $y^* < 0.3$ is used to select higher p_T jets at a given invariant mass and provides a mass distribution which is unbiased by the leading jet selection [7].

3.3. Monte Carlo Event Generation

Sophisticated Monte Carlo (MC) simulations are used to study the properties of physical processes and to make theoretical predictions for signal and background processes, They allow for direct comparisons between the expectation of theoretical models and measured data. In the following, a short summary of the involved steps of a MC simulation is given, based on a similar description in reference [62]. All steps mentioned in the discussion of Figure 3.1 have to be included in a MC




Figure 3.7.: Measurements of the dijet cross section as a function of m_{jj} and y^* compared to NLO (next-to-leading-order) predictions, from reference [61].

simulation. In the first step, matrix elements for the hard scattering process are calculated to a certain order in perturbation theory. These are then convoluted with the PDFs of the incoming partons, where choices for μ_R as well as for μ_F have to be made. The resulting differential cross sections can be used to generate four-momentum vectors for the final state particles using pseudo random number generators. Due to the random nature of this approach, the term *Monte Carlo Event Generator* has been established for programs applying this procedure.

Afterwards, radiation of additional partons via ISR or FSR is performed in the simulation of the parton showers. The shower development is described using an evolution process in momentum transfer starting from the initial momentum scale of the hard process down to the scale of hadronization. The evolution is performed using a step-by-step approach based on the DGLAP equations. If the matrix element was calculated using higher orders of perturbation theory, a matching of the parton showers to the matrix elements is necessary to avoid double counting, as high order corrections to matrix elements can also contain radiation of hard partons. After the shower development has stopped at the hadronization scale, the transformation of low-momentum partons into hadrons is done using phenomenological fragmentation models. A frequently used model is the so-called *string model*, where the potential between two quarks is described by gluonic color strings similar to Equation (2.20). Eventually, decays of unstable final state hadrons are simulated.

The evolution of the underlying event is included via the simulation of two effects. Additional interactions of partons in the proton not participating in the hard scattering can occur, so-called multiple parton interactions. In addition, the proton remnants can be left in colored states leading to the creation of additional low energy hadrons. These non-perturbative effects are simulated using phenomenological models whose parameters are usually determined from fits to data.

A variety of MC generators are available, implementing all or parts of the steps introduced in this section. In the following, the generators used in the scope of this thesis are introduced briefly.

3.3.1. Madgraph

Madgraph [63] is a program which can automatically generate tree level and one-loop level matrix elements for arbitrary processes. It does so by applying a technique for the algorithmic evaluation of renormalized one-loop amplitudes as well as special formalism for the treatment of infrared divergences. The user simply has to specify the initial and the final states for the process of interest. Feynman diagrams for the desired process are then generated by Madgraph. Using Feynman rules, the matrix elements are evaluated and the results can be used for the prediction of observables, for example cross sections, or for event generation using automated phase space integration. Madgraph also allows putting restrictions on the considered phase space, which is particularly important to reduce statistical uncertainties in phase space regions with low cross sections. While a simulation of parton showering and hadronization is not performed in Madgraph, it is possible to use the provided interface to external programs, for example Pythia.

3.3.2. Pythia

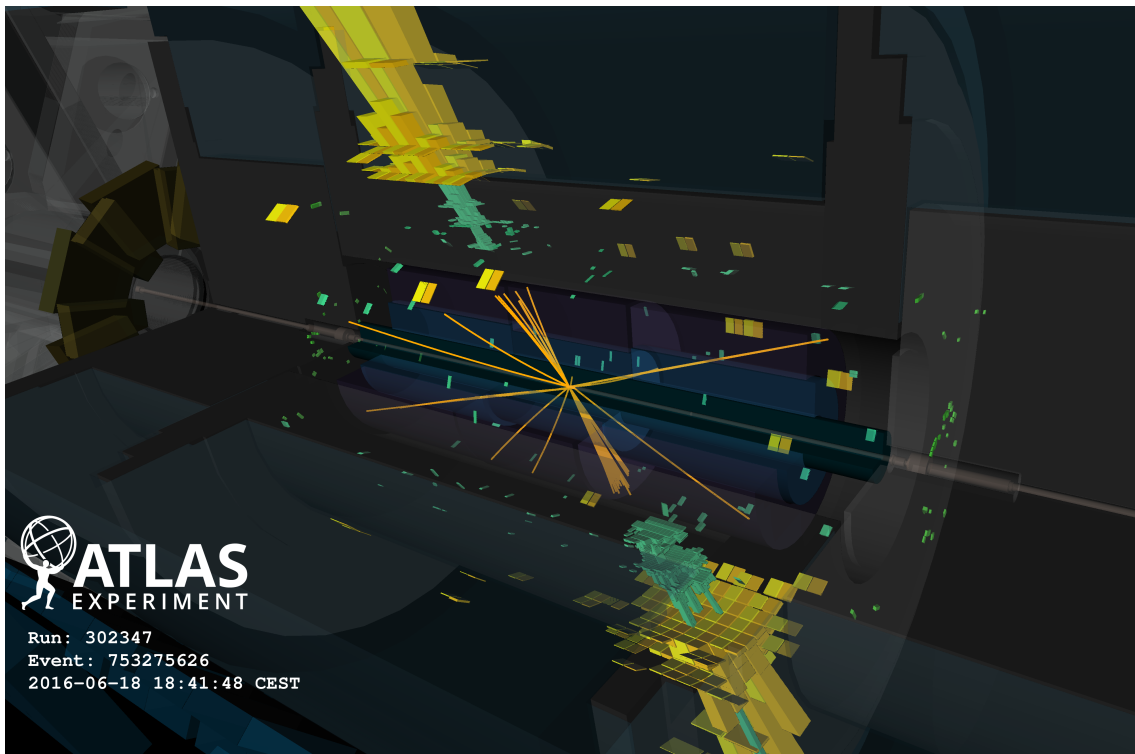
Pythia [64] is a general purpose event generator with a wide range of applications in high energy physics. It provides support for the generation of many hard coded processes, but no automatic generation of new processes. Instead, interfacing with general-purpose matrix event generators like Madgraph is possible. Pythia simulates the parton showering, multiple parton interactions, interaction of beam remnants and the underlying event. It also models hadronization using the Lund string model [65] and decays of the resulting hadrons.

Part II.

The Experiment

Technology frightens me to death. It's designed by engineers to impress other engineers. And they always come with instruction booklets that are written by engineers for other engineers – which is why almost no technology ever works.

— John Cleese



An event display of event 753275626 in run 302347 showing a dijet event recorded with ATLAS is shown. The two central p_T jets have transverse momenta of 3.13 TeV and 2.98 TeV, with a total invariant mass of 7.5 TeV. For both jets, tracks of the contained charged particles are visible in the inner detector (yellow trajectories), as well as high energy depositions in the electromagnetic calorimeter (green blocks) and in the hadronic calorimeter (yellow blocks). From reference [66].

Technological Foundations

4.1. Introduction

In experiments in high-energy physics, high data rates and small latency budgets make it impractical to use common hardware like CPUs or GPUs for the initial stages of event processing. Both CPUs and GPUs cannot provide the desired low latencies of $\mathcal{O}(100\text{ ns})$ for fast, real-time data processing and are limited in their Input/Output (I/O) capabilities, where data rates in the Gbit range are required. While CPUs also lack processing power, especially due to their limited parallelization capabilities, GPUs are lacking in terms of generality and customizability for problems with limited parallelizability. For these reasons, data processing hardware is mostly designed based on custom-built electronics utilizing ASICs (Application-Specific Integrated Circuits) and FPGAs (Field-Programmable Gate Arrays). FPGAs are integrated circuits which can be flexibly reconfigured after manufacturing. The first FPGAs were developed as an alternative to ASICs, which contain hard-wired logic that cannot be redesigned after manufacturing. Due to the possibilities of reconfiguration also during development, FPGAs provide a huge improvement in terms of flexibility as well as production costs. Therefore, they provide the ideal functionality for data acquisition systems in high energy physics, where fast processing of data is essential as well as flexibility and the ability to adapt to evolving data-taking conditions and requirements. FPGAs implement this versatility by combining programmable logic blocks with reconfigurable instead of hard-wired interconnections. In addition to these basic blocks, FPGAs contain specialized blocks for individual tasks, for example for high-speed multi-gigabit I/O, dedicated memory blocks, support for advanced arithmetic operations, and sophisticated devices for efficiently distributing clocks and clock networks. Figure 4.1 shows a schematic depiction of the structure of an FPGA with some typical logic blocks.

In the following sections, the blocks most relevant for the contents of this thesis are described in more detail. Since the FPGAs used in the scope of this thesis are designed and vended by Xilinx, some information presented is specific to Xilinx FPGAs. The FPGA that is targeted in this thesis is from the Virtex-7 family [68]¹. Section 4.2 introduces the logic blocks available on these Xilinx FPGAs. In Section 4.3, the IPBus protocol is introduced, which plays an important role in module communication. The programming of FPGAs is done in a Hardware Description Language (HDL), which allows a detailed and formal description of electronic circuits. A brief introduction into programming of FPGA firmware using VHDL is given in Section 4.4, which also introduces *Vivado*, the FPGA development environment provided by Xilinx. Most of the information presented in the following is extracted from the respective Xilinx manuals and indicated in the text, some additional information is taken from previous theses on FPGA firmware development [69, 70].

¹The exact model number is XC7VX690TFFG1927-3.




Figure 4.1.: A simplified schematic of an FPGA is shown, containing programmable interconnections, I/O blocks, and logic blocks (CLB and DSP), from [67]. The logic blocks are arranged in columns on the device.

4.2. Setup of an FPGA

4.2.1. Configurable Logic Blocks

Configurable Logic Blocks (CLBs) [71] are the primary FPGA resource used to implement combinational logic. A CLB can consist of one or multiple slices², which can be understood as a group of directly connected logic resources, and of the corresponding configurable interconnecting routing resources. The logic elements in a slice are lookup tables (LUTs) and flip-flops. A LUT implements a truth table by storing the table values in an array in memory. This functionality is used to replace an arithmetic or logic operation by a simpler lookup operation in the corresponding truth table. Any n -input logic function can be encoded by an n -bit LUT by storing the corresponding truth table in the FPGA. Due to this, LUTs are the most versatile logic blocks and can be applied to implement a wide range of designs. A complete reconfiguration of their functionality is possible by simply changing the memory values. Flip-flops are clock-triggered circuits which can be used to store small amounts of data for an arbitrary amount of time. In the considered device, a CLB contains eight six-input LUTs and sixteen flip-flops. Each six input LUT (LUT6) consists of two five-input LUTs (LUT5) with shared inputs. Since one of the LUT outputs is connected to a multiplexer, the possibility to either use shared or distinct outputs exists, as shown in Figure 4.2. This setup can be used to implement any six-input Boolean function, two independent five-input Boolean functions with shared inputs or two completely independent functions with three and two inputs or fewer. The LUT outputs can either be directly connected to slice outputs, or registered in a flip-flop storage element. Higher-order Boolean functions can be implemented by connecting two or more LUT6 via dedicated multiplexers [72]. In addition, carry logic is provided

²For the Virtex-7 FPGA a CLB contains two slices.

to implement arithmetic functions, for example adders, counters or multipliers. Two different type of slices exist, the SLICEL (“Logic”) and the SLICEM (“Memory”). While the SLICEL contains only the LUTs and storage elements shown in Figure 4.3, the SLICEM has additional functionality to run it as distributed 64 bit Random-Access Memory (RAM). This functionality allows the usage of SLICEM blocks as additional storage elements if this is required by the design.




Figure 4.2.: Six input LUT containing two five input LUTs (shown in red), modified from [72]. In dependence on the value of the input I5, the output O6 is set to either output of the five input LUTs via a multiplexer (shown in green).

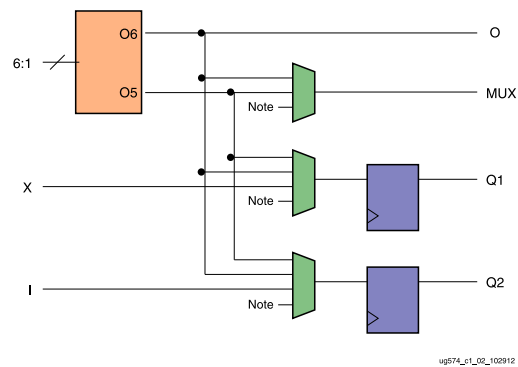


Figure 4.3.: Schematic of a LUT (orange) connected to slice I/O and to flip-flops (blue). The multiplexer (green) inputs include carry logic and wide multiplexers (not shown). In total, eight of these blocks form a CLB. Modified from [73].

4.2.2. DSP Slices

The approach to implement n -input logic functions via the combination of LUTs requires an exponentially increasing amount of LUT resources, as a 7-input LUT requires the combination of two 6-input LUTs, an 8-input LUT already four 6-input LUTs, and so on. Therefore, this method gets increasingly inefficient for the implementation of operations with wider inputs. Two common operations which suffer from this limitation are addition and multiplication. For this purpose, Xilinx FPGAs contain dedicated optimized addition and multiplication blocks, which are implemented inside so-called Digital Signal Processor (DSP) slices [74]. A schematic view of a DSP48E1, as used on the 7-Series FPGA devices, is shown in Figure 4.4. The DSP in the 7 series devices contains a 25 bit pre-adder, a 25×18 bit multiplier and an 48 bit Arithmetic Logic Unit (ALU), which can be used to implement an additional \pm operation or bitwise logic operations on its inputs, either the input C or the previous result P. All the flip-flops included in the DSP as well as the pre-adder are optional, and can be excluded from the data flow if not required. This allows the flexible implementation of various operations combining additions, subtractions and multiplications. DSP functionality for wider inputs or more complex operations are supported by cascading multiple DSP slices.

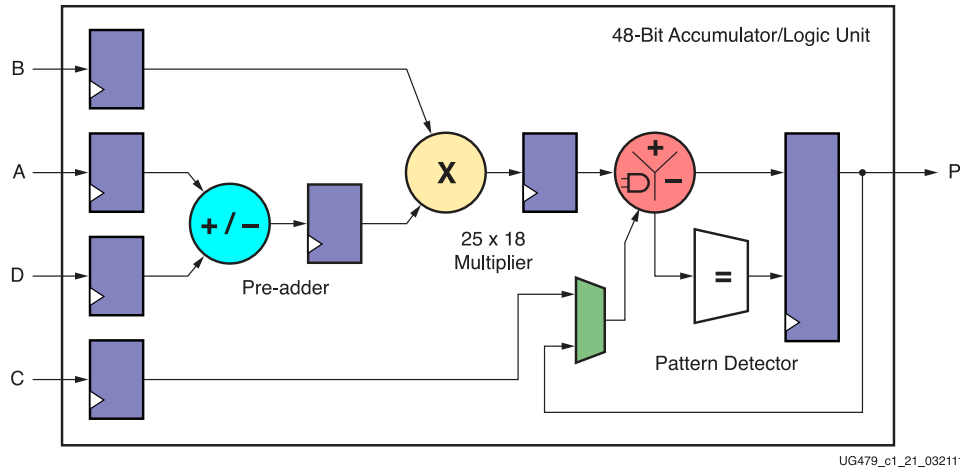


Figure 4.4.: Simplified architecture of a DSP, modified from reference [74]. The most relevant parts are colored for a better overview: flip-flops in dark blue, multiplexers to choose between inputs in green, the pre-adder in cyan, the multiplier in yellow and the ALU in red. The most general arithmetic operation that a DSP can implement is given as $P = (A \pm D) \cdot B \pm C$.

4.2.3. Block RAM

In order to store large amounts of data inside an FPGA, dedicated RAM blocks exist, so-called Block Memories (BRAMs) [75]. Compared to SLICEM, BRAMs provide an increased memory capacity as well as a higher memory density on the device. The BRAMs in 7 Series FPGAs can store up to 36 Kibit³ and can be configured as either two independent 18 Kibit RAMs, or one 36 Kibit RAM. Each BRAM has two independent write and read ports, therefore the BRAM can be run in various configurations. Besides running each 18 Kibit memory independently with a single read or write port, it is also possible to use the different ports independently to access the same memory block in the so-called dual-port mode. Two different dual-port operation modes are possible: in the simple dual-port mode, one write and one read port are used, while two bi-directional ports are used in the true dual-port mode. All read and write operations are synchronous, clock-triggered operations and require therefore a clock edge. Different configurations of a BRAM are also possible in terms of memory width and memory depth, where the width describes the number of bits in a single data word, while the depth describes the total number of words that can be stored. Both of the ports of each 18 Kibit BRAM can be configured independently as 36 bit wide and 512 elements deep memories, but different configurations like 32×576 , 18×1024 , 4×4096 or 1×16384 are also possible. Analogously, the usage as a single memory unit allows for configurations of 72×512 down to 1×32768 . Support for other, compatible operation modes exists as well. Larger memories can be obtained by combining multiple BRAMs, as these blocks are cascable. Figure 4.5 shows a schematic of the BRAM setup and its connectivity in a 7 Series device in the true dual-port mode.

4.2.4. Multi Gigabit Transceivers

To provide the necessary functionality to handle the high bandwidths required for data transmission, FPGAs are equipped with so-called Multi Gigabit Transceivers (MGTs) [76]. Line rates from

³The binary prefix kibi denotes 2^{10} .

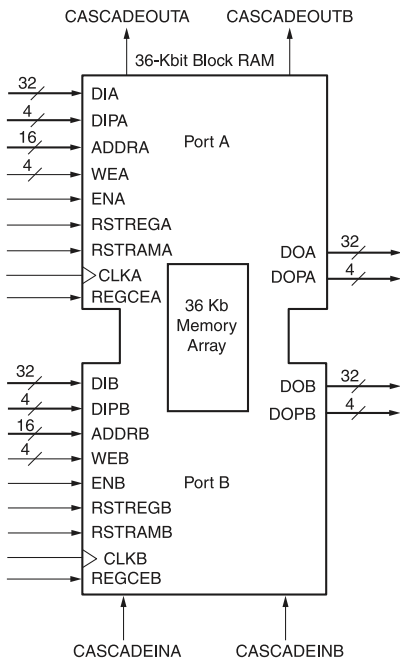


Figure 4.5.: Data flow for a dual-port BRAM, from reference [75]. The 32 bit wide data input (DI), the 16 bit memory address (ADDR), the clock signal (CLK) and the 32 bit wide data output (DO) buses are shown for both ports A and B. The bits of the parity ports (DIP, DOP) can be used for additional data inputs/outputs and are used for port widths of $\times 9$, $\times 18$, and $\times 36$. To enable or disable writing, the write-enabled (WE) port is used, the enable port (E) can be used to enable or disable a port completely. The CASCAD EIN and CASCAD EOUT ports are provided to combine multiple BRAMs to form larger memories. Other shown ports are not relevant in the scope of this thesis.

500 Mbs⁻¹ up to 13.1 Gbs⁻¹ are supported by the transceivers of the Virtex-7 FPGAs. An MGT can be used either as a receiver, which deserializes the incoming data for parallel data usage in the FPGA, or as a transmitter, which serializes parallel data from the FPGA for the outgoing data transmission. MGTs transmit or receive data as a serial stream of bits, providing the same speed at lower costs compared to parallel data interfaces. Besides serialization and transmission, many additional functionalities are supported, for example different data encodings, tools to recover the clock from incoming data, or tools to convert the parallel data clock to a higher frequency serial clock. Data consistency checks for error detection can be performed via the calculation of check sums, for example using the Cyclic Redundancy Check (CRC) [77].

4.3. The IPBus Protocol

Most electronic systems within recent high-energy physics experiments are based on the ATCA (Advanced Telecommunications Computing Architecture) standard, which provides an open, multi-vendor architecture for modern high-speed interconnect technologies [78]. For the purpose of providing a protocol to read and modify the memory spaces of ATCA boards using external software applications, a dedicated protocol has been developed for high energy physics, the IPBus protocol [79]. It is a packet-based protocol to control hardware like FPGAs via IP (Internet Protocol) over Ethernet using UDP (User Datagram Protocol). The IPBus protocol defines several operations, two of them listed below

- **Read** A read command with user-definable depth. It is possible to read multiple continuous registers in address space (incrementing read) and single ports (non-address-incrementing).
- **Write** A write command with user-definable depth. As with read commands, both incrementing and non-incrementing writes are available.

Both read and write operations require the address of the memory space that the operation interacts with. For each IPBus operation, the client sends a request to the IPBus device, which then sends back a response message with an error code indicating whether the transaction was successful, followed by the returned data in case of a successful read operation. Both the address and the data word width are 32 bit. To reduce overhead and make the transactions faster, multiple transactions can be concatenated into a single IPBus packet. The IPBus suite consists of three major components: a firmware implementation of an IPBus UDP server, which is written to run alongside the main processing firmware on the same FPGA and requires no further external resources, the ControlHub software application, which mediates simultaneous accesses from multiple clients and implements the IPBus reliability mechanism, and the μ HAL C++/Python programming interface for read and write requests, which is used to develop the corresponding software. Due to its simplicity, UDP has been chosen as the transport protocol in order to minimize the FPGA resource usage. The additional reliability mechanism allows the client to correct for packet loss, duplication or reordering of packages.

The latency and block transfer throughput of IPBus requests have been measured [79], as these are crucial parameters in a real-time application with possible latency constraints. Figure 4.6 shows the measurement results for one software client controlling one IPBus device via the ControlHub. For several hundred words, the latency is in the magnitude of a few hundred μ s. The throughput is given as several hundred MBit/s and slightly above 0.5 GBit/s for payloads larger than 1 MB, providing both low latency and high throughput required for real-time applications.




Figure 4.6.: The median write/read latency (left plot) and throughput (right plot) as a function of the number of words for one software client controlling one IPBus device is shown, from [79].

4.4. Firmware Development

To design firmware for FPGAs, HDLs are commonly used, of which VHDL (Very High Speed Integrated Circuit Hardware Description Language) [80] and Verilog [81] are the most common choices. In these languages it is possible to describe circuits, their combinational logic, registers and wiring on a very abstract level. Besides that, the possibility for *high level synthesis* exists, where code from a high level language like C/C++ is translated into FPGA designs. Choosing the correct tool is a matter of compromise: while offering advanced features of a high-level language,

high level synthesis approaches provide less flexibility and control over the resulting design, as the design is described on a more abstract level. In this thesis, all firmware development is done using VHDL.

In VHDL, all designs are hierarchically built up from smaller blocks, which are connected with each other. Every design entity consists of three parts: the interface definition, which is specified in the *entity*, the actual behavioral description in the *architecture*, and the configuration. The entity declaration describes a list of *ports*, which defines the input and output signals of the entity, and a list of *generics*, which are configurable parameters of the entity. The behavioral description contains all combinational logic, registers and their connections inside the entity. It is possible to instantiate other entities inside an architecture, which is essential for structuring more complex designs. The configuration specifies the actual wiring and defines which external signals are connected to which ports, where signals can be understood as the wires between the ports of an entity or between two logic blocks. Each FPGA design has a *top module entity*, which specifies the connections of the physical pins of the FPGA, i.e. the connections to external resources or devices. Inside this top module, all other entities are instantiated and connected with each other. To implement the desired functionality, it is possible to either directly instantiate design primitives, for example a DSP, or to implement the operations directly on a more abstract level and let the design tools infer the required resources. The former method allows for full control over the exact placement of the blocks, while the latter provides more flexibility and the possibility to optimize the design, in terms of performance or the utilization of the different resources.

For each entity port and each signal, a data type has to be specified. The basic data type is a `std_logic` representing a single bit, which can in total take one of nine possible values, among those '0' and '1' for logical values, or special values indicating uninitialized or undefined values. A `std_logic_vector` is an array of `std_logic`, representing multiple bits and the data type which is used the most. For arithmetic operations, it is possible to interpret a `std_logic_vector` in one of two possible ways, as an *unsigned* binary numbers or as a *signed* two's complement binary number. In addition, it is possible to define arrays of `std_logic_vector` or their signed/unsigned interpretations, as well as records, which group several `std_logic_vector` in one structure.

4.5. Xilinx Vivado

Xilinx Vivado [82] is an integrated development environment for FPGA firmware development. Using Vivado, it is possible to program, assemble, implement and validate an FPGA design using either a Graphical User Interface (GUI) or a shell in the Tcl scripting language. In the context of Vivado design flows, the design specified by VHDL code is referred to as a Register Transfer Level (RTL) design, as the code describes the data flow between registers. In the following, a summary of the steps involved in creating the FPGA configuration file from a RTL design is given: *synthesis* and *implementation*. Afterwards, a discussion of various design characteristics categorizing the quality and success of a completed implementation is given. The section is closed with an introduction of IP Cores, predefined code blocks shipped with Vivado which implement commonly required functionality, and a short introduction to waveform simulations.

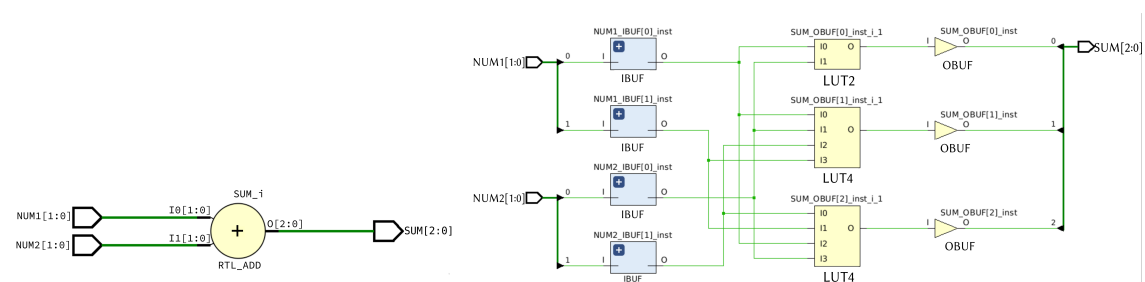
4.5.1. Synthesis and Implementation

In the first synthesis step, an elaborated design is created from the RTL design by translating the logic described in VHDL code into a description of generic behavioral components and their relations. This is done by translating the hardware description into Boolean functions which are

optimized with certain algorithms and device-independent netlists are created. Afterwards, the blocks of the elaborate design are translated to the logic blocks of the device. Figure 4.7 shows a simple example of how the elaborated RTL design compares to the synthesized design for a 2-bit adder implementing the operation $\text{sum} = \text{num1} + \text{num2}$. While the RTL design still contains the “+” operation, the synthesized design has mapped the addition to three LUT blocks, one for each of the three bits of the summation result. Several optimization steps are already performed during synthesis, where unused components are removed and constant signals are propagated. After the synthesis has finished, an estimation of the required logic resources is given. However, since the design has not been adapted to the specific device yet, the final resource utilization can deviate from this estimate. In addition, no reliable latency information is available.

This information is provided after the next step, the implementation. It contains three major steps, namely physical design optimization, placement and routing. Before these steps are executed, possibly specified constraints on clocks or logic block placements are imposed on the design. During design optimization various steps are involved to improve the design, for example merging of instances with equivalent logic or the removal of unrequired logic blocks. Depending on the exact strategy, optimization is executed with a different goal, for example making it easier to fit the design on the given device or to minimize the power demands of the design. In placement, the optimized design is now actually placed on the device, the required resources as determined by the synthesis are allocated to the physically available logic blocks of the device. The placement step is a complex and time-consuming task, especially for larger, complex designs, as it is desired to have all signal paths as short as possible to minimize latency and power consumption.

Afterwards, the routing step fixes the necessary signal paths by connecting the logic blocks of the design with the configurable interconnections. Short paths are desired if possible to meet the timing constraints, as discussed later in more detail in Section 4.5.2. However, in large, complex designs local network congestions and path complexity have to be taken into account, making the routing a non-trivial and time-consuming task. Depending on the chosen strategy, additional optimization steps can be executed at various points, for example after placement or after routing. After the implementation is finished, Vivado reports the final resource utilization and checks whether the timing constraints are met. If everything was successful, the last remaining step is the *bitstream* generation. The bitstream is a binary file which can be directly loaded onto the FPGA and specifies the final configuration of the FPGA.



(a) RTL adder in the elaborated design, (b) Synthesized adder, where the addition has been translated where the operation is given as RTL_ADD. into logic blocks available on the device, in this case three LUTs. The IBUF/OBUF blocks are used for connections to the physical pins of the FPGA.

Figure 4.7.: Comparison of the elaborated design representing the VHDL code (a) and the synthesized design (b) for a simple 2-bit adder.

4.5.2. FPGA Design Characteristics: Resource Utilization and Timing

In real time applications, there are often constraints on the total number of clock cycles that a design can use. In order to perform complex operations in these timing constraints, it is desirable to perform as much combinational logic between two flip-flops as possible. At the same time, the time a signal requires on its path between two registers is not allowed to be higher than the clock period. If this is not the case, data will arrive too late at a register, leading to errors as the data is now out of sync. In addition, the amount of available logic blocks is limited, putting constraints on the total amount of possible parallel logic operations in a clock cycle. Therefore, multiple quantities characterizing the design quality and performance are reported after the synthesis and implementation step. These quantities show whether the design is appropriate for the target device and whether it is expected to work correctly when loaded on the FPGA. Two groups of characteristics are important, the first one quantifying *resource utilization*, while the second one quantifies *timing*. Resource utilization indicates whether the logic resources available on the FPGA are sufficient for the design requirements. When evaluating the designs in this thesis, the resource utilization of LUTs, DSPs, flip-flops and BRAMs are reported if appropriate. The total resources provided by the Virtex-7 FPGA are listed in Table A.1. For the timing, Vivado calculates how long the signals propagate on the paths between registers, and reports whether these durations conform to the timing constraints given for the design, i.e. the clock frequency.⁴ Vivado summarizes the timing of a design in terms of multiple timing characteristics, among them the negative slack and the hold slack. In general, the negative slack of a path corresponds to the difference of the time when the rising edge of a clock arrives and the time when a signal update arrives at the register. Hence, a positive value indicates that a signal has arrived early enough, while a negative value means that the data has arrived too late. Since data values which propagate very fast from one register to the next can be caught by the same clock edge and overwrite the previous data, a minimum time is required for each path. The hold slack corresponds to the difference of the time when data arrived and the minimum time that was required for the data. In this case, a positive value means that more than the minimum time was required, while a negative value means a hold violation. To meet the timing constraints, all paths of the design must meet the criteria from both negative slack and hold slack.

4.5.3. IPCores

Vivado is shipped with a library of Intellectual Property Cores (IP Cores), which implement specific and often required functionalities. An IP Core can be understood as a predefined entity, which can be reused in multiple places or designs. Their main advantage is that professionally written and well tested code can be reused instead of reimplementing complex functionality, leading to shorter development cycles. IP Cores are offered for a wide range of functionalities, for example for peripheral controllers like Ethernet controllers, generation of multiple clock frequencies, complex arithmetic operations like Fast Fourier Transformations or for the efficient and explicit allocation of FPGA resources. IP Cores are used in this thesis primarily for the allocation of BRAM resources [83].

⁴Throughout this thesis, the clock frequency is given by an integer multiple of the LHC frequency, which is 40.08 MHz, corresponding to a time period of 24.95 ns.

CERN and the LHC

In order to validate theoretical predictions with high precision measurements and to investigate possible new physics beyond the Standard Model, experiments at previously inaccessible collision energies are necessary. To reach these energies as well as very high beam intensities, the Large Hadron Collider (LHC) was designed, providing the world's highest center-of-mass energies of up to 14 TeV [85]. The LHC is a 27 km long synchrotron and can be operated to accelerate protons as well as beams of heavy ions. It is located in a tunnel up to 175 m beneath the Geneva region at CERN, the European Organization for Nuclear Research. At CERN, fundamental research in the fields of nuclear and particle physics is realized using accelerators and particle detectors. The organization was founded in 1954 and consists of 23 member states as of October 2019 [86].

This chapter starts with a brief overview of the entire proton accelerator complex at CERN in Section 5.1, followed by an introduction to the experiments at the LHC in Section 5.2. Afterwards, Section 5.3 presents the key parameters characterizing the LHC performance. This chapter concludes with a discussion of LHC filling schemes and bunch groups in Section 5.4.

5.1. The CERN Accelerator Complex

The protons at CERN are not only accelerated by the LHC, but by a whole chain of accelerators with increasing energies. The entire acceleration complex is shown in Figure 5.1. The acceleration process starts at a hydrogen bottle, where hydrogen atoms are ionized and the resulting protons are injected into the accelerator chain. They are accelerated linearly to an energy of 50 MeV in the LINAC2 and are then injected into a sequence of circular accelerators, starting with the PS Booster. The proton beam, which now contains protons with energies of 1.4 GeV, is then fed to the Proton Synchrotron (PS) and the protons are accelerated to 25 GeV. In addition, the PS prepares the proton bunch structure for the LHC, for example with 72 bunches and a bunch spacing of 25 ns [88]. Afterwards, the protons are transferred to the Super Proton Synchrotron (SPS), where usually two to four PS batches of 72 bunches are accumulated and proton energies of 450 GeV are reached. Finally, the proton beam is injected into the two beam pipes of the LHC in a clockwise and an anti-clockwise direction. Currently, each proton can reach energies up to 6.5 GeV in the LHC. For Run-3 of the LHC, which will start in 2021, an increased maximum beam energy of 7 TeV is foreseen. In order to bend the charged protons on their circular path along the accelerator, 1232 superconducting dipole magnets are used to provide the necessary magnetic fields of more than 8 T. Additional transverse beam focusing is provided by 392 quadrupole magnets.

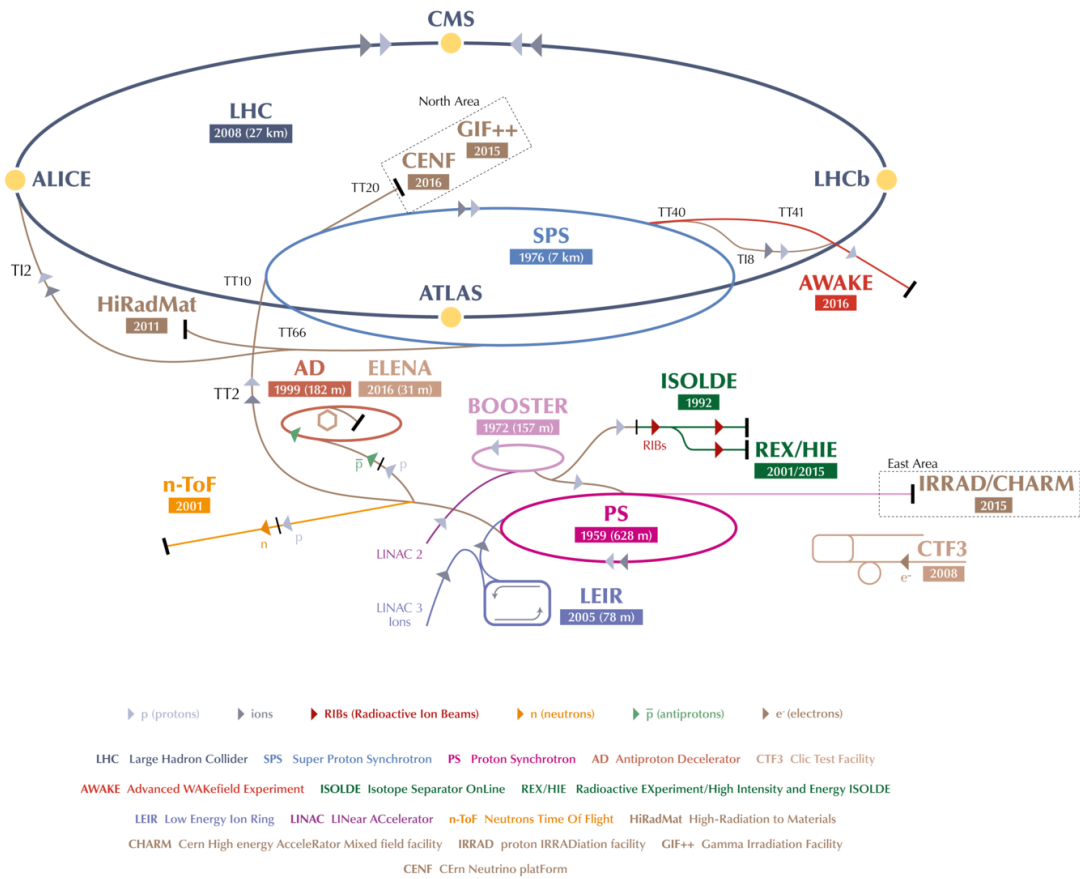


Figure 5.1.: Overview of the CERN accelerator complex including the LHC accelerator, from reference [87]. The different accelerators are shown with their spatial dimensions and year of start-up. The four main experiments at the LHC are shown, namely ATLAS, CMS, LHCb and ALICE.

5.2. The Experiments at the LHC

As shown in Figure 5.1, four main experiments are located at different interaction points of the LHC. Two experiments, the ATLAS [89] and the CMS [90] experiments, are built as multi-purpose detectors and cover a wide variety of physics searches and measurements. Both experiments, while sharing the same scientific goals, have been designed independently in order to make cross-confirmation of new discoveries possible [91]. The LHCb experiment [92] is designed to measure physics involving bottom quarks. Precision measurements of rare B meson decays and CP violation are performed to test the Standard Model. The ALICE experiment [93] is designed to primarily study heavy-ion collisions. During these collisions the energy density allows for the creation of quark-gluon plasma, a state of matter in which quarks and gluons are freed at very high temperatures comparable to the early universe.

5.3. LHC Performance

The performance of the LHC depends not only on the beam energy, but also on the number of colliding protons, the transverse beam dimensions and other parameters. A key parameter to quantify the performance of an accelerator in dependence on these parameters, the *luminosity*, is introduced in the following. The instantaneous luminosity \mathcal{L} relates the cross section σ_{process} of a physical process to the event rate \dot{N}_{process} :

$$\dot{N}_{\text{process}} = \frac{dN_{\text{process}}}{dt} = \mathcal{L} \cdot \sigma_{\text{process}}. \quad (5.1)$$

Equation (5.1) writes the event rate of a process as the product of the cross section σ_{process} , which contains all information about the physical process, and the luminosity \mathcal{L} , which only depends on properties of the collider and beam parameters. In order to increase the event rate for a given process, a higher luminosity is required. Assuming two identical Gaussian beam distributions in the transverse plane, the luminosity can be written as a function of multiple beam parameters [94]:

$$\mathcal{L} = \frac{N_p^2 n_b f_r}{4\pi \Sigma_x \Sigma_y} F. \quad (5.2)$$

N_p describes the number of protons per bunch and n_b the number of bunches per beam, $f_r = 11.245$ kHz the beam revolution frequency, Σ_x the horizontal and Σ_y the vertical dimension of the beam in the transverse plane. F describes a geometrical correction accounting for the crossing angle with which the beams are colliding at the interaction point. While the numerator in Equation (5.2) corresponds to the number of proton-proton crossings per time interval, the denominator describes the transverse beam dimensions. Hence, a high luminosity corresponds to a high interaction rate at an interaction point where the transverse beam dimensions are as small as possible. Table 5.1 shows the values of multiple beam performance parameters for the LHC runs of 2017 and 2018 compared to the predictions for Run-3.

Table 5.1.: LHC parameters for the 2017 and 2018 runs [95, 96] compared to the plans for Run-3 [96]. The quantity $\beta^* \sim \Sigma^2$ is the beta function at the collision point and directly related to the transverse size of the beam, therefore a smaller β^* corresponds to smaller transverse beam dimensions.

Year	E_{beam} [TeV]	N_p [10^{11}]	n_b	β^* [cm]	$\mathcal{L}_{\text{Peak}}$ [$10^{34} \text{ cm}^{-2} \text{ s}^{-1}$]
2017	6.5	1.25	1900,2556	40 to 30	2.09
2018	6.5	1.1	2556	30 to 25	2.10
Run-3	7	≤ 2.1	≤ 2748	30 to 25	~ 2

Due to the high number of protons in every proton bunch and due to the beam focusing at each interaction point, every Bunch Crossing (BC) can lead to multiple proton-proton interactions. These additional inelastic pp interactions are called *pile-up*. While a higher luminosity increases the rate of interesting events, it also leads to increasing noise levels and backgrounds due to pile-up effects, since particles from multiple interactions overlay in the detector at the same time. To quantify the level of pile-up, the mean number of interactions per bunch crossing $\langle \mu \rangle$ is specified. This number is not constant, but depends on the exact run conditions, as can be seen




Figure 5.2.: The distribution of $\langle\mu\rangle$, the mean number of interactions per bunch crossing in the ATLAS detector, is shown for the years of Run-2 [97].

in Figure 5.2, which shows the profile of $\langle\mu\rangle$ for the four years of data-taking in Run-2. In general, the level of pile-up has risen with the past years, where values of over 70 interactions per crossing have been reached in 2017 and 2018. Though, the pile-up was leveled at $\langle\mu\rangle \approx 40$ for most of the data taking in 2018, since the experiments have constraints on the amount of pile-up events that can be processed. It is possible to express the instantaneous luminosity \mathcal{L} as a function of $\langle\mu\rangle$ and the inelastic pp cross section [98]:

$$\mathcal{L} = \frac{\langle\mu\rangle n_b f_r}{\sigma_{\text{inelastic}}}. \quad (5.3)$$

$\langle\mu\rangle$ is therefore directly proportional to the instantaneous luminosity. Increasing the event rates by raising the luminosity is therefore inevitably connected with higher pile-up.

While the instantaneous luminosity \mathcal{L} relates the rate of events and the cross section, a similar quantity is used as a measure to specify the amount of data produced by a collider in a certain time interval, the *integrated luminosity*

$$L = \int \mathcal{L} dt. \quad (5.4)$$

It relates the observed number of events and the cross section of a certain process:

$$N_{\text{process}} = L \cdot \sigma_{\text{process}} \quad (5.5)$$

In the year 2018, peak luminosities of $2.1 \cdot 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and a recorded integrated luminosity of 60.6 fb^{-1} were achieved, compare Figures 5.3(a) and 5.3(b), which show both the instantaneous and total luminosity delivered at ATLAS over time.

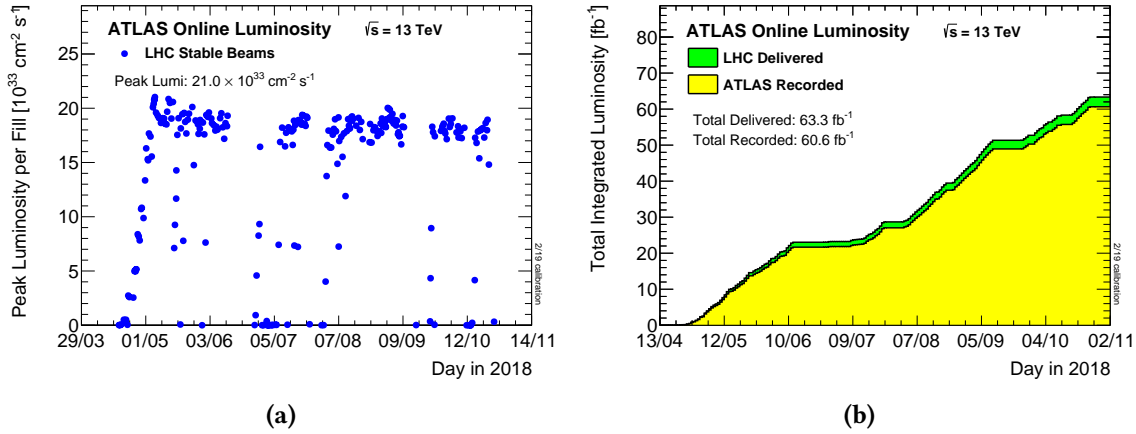


Figure 5.3.: Peak instantaneous luminosity (a) and integrated luminosity (b) at the ATLAS experiment in 2018 [97].

5.4. LHC Bunch Structure and Filling Schemes

With a revolution frequency of 11.245 kHz and a bunch frequency of 40.08 MHz, a total number of 3564 possible bunch positions exist in the LHC [99]. However, not every available bunch position actually contains protons. This is caused by the properties of the accelerator chain, since the smaller accelerators contain fewer bunches than the following larger accelerators, which accumulate the bunches from the previous accelerator to form longer *bunch trains*. When the proton bunches are transferred from a smaller to a larger accelerator, so-called injection kicker magnets are required to inject and bend the particles onto the correct trajectory. These magnets are designed to have very fast rise and fall times, though they still take longer to start up than the bunch spacing of 25 ns. In the time of the ramp up no particles should pass through these magnets, therefore short gaps are left between the bunch trains. The exact contents of all available bunch positions as well as the bunch spacing, the number of bunches colliding at the interaction points and other parameters are defined in so-called filling schemes [100]. Analogously, a total of 3564 bunch crossings per beam revolution exist for each interaction point. To identify each crossing, a unique Bunch Crossing Identifier (BCID) between 0 and 3563 is assigned to every bunch crossing in a revolution. In order to put each event in coincidence with the corresponding bunch contents, the ATLAS experiment categorizes each BCID in dependence on its bunch contents and groups them into so-called bunch groups. The collection and configuration of all bunch groups is called a bunch group set, Figure 5.4 shows an exemplary bunch group set from run 351698. The paired category is the most important one for physics, as it corresponds to a bunch crossing with protons in both beams. Calibrations of subdetectors *during* a run are performed in the CalReq category at certain BCIDs with no proton bunches at the end of the scheme. Using the Unpaired categories, which contain only a single beam, non-collision beam backgrounds can be studied. Other categories like empty bunch crossings directly before or after collisions are used for monitoring backgrounds and for slow particles searches. As the filling conditions of the LHC can change, ATLAS defines multiple bunch group sets, which can be used in dependence on the filling scheme. In ATLAS, this information is utilized in the Central Trigger Processor, where the trigger decision is set in coincidence with the corresponding bunch group.

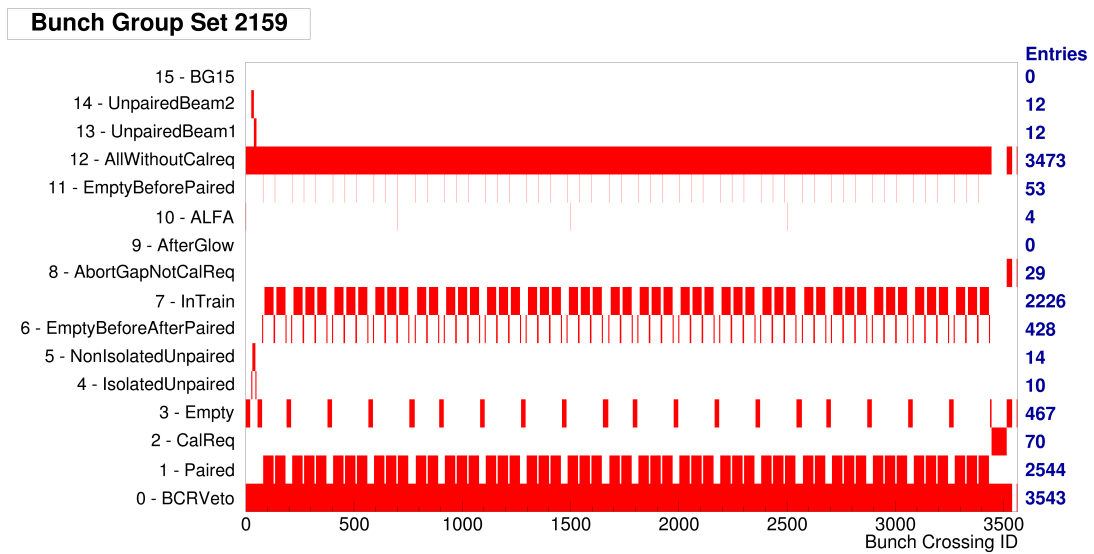


Figure 5.4.: Bunch group set 2159, as used for example in run 351698. The filled area indicates which bunch group is active at which BCID. The LHC bunch train structure can be recognized in the Paired category. The figure was created using the ATLAS run query tool [101].

The ATLAS Experiment

In this chapter, the ATLAS experiment, a large general-purpose particle detector at the LHC, is introduced. An overview of the general design of the detector is given in Section 6.1, including the coordinate system used by ATLAS and important kinematic variables. The parts of the ATLAS detector used for the measurement of particles, namely the tracking system, the calorimeters and the muon system, are introduced in Sections 6.2 to 6.4. In Section 6.5, an introduction to the detectors used for luminosity measurements is given. The trigger system is discussed in Section 6.6 with a particular focus on the level-1 trigger. Finally, an overview of detector simulation is given in Section 6.7. Most information presented in this section is taken from reference [89], additional sources are indicated when used.

6.1. General Detector Overview

The ATLAS detector is a general purpose detector and is used to identify the particles which are created in particle collisions at the LHC. It provides large angular acceptance, precise measurements of particle momenta and energies as well as the possibility to identify different kind of particles. It has a cylindrical geometry and is constructed in several layers around the beam axis. The detector is segmented in three regions: the central *barrel* region, where the detector elements are arranged in a cylindrical pattern around the beam axis, the *end-cap* region at each end of the barrel, where the detector elements are located on disks perpendicular to the beam direction, and the *forward* region, which is contained inside the end-cap region. Due to this arrangement, the detector has a forward-backward symmetry as well as a cylindrical symmetry. With an approximate weight of 7000 t and dimensions of 25 m in height and 44 m in length, the detector is one of the largest detectors at particle colliders in the world. A solid angle of almost 4π is covered by the detector, which is only interrupted by the beam pipe. Figure 6.1 shows a schematic overview of the detector.

6.1.1. Coordinate System and Kinematic Variables

ATLAS uses a right-handed coordinate system, with the interaction point in the center of the detector as its origin. The z -axis corresponds to the beam direction, whereas the x - y plane is transverse to the beam. To completely define the coordinate system, the positive x -axis is required to point from the interaction point in the detector towards the center of the LHC ring, while the positive y -axis is defined as pointing upwards. The position of a particle in the detector is described by spherical coordinates (r, θ, ϕ) , where ϕ is the azimuth angle in the transverse plane and θ is the angle relative to the beam direction. Since the momenta of the colliding partons are approximately




Figure 6.1.: Sectional view of the ATLAS detector, from [102]. The inner detector, formed by the pixel detector, semiconductor tracker and transition radiation tracker, is enclosed by the calorimeter system. The calorimeter system consists of the electromagnetic calorimeter and the hadronic calorimeter and is surrounded by the muon system.

collinear to the z -axis, longitudinal momentum components of detected particles can be remnants of the momenta of the incoming partons. In contrast to this, transverse momentum components are very likely originating from the hard physics process itself. It is therefore useful to define several variables in the transverse plane. A commonly used variable is the *transverse momentum* p_T , which is defined as

$$p_T = \sqrt{p_x^2 + p_y^2}. \quad (6.1)$$

The corresponding *transverse energy* of a particle with mass m is defined as

$$E_T = \sqrt{p_T^2 + m^2}. \quad (6.2)$$

Momentum conservation requires the vector sum of all transverse momenta to vanish before and after the collision. In case there are invisible particles like neutrinos, the total vector sum of measured transverse momenta is not equal to zero. The missing part of the transverse momentum is described by E_T^{miss} :

$$E_T^{\text{miss}} = - \sum_{i \in \text{measured}} \mathbf{p}_{T,i} \quad (6.3)$$

The magnitude of Equation (6.3) is the *missing transverse energy* E_T^{miss} . Due to the possibly different momenta of the incoming partons, the center-of-mass frame is moving along the z -axis relatively to the laboratory frame. The two frames are therefore related via a Lorentz boost along the z -axis. Since differences in the angle θ are not invariant with respect to these transformations, a different variable is used for position specifications, the *pseudorapidity* η :

$$\eta = - \ln \tan \left(\frac{\theta}{2} \right) = \frac{1}{2} \ln \left(\frac{|\mathbf{p}| + |p_z|}{|\mathbf{p}| - |p_z|} \right). \quad (6.4)$$

For relativistic particles, i. e. $E \gg m$, the pseudorapidity is equivalent to the rapidity y , which has the property of invariant differences $|y_1 - y_2|$ with respect to Lorentz boosts along the z -axis. Compared to θ , η has the additional advantage that the particle flux from hadron-hadron collisions is approximately constant per pseudorapidity interval. For massless particles, the four-momentum vector can be written as a function of E_T , η and ϕ only:

$$p = (E_T \cosh \eta \quad E_T \cos \phi \quad E_T \sin \phi \quad E_T \sinh \eta) \quad (6.5)$$

Commonly used quantities to denote distances in the angular η - ϕ plane between two particles with coordinates (η_1, ϕ_1) and (η_2, ϕ_2) are listed below:

$$\Delta\eta = |\eta_1 - \eta_2| \quad (6.6)$$

$$\Delta\phi = \begin{cases} |\phi_1 - \phi_2|, & \text{if } |\phi_1 - \phi_2| \leq \pi \\ 2\pi - |\phi_1 - \phi_2|, & \text{if } |\phi_1 - \phi_2| > \pi \end{cases} \quad (6.7)$$

$$\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2} \quad (6.8)$$

For relativistic particles, the invariant mass of two objects can be expressed in terms of angular and transverse quantities only using Equation (6.5):

$$M^2 = (p_1 + p_2)^2 = 2E_{T,1}E_{T,2} (\cosh \Delta\eta - \cos \Delta\phi). \quad (6.9)$$

6.1.2. Particle Measurements with ATLAS

When traversing the detector, particles undergo various interactions with the detector material, leading to typical, distinct signals for each class of particles. Particles created at the interaction point travel through the different layers of the detector. First, they encounter the inner detector, which consists of semiconductor- and gaseous-based tracking detectors. Charged particles, like the muon, electron or proton, ionize the material along their trajectory, which allows for a reconstruction of their tracks with high resolution. The inner detector is enclosed by a solenoid magnet, creating a magnetic field which bends the charged particles' tracks to measure their momenta and sign of charge. To determine the energies of the traversing particles, calorimeters are used, which make up the next detector layer. The Electromagnetic Calorimeter (ECAL) is designed to absorb low-mass electromagnetically interacting particles like photons or electrons, which interact with the detector material resulting in a shower of electromagnetic particles. The Hadronic Calorimeter (HCAL) is designed to absorb hadrons due to their on nuclear interactions, resulting in hadronic showers. Finally, the outermost layer, the muon spectrometer, allows for a measurement of muon tracks which are bent in the magnetic field of the toroid magnets. In the following, the components of the ATLAS detector introduced in this section are described in more detail.

6.2. The Inner Detector

The inner detector is contained within a cylindrical envelope of length ~ 3.5 m and radius 1.15 m within a superconducting solenoid magnet producing a magnetic field of 2 T. Its main tasks are the reconstruction of trajectories of charged particles as well as the precise determination of primary and secondary vertices. To achieve track reconstruction, information from three different subdetectors is combined: the pixel detector, the Semiconductor Tracker (SCT) and




Figure 6.2.: Components and layers of the inner detector in the barrel region, from [103].

the Transition Radiation Tracker (TRT). Using the track curvature of a particle in the magnetic field of the solenoid, the momentum and charge of the particle are reconstructed. A computer generated, sectional image of its setup is shown in Figure 6.2. The pixel detector is the innermost part of the inner detector and covers a region of $|\eta| < 2.5$. It is composed of four barrel layers of silicon pixel sensors including the insertable *B*-layer (IBL) which was added for Run-2, and two end-caps of three silicon pixel sensor disks each [104]. The SCT is a multi-layered silicon strip detector which follows on the pixel detector. It is composed of four double-layers of silicon strips arranged in a barrel structure, and of eighteen disks in the end-caps. In contrast to the other components, the TRT encloses only an area of $|\eta| < 2.0$. Gas-filled drift tubes interleaved with polypropylene provide numerous hits for track reconstruction as well as the ability to distinguish between electrons and tracks from hadrons.

6.3. The Calorimeter System

The calorimeter system covers a pseudorapidity region of $|\eta| < 4.9$ and encloses both the inner detector and the solenoid magnet. It has a total length of approximately 12 m and an outer diameter of 8.5 m. Energies of all charged and neutral particles, except for neutrinos and muons, are measured with sampling calorimeters. These consist of two components: the passive material, a dense and ideally high-*Z* absorber material, and the sensitive detector elements of the active material. Incoming particles lose their energies through interactions with the passive material, resulting in a conversion of the energy into a shower of secondary particles, which then cause signals in the sensitive detector elements via ionization or scintillation. Detector material and active material alternate many times within the layers of the calorimeter, allowing the reconstruction of energies of the traversing particles from the deposited energies of these showers. An overview of the calorimeter system of the ATLAS detector is shown in Figure 6.3. The design of the ATLAS calorimeter was driven by the properties of the expected particle showers, especially their shapes and spatial extensions. Due to the different kinds of involved interactions, elec-




Figure 6.3.: Illustration of the calorimeter system which encloses the inner detector, taken from [89].

tromagnetic showers resulting from high energy electrons, positrons or photons, and hadronic showers resulting from high energy hadrons are different in size and shape. Therefore, distinct parts of the calorimeter system are used to detect them, the ECAL and the HCAL. The calorimeter parts which are closest to the beam pipe, both electromagnetic barrel and end-caps as well as the hadronic end-caps and forward calorimeter, use liquid argon as active material, which provides a stable, linear response as well as radiation hardness. This part of the detector is usually referred to as the *Liquid Argon (LAr) Calorimeter*. The calorimeter parts further away from the beam pipe need to handle a less dense particle influx and lower radiation levels, therefore a different active material was chosen. The so-called *Tile Calorimeter* is made of more cost-efficient steel absorbers alternating with scintillating tiles. In contrast to liquid argon, these tiles can be operated at room temperature.

The resolution of such a calorimeter as a function of energy is described by the addition in quadrature of three η -dependent terms:

$$\frac{\sigma_E}{E} = \frac{N}{E} \oplus \frac{S}{\sqrt{E}} \oplus C, \quad (6.10)$$

where the noise term N describes detector noise, electronic noise and pile-up effects, the stochastic term S the fluctuations due to the counting of particles in the shower, and the constant term the energy loss in the passive material. Since the noise term is independent on the energy, it contributes most at low energies. The stochastic term grows with \sqrt{E} and dominates in an intermediate energy range, while the constant term dominates at high energies.

6.3.1. The Electromagnetic Calorimeter

The ECAL measures the energies of particles based on electromagnetic interactions. In the GeV range, electrons and photons repeatedly lose their energy primarily via bremsstrahlung and pair-production, resulting in a growing cascade of secondary particles of smaller energies. When the energies of the secondary particles reach the MeV range, ionization and the photoelectric




Figure 6.4.: Sketch of an ATLAS ECAL barrel module, the granularities of the different layers in η and ϕ are specified. The size of a trigger tower is indicated which combines several calorimeter cells. It is commonly specified as 0.1×0.1 in $\Delta\eta \times \Delta\phi$, even though it actually corresponds to $0.1 \times 2\pi/64 \approx 0.1 \times 0.0982$. From [89].

effect become the dominant processes, and the development of the shower is stopped when the resulting low-energetic particles are absorbed. The energy of the initial particle is determined from the signals these secondary particles cause in the active material. In order to correct for the energy deposited in the passive material, additional scaling factors are applied. The energy is thereby calibrated to the so-called electromagnetic scale (EM scale), as this corresponds to the correct energy for electromagnetically interacting particles.

The ECAL is composed of a central component in the region $|\eta| < 1.475$ (LAr electromagnetic barrel) and two end-caps in the region $1.375 < |\eta| < 3.2$ (LAr electromagnetic end-caps). Over its full coverage, accordion-shaped lead plates, which are interleaved with Kapton electrodes, are used as the absorber material. Varying thicknesses of the lead absorbers have been chosen to achieve at least 24 (26) radiation lengths X_0 in the barrel (end-caps) while still providing high energy resolution. The active material liquid argon is floating in the gaps between the electrode plates. In the central region of $|\eta| < 1.8$, a presampler is used to correct for the energy loss upstream of the calorimeter, consisting of thin layers of LAr without absorber material. Granularities in $\Delta\eta \times \Delta\phi$ between 0.025×0.025 in the barrel region and 0.1×0.1 in the end-caps are provided by the ECAL. Over the central region ($|\eta| < 2.5$) the electromagnetic calorimeter is segmented in three layers in depth, while only two layers are available in the end-cap region. Figure 6.4 shows a module in the barrel region, showing the various layers in depth as well as the accordion geometry.

6.3.2. The Hadronic Calorimeter

Due to their higher masses compared to electrons, hadrons lose only a small fraction of their total energy via bremsstrahlung. In addition, uncharged hadrons like the neutron do not interact electromagnetically at all. A different part of the calorimeter is therefore required for an energy

measurement of hadrons. Figure 6.5 shows a cascade caused by a high energetic hadron interacting with the detector material.




Figure 6.5.: An incoming hadron scatters inelastically with a single nucleon of a nucleus, resulting in a cascade of secondary particles. Adapted from [105].

Inelastic interactions with a nucleus cause the creation of secondary particles, preferably in forward direction. Secondary particles with sufficient energies can again perform inelastic interactions with other nuclei, leading to a cascade of secondary particles. The nuclei are left in highly excited states and release their energy via spallation processes by radiating nucleons and nuclear fragments. When the remaining energy reaches the binding energy of a nucleon, it is dissipated via an evaporation or fission process [12, 105]. In each of these inelastic hadronic interactions, parts of the energy are removed from further hadronic interactions due to the production of secondary π^0 and η mesons which decay almost exclusively into photons resulting in electromagnetic showers. In addition, energy carried by neutrinos and muons from charged hadron decays is undetected and lost as well as energy required to overcome the binding energies of nuclei. The variety of the involved processes cause fluctuations in the development of a hadronic shower and the different energy components. Therefore, the energy resolution of hadronic showers is worse compared to the resolution of electromagnetic showers. In general, the energy for hadronic sources is underestimated due to the invisible and escaped energy in hadronic interactions, therefore the ATLAS calorimeter has a lower *response* to hadronic showers than to the electromagnetic component. Calorimeters with this property are called *non-compensating*. Additional calibrations are therefore required when hadronic objects are reconstructed from the calorimeter signals. Analogously to the radiation length X_0 , the hadronic interaction length λ_I is defined. As λ_I is in general much larger than X_0 , hadronic showers are longer and broader than electromagnetic ones, requiring the HCAL to be thicker than the ECAL.

The HCAL of ATLAS consists of three different parts, each covering a different detector region. The region $|\eta| < 1.7$ is covered by the Tile barrel and Tile extended barrel calorimeters, where steel is used as the absorber and scintillating tiles made of polystyrene doped with para-terphenyl as the active material. Photomultiplier tubes (PMTs) with wavelength shifting fibers are used to readout the light signals produced in the tiles. The tile calorimeter contains a central barrel with a length of 5.8 m at $|\eta| < 1.0$ and two extended barrels with lengths of 2.6 m in the region $0.8 < |\eta| < 1.7$. Its total radial depth is given as $7.4\lambda_I$, each barrel consists of 64 modules of size 0.1 in $\Delta\phi$. The geometry of a Tile module is shown in Figure 6.6. Due to the high rate of particles as

well as the challenging radiation conditions, the design choices in the outer regions differ from the central barrels. The *LAr hadronic end-cap* covers a range of $1.5 < |\eta| < 3.2$ and overlaps slightly with both other systems. It consists of two wheels in each end-cap, liquid argon is used as the active material, copper plates serve as the absorptive material. The *LAr forward calorimeter* covers the region $3.1 < |\eta| < 4.9$ and consists of three modules in each end-cap. They are sharing the same cryostats as the end-cap calorimeters. The first module is made of copper and optimized for electromagnetic measurements, while the other two are made of tungsten and measure primarily the energy of hadronic interactions. A depth of approximately $10\lambda_I$ is provided.




Figure 6.6.: Module of the tile calorimeter, showing the arrangement of active and passive material as well as the optical readout assembly with front-end electronics, from [89]. Particles from the interaction point enter the module from the bottom.

6.4. The Muon System

Due to their mass, which is roughly 200 times higher than the electron mass, and their abstinence from the strong interaction, muons are able to pass the calorimeter system without considerable energy loss. Muons are therefore detected using a large spectrometer enclosing the calorimeter as the outermost detector, which covers a pseudorapidity range of $|\eta| < 2.7$. Large superconducting air-core toroid magnets provide a magnetic field of 2 T to 8 T, allowing for the reconstruction of muon momenta based on their track curvature measured in high-precision tracking chambers. A cut-away view of the ATLAS muon chambers and the toroid magnets is shown in Figure 6.7. In the central region of $|\eta| < 2.0$, three cylindrical layers of Monitored Drift Tubes (MDTs) provide track measurements. At larger pseudorapidities of $2.0 < |\eta| < 2.7$, multi-wire proportional chambers are used, so-called Cathode Strip Chambers (CSCs). In addition to the high-precision tracking chambers, trigger chambers which provide fast signals in $|\eta| < 2.4$ are present. The trigger chambers measure the angular (η, ϕ) coordinates of the track and well-defined p_T thresholds. They deliver fast signals in a time window of 15 ns to 25 ns after the passage of a particle and can therefore provide muon information for the trigger system as well as the ability to tag the corresponding bunch crossing.




Figure 6.7.: A detailed overview of the muon system is shown. The different kinds of muon chambers and the toroid magnets are marked. The figure can be found in [89].

Different designs have been chosen for the trigger chambers in the barrel region $|\eta| < 1.05$ and for the end-cap $1.05 \leq |\eta| \leq 2.4$. This is motivated by the fact that at higher values of $|\eta|$, the muon momentum p is much larger at a given p_T compared to central values $\eta \approx 0$, requiring for an η -dependent increasing granularity in the end-cap trigger system. Gaseous parallel electrode-plate detectors, so-called Resistive Plate Chambers (RPCs), are used in the barrel region. They have no wires, making them easy to install while providing sufficient resolution. Thin Gap Chambers (TGCs) are multi-wire proportional chambers and are used in the end-cap region. Their spatial resolution is determined by the readout channel density and can therefore be adjusted to the requirements in this region. The usage of the trigger chambers for the muon trigger system is explained in more detail in Section 6.6.1.2.

6.5. Luminosity Detectors and Luminosity Measurement

In order to determine the instantaneous and integrated luminosity, ATLAS utilizes dedicated luminosity detectors [106]. The most important luminosity detector is LUCID (Luminosity measurement using Cerenkov Integrating Detectors) [107], complemented by additional measurements, for example from the BCM (Beam Conditions Monitor) system. LUCID is a forward detector located at both sides of the interaction point at distances of ± 17 m. Each of the detector arms contains 16 PMTs. Protons traversing the PMTs create Cherenkov light in their quartz windows, providing raw hit counts for each Bunch Crossing (BC). BCM consists of four diamond detectors arranged in cross patterns on each side of the interaction point at ± 1.84 m [108]. Its horizontal and vertical sensors can be read out separately, leading to the possibility of combining two independent measurements.

To determine the luminosity, the measured raw hit counts for each BC are converted to a visible interaction rate μ_{vis} . Using Equation (5.3), the following expression can be determined for the instantaneous luminosity per bunch:

$$\mathcal{L}_b = \frac{\mu_{\text{vis}} f_r}{\sigma_{\text{vis}}} \quad (6.11)$$

The total instantaneous luminosity corresponds to the sum of \mathcal{L}_b over all colliding bunch pairs. The quantity σ_{vis} is a calibration constant representing the absolute luminosity calibration of a luminosity algorithm, and is experimentally determined in so-called *Van der Meer* scans [98], which are performed in dedicated LHC runs with special run conditions. In these scans the beams are displaced in different steps in x and y direction, and the interaction rate is measured as a function of the displacement. By combining Equation (6.11) and Equation (5.3) for a single bunch an expression for σ_{vis} is derived:

$$\sigma_{\text{vis}} = \mu_{\text{vis}}^{\text{max}} \frac{2\pi\Sigma_x\Sigma_y}{n_1n_2} \quad (6.12)$$

where $\mu_{\text{vis}}^{\text{max}}$ is the visible interaction rate per bunch at the peak of the displacement scan curve, where both beams overlap completely. The transverse beam dimensions Σ_x and Σ_y are determined from the decrease of the interaction rate with the displacement. With this calibration, Equation (6.11) is used to directly convert the hit counts into the instantaneous luminosity.

6.6. The Trigger and Data Acquisition System

In order to channel the collision data from the detector to storage, ATLAS uses a specialized Data Acquisition (DAQ) system. Bunch crossings occur every 24.95 ns at the LHC interaction points, resulting in a rate of 40.08 million events each second with instantaneous luminosities in the order of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. Since a fully recorded event contains all final state information from the complete detector, it has a size of about 1.5 MB and takes several seconds of reconstruction time [109]. This makes it evident that it is neither possible to permanently store all collision data nor to provide enough computation power for the reconstruction of all events. However, the cross sections for all processes of interest are orders of magnitude below the total pp cross section, as can be seen in Figure 6.8. Hence, only a small fraction of all events is of interest for physics analyses. The goal is therefore to select only those relevant events, as this allows to reduce the data rate by many orders of magnitude and makes it possible to permanently store the remaining collision data. For this purpose, ATLAS uses a dedicated event selection system, the Trigger system. Figure 6.9 shows an overview of the architecture of the Trigger and Data Acquisition (TDAQ) system. The ATLAS trigger system consists of two levels: a hardware based level-1 trigger (L1) and a software based trigger, the High Level Trigger (HLT). It reconstructs events *online* parallel to the run and decides whether the event is interesting and worth saving in only a few seconds. The level-1 trigger is based on custom-built electronics and uses signals from the muon system and the calorimeters to find events with interesting signatures and to reduce the initial event rate of 40 MHz to about 100 kHz. The HLT works with the trigger decision from the first level and refines it by applying additional selection criteria utilizing the full detector information. It is based on software algorithms running on a processor farm. In total, the event rate is reduced down to a level of 1 kHz of interesting events. In this context, interesting means that the event fulfills the conditions of one of several so-called trigger items, which are listed in the *trigger menu*. A single trigger item defines specific conditions which an event has to fulfill, these are usually motivated by the experimental searches and measurements, for example a requirement on the particle content of the event, kinematic thresholds or topological requirements between particles. The trigger menu comprises the list of full L1 to HLT trigger items.

While the trigger decision is formed, it is necessary to buffer all detector information. Dedicated front-end pipeline buffers on the subdetectors are responsible for this while the L1 decision is formed. When an event is accepted by the level-1 trigger, the Readout Driver (ROD) cards push the




Figure 6.8.: A selection of Standard Model total production cross section measurements compared to the respective theory predictions is shown [2]. The dominant cross section is the total pp cross section. All other cross sections are smaller by a factor between 6 (e.g. W boson production) and 11 ($t\bar{t}Z$) orders of magnitude.

information from the pipeline buffers to the Readout System (ROS). The ROS buffers all data while the HLT forms its decisions. Upon requests, it provides the required data to the HLT algorithms. When the HLT decision is formed, the buffered information on the events is sent to permanent storage in case of an accepted event or discarded otherwise.

Monitoring of all modules and applications running in the TDAQ system is provided by several applications, one of them being the Information Service (IS) [111]. Every application can make information publicly available for other modules via the IS. It is a scalable software-based system which provides means to share control and monitoring information to other applications. When updated information is available, notifications on the updates are provided to all clients within a few ms of time. In addition, it is possible to subscribe to automatically receive updated information from specified sources. For example, information on whether ATLAS is running or whether the LHC has declared stable beams are published via the IS. A related application is the Online Histogramming Service (OHS), which allows monitoring applications to exchange data in form of histograms. For example, the accept rates of the various trigger items are published this way. These monitoring histograms are created continuously and archived to tape at the end of a run for further analyses using the Monitoring Data Archive (MDA).

6.6.1. The Level-1 Trigger

The ATLAS level-1 trigger is a fixed-latency, pipelined system based on fast custom-built electronics, utilizing mostly FPGAs and ASICs. It consists of the Level-1 Calorimeter Trigger (L1Calo), the Level-1 Muon Trigger (L1Muon), the Level-1 Topological Processor (L1Topo) and the Central

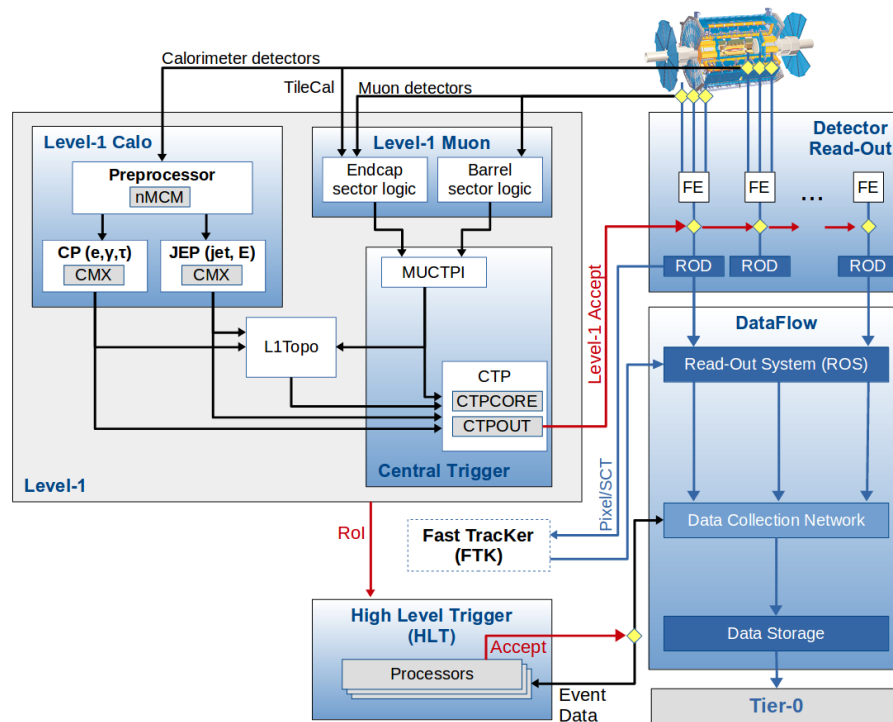


Figure 6.9.: Overview of the TDAQ system. The trigger system with the level-1 trigger and the HLT (High Level Trigger) is shown as well as the structure of the readout and data storage infrastructure. From reference [110].

Trigger Processor (CTP). Coarse granularity information from the calorimeters and the muon system is used to perform L1 trigger decisions based on multiplicities of jets, electrons, photons, muons, or τ leptons exceeding predefined energy thresholds as well as thresholds of global quantities like $\sum E_T$ or E_T^{miss} . In addition, topological quantities are calculated for the trigger decisions. However, no tracking information from the inner detector is available for the decisions in the first trigger level. The level-1 trigger also identifies Region of Interests (RoIs), detector regions containing objects which contributed to the trigger decision, these can be used to seed the HLT algorithms. The maximum capacity of the pipeline buffers lead to latency requirements defining a total latency budget of $2.5 \mu\text{s}$ including cabling to form the level-1 decision.

6.6.1.1. The Level-1 Calorimeter Trigger

L1Calo is a pipelined system which provides trigger signals in real-time to the CTP using signals from the LAr and Tile calorimeter systems. Its input data consists of 7200 trigger towers, which are analog sums over calorimeter cells with a granularity of 0.1×0.1 in $\Delta\eta \times \Delta\phi$ for most trigger towers. The exact granularity of all trigger towers is shown in Figure 6.10. Separate sets of trigger towers exist for the ECAL and HCAL. The trigger tower signals are then digitized, assigned to a certain bunch crossing and noise-filtered by the L1Calo PreProcessor Modules (PPMs). Afterwards, the PPMs convert the digitized input signals into E_T values using LUTs. This trigger tower E_T information is then forwarded to two feature processors: the Cluster Processor (CP) and the Jet/Energy-Sum Processor (JEP), both identifying localized energy depositions using sliding window algorithms. The CP consists of a total of 48 Cluster Processor Modules (CPMs) and




Figure 6.10.: Trigger tower granularity for $\eta > 0$ and one quadrant in ϕ , the figure is adopted from reference [112]. The trigger towers at $|\eta| > 2.5$ are irregular and larger than in the central region. The system is symmetric both in η and ϕ .

identifies electron/photon and τ candidates satisfying programmable E_T thresholds in the region of up to $|\eta| < 2.5$, as this is the limit of the inner detector and high-granularity electromagnetic calorimetry. The JEP identifies jets and produces global sums of total, missing and jet-sum E_T . It consists of two crates of 16 Jet/Energy Module (JEM) modules each, every JEM covers a small calorimeter region in $\eta \times \phi$. The inputs for the JEP algorithms are *jet elements*, which are 2×2 sums of trigger towers in depth over both calorimeters, with a granularity of 0.2×0.2 in most cases. Due to the irregular size of the trigger towers for $|\eta| > 2.5$, some jet elements in the outer regions are larger in η . The jet elements along with the calorimeter coverage of the JEMs is shown in Figure 6.11 for one quadrant of ϕ and $\eta > 0$. A description of the algorithms used by L1Calo to reconstruct hadronic or electromagnetic objects is given in Chapter 7.

The information on reconstructed electromagnetic objects, jets and energy sums is then sent to the merging modules, the Common Merger eXtended (CMX). The CMX combines the object information into so-called Trigger Objects (TOBs), which contain the angular coordinates (η, ϕ) , the transverse energy E_T and information about the type for each object. Each CMX module then sorts its TOBs by transverse energy, and provides these TOBs as inputs for the L1Topo system. The CMX also counts object multiplicities above configurable E_T thresholds and provides this information to the CTP. In addition, it also includes the current BCID in its data for L1Topo.




Figure 6.11.: Jet Elements and calorimeter coverage of the JEMs of the JEP. Each colored region depicts the coverage region of a single JEM. The JEMs in the forward region cover a larger η range.

6.6.1.2. The Level-1 Muon Trigger

The L1Muon subsystem is based on signals from the muon trigger chambers RPCs and TGCs. FPGA based sector logic modules process the detector signals to determine a total of six configurable p_T thresholds using dedicated coincidence logic. This information is then provided to the Muon

CTP Interface (MUCTPI), which extracts coarse topological information with a resolution of $\Delta\eta \times \Delta\phi \approx 0.3 \times 0.1$ from all RPCs and TGCs. It then combines the information with the associated thresholds to create TOBs, which are provided to L1Topo via the MUCTPIToTopo [113]. In addition, the MUCTPI calculates muon multiplicities from the estimated p_T thresholds, taking possible double counting due to overlapping regions in the different trigger sectors into account. Again, these are sent to the CTP.

6.6.1.3. The Level-1 Topological Processor

In order to cope with rising luminosities and the corresponding increase of trigger rates, L1Topo has been introduced as a new module to the ATLAS level-1 trigger for Run-2 [6, 114]. The L1Topo system consists of an ATCA crate and two L1Topo modules, each containing two processor FPGAs. L1Topo gets inputs from the calorimeter trigger L1Calo as well as from the muon trigger L1Muon. All reconstructed TOBs are provided for each event, meaning that the *complete event* is available on a single board in the level-1 trigger. The event information is then used to perform trigger decisions based on topological algorithms by combining objects from the different subsystems of the L1 trigger. Topological algorithms can therefore be applied on combinations of electrons/photons, taus, jets, energy sums and muon particle candidates. Examples for topological algorithms which can be calculated on L1Topo are shown in Figure 6.12. For each algorithm a decision bit and an overflow bit are formed, the former indicating whether the kinematic or angular selection of the algorithm has been passed, while the latter indicates energy saturation or other sources of overflow. These bits are provided to the CTP via electrical cables, where they are used for the final level-1 trigger decision. Details on the hardware design, firmware development and the physics performance are given in a dedicated chapter in Part III.




Figure 6.12.: Examples of topologies calculated on L1Topo. Angular variables, for example $\Delta\phi$, $\Delta\eta$ or ΔR are included as well as object clustering and invariant or transverse masses.

6.6.1.4. The Central Trigger Processor

The CTP is the component of the L1 trigger where the final L1 trigger decision is performed [115, 116]. All decision bits and multiplicity counts from the other parts of the L1 trigger are inputs to the CTP, where they are either considered individually or combined using logical or and operations, resulting in up to 512 trigger items as configured by the trigger menu. The trigger items are put in coincidence with the corresponding bunch groups as introduced in Section 5.4. Afterwards each item is scaled by an adjustable *prescale* factor $p \geq 1$, allowing for a possible reduction of rates of trigger items. In this case, the CTP accepts at random $1/p$ events passing the trigger selection. To correct for the loss of events, the accepted event is weighted by a factor of p . Finally, the so-called L1 Accept (L1A) is formed by taking the logical or of all trigger items,

indicating whether the event is forwarded to the higher trigger stages and to the ROS. In addition, the CTP is responsible for the timing and synchronization signals which are distributed to the various subdetectors of ATLAS, for example the clock signal synchronized to the LHC bunches or the BCID information. This is done via the so-called Timing, Trigger and Control (TTC) system. As parts of this information is relevant for multiple experiments, parts of the TTC have been developed as a joint project between the experiments and the LHC. In addition, the CTP also defines the length of each *luminosity block* (*lumiblock*). A lumiblock is the smallest time interval for which quantities like the integrated luminosity can be determined correctly. The duration and the number of each lumiblock are provided as a part of the readout, allowing to label all data in ATLAS with a run number and a lumiblock number. In addition, lumiblocks with detector failures or other unwanted recording conditions can be dropped out. In order to avoid the unnecessary loss of data when lumiblocks are dropped, the interval should be as small as possible. On the other hand, the time interval needs to be large enough so that the statistical uncertainties on the luminosity determination can be neglected compared to the systematic uncertainties. For ATLAS, the length of most lumiblocks is in the order of one minute.

6.6.2. The High Level Trigger

The HLT is a purely software-based trigger which is running on a farm of computing nodes. All events accepted by the L1 trigger are processed by the HLT using the full detector information at highest granularity, including tracking information from the inner detector. Various trigger algorithms can be performed in the RoIs identified by the first trigger level as well as for the full detector. Most HLT triggers use a two-stage approach in order to reduce the required processing time, consisting of a fast first-pass reconstruction to reject most of the events followed by a more precise reconstruction for the remaining events. The total processing time per event is given as ≈ 200 ms, depending on the instantaneous luminosity and the trigger menu configuration [117]. The average output rate of the HLT is at about 1 kHz and mainly determined by data storage constraints [118].

6.7. Detector Simulation

As discussed in Section 3.3, MC simulations are frequently used in high-energy physics. The output of these simulations are given in terms of *truth* information, which contains all information available at the MC generator level, for example the four vectors of all particles and information on their interactions and decay chains. However, this output does not correspond to the format of data recorded with the ATLAS detector, which consists of the digitized detector signals created by particles traversing the detector. In order to be able to truly compare the MC events to recorded data, it is necessary to simulate the response that the particles in the simulation would cause to the real detector as well as the digitization of detector readout signals. Afterwards, it is possible to run the same event reconstruction on both.

In this thesis, two different types of detector simulation are used. The first one is widely used by the ATLAS collaboration [119] and based on the GEANT4 framework [120]. It contains a detailed model of the ATLAS detector, including its whole geometry and detailed material information. All possible interactions of the truth particles from the MC simulation with the detector material are simulated, additional particles created from bremsstrahlung or in showers are also propagated. Finally, the digitization of the analogue detector signals and the decisions of the trigger system are modeled. The presence of additional pile-up collisions is also simulated by

overlaying the events with hits from other events. Afterwards, the simulated events are stored in the same format as the recorded data, allowing to run the identical event reconstruction on both. However, the simulation of the detector is a time-consuming task and can require several minutes of processing time for a single event. Since large computing resources are required to accurately model the detector geometry and all interactions, alternative ways of simulation have been developed to achieve the required simulated statistics for a variety of physics studies [119]. These are commonly referred to as *fast simulation*, as opposed to the GEANT4 based *full simulation*. In this thesis, fast simulations based on the Delphes [121] toolkit are used. Delphes is based on a modular design and allows for flexible combinations of tracking detectors, calorimeters and muon systems, but also includes a predefined configuration for the ATLAS detector. Using Delphes it is possible to use flexible configurations for phenomenological studies by modifying the detector granularities or resolutions. Instead of fully simulating the interactions of particles with the detector material, the responses of the detector parts are modeled in a parametrized fashion to replace the time-consuming MC simulation. For example, particles reaching the calorimeter are assumed to deposit a fraction f of their energy in exactly one ECAL cell and the remaining $1 - f$ in one HCAL cell. The energy resolutions of the calorimeters are then modeled via parametric functions of E and η independently for the ECAL and the HCAL. Jets are reconstructed from the calorimeter information using the FastJet package [122]. This approach significantly speeds up the simulation to ~ 0.01 s to 1 s per event. Though, some trade-offs need to be made: since a simplified geometry is assumed, effects like dead material are not considered. In addition, more complex physical effects like bremsstrahlung are also not included, limiting the precision of the fast simulation.

Jet Reconstruction in ATLAS

As discussed in the previous chapter, different particles leave distinct signals in the ATLAS detector. For physics analyses as well as for trigger decisions it is necessary to identify these signals with physical objects by reconstructing objects based on detector signals. This chapter presents an overview of jet reconstruction with the ATLAS level-1 trigger in Section 7.1. The tight latency budget of the level-1 trigger requires a trade-off between simplicity and performance when designing the object reconstruction algorithms. To fully utilize the capabilities of the underlying hardware, level-1 algorithms should be highly parallelizable. In contrast to this, higher trigger stages and offline reconstruction are able to utilize more sophisticated algorithms than the L1 trigger, for example based on iterative procedures. An overview of offline jet reconstruction in ATLAS is presented, starting with a short introduction to tracks and vertices in Section 7.2.1 and then continuing with a description of jets in Section 7.2.2. To further improve the quality of reconstructed jets, a jet calibration procedure and a jet cleaning procedure are applied, these are discussed in Section 7.2.3 and in Section 7.2.4, respectively.

7.1. Jet Reconstruction in the Level-1 Trigger

The level-1 jet algorithm is based on energy depositions in jet elements with a granularity of 0.2×0.2 in the central region, while their granularity is irregular for higher values of $|\eta|$. These jet elements are created by combining trigger towers of sizes of mostly 0.1×0.1 . Since the jet element energies are calculated from the corresponding trigger tower energies, it is necessary to understand how E_T values are assigned to each trigger tower based on the measured calorimeter signals. This conversion is performed in the PPMs of L1Calo and is described in the following. In particular, methods to reduce the impact of pile-up on reconstructed objects are already applied during this step. Two kinds of pile-up have to be distinguished, in-time pile-up and out-of-time pile-up. The former refers to effects which occur in the same BC as the hard scattering, the latter to phenomena due to signals spanning over multiple BC.

Jet Element Energies

The calorimeter front-end electronics combine calorimeter cells into trigger towers by grouping multiple cells in predefined η - ϕ regions. The analogue signals of all these cells are summed up, corrected from E to E_T and signals from different detector regions are calibrated to a common scale. These trigger towers are then forwarded to the L1Calo PPMs. Depending on the calorimeter region and used technology, the shapes of the input signals can be different. Figure 7.1 shows exemplary pulse shapes of trigger tower signals from different calorimeter regions. Since the




Figure 7.1.: Normalized signal pulse shapes from different parts of the calorimeter. “FCAL” denotes the forward calorimeter. In general, the pulses are much longer than the bunch spacing of 25 ns. The pulses from LAr are bipolar with a large positive peak followed by a negative undershoot. Tile pulses have a unipolar shape. From reference [123].

pulses span over multiple BC, it is important to correctly assign a pulse to the corresponding bunch crossing. In addition, pulses from different BC can overlap, resulting in out-of-time pile-up effects which need to be handled accordingly. The observed pulse shapes vary in η and with the different detector regions, however only little dependencies on ϕ are observed. Therefore, trigger towers with the same η coordinate are treated equally in the PPM algorithms. To calculate the trigger tower energies, the analogue calorimeter signals are first digitized by analog-to-digital converters (ADCs) to ADC samples A_i . Afterwards, a peak-finder algorithm is used to assign the correct BC to each input signal. At each BCID i , the peak finder calculates the filter f_i as the convolution of the last five digitized ADC samples A_{i-k} with filter coefficients c_k :

$$f_i = \sum_{k=0}^4 c_k A_{i-k}. \quad (7.1)$$

The filter coefficients are optimized for the LHC operation conditions and change between runs, as both the bunch train structure and the pile-up of single BCs are taken into account. Afterwards, the *pedestal correction* is applied to obtain the corrected filter output f_i^{corr} . The pedestal correction aims at correcting the interplay of the LHC bunch train structure and the long calorimeter pulses. Within a train, the overlapping of pile-up from many subsequent BC lead to cancellations due to overlap of peaks and negative undershoots.¹ However, this is not true at the beginning or end of a bunch train, leading to a shift of the signal baseline, the so-called pedestal. This effect is depicted in Figure 7.2. If left uncorrected, this shift could result in changes in the signal amplitudes and lead to an overestimation of the deposited energies for certain BC. This is compensated by the

¹No negative undershoots exist in the Tile signals, as the typically low-momentum pile-up effects are mostly absorbed in the ECAL, leaving the HCAL unaffected [124].




Figure 7.2.: Mean shift of the pedestal for different values of $\langle\mu\rangle$. Large deviations from a flat pedestal are observed at the start or end of a bunch train. From reference [123].

following BCID dependent correction:

$$f_i^{\text{corr}} = f_i - (\bar{f}_i - f_i^{\text{ped}}), \quad (7.2)$$

where f_i^{corr} is the corrected filter output, \bar{f}_i the average filter output, and f_i^{ped} the configured target value. The average filter is calculated and updated frequently by the firmware by averaging for about 5.9 s [123], the target value corresponds to the expected flat pedestal in the absence of any beam. Afterwards, the corrected filter output is then checked for the peak condition²

$$f_{i-1}^{\text{corr}} < f_i^{\text{corr}} \geq f_{i+1}^{\text{corr}}. \quad (7.3)$$

If a trigger tower fulfills the peak condition in a certain BC, an E_T value is obtained by forwarding the filter value f_i^{corr} to the JEP LUTs. The LUTs assign an energy value of 0 GeV to 255 GeV with a resolution of 1 GeV to each trigger tower. Since statistical fluctuations in the pile-up noise can result in fake E_T assignments, noise cuts are applied in the LUTs to improve the object reconstruction and performance of L1Calo triggers. The noise cuts are determined based on the occupancy of trigger towers as a function of η , which means that the noise cut values are tuned such that only the highest 0.5 % of filter outputs in each η bin are set to non-zero E_T values [123, 125]. Since the occupancy depends on the beam conditions as well as on the LHC filling scheme, the noise cuts have to be recalibrated after major changes in the run conditions. The resulting noise cuts are shown in Figure 7.3 for a few different scenarios. Noise cuts are applied, which range from 1 GeV to 2 GeV in the central region and up to 9 GeV in the forward region. Before forwarding the energy information to the JEP, the trigger tower energies are summed up in depth over both ECAL and HCAL to form the jet elements. The exact way in which the trigger towers are combined to form jet elements is depicted in Figure 7.4. While the combination is obvious for

²A modified peak-finder algorithm is used for saturated pulses, where multiple neighboring BC can have the same ADC values.




Figure 7.3.: JEP energy LUT noise thresholds for the electromagnetic trigger towers (a) and hadronic trigger towers (b) for different values of $\langle\mu\rangle$. The term 8be4 denotes a special filling scheme used by the LHC in 2017. From reference [123].

$|\eta| < 2.4$, the jet element sizes vary for $2.4 \leq |\eta| < 3.2$. In the forward region $3.2 \leq |\eta| < 4.9$, only 16 trigger towers are available in ϕ . To match the granularity of the inner regions, every tower is artificially split into two towers, each carrying half of the energy of the original tower. This has the advantage that the jet elements now form a grid of 32×32 elements in $\eta \times \phi$, allowing to run the same algorithm on all JEMs by formulating it in terms of local jet element numbers and independent on the actual coordinates.

Jet Finding

Based on the jet elements formed in the PPMs, a jet algorithm is run in the JEP. In its first step, the jet algorithm calculates the corresponding transverse energy E_T for each possible window of 2×2 jet elements as the sum of the individual energies. These sums are the so-called seeds of the algorithm. For each seed the energy is then compared to the energies of all possible surrounding and overlapping seeds in a window of 4×4 jet elements, as shown in Figure 7.5(a). In order to define a jet candidate, the central seed needs to be a local maximum, requiring its energy to be larger or larger equal than its neighbors with the exact condition being shown in Figure 7.5(b). A mixture of $>$ and \geq is used to introduce a preferential direction in case of neighboring seeds with the same energy. This algorithm is also referred to as a Sliding Window (SLW) algorithm.³ A jet candidate is then defined by two components, a 0.4×0.4 seed which is a local maximum and defines the position of the jet candidate as its geometrical center, and a jet window, which is used to measure the E_T of the jet. Jet windows in three different sizes, 0.4×0.4 , 0.6×0.6 and 0.8×0.8 , are implemented, though only the two windows with sizes 0.4×0.4 and 0.8×0.8 are used for trigger algorithms in the CMX and L1Topo. The jet windows are shown in Figure 7.6 together with the position of the corresponding jet seed. In addition to the local maximum criterion, a threshold condition on the minimum energy in the jet window is required. In summary, the requirements for a jet TOB are as follows:

³The algorithm name is inspired from a sequential software implementation, where a single search window is processed and subsequently moved in η and ϕ . In the trigger implementation on FPGAs all search windows are evaluated in parallel, making the naming rather counter-intuitive.




Figure 7.4.: Trigger towers (a) and the corresponding jet elements (b) for one quarter in ϕ and $\eta \geq 1.6$. Color-shading indicates which trigger towers are combined to form a jet element. For $|\eta| \leq 2.4$ the granularity is uniform, for higher values of η the jet elements differ in size. For $|\eta| \geq 3.2$, trigger towers cover 0.4 in ϕ and are split in half for the jet elements (purple). Based on the information received in [126].

1. The jet seed must be a local energy maximum compared to its neighbors, see Figure 7.5(b).
2. The total E_T in the jet window must exceed a configurable threshold. For Run-2, a value of 12 GeV was used for the 2×2 and 4×4 windows.

7.2. Offline Jet Reconstruction

The objects described in this section are reconstructed in recorded data utilizing the full detector information. At the HLT and in the offline reconstruction, more sophisticated reconstruction algorithms using finer-granularity detector information can be used, as significantly more processing time is available. In contrast to the first trigger level, tracking information from the inner detector is available at this level.

7.2.1. Tracks and Vertices

The goal of the track reconstruction is to reconstruct the path of a charged particle through the inner detector. Hits of particles in the inner detector are combined to obtain three-dimensional space-point information. Beginning with seeds built from high-precision information in the inner layers of the pixel detector and SCT, additional space-points compatible with the track are subsequently added by track fitting algorithms. Detailed information on various track fitting algorithms used in Run-2 is given in reference [128]. Reconstructed tracks are required to pass certain quality criteria, for example on their minimum p_T or the number of associated hits in the different subdetectors.

After all tracks are fitted, a vertex finder algorithm associates tracks to vertices [129]. Starting at the position of the beam spot, an iterative procedure is used to determine the vertex position. Each track is assigned a weight which denotes the probability of the track to belong to the vertex and indicates its influence on the vertex position. At the end of the iteration procedure, incompatible tracks are removed and used for the next vertex finding iteration. From the set of all vertices, the

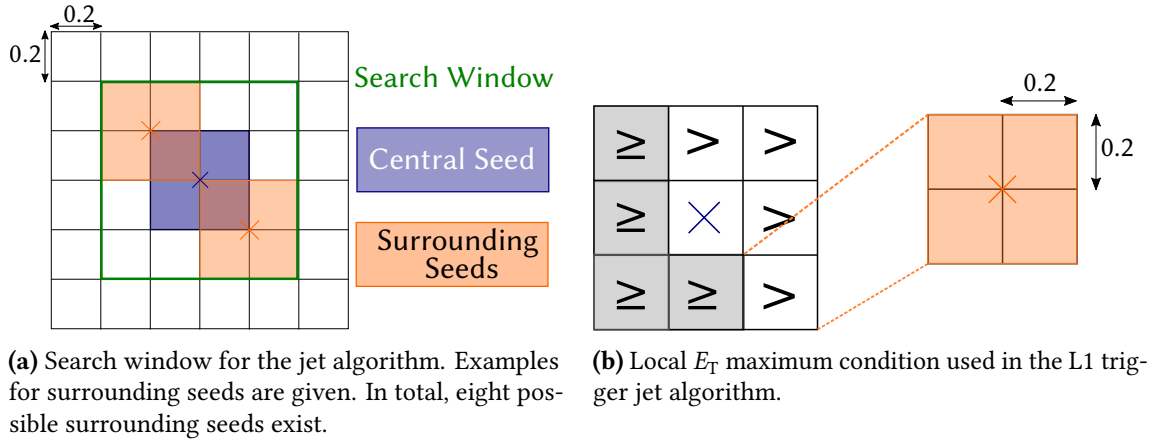


Figure 7.5.: Local maximum identification of the L1 trigger jet algorithm. Based on a figure in [127].

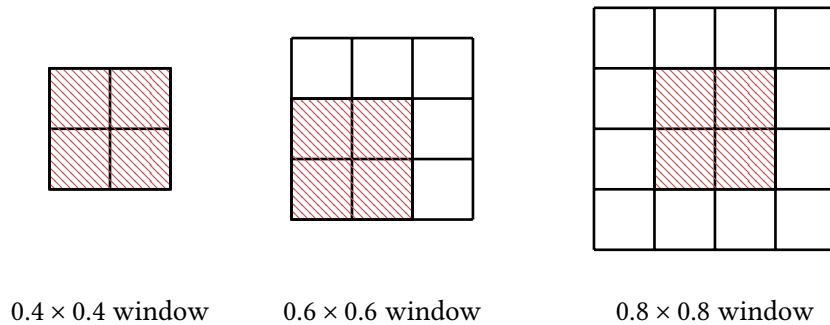


Figure 7.6.: The window sizes used in the jet algorithm are shown, based on 0.2×0.2 jet elements. Shaded regions show the jet seed and define the seed position. In case of the 0.6×0.6 size, four different positions are possible for the seed. Modeled after a figure in reference [112].

one belonging to the hard scattering process is of particular interest. To identify this vertex, the primary vertex is defined as the vertex with the highest value of $\sum p_T^2$.

7.2.2. Jets

In the following, the offline jet reconstruction is discussed briefly. Instead of relying on coarse trigger tower information, the full granularity information on calorimeter cells is available. The energy depositions caused by hadronic showers in the calorimeter cells are used to create so-called *topological clusters*, three-dimensional combinations of topologically connected calorimeter cells with large energy depositions significantly above background noise level [130]. These clusters then serve as input to a jet algorithm. In ATLAS, so-called sequential recombination algorithms are frequently used. These algorithms use an iterative procedure to combine pairs of topological clusters (i, j) with transverse momenta p_{Ti} and p_{Tj} using two distance measures

$$d_{ij} = \min \left(p_{Ti}^{2p}, p_{Tj}^{2p} \right) \frac{\Delta R_{ij}^2}{R^2} \quad (7.4)$$

$$d_{iB} = p_{Ti}^{2p},$$

where R is the radius parameter of the jet algorithm and p defines different combination schemes, of which only the choice $p = -1$ corresponding to the anti- k_t algorithm [131] is relevant for this thesis. The jet algorithm then performs the following steps:

1. For each cluster i , identify the smallest distance of all d_{ij} and d_{iB}
2. If the minimum is a d_{ij} , combine the clusters (i, j)
3. If the minimum is d_{iB} , call it a jet and remove it from the procedure
4. Repeat 1.-3. until no cluster are left

It preferably combines nearby clusters with high transverse momenta, and is the standard jet clustering algorithm used in ATLAS during Run-2, with jet radius parameter of $R = 0.4$ or $R = 1.0$. The four-momentum of the jet is then taken as the vector sum of its constituent clusters.

7.2.3. Jet Calibration

The reconstructed energy of such a jet does not take the limited response of the hadronic calorimeter into account. In addition, effects like dead detector material, for example due to supporting structure, the possibility of jet constituents falling outside the jet radius, cell energies not passing noise cuts or parts of the jets leaking outside the calorimeters are not considered. Therefore, the reconstructed jet energy will in general not correspond to the true energy of the original particle. In ATLAS, this is corrected by a dedicated *jet calibration* procedure. It corrects the originally measured jet energy at the *EM scale* to the Jet Energy Scale (JES) [132, 133]. While the whole procedure is often referred to as the JES correction, it actually calibrates the whole jet four-momentum and not just its energy. The steps performed in the calibration of an EM scale jet are summarized in Figure 7.7 and briefly discussed in the following. Parts of the calibration procedure rely on tracking information from the inner detector as introduced in Section 7.2.1. Tracks associated to charged particles included in a jet can be associated to a jet using *ghost association*. The tracks are added as additional inputs to the jet clustering algorithm with infinitesimal p_T , and are therefore referred to as ghosts. This way, they do not change the four-momentum of the jet. If the tracks are included in a jet in this clustering, they are associated to jet.

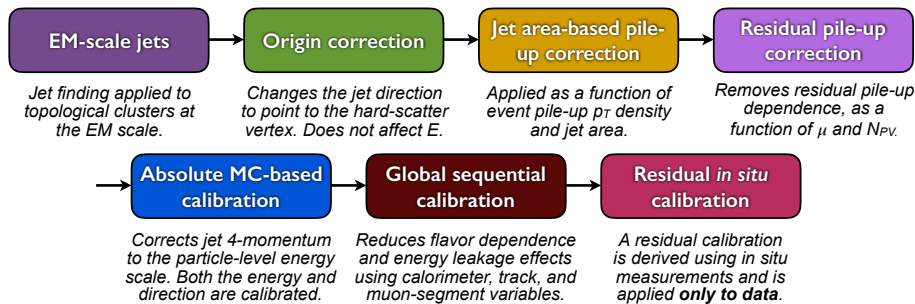


Figure 7.7.: Overview of all calibration stages applied to EM-scale jets, from reference [133].

Origin Correction

The coordinates of the topological clusters used as inputs for the jet algorithm are defined with respect to the center of the ATLAS detector in the origin of its coordinate system. In this step, the

jet four-momentum is adjusted to point towards the primary vertex instead while keeping the jet energy constant.

Pile-Up Correction

In general, the measured energy of reconstructed jets is affected by pile-up. The correction is modeled in three pile-up related variables, namely the median p_T density in the event ρ , the number of primary vertices N_{PV} and the mean number of proton-proton interaction μ , and is given as

$$p_T^{\text{corr}} = p_T^{\text{reco}} - \rho \cdot A_i^{\text{jet}} - \alpha \cdot (N_{PV} - 1) - \beta \cdot \mu. \quad (7.5)$$

ρ is calculated from the median p_T density of all jets without any p_T selection within $|\eta| < 2.0$:

$$\rho = \text{median} \left(\frac{p_{T,i}}{A_i^{\text{jet}}} \right) \quad (7.6)$$

where A_i^{jet} is the jet area⁴. The first term of Equation (7.6) is called the *jet area based pile-up correction*. Residual dependencies of the jet p_T on N_{PV} and $\langle \mu \rangle$ are subtracted using coefficients α and β determined from MC. Since truth jets reconstructed from truth particles in MC simulations are not sensitive to pile-up, the geometrical matching of reconstructed and truth jets allows to determine the pile-up dependence of reconstructed jets on these variables. The jet four-momentum is then corrected by a scaling factor given by the ratio $p_T^{\text{corr}}/p_T^{\text{reco}}$.

Absolute MC-based Calibration

The MC-based calibration compares reconstructed jets to truth jets and corrects their energy and direction to the particle-level. This is achieved by geometrically matching reconstructed and truth jets, and by calculating the energy response

$$\mathcal{R} = \frac{E_{\text{reco}}^{\text{EM}}}{E_{\text{truth}}} \quad (7.7)$$

as a function of E_{reco} in bins of η_{detector} . The binning in η_{detector} is required since the various sub-detectors of ATLAS have different responses and therefore require different corrections. For every bin of η_{detector} and E_{truth} , the mean response $\langle \mathcal{R} \rangle$ is determined from a Gaussian fit. Figure 7.8 shows the mean response in bins of η_{detector} for different truth jet energies. The response is always below 1 and shows a strong η dependence, with a particular bad response in the transition regions between barrel and end-cap region at $|\eta_{\text{detector}}| \approx 1.4$ and between the end-cap and forward region at $|\eta_{\text{detector}}| \approx 3.2$. Moreover, the response gets higher when higher truth jet energies are considered. The dependence of $\langle \mathcal{R} \rangle$ on $E_{\text{reco}}^{\text{EM}}$ in each η bin k is fitted with an empirical function of the following form

$$\mathcal{F}_k(E_{\text{reco}}^{\text{EM}}) = \sum_{i=0}^N a_{i,k} (\log E_{\text{reco}}^{\text{EM}})_k^i \quad (7.8)$$

⁴The jet area is defined by adding additional, infinitely soft ghost particles to the jet clustering. The extent of the region where these ghosts are clustered with a given jet gives a measure of the jet area [134].




Figure 7.8.: The average energy response of jets simulated by ATLAS is shown as a function of the detector region and different matched truth jet energies. From reference [133].

where the $a_{i,k}$ are free parameters of the fit, their number depends on the model hyperparameter $1 \leq N \leq 6$. A correction is then derived as the inverse of the obtained response function:

$$E_{\text{reco}}^{\text{EM+JES}} = \frac{E_{\text{reco}}^{\text{EM}}}{\mathcal{F}(E_{\text{reco}}^{\text{EM}})_{\eta_{\text{detector}}}} \quad (7.9)$$

Afterwards, a dedicated jet η calibration corrects for differences in η_{reco} and η_{truth} without altering the jet energy.

Global Sequential Correction

A global sequential correction is applied after the MC-based calibration and is used to remove the dependence of the response on further variables without changing the average jet energy [135]. The corrections are applied similarly to the MC-based calibration by parametrizing the response as a function of the jet p_{T} and the corresponding variable x in bins of η_{det} , and are applied sequentially for one variable at a time. Correction factors are derived as $c(x) = k / \langle \mathcal{R}(x) \rangle$, where k is chosen to ensure that the average jet energy is not changed by the calibration.

As the ATLAS calorimeter is non-compensating, a large deposited energy fraction in the hadronic layers indicates a lower calorimeter response of this jet on average. Due to this, the first step of the global sequential correction uses variables characterizing the longitudinal energy-profile of the jet. For this, the energy fraction deposited in certain calorimeter layers $f_{\text{layer}} = E_{\text{layer}}^{\text{EM}} / E_{\text{jet}}^{\text{EM}}$ is used for the third electromagnetic layer LAr3 and the first hadronic layer Tile0. The next corrections are done using track-based variables, namely the number of tracks associated to a jet n_{trk} and the p_{T} -weighted average distance between the jet axis and the tracks associated to the jet \mathcal{W}_{trk} . By this, flavor dependencies are mitigated, as jets originating from a gluon are on average wider and contain more low- p_{T} particles compared to quark jets. The final correction targets jet leakage beyond the calorimeter by inspecting the activity in the muon chambers behind the calorimeter. This is characterized by the number of partially reconstructed tracks in the muon spectrometer behind the jet.

Residual In Situ Calibration

Since the process of jet formation, the detector response on jets and pile-up effects are not perfectly modeled in MC, some residual differences are expected between jets in data and simulation. A final residual *in situ* calibration corrects these differences. The *in situ* calibrations are based on the balance of the p_T of a jet against well measured reference objects in both data and MC. The response is then calculated as the ratio of the jet momentum p_T^{jet} and the momentum of the well-balanced reference object p_T^{ref} :

$$\langle \mathcal{R} \rangle = \left\langle \frac{p_T^{\text{jet}}}{p_T^{\text{ref}}} \right\rangle, \quad (7.10)$$

with correction factors defined as

$$c = \left(\frac{\langle \mathcal{R}_{\text{MC}} \rangle}{\langle \mathcal{R}_{\text{data}} \rangle} \right). \quad (7.11)$$

Using dijet events, an η -intercalibration is performed, where the average p_T for non-central jets ($0.8 \leq |\eta_{\text{det}}| \leq 4.5$) is equalized to the p_T of balancing central jets. This ensures a uniform response over the whole detector, and allows to apply the next corrections for all values of $|\eta|$. Afterwards, the balance of photons and Z bosons recoiling against jets with $|\eta|_{\text{det}} < 0.8$ is used to calibrate jets with p_T up to 950 GeV. In addition, several low- p_T jets are used as the reference objects to calibrate high- p_T jets up to 2 TeV. Since the Z/γ +jets and multijet balance overlap, a combined correction is derived.

The resulting jet energy scale uncertainty is shown in Figure 7.9. It is known at a level of about 5% for momenta of ~ 20 GeV, down to the sub-percent level for 400 GeV, and up to 3.5% for 2 TeV. A total of 80 systemic uncertainty terms propagated from the individual calibration steps are included, where the majority comes from the *in situ* calibrations. The full list of systematic uncertainties are summarized in reference [133].




Figure 7.9.: Combined jet energy scale uncertainty as a function of the jet p_T for jets with $\eta = 0$, from reference [136].

7.2.4. Jet Cleaning

Besides being the result from partons arising from the hard scattering process, jets can also be the result of misidentifications of multiple non-collision origins. The so-called jet cleaning is targeted at removing these fake jets [137, 138]. Three main sources of fake-jet background are distinguished and summarized in the following. Noise bursts in the Hadronic End Cap (HEC) can result in particularly noisy cells, resulting in energy depositions reconstructed as a jet. These fake jets can be identified by their properties, in particular a large fraction of the energy deposited in the HEC, where in addition the energy is deposited in only a few calorimeter cells. Furthermore, the signal shape differs from what would be expected for a real jet. Another source of fake jets is coherent noise in multiple calorimeter cells. Jets built from these cells have a high fraction of their total energy deposited in the ECAL, and an unusual amount of associated cells with negative energy. Besides jets from calorimeter noise, there are also jets originating from non-collision backgrounds. For example, high-energy cosmic-rays can lead to additional particles being produced in the ATLAS detector. In addition, protons lost from the beam interacting with the collimators of the detector or protons interacting with residual gas in the beam pipe can result in particle showers in the detector which can be reconstructed as jets. While jets from the hard interaction usually extend from the primary vertex through many layers in the detector, fake jets from non-beam backgrounds are more localized in the detector and are not associated to the primary vertex and can therefore be identified. The exact criteria applied for the jet cleaning procedure are listed in the Appendix A.5.

Both jet calibration and jet cleaning are standard procedures used in all analyses relying on reconstructed jets. However, no such calibration or cleaning procedures have been applied to level-1 jets during Run-2 for multiple reasons. Since level-1 jets are used to form trigger decisions in real-time, such a calibration would have to be performed online on the modules of L1Calo, which would be challenging due to the tight latency constraints. In addition, it is likely that parts of the calibration procedure like the pile-up correction will depend on the exact running conditions, requiring frequent adjustments. Furthermore, some calibration steps utilize track-based inner detector information, which is not available on the first trigger level.


When using level-1 jets in an analysis, the effects of the missing jet calibration need to be studied. The feasibility of adapting the jet energy calibration to the level-1-trigger is discussed in Section 12.2.1. In addition, the effects that the jet cleaning procedure would have on level-1 jets is discussed in Section 12.5.

Part III.

The Level-1 Topological Processor at ATLAS

One of the universal rules of happiness is: always be wary of any helpful item that weighs less than its operating manual.

— Terry Pratchett, *Jingo*



Visualization of event topologies calculated on L1Topo used to trigger events at ATLAS. The angular distance $\Delta\phi$ between the missing transverse energy and a lepton can be used to target W boson-like decays. Angular distances like ΔR help targeting boosted signal regions of resonances with huge backgrounds with a similar detector signature, for example $H \rightarrow \tau\tau$ decays. By summing up objects inside a ΔR cone, L1Topo makes it possible to utilize substructure information already on the first trigger level, which is used for large- R jet triggers. Invariant mass triggers help resonance searches, in particular low mass resonances with otherwise overwhelming background, for example searches for rare meson decays like $B_s^0 \rightarrow \mu^+ \mu^-$.

The Level-1 Topological Processor

In the following, an overview of the L1Topo system as used during Run-2 is given. The L1Topo hardware has been introduced to the level-1 trigger during the LHC shutdown between Run-1 and Run-2 from 2013 to 2015 [6]. After the commissioning of the system, L1Topo triggers have been used for physics from 2016 until the end of Run-2 in 2018. An upgraded L1Topo system will continue to provide trigger decisions in Run-3 [139]. The chapter starts with an overview on the hardware of L1Topo as used in Run-2 and gives a brief introduction to its functionality in Section 8.1. Afterwards, the firmware structure of L1Topo is introduced in Section 8.2, starting with the Real-Time Data Path (RTDP) in Section 8.2.1, the control path and infrastructure in Section 8.2.2 and the readout path in Section 8.2.3. A particular focus is placed on the algorithm firmware, the central part of the RTDP, in Section 8.3. The configuration of the L1Topo firmware, its implementation and resource utilization are introduced in Section 8.4. Afterwards, the simulation and validation of the firmware is described in Section 8.5. The section concludes with a summary of the physics performance of L1Topo triggers in Section 8.6.

8.1. Hardware Overview

The L1Topo system consists of an ATCA crate [78] and two L1Topo modules. Each of the modules contains two Xilinx Virtex-7 processor FPGAs for algorithm processing and one Kintex-7 FPGA for module control and communication logic via external software applications using the IPBus protocol. A photo of the front panel of the ATCA crate with both L1Topo modules with all connections is shown in Figure 8.1, a photo of an unplugged module can be seen in Figure 8.2. Input data are received via a total of 160 optical fibers with data rates of 6.4 Gb/s per fiber to the backplane from L1Calo and to the front panel from L1Muon. The optical signals are converted into electrical signals in Avago MiniPOD receivers [141] and forwarded to the MGTs of the processor FPGAs, where the data is deserialized for the processing in parallel algorithms. Both processors operate in parallel and are supplied with the full data, but are otherwise completely independent of each other. Every processor can provide a maximum of 32 trigger decisions, resulting in up to 128 possible trigger algorithms. The real-time data output to the CTP uses 32 low latency electrical channels providing speeds of up to 160 Mb/s, an additional 24 channels of optical fibers with 6.4 Gb/s are also available. An additional fiber for the output to the DAQ and RoI interface is provided. To ensure synchronous real-time operation with the trigger system, most of the FPGA internal clocks need to be derived from the LHC clock, which is received via the TTC system. For this purpose, a so-called TTCdec mezzanine card, which is a L1Calo specific mezzanine module to connect to the TTC system, is carried on each module. In addition, a 40.08 MHz crystal clock is available on each L1Topo board as an additional clock source for standalone tests.




Figure 8.1.: The Run-2 L1Topo system inside the ATCA crate. From reference [140].




Figure 8.2.: An unplugged L1Topo module is shown, a few important parts have been marked. Modified from [6].

8.2. Firmware Overview

The actual functionality of the topological processor, including the trigger algorithms, is configured in the firmware of L1Topo, which is written in VHDL. This section discusses the firmware of the processor FPGAs and the algorithmic firmware. From a structural point of view, the processor firmware of L1Topo is partitioned into three large parts: The RTDP, which covers the whole process from data reception over algorithm processing and results output, the infrastructure firmware, responsible for module control and general infrastructure, and the readout path, which provides the input objects and algorithm results of L1Topo to the DAQ system.

8.2.1. The Real-Time Data Path

The RTDP encompasses the whole process of data reception, algorithmic processing of data and the transmission of the algorithm results. Figure 8.3 shows a block diagram of the L1Topo realtime data path firmware and the data flow. The RTDP contains a data reception block, which includes dedicated logic for both L1Calo and L1Muon links, an algorithm block and a block for the output links to the CTP. All blocks contain configurable logic, which can be controlled via IPBus applications. The data reception block processes the received input data of every MGT, outputting 32 bit words at each rising edge of a 160 MHz clock. Afterwards, an alignment step is performed to assign all links to the same LHC BCID by delaying every link with so-called coarse delays. By additional adjustable fine delays, each 32 bit subword is shifted to its expected position in the final 128 bit data word.

After the deserialization took place, a multiplexer is placed to switch between RTDP data from the MGTs and data coming from the playback memories. The playback memories are implemented using the FPGA BRAM resources and are used to playback data consisting of test patterns from an external software source, which makes it possible to inspect the firmware behavior for well-defined inputs. In addition, a configurable branch-off to the spy memories is placed directly after the deserializer. The spy memories allow checks on the deserialized inputs of L1Topo using software applications. Behind this branch-off, Cyclic Redundancy Checks (CRCs) are performed to ensure data consistency. In case of a faulty transmission, data is zeroed in order to avoid processing of erroneous data in the algorithms.

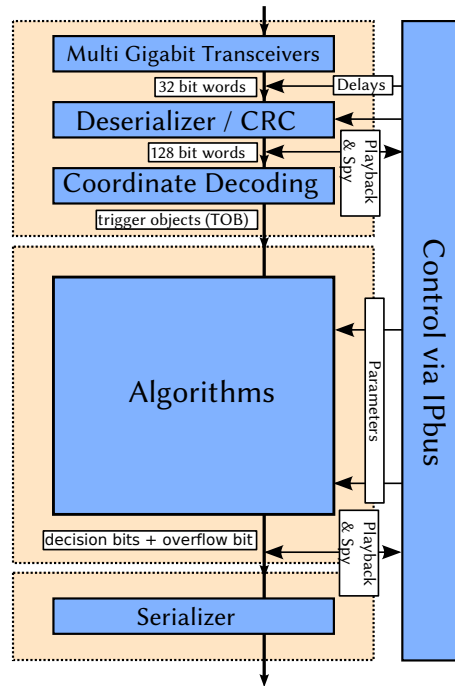


Figure 8.3.: Schematic view of the L1Topo RTDP path. Modified from a similar figure in reference [142].

In the last step of the data reception block, the coordinate decoding is done. The 128 bit data words are converted into the Trigger Objects (TOBs) of L1Topo according to their data formats, which are described in more detail in Section 8.3.1. Electron/photon and tau candidates are received from the CP modules via the CMX system, while jets and missing transverse energy information are provided by the JEP via the CMX, and muons from the muon sector logic via the MUCTPIToTopo. Documentation on how these objects are built at the level-1 trigger is given in Appendix A.2. These trigger objects are passed to the algorithm firmware, where topological algorithms are applied to combinations of the input objects. Configurable parameters, for example decision thresholds of the algorithms, are set from software applications using IPBus. The output of the algorithm block are decision and overflow bits. Additional playback and spy memories are placed after the decision bits are formed and are used to either readout or manually set the algorithm decision bits for debugging purposes. A logical or operation between the decision bits and overflow bits is then performed to create 32 result bits. These result bits are then serialized and provided to the CTP via the electrical channels, which are directly connected to the pins of the FPGA.

8.2.2. The Control Path and Infrastructure

The correct operation of the MGTs and FPGA internal clock circuitry requires a dedicated infrastructure firmware. Since some functionality, for example the deserialization of data or some of the more complex algorithms, runs at higher clock frequencies, several clock signals running at multiples of the 40.08 MHz input clock frequency are created using dedicated clock buffers. In addition, module control via external software requests and inter-module communication is controlled by the firmware of the control path. For this purpose, the standard IPBus firmware as

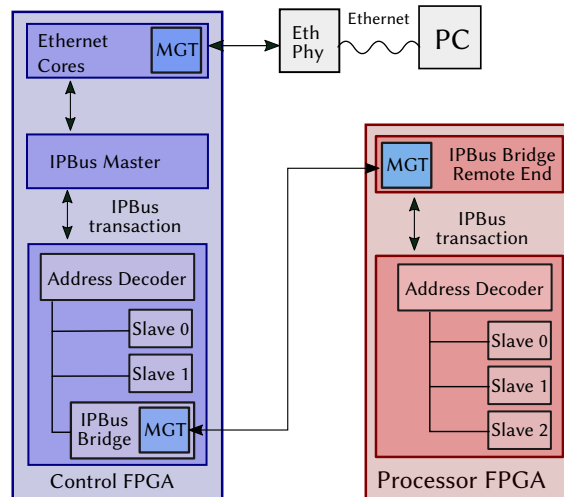


Figure 8.4.: IPBus firmware schematic. An IPBus request is sent via Ethernet, the analog Ethernet signal is digitized using an EthPhy Chip. It is then forwarded to the IPBus master, which uses the address decoder to address the desired slave. Requests for a slave on the processor are forwarded via the IPBus bridge. It exists on both controller and processor side, where the processor block serves as a slave to the controller block. Created from reference [143].

introduced in Section 4.3 has been implemented and adapted to the functionality of L1Topo. In general, the IPBus firmware is implemented in a master/slave model. The IPBus master is located in the firmware of the controller FPGA and contains an implementation of the required protocol stack, namely Ethernet, IP, UDP and IPBus. When an IPBus packet is received via the external Ethernet connection, it is first checked whether multiple IPBus requests have been bundled into a single packet. In the case of multiple requests these are stripped down to numerous single transactions, which are then forwarded to the so-called IPBus slaves. An IPBus slave can be understood as anything that can be addressed with an IPBus request, for example the playback and spy memories introduced in Section 8.2.1, or the parameters of the topological algorithms. Every IPBus slave is identified by a unique memory address. At the beginning of the IPBus slave block, a dedicated address decoder maps the memory address of the IPBus request to the corresponding slave. If the request addresses an IPBus slave living in the processor firmware, the request is forwarded to the so-called IPBus bridge, which transmits requests to the correct processor FPGA via MGT connections. The IPBus bridge itself is an IPBus slave and simply forwards the request. On the processor side, the IPBus bridge remote end deserializes the incoming requests. The remote end serves again as an IPBus master to another address decoder, which then forwards the request to the desired slave on the processor side. The IPBus bridge is completely transparent to both slaves and master, as they are not aware of the inter-FPGA communication. This way, no further distinction between slaves on different FPGAs on the same module is required. A schematic depiction of the IPBus firmware is shown in Figure 8.4.

In total, a single IPBus request from the controller to the processors contains 66 bit, including 32 bit of write data, 32 address bits, 1 bit writeEnable indicating whether the request is a write request and a 1 bit IPBus strobe signal. A change in the strobe signal indicates an incoming IPBus request to the firmware. The total provided data speed for the IPBus bridge is given as $32 \text{ bit} @ 40.08 \text{ MHz} = 1.6 \text{ Gb/s}$. IPBus responses from the processor to the controller contain 32 bit

of read data, an additional acknowledge bit indicates that the slave received the request and an IPBus error bit indicates whether the transaction was successful or an error has occurred. The coherence between the software and firmware view of IPBus slaves and their addresses is ensured by using a common XML description.

8.2.3. The Readout Path

After sending out the algorithm results, the input TOBs used by L1Topo and the trigger decisions need to be buffered for every event as long as the corresponding L1 trigger decision is not formed yet. In case of an L1A for a particular event, the corresponding event data needs to be retrieved from the buffer memories and forwarded to the ROS. A dedicated firmware part, the readout firmware, is responsible for managing the data flow from the RTDP and the algorithmic firmware to the ROD [144]. The process of transporting the data from the RTDP to the ROD is divided in multiple parts. In the first stage, all relevant data from L1Topo is collected and buffered in circular buffers, awaiting a possible L1A from the CTP. For these buffers, the BRAMs are used in dual-port mode with a total memory depth of 256 words to provide enough size to wait for the delayed L1 trigger decision. Individual buffers exist for every data source, e.g. for the inputs of L1Topos from L1Calo and L1Muon system, as well as for the result and overflow bits. An additional FIFO (First In, First Out) buffer saves pointers to the positions of the data to properly access the data from the correct buffers when required. If a L1A is received, the corresponding data slices are fetched from the buffers. Afterwards, the data from different sources is merged, and formatted correctly by adding the appropriate headers and trailers for readout data. The data is now ready for transmission and can be transmitted to the ROD via the backplane Gbit DAQ fiber.

8.3. The Algorithmic Firmware

The algorithmic firmware is the central component of the L1Topo firmware, as it performs the topological algorithms used for the L1 trigger decision. A block diagram of the algorithmic firmware is shown in Figure 8.5. A total latency of 3 BC is available for the algorithm processing. The real-time input of L1Topo consists of 120 electromagnetic clusters and 120 tau objects, 32 muons, 64 jet objects and missing energy information. These numbers fully utilize the available bandwidth and allow each trigger module to provide most relevant objects even for events with extreme amounts of particle activity. However, in almost all cases, the number of objects present in an event is significantly lower than what can be provided via the available bandwidth, in these cases most of the provided objects are equal to zero. Since the positions of the non-zero objects are not known a priori, it is important that all provided objects are considered. In principle, it would be possible to calculate topological quantities, for example invariant masses, directly from the lists of all particles. However, most topological algorithms work on the combinations of pairs of TOB from different input lists. Due to this, the amount of required operations scales quadratically with the length of the input lists, requiring a rapidly growing amount of the available FPGA resources. To avoid this, dedicated algorithms can be used to identify the most significant, high-energetic TOBs, as physics triggers are usually interested in those objects. Therefore, the first stage of the topological algorithm firmware is used to reduce the number of objects by choosing only the potentially interesting TOBs. Two algorithm types are applied to reduce the number of input objects without losing important event information. The first one are *sort algorithms*, which create lists of the six leading particles for a certain input type. The second type are *select algorithms*, which create lists of ten particles above a configurable E_T threshold. For

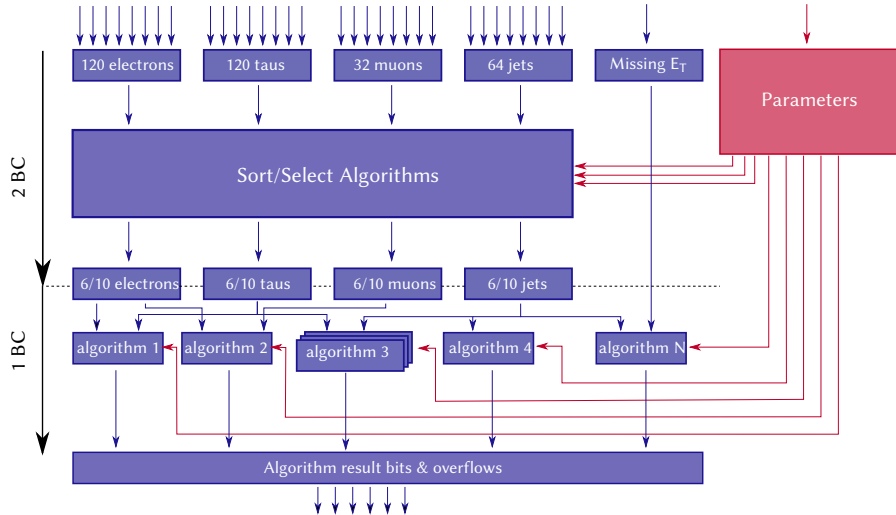


Figure 8.5.: L1Topo algorithm firmware structure. Two stages of algorithms are implemented, the sort/select algorithm stage and the decision algorithm stage. The sort/select algorithms create short lists of input objects, which are then combined in the decision algorithms to create result and overflow bits. Algorithms can be configured by multiple parameters.

both algorithms, additional selections on $|\eta|$ or on object specific information like isolation for electrons or taus can be performed. In addition, the desired jet energy window of either 2×2 or 4×4 jet elements has to be selected at this stage. The sort and select algorithms are relatively resource intensive and have a fixed delay of 2 BC (≈ 50 ns).

All sort/select algorithms convert their input TOBs into a global data format, the so-called *GenericTOB*. This is necessary since the different subsystems of the L1 trigger provide TOBs with different spatial granularities as well as different energy scales. For example, the objects reconstructed by the CP are built from 0.1×0.1 trigger towers, and therefore have a higher granularity and different data format than the jets reconstructed by the JEP based on jet elements of size 0.2×0.2 . The conversion into the *GenericTOB* format utilizes dedicated LUTs for each kind of input TOB, which map the object specific coordinates to common global coordinates. Some object specific information, like isolation for the CP objects, cannot be saved in a generic format and is therefore used in the beginning of the sort/select algorithm stage. Having all TOBs available in the same data format means that the identical algorithm logic can be reused independent on the input objects, significantly reducing code duplication and the potential for bugs. In addition, this approach makes it easy to combine TOBs from different subsystems in the same algorithm logic.

The output of the first algorithm stage consists of shortened lists of *GenericTOBs*, which can then be flexibly combined into the *topological decision algorithms*. Their latency is limited to one BC. A wide variety of topological algorithms is available, examples will be discussed in Section 8.3.3. Their outputs are one or more decision bits, indicating whether the trigger condition is fulfilled, and overflow bits, indicating several sources of overflows, for example energy saturation or integer overflows of internally calculated quantities. Sort/select algorithms as well as decision algorithms depend on several parameters, which are always implemented to be configurable via IPBus. These parameters encompass everything that can be changed without affecting the actual algorithm logic, for example isolation requirements, E_T and η cuts of the select algorithms or

decision thresholds of the topological algorithms. In the following, some aspects of the firmware are discussed in more detail, namely the data formats in Section 8.3.1, the sort/select algorithms (Section 8.3.2) and the decision algorithms (Section 8.3.3).

8.3.1. Data Formats

In order to give a good overview of the most important aspects of the internal data formats, the jet data format and its conversion to the GenericTOB is discussed as an example. An in-depth description of all the different formats has been documented for example in reference [145]. In order to utilize the available bandwidth as good as possible, the TOBs provided to L1Topo contain only necessary information that cannot be inferred implicitly. Coordinates like η and ϕ are sent as local bin numbers in the coordinate system of the respective submodule. The global coordinates are then inferred from information on the calorimeter coverage of the module the object was found on. For example, each JEM provides coordinates only in its local coordinate system, compare Figure 6.11 for the JEM coverage. Each JEM then simply enumerates all of its covered jet elements to define a local coordinate system. The format of a single jet TOB found by the JEP and provided by the CMX is shown in Figure 8.6. The 10 bit energy is the energy inside the 4×4 window of jet elements, and the 9 bit energy inside the 2×2 window. The local position information is encoded in the 3 bit frame information, which enumerates eight different windows of 2×2 jet elements inside the coverage of a single JEM. An additional JEM local coordinate information enumerates the local position inside the window of 2×2 jet elements. In a first step, this information is combined to calculate local bin numbers for both η and ϕ , enumerating the jet elements inside the coverage of a single JEM from 0 to 3 in η and 0 to 7 in ϕ . To find the corresponding global bin number, the corresponding JEM can then be uniquely identified using the provided JEM number information and the corresponding crate number, which is inferred from the fiber the object is transmitted on. For each JEM, individual offsets are calculated to give each jet element a unique pair of global η and ϕ bin numbers. For example, the JEM which is marked green in Figure 6.11 gets an η offset of 16 and a ϕ offset of 0. This procedure results in 5 bit η and 5 bit ϕ global bin numbers, enumerating the jet elements starting at the lowest η and lowest ϕ possible. The data format for jets in L1Topo after decoding of the bit words from the CMX system is given as follows:

- Jet TOB: 9 bit E_{T1} , 10 bit E_{T2} , 5 bit η , 5 bit ϕ

In order to correctly calculate topological quantities like $\Delta\phi$, ΔR or m_{jj} , these bin numbers have to be converted to global coordinates. Since the physical positions of the jet elements are known and

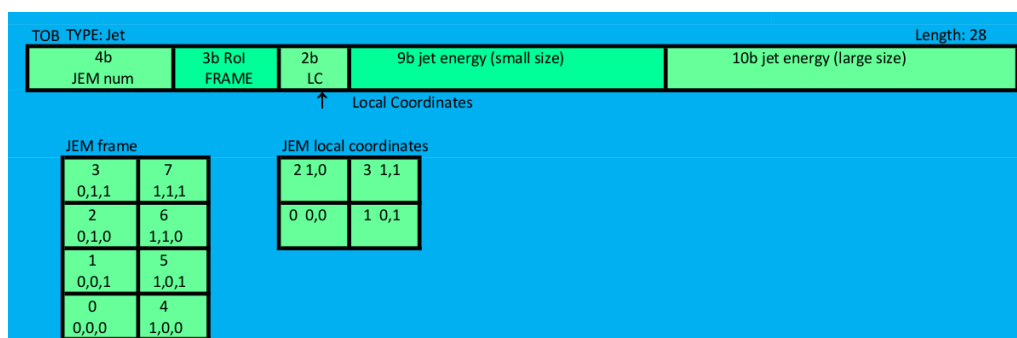


Figure 8.6.: Jet TOB data format from the CMX. Taken from reference [146].

fixed, predefined LUT are used for this purpose. In this context, every type of TOB is converted into the GenericTOB. The GenericTOB format is designed to provide sufficient granularities for all different input formats and is given as

- GenericTOB: 10 bit E_T , 7 bit η , 6 bit ϕ 1 bit overflow.

The 10 bit E_T are used to represent unsigned values of 0 GeV to 1023 GeV with a granularity of 1 GeV, an additional overflow bit is used to indicate an energy value of larger than 1023 GeV. For E_T the conversion is straight forward, as the jet energy is already given at a scale of 1 GeV. The 7 bit η are used to represent values between -6.4 and 6.3 using integer numbers between -64 and 63 . Due to the non-uniform size of the jet elements for higher η , the conversion for η is non-linear and therefore requires a LUT, which is shown in the appendix in Table A.2. For jets, possible η values lie between -3.9 and 3.9 as given by the possible jet element positions. The 6 bit ϕ are used to represent values of $0 \leq \phi < 2\pi \approx 6.4$ using integer number between 0 and 63. The conversion for ϕ is linear and given as follows:

$$\phi_{\text{generic}} = (2 \cdot \phi_{\text{jet}} + 2) \bmod 64, \quad (8.1)$$

where the approximation $\phi_{\text{generic}} = 32 \frac{\phi_{\text{float}}}{\pi}$ is made. To avoid FPGA resource and latency expensive floating-point arithmetic, all coordinates and energies of the generic format are represented as integer values. While implementation of floating-point IP Cores exist, they require more resources and are slower than comparable integer operations, especially due to the increased number of corner cases which have to be handled, for example NaN values or shifters to align mantissas of different numbers. Due to this, no floating-point operations are performed in the scope of this thesis. Event topologies like $\Delta\eta$, $\Delta\phi$ or ΔR^2 between pairs of objects are now simple arithmetic operations involving integers only.

8.3.2. Sort and Select Algorithms

Due to the tight latency budget of only 2 BC, the algorithm implementations of the sort/select stage need to be as parallelizable as possible. Dedicated algorithm exists for every type of input TOB, taking object specific information and the varying number of objects into account. The conversion of the TOBs to GenericTOB is the first step of these algorithms, as this allows to apply the identical sorting or selection logic to each list of input TOBs. TOBs failing isolation criteria are set to zero during the conversion. Afterwards the algorithms are based on two stages, which utilize multiple sort and select blocks. These are the most basic blocks and are therefore introduced in the following. The sort block is based on performing many comparison operations in parallel to construct a quadratic comparison matrix, which is shown in Figure 8.7. The number of columns and rows corresponds to the number of objects that have to be sorted. For each field below the diagonal, the i^{th} and the j^{th} input TOB are compared, where i corresponds to the row and j to the column of the matrix. If the energy of the i^{th} TOB is larger than or equal to the energy of the j^{th} input TOB, the corresponding matrix entry is set to 1, else it is set to 0. Due to the anti-symmetry of the matrix, the upper half of the matrix can be derived from the lower half. The number of ones in the i^{th} column determines the position of the i^{th} input object in the output list: when n objects are sorted, the column of the leading object will contain exactly n ones, the column of the subleading object $n - 1$ ones, and so on.

Furthermore, the sorting matrix approach allows for the utilization of presorting. Presorting means that equal-sized subblocks of the input list have already been sorted, for example by the




Figure 8.7.: Example of the sorting matrix for 6 input channels, modified from reference [70].

module which identified these objects. If the input of n objects consists of n/k blocks of k presorted objects, the presorting information can be used to pre-fill blocks of size $k \times k$ around the diagonal of the matrix, this is illustrated in Figure 8.8 for $n = 6$ and $k = 2$.

The logic of the selection block is similar to the logic of the sorting block, though the energy of every TOB is only compared once against a threshold value instead of all other objects. A single bit is used to indicate whether this threshold is passed or not and TOBs failing the selection are zeroed. In order to build the output array, the remaining TOBs are then left-aligned. The resulting mapping of the input channels of a select module to the output channels are shown in Figure 8.9.

The sorting algorithms are based on these two blocks and are performed in two stages as shown exemplary in Figure 8.10 for the jet sorting algorithm. For other input objects the algorithms are designed similarly, though the number of input objects or the number of presorted blocks can be different. Before the sorting stages, the jet TOB objects are converted into GenericTOB objects. In the first stage, the 64 inputs are divided into two smaller lists, as the CMX system provides two presorted list of 32 jets. Both smaller input lists are then forwarded to select modules, which are configured to select the six most energetic objects of the presorted block. If no η cuts are

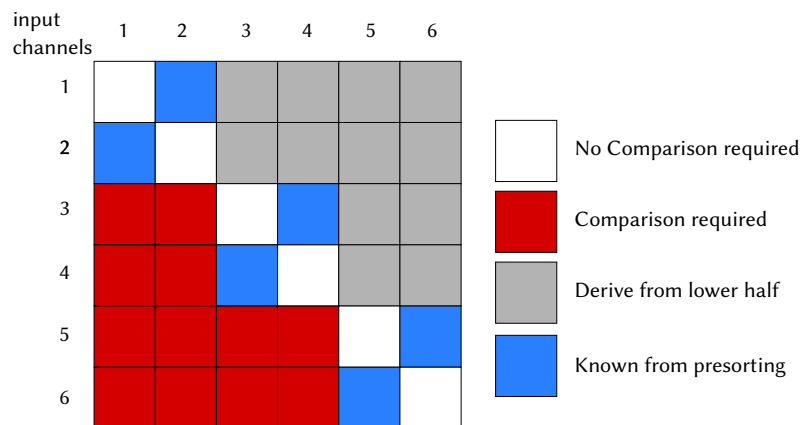


Figure 8.8.: Example of the sorting matrix for 6 input channels utilizing presorting information, where blocks of 2×2 objects are already sorted.




Figure 8.9.: Scheme of the L1Topo select block, adapted from [70]. In this example, six objects are selected out of ten objects. Accepted objects are depicted in green, object below the threshold in red. The output list is left-aligned, additional objects above the output length are lost.

required this procedure is trivial, with η cuts an additional alignment step is required. In the next BC, the 12 outputs from the first stage are sorted with the described sorting matrix logic, utilizing the fact that the list consists of two presorted blocks of six TOBs. In this case, the utilization of presorting reduces the amount of required comparison operations by about 45%. The FPGA resource utilization of the JetSort algorithm is shown in Table 8.1. While a utilization of 3.48% sounds acceptable, it is important to consider that multiple instances of this algorithm can be required in a firmware, for example configured with and without η cuts for both small and large energy window. In addition, sorting algorithms for muons, electron/photons, or tau objects have to be implemented as well. Therefore, these algorithms dominate the FPGA utilization. The select algorithms are also based on two stages similar to the sort algorithm, but with a modified second stage which only merges the first stage outputs instead of sorting them.

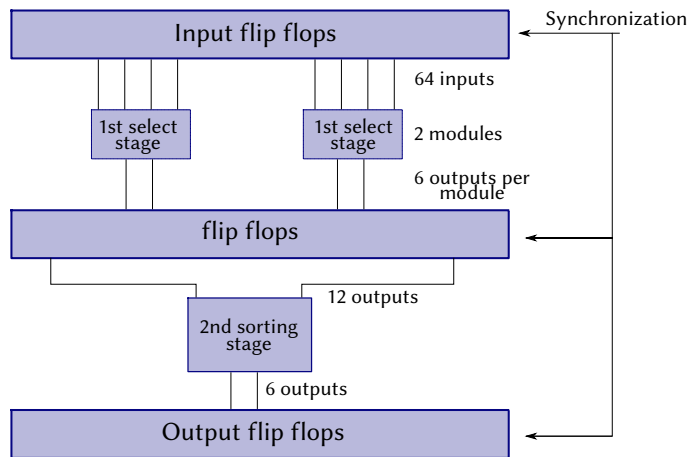


Figure 8.10.: Scheme of the L1Topo jet sorting algorithm, adapted from reference [70]. 64 inputs in blocks of two presorted objects are sorted down to six objects.

Resource	Required	Utilization [%]
LUT	15087	3.48
Flip-Flops	4040	0.47

Table 8.1.: Resource utilization of the JetSort algorithm. A high amount of LUT resources is required for this algorithm, while no DSP or BRAM logic is required.

8.3.3. Topological Decision Algorithms

A large variety of topological algorithms is implemented, and has been documented in detail, for example in reference [147]. Since an in-depth description of all algorithms would go far beyond the scope of this thesis, the discussion in the following is concentrated on the general algorithm structure. Possible topological algorithms encompass all possible combinations of E_T , η and ϕ of multiple objects, given that enough resources and latency is available.

In general, the firmware of the topological algorithms is designed in a modular fashion, meaning that dedicated internal entities are implemented to calculate topological quantities which are required in multiple algorithms. For example, the *DeltaEtaCalc* entity to calculate $\Delta\eta$ or the *DeltaPhiCalc* to calculate $\Delta\phi$ can be reused inside multiple algorithms. This avoids code duplication and reduces the number of possible sources for firmware bugs. Algorithms which combine pairs of objects often exist in two versions: an *Inclusive1* algorithms uses all combinations of TOBs from a single list, an *Inclusive2* algorithm uses all combinations of TOBs from two different input lists. Every topological algorithm takes one or more lists of GenericTOB as inputs. An overview table with examples of topological algorithms is shown in Table 8.2, summarizing their general functionality. As an example and due to its importance in the dijet analysis presented in Part IV, the functionality of the *InvariantMass* algorithm is explained in detail in the following. Implementation details for the other algorithms whose trigger performance is discussed in Section 8.6 are given in Appendix A.4.2 for ΔR^2 and in Appendix A.4.3 for SimpleCone.

Table 8.2.: Examples for topological algorithms performed on L1Topo.

Algorithm Overview		
Algorithm name	Quantity	Decision
DeltaEta	$\Delta\eta = \eta_1 - \eta_2 $	$\eta_{\min} \leq \Delta\eta \leq \eta_{\max}$
DeltaPhi	$\Delta\phi = \phi_1 - \phi_2 $	$\phi_{\min} \leq \Delta\phi \leq \phi_{\max}$
DeltaRSqr	$\Delta R^2 = \Delta\eta^2 + \Delta\phi^2$	$\Delta R_{\min}^2 \leq \Delta R^2 \leq \Delta R_{\max}^2$
InvariantMass	$2E_{T,1}E_{T,2}(\cosh(\Delta\eta) - \cos(\Delta\phi))$	$M_{\min}^2 \leq M^2 \leq M_{\max}^2$
TransverseMass	$2E_{T,1}E_T^{\text{miss}}(1 - \cos(\Delta\phi))$	$m_{T,\min}^2 \leq m_T^2 \leq m_{T,\max}^2$
Disambiguation	$\eta, \phi, \Delta R$	$(\phi_1 \neq \phi_2 \wedge \eta_1 \neq \eta_2) \vee \Delta R^2 > R_{\max}^2 \vee \text{TOBs}$
JetHT	$H_T = \sum_{\text{TOBs}} E_T$	$H_T > H_{T,\min}$
SimpleCone	$E_{T\text{sum}} = \sum_{\Delta R \leq \Delta R_{\text{cone}}} E_T$	$E_{T\text{sum}} > E_{T\text{min}}$

8.3.4. The Invariant Mass Algorithm

In general, the invariant mass is implemented using Equation (6.9), as this expresses the invariant mass using only (E_T, η, ϕ) , i.e. the information present in the TOBs of the level-1 trigger. For each pair of TOBs, the topological quantities $\Delta\eta = |\eta_1 - \eta_2|$ and $\Delta\phi$ are calculated via the dedicated entities, where $\Delta\phi$ is calculated as

$$\Delta\phi = \begin{cases} 64 - |\phi_1 - \phi_2|, & \text{if } |\phi_1 - \phi_2| > 32 \\ |\phi_1 - \phi_2|, & \text{else} \end{cases} \quad (8.2)$$

The main difficulty in a firmware implementation is the calculation of the angular difference term $\cosh \Delta\eta - \cos \Delta\phi$, as both operands take on non-integer values. Instead of using floating-point numbers, the required decimal precision is defined beforehand, which allows for integer arithmetic using *fixed-point* numbers. In these numbers, a fixed number of bits is allocated to represent the numbers before and after the radix point. Operations on binary patterns in fixed point representation can be treated equally to operations with integer numbers. However, some details have to be considered when performing operations with fixed-point numbers. The number of bits allocated for decimal bits needs to be identical for both operands in an addition or subtraction. When multiplying two fixed point numbers, the numbers of required integer and decimal bits for the result are equal to the sum of the number of bits of the operands. For the invariant mass, 7 bits have been chosen for the decimal places of $\cosh \Delta\eta$ and $\cos \Delta\phi$, as this is the minimum number of bits where no duplicate values occur due to rounding imprecisions with the given input granularities of η and ϕ . One integer and one sign bit are required in addition for the cos operation, while 11 bit integer are required for the cosh operation, as this is sufficient to represent the highest possible value of $\cosh 7.8$.¹ Given a floating-point number z , the corresponding fixed-point representation using 7 bit for the decimal places is found as follows

$$z_{7\text{ bit}} = \frac{1}{2^7} \lfloor z \times 2^7 \rfloor, \quad (8.3)$$

where the operation $\lfloor \cdot \rfloor$ denotes rounding to the closest integer value. Due to the limited granularity of L1Topo inputs, only 79 different possible values exist for $\Delta\eta$, ranging from 0 to 7.8 in steps of 0.1. In addition, only 33 possible different values exist for $\Delta\phi$, compare Equation (8.2). Therefore, the cos and cosh function calls can be replaced by LUT operations, where the values for all possible inputs are pre-computed and saved in the LUT memories of the FPGA. These LUTs return the fixed point representations $\cosh \Delta\eta_{7\text{ bit}}$ ($\cos \Delta\phi_{7\text{ bit}}$) multiplied by 2^7 , which effectively shifts the decimal point by seven places to obtain integer values. As an example to illustrate the steps involved when calculating an invariant mass, consider the case of two jets, the first one represented by a TOB of $(E_T \ \eta \ \phi) = (35 \text{ GeV} \ 20 \ 10)$, the second one by $(21 \text{ GeV} \ -9 \ 22)$. In the first step $\Delta\eta$ and $\Delta\phi$ are calculated as

$$\Delta\eta = |20 - (-9)| = 29, \quad \Delta\phi = |10 - 22| = 12. \quad (8.4)$$

The angular functions $\cosh \Delta\eta$ and $\cos \Delta\phi$ are then calculated using LUTs as follows

$$\begin{aligned} \cosh(2.9) &\approx \cosh(2.9)_{7\text{ bit}} = 9.1171875 = \frac{1}{2^7} 1167 \Rightarrow \text{LUT}_{\cosh}(29) = 1167 \\ \cos(1.2) &\approx \cos(1.2)_{7\text{ bit}} = 0.3828125 = \frac{1}{2^7} 49 \Rightarrow \text{LUT}_{\cos}(12) = 49 \end{aligned} \quad (8.5)$$

¹The maximum value of $\Delta\eta = 7.8$ is reached if one jet has an η of -3.9 and the other one of 3.9 .

The intermediate value of the squared invariant mass is calculated by multiplying the angular difference with the energy term $2E_{T,1}E_{T,2}$

$$m_{jj,\text{int}}^2 = 2 \cdot 35 \text{ GeV} \cdot 21 \text{ GeV} \cdot (\text{LUT}_{\cosh}(29) - \text{LUT}_{\cos}(12)) = 1643460 \text{ GeV}^2. \quad (8.6)$$

The final invariant mass is then calculated by performing the division by 2^7 using a simple bit-shift by 7 bits

$$\begin{aligned} m_{jj}^2 &= \frac{1}{2^7} m_{jj,\text{int}}^2 = \frac{1}{2^7} 1643460 \text{ GeV}^2 = \frac{1}{2^7} 11001000100111 \underbrace{1000100}_{\text{Drop}} |_2 \text{ GeV}^2 \\ &\approx 11001000100111 |_2 \text{ GeV}^2 = 12\,839 \text{ GeV}^2, \end{aligned} \quad (8.7)$$

where the notation $|_2$ indicates the binary representation of an integer number. Even though an implementation of a square root algorithm exists, it is usually not applied in the InvariantMass algorithms as it is not strictly necessary. Instead, all trigger decisions are performed using the squared invariant mass. To compare the result with the float value of

$$\sqrt{2 \cdot 35 \text{ GeV} \cdot 21 \text{ GeV} \cdot (\cosh(2.9) - \cos(1.2))} = 113.43 \text{ GeV}, \quad (8.8)$$

the square root of the result of Equation (8.7) is performed:

$$\Rightarrow m_{jj} = \sqrt{12\,839 \text{ GeV}^2} = 113 \text{ GeV}, \quad (8.9)$$

which shows that the desired result of the closest integer number is achieved. An analysis of the precision of this method for all possible input values shows that the median of the relative error introduced by this method is below 10^{-4} and negligible. However, in a few cases where $\cosh \Delta\eta \approx 1$ this relative error reaches the percent level. A more detailed discussion of the precision of the invariant mass algorithm is given in Appendix C.

To investigate how the invariant mass algorithm is synthesized and implemented, the elaborated design for this algorithm can be checked and is shown in Figure 8.11 for the simple case of a single invariant mass between two TOBs. After firmware synthesis, the calculation of a single invariant mass requires a total of 170 LUTs and 2 DSPs. The first DSP is required for the multiplication of the two energies and the second one for the multiplication of the energy and angular parts.




Figure 8.11.: Schematic of the invariant Mass algorithm created by Vivado. The DeltaEtaCalc and DeltaPhiCalc are instantiated first, afterwards the cos and cosh entities are placed. Two multiplications and a subtraction are combined to calculate the invariant mass.

8.4. The Trigger Menu and Firmware Implementation

The available sort/select and decision algorithms can be combined into hundreds of different unique trigger items. The actual configuration is driven by the demands of the ATLAS physics program, and can change several times a year. To ensure a consistent configuration, the trigger items running on each processor FPGA are defined by the trigger menu, where the full algorithm chain, including sort/select, decision algorithm and parameters is listed for each used output bit of the FPGAs. Therefore, the algorithm top modules in the firmware are derived from the configuration in the trigger menu. Changes in the menu which require changes of algorithm logic or the introduction of new algorithms require rebuilding of the firmware. However, parameter changes are fully handled by the corresponding trigger online software. The menu configuration is provided in a parsable XML format. From this file, the required VHDL algorithm top module files and IPBus addresses for the algorithm parameters are generated automatically, ensuring a consistent configuration of the firmware by avoiding manual adjustments. At the end of 2018, 113 triggers have been configured in total from the trigger menu. The resources required by this configuration are shown in Figure 8.12 for the four processor FPGAs of the two modules, which are called Topo0 U1 and Topo0 U2 for the first module and Topo1 U1 and Topo1 U2 for the second module. The most utilized resource are LUTs, as these are the most general-purpose resource and required for the logical comparisons in the resource intensive sort- and select algorithms. BRAM resources are mainly utilized by the readout firmware to implement the required data buffers. The DSP resources are required by all algorithms containing multiplications of longer numbers, most prominently the invariant mass algorithms.




Figure 8.12.: Resource utilization of the final Run-2 L1Topo firmware in 2018. The two Topo modules are named Topo0 and Topo1, each containing two processor FPGAs named U1 and U2.

8.5. Commissioning, Simulation and Validation

Commissioning of the system and validation of the trigger logic are essential steps on the way to routine usage of L1Topo triggers for physics. In general, multiple commissioning steps are




Figure 8.13.: The L1Calo test rig.




Figure 8.14.: The L1Topo test rig module.

performed, starting with basic low-level component and hardware system tests and moving up to more and more complex tests involving multiple subsystems and advanced logic [148]. The most basic algorithm validation steps are performed using the simulation engine provided by Vivado as introduced in Section 4.5.4. A simulation test bench is written for each algorithm, and basic functionality like the timing behavior is checked by utilizing the waveform window. In a next step, the VHDL simulations are automatized to allow checking the algorithm logic for a larger amount of possible inputs. For this purpose, a corresponding implementation in the high-level programming language python is made for each sort, select and decision algorithm. To emulate the real-time inputs, text files of bit words corresponding to the different data types of L1Topo are generated. This way, it is possible to automatically run tests using random lists of TOBs that the various algorithms expect as inputs. These bit words are read in and processed in the python implementation of the algorithm. The resulting TOBs or decision bits are then written to an output text file. At the same time, the TOBs from the input text files are read in VHDL and the corresponding algorithm simulation is run. The results are again written to a text file and the results of python and VHDL are compared afterwards. Every time when the VHDL code of an algorithm changes, its logic can be checked for numerous inputs using the python implementation. However, this procedure only implements unit tests for the logic of individual, single algorithms. While those are very useful to catch even subtle flaws in the logic of an individual algorithm, not all corner cases can be covered. On the one hand, it is very difficult to cover the whole phase space of possible input objects, on the other hand, the integration within the actual hardware and all required infrastructure cannot be tested this way. Integration tests of the behavior of algorithms integrated into the control and readout firmware can be performed using the playback and spy memories, where well-prepared inputs are written to the playback memories at the beginning of the RTDP, and intermediate algorithm result or the final decision bits are spied upon using the spy memories. To test the algorithms using real inputs from L1Calo modules and to test the integration of hardware modules and firmware, the L1Calo test rig is available. The test rig contains a simplified, reduced version of the L1Calo setup at ATLAS, including a few JEMs, CPMs, PPMs and one L1Topo board. Photographs of the test rig can be seen in Figures 8.13 and 8.14. It is useful for a first steps of firmware validation using the actual hardware. In tests in the test rig, it is possible to create simple patterns from the JEP or CP system, e.g. jets with monotonically increasing E_T , and check the firmware behavior for individual events. The final stages of validation are performed using the full ATLAS system. For this purpose, a complete bitwise simulation of all algorithms has been implemented and is performed for real events. If

an event has been accepted by the L1 trigger, all the inputs of L1Topo as well as its decision and overflow bits are provided to the ROS via the readout firmware of L1Topo. The bitwise simulation is then run in the HLT on these input objects, and the results of hardware and simulation are compared for each accepted event and monitoring histograms of the mismatch rates are provided during a run. These histograms are written out via the MDA and are then available offline for further inspections. This information allows to compare the hardware and simulation decisions for all events accepted by the L1 trigger. Eventual mismatches are further investigated. When the mismatch rates reach a satisfying level, the triggers are enabled for the data taking. During ATLAS data taking, the comparison between the hardware and simulation decision is performed in real-time in an online monitoring system to ensure the integrity of the recorded data and to react promptly in case of any issues. Figure 8.15 shows an example of recorded mismatch rates for a run taken on September 24, 2018.




Figure 8.15.: The mismatch rates between the L1Topo hardware decisions and the simulation decisions are shown for a selection of algorithms. [149]. Mass and energy selections are shown in units of GeV. Indices “C” and “F” denote a selection of central ($|\eta| < 2.5$) and full range objects, respectively. For all shown algorithms, the agreement is at least at the per mille level.

8.6. Trigger Performance

When an L1Topo trigger is completely validated, it can be enabled for physics, which means that data is actually recorded using the decisions of the corresponding trigger items. In this section, the performance of some trigger algorithms of L1Topo is discussed, presenting some of the most important applications. For all use cases, significant improvements of background rejection and improved acceptance of physics signal events are provided even with rising instantaneous luminosities, making L1Topo indispensable for many analyses.

Hadronic τ Lepton triggers for $H \rightarrow \tau\tau$ One of the largest positive impacts of L1Topo is for analyses with ditau final states, which occur prominently in $H \rightarrow \tau\tau$ decays, where one or both taus decay hadronically [150]. The non-topological selection requires two isolated tau objects

identified by the CP, with the first one having a transverse energy of at least 20 GeV and the second one of at least 12 GeV. However, this low- E_T selection suffers from a large background of QCD multijet events with a similar detector signature. To improve this, a L1Topo trigger is designed which applies selections made by the offline analysis already on trigger level. For background suppression, a boosted signal region with higher Higgs momentum is considered, where both τ leptons tend to be closer together in ΔR . At the same time, QCD multijet events with the same signature will have on average two taus created back-to-back. By using the L1Topo DeltaRSqrIncl1 algorithm, a requirement on the angular distance of $\Delta R(\tau_1, \tau_2) < 2.9$ is introduced, which significantly reduces the trigger rate by a factor of four, compare Figure 8.16(a). At the same time, no signal efficiency is lost by requiring the ΔR selection, as shown in Figure 8.16(b). An additional rate reduction is achieved by requiring the presence of an additional jet with at least 25 GeV momentum to balance the boosted Higgs.




Figure 8.16.: Application of an L1Topo ΔR algorithm on $H \rightarrow \tau\tau$: A large rate reduction is achieved (a), without any efficiency loss (b), from reference [150].

B physics triggers Many of the ATLAS B -physics and light states analyses are based on the identification of B mesons or other low mass resonances via their decay products, often containing pairs of leptons. Due to the low meson masses in the range of a few GeV, the final state leptons tend to have very low momenta. For example, the non-topological trigger requires two muons of 6 GeV for dimuon final states. However, this also accepts many background events and results in a very high rate. Using a combination of the DeltaRSqrIncl1 and the InvariantMassInclusive1 algorithms, additional requirements on the muon invariant mass $m_{\mu_1\mu_2}$ and the angular separation $\Delta R_{\mu_1\mu_2}$ are imposed. By requiring $m_{\mu_1\mu_2}$ to be in the range of 2 GeV to 9 GeV and $0.2 < \Delta R_{\mu_1\mu_2} < 1.5$, the trigger rate is reduced by a factor of 3.5, compare Figure 8.17(a), with just small losses in the efficiency, compare Figure 8.17(b), and without introducing a bias in the mass distribution. This allows to keep the required low- p_T muon thresholds and keeping the rates at a manageable level at the same time.

Jet Substructure In boosted topologies, for example in decays of high energetic top quarks with hadronically decaying W bosons, it frequently occurs that multiple jets are very close together in

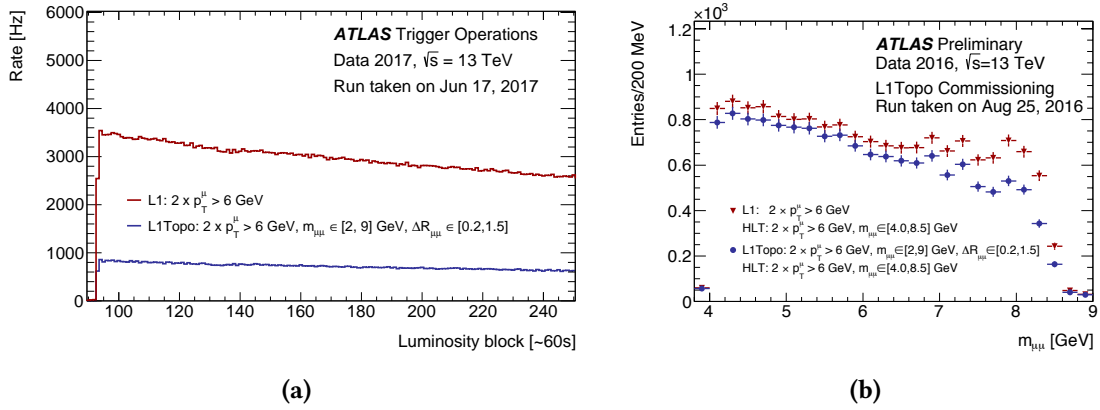


Figure 8.17.: Large rate reduction for topological dimuon triggers (a), from reference [151] and only small efficiency loss (b), from reference [152].

ΔR . In these cases, the level-1 sliding window algorithm is not able to correctly identify the three single jets that are actually part of one large radius, boosted jet. Analyses relying on these boosted jet signatures often use larger radius parameters of $R = 1.0$ for the offline jet reconstruction. However, even the widest jet window provided by the JEP is too small to capture the complete energy of these jets, having only a size of 0.8×0.8 . This effect gets worse when large jets with multiple subjets are considered, where it is more probable that parts of the energy are not captured correctly by the level-1 jet algorithm. To regain the efficiency lost in the first trigger level, the SimpleCone algorithm is used and configured to sum up all jets with energy $E_T > 15$ GeV within a cone of size $\Delta R = 1$. The resulting energy sum is required to be larger than 111 GeV, as this equals the rate of the non-topological trigger running at 100 GeV. When the number of subjets gets higher, the efficiency of the non-topological trigger degrades, as can be seen from the open points in Figure 8.18. The filled points indicate the efficiency of the topological trigger, which regains efficiency for jets with at least two subjets, providing a suitable large-jet trigger on L1.





Figure 8.18.: Efficiencies for two different L1 triggers are shown for large- R jets with different numbers of subjets as a function of the offline leading jet p_T . Taken from reference [153].

Part IV.

The Analysis

Das haben wir noch nie probiert, also geht es sicher gut.

— Pippi Langstrumpf



Event visualization of a leading order $pp \rightarrow Z' \rightarrow qq$ event generated with Madgraph and Pythia. The visualization was created using `mcviz` [154]. While the upper left part of the visualization shows primarily soft QCD interactions of the underlying event, the lower right part shows the decay of a Z' boson into two high-momentum charm quarks, resulting in two collimated jets of high-energetic particles.

L1Topo Histogramming

9.1. Motivation for a Triggerless Analysis

Due to the high cross section of inelastic QCD processes, low- p_T single jet triggers are required to have high prescales to keep the trigger rates manageable for the DAQ system. Figure 9.1 shows the evolution of prescales of a level-1 trigger item requiring at least one L1 jet with at least 25 GeV E_T in a run in 2018.




Figure 9.1.: Evolution of prescales of the L1_J25 trigger, which requires one level-1 jet with at least 25 GeV E_T . The run was taken on August 15, 2018.

The prescales range from 175 000 at the start of the run to about 50 000 at the end of the run, meaning that only 1 in about 100 000 events fulfilling the trigger condition is saved. For a correct normalization of the recorded data, every accepted event has to be weighted by the corresponding prescale p . However, the statistical uncertainties also scale with a factor \sqrt{p} compared to data recorded with an unprescaled trigger. Due to the increased uncertainties, the sensitivity of a search relying on such a low-energy trigger is severely reduced. This results in statistical limitations for dijet resonance searches, which therefore resort to using unprescaled triggers at higher p_T values. Figure 9.2 shows the invariant mass spectrum of the ATLAS offline dijet resonance search, where only masses above 1.1 TeV are included in the analysis [155] due to the single-jet trigger prescales. This in turn limits the sensitivity in the low mass region, as can be seen in Figure 9.3,




Figure 9.2.: Invariant mass spectrum of the offline analysis. The blue curve counts events accepted by any single-jet trigger, the red curve contains additional prescale corrections. From reference [7].




Figure 9.3.: Limits on the Z' mediator coupling to baryons g_B as a function of mass [156]. Some initial limits from the ATLAS and CMS searches are shown based on Run-1 data.

where the exclusion limits on the vector coupling g_q as a function of the mediator mass $m_{Z'}$ are shown for multiple collider experiments. Improvements in the center of mass energies of modern colliders allow to push the discovery potential to larger masses. However, no more improvement can be gained in the region of $m_{jj} < 1000$ GeV even with increasing luminosities. As the trigger thresholds have to be set higher with increasing luminosities, the exclusion limits obtained at ATLAS and CMS are dominant only at higher masses than those of older experiments like CDF. After Run-1 of the LHC, the best limits for this region were therefore still those from the experiments at the SPS or the Tevatron collider.

In order to improve the sensitivity of dijet searches in the low mass regions, multiple strategies have been developed in Run-2 to circumvent the prescale problem. The used strategies can be divided into two classes: the first class focuses on recording only partially reconstructed event data, which in turn allows the recording of more events at the same data rate, the other approach requires additional particles in the final state to discriminate against QCD multijet background. An example of the first approach is the *Trigger Level Analysis (TLA)*, which relies on partially recorded events where only event information used and calculated by HLT jet algorithms is recorded [7]. An example of the second approach is the dijet+photon search, where boosted dijet systems recoiling against a high- E_T ISR photon are considered [8]. Figure 9.4 shows the status of the ATLAS limits in 2017 including these searches. Compared to Figure 9.3, the limits have been extended to lower masses. However, they are noticeably weaker in the region below 450 GeV. While being much more sensitive to small couplings at lower masses than the offline dijet search, both presented approaches have their disadvantages. As the TLA runs on the HLT, it is still statistically limited by the rate of level-1 accepted events of 100 kHz, limiting the sensitivity of the approach below $m_{jj} = 450$ GeV, where no limits are set by the TLA. The dijet+photon search can target even lower masses, but the required additional photon results in lower cross sections and reduces the signal efficiency, meaning that the limits achieved using this method are weaker than the ones from the TLA. The obtained limits do not cover Standard Model-like couplings of $\mathcal{O}(0.01)$ at the electroweak scale of a few hundred GeV. It is therefore important to improve the search strategies targeting this region. Statistical limitations could be circumvented if it was possible to develop a strategy where every single event can be analyzed before being dropped due to the




Figure 9.4.: Summary plot of ATLAS limits for the dark matter mediator mass and coupling using 2015 and 2016 Run-2 data, comparing the offline search to approaches targeting the lower mass regions [157].

limited level-1 trigger rate. This in turn would require an analysis of events online on the first trigger level before the trigger decision is formed. The only place in ATLAS where this is possible is the L1Topo module of the level-1 trigger, since L1Topo receives all reconstructed objects for every event and provides the capability to calculate the topological quantities required for an analysis. Using the jet TOBs identified by the JEP and the invariant mass algorithm implemented on L1Topo, it could be possible to perform such an analysis directly online on trigger level without any statistical limitations. In the following, this innovative new idea enhancing the approach taken by the TLA using L1Topo is presented and developed. In order to evaluate the feasibility of this approach, the statistical significance that could be achieved using such an approach has been studied in a previous thesis [1]. By estimating the invariant mass spectrum of L1 jets from a fit to recorded data with a luminosity of 3.4 fb^{-1} , it was found that the exclusion limit on the cross section of a 500 GeV signal is comparable to the limit obtained by the TLA for the same invariant mass. This means that a comparable statistical sensitivity in the mass region covered by the TLA can be achieved. Compared to the TLA, which stops at masses of 450 GeV, the approach on L1Topo has the additional advantage of not having this lower mass limitation. Due to these promising results, it was decided to develop a concept to implement this analysis in the firmware of L1Topo, which is presented in the following.

9.2. The Histogramming Concept

Since L1Topo is part of the L1 trigger, it processes all events with a frequency of 40 MHz. As described in Chapter 8, L1Topo is provided with all objects reconstructed in the L1 trigger for each event, including all reconstructed jet TOBs from the JEP system. As the jets from the hard scattering process are on average expected to have higher transverse momenta than background jets in pile-up events, the two jets used for the mass calculation can be chosen by determining the two most energetic jets in the event. Using the jet sort algorithm described in Section 8.3.2, it is possible to identify the leading and subleading jet in every event. The invariant mass algorithm

presented in Section 8.3.3 can then be used to calculate the dijet invariant mass directly on L1Topo. In this case, a square root operation is performed to directly calculate m_{jj} from m_{jj}^2 . Ideally, the invariant mass could be read out event-by-event, for example as part of the trigger decision bits. However, even writing out a single 32 bit quantity would lead to data rates in the Gbit s^{-1} regime of about $32 \text{ bit} \cdot 40 \text{ MHz} = 1.28 \text{ Gbit s}^{-1}$, meaning that additional optical fibers would be required to transfer such an enormous amount of data. When considering this, it is important to remember that L1Topo provides crucial functionality to the ATLAS trigger system. As modifications on critical parts of the trigger system infrastructure should be avoided if possible, the installation of additional fibers to the hardware was not foreseen. The usage of existing fibers is not possible, as this would block large parts of the bandwidth for other required operations. A further reduction of the amount of readout data is therefore required for this approach. Since the size of the readout data cannot be further reduced, the only possibility is to reduce the frequency with which the data is read out. This can be achieved by creating *histograms* of the calculated invariant masses directly on L1Topo. Instead of reading out event level data, a histogram of the dijet invariant mass spectrum is recorded directly on L1Topo and read out in much larger time intervals of $\mathcal{O}(\text{s})$. For an exemplary histogram with about 200 bins, each containing 32 bit of event counts, the resulting data rate is significantly lower at about 6.4 kbit/s. Instead of requiring additional optical fibers, a readout via software using the IPBus connections of L1Topo is therefore possible. This way, an entirely triggerless analysis can be implemented directly on hardware, recording data from dijet events utilizing the full data rate of 40 MHz. Due to the utilization of histograms, this approach will be referred to as the L1Topo *histogramming*. Figure 9.5 shows a schematic view of the general histogramming concept. The histogramming block receives the identical inputs as the RTDP algorithms with a frequency of 40 MHz. In order to avoid any interference with the trigger algorithms of the RTDP, the histogramming runs in an independent block of the firmware. In a first step, the jets are sorted to identify leading and subleading jets. For both jets, the larger energy window of 4×4 jet elements is chosen, as this window corresponds better to a reconstructed offline jet with a radius of $R = 0.4$ than the smaller window size. The $R = 0.4$ anti- k_T jets are used for most of the dijet resonance searches at ATLAS, the 4×4 window of the SLW algorithm therefore allows for better comparisons of the results. In cases where there are only one or zero jets in the event, the subleading or the leading jet are equal to zero. After the selection of the jet pair,

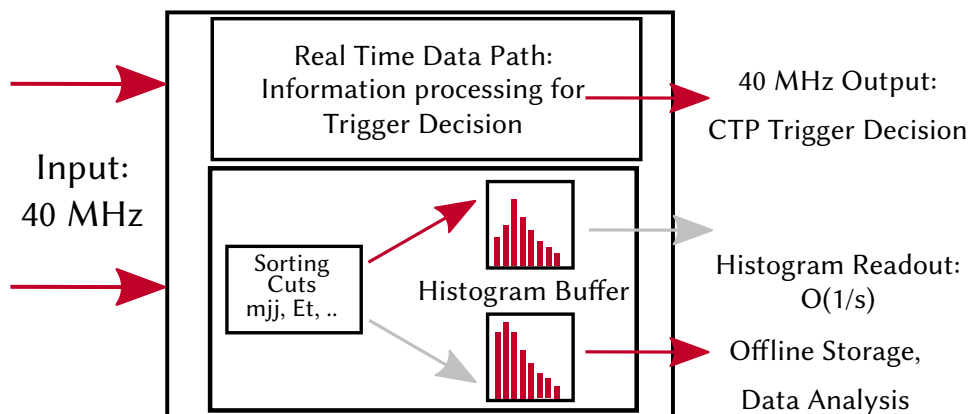


Figure 9.5.: Overview of the L1Topo histogramming scheme running parallel to the RTDP. Invariant masses are filled into one of two histogram buffers, the other buffer is read out for offline storage in the meantime.

various kinematic selections can be applied to define multiple kinematic regions. In addition, it is possible to veto events at this stage based on several criteria, for example overflow information or certain BCIDs. Then the histogrammed quantity, for example the invariant mass m_{jj} or the jet E_T , is calculated and filled into a histogram buffer. After a certain time has elapsed, the filled histogram buffer is read out. To ensure that the contents of the histogram are not modified or overwritten while it is read out, a second histogram buffer is implemented which is filled in the meantime. The read out histograms are then permanently stored offline, and can later be accessed in a data analysis. Since the information available for this analysis is limited to the information stored in histograms, the choice of implemented histograms is decisive for the analysis and is discussed in the following.

9.2.1. Kinematic Selection and Histograms

Since the invariant masses are calculated and histogrammed online on L1Topo, no event level information is available after the event has been processed. As a consequence, no event selections can be performed after the histograms are recorded. Therefore, all kinematic selections have to be made when filling the histograms and have to be fixed before the firmware implementation. The simplest case possible is an inclusive histogram of the invariant mass for all events without any kinematic selections. In the following, this histogram will be referred to as the *total* invariant mass histogram. For the kinematic selections, three variables are used which play important roles for dijet resonance searches, namely y^* as introduced in Equation (3.7) and the pseudorapidities of the jets η_1 and η_2 . As discussed in Section 3.2.2, the rapidity difference y^* between the final state jets in the laboratory system can be used to discriminate between possible signal events and QCD multijet background. In order to allow for flexible combinations of the different regions and avoid double counting, the regions are defined mutually exclusive. Three y^* regions are introduced for the analysis: a region with $y^* < 0.3$, which is mainly chosen for an easy comparison with the signal region of the TLA analysis, a region with $0.3 \leq y^* < 0.6$ offering the best background rejection combined with the first region and a region with $y^* \geq 0.6$ to cover the remaining phase space. In addition, the histograms are also partitioned into jet η slices, defining three different detector regions. Since the granularity of the jet elements used for the L1 jet algorithm degrades for higher values of $|\eta|$, see for example Figure 6.10, it can be useful to restrict the analysis to more central detector regions where the best granularity is available. The *central* region was chosen as $0 \leq |\eta| \leq 2.0$, since this guarantees that all jets are fully contained in $|\eta| < 2.4$ and avoids jet element granularity transitions. A second region, the *extended* region, is chosen as $2.0 < |\eta| \leq 3.0$, and the *forward* region is defined as $|\eta| > 3.0$. Similar to the y^* categories, these are chosen mutually exclusive to allow for flexible combinations of the kinematic regions. Since it is possible that the two jets fall in different $|\eta|$ categories, a total of six jet η categories exist. When combining y^* and η categories, a total of 18 invariant mass categories are defined, which are summarized in Table 9.1. After the kinematic selections for the inputs of every histogram have been fixed, it is necessary to decide on the histogram binning by defining both histogram ranges and bin widths. In general, the bin widths will be non-uniform depending on m_{jj} , larger bin widths can be chosen for higher masses for two reasons. Since the dijet cross section drops over many orders of magnitude as a function of m_{jj} , compare Figure 3.7, fewer events are expected at higher masses. In addition, the absolute mass resolution $\sigma_{m_{jj}}$ is expected to grow with m_{jj} as a calorimeter based quantity. An upper boundary for the invariant mass is given by available center-of-mass energy $\sqrt{\hat{s}}$, limiting the maximum possible value. The bin widths at low masses are kept small on purpose to keep as much information as possible. For a later statistical analysis, a coarser binning which

Table 9.1.: All 18 available categories of implemented histograms are listed. For each category of jet pseudorapidities, three different y^* categories are available.

η_1	central	central	extended	extended	forward	central
η_2	central	extended	extended	forward	forward	forward
y^*	$y^* < 0.3$	$y^* < 0.3$	$y^* < 0.3$	$y^* < 0.3$	$y^* < 0.3$	$y^* < 0.3$
	$0.3 \leq y^* < 0.6$	$0.3 \leq y^* < 0.6$	$0.3 \leq y^* < 0.6$	$0.3 \leq y^* < 0.6$	$0.3 \leq y^* < 0.6$	$0.3 \leq y^* < 0.6$
	$0.6 \leq y^*$	$0.6 \leq y^*$	$0.6 \leq y^*$	$0.6 \leq y^*$	$0.6 \leq y^*$	$0.6 \leq y^*$

Table 9.2.: Invariant Mass histogram overview, showing the bin widths, the ranges, the number of bins n_{bins} and the resulting amount of data for the different regions.

m_{jj} region	bin width [GeV]	range [GeV]	n_{bins}	bits per bin	total bits
1	2	0–200	100	32	3200 bit
2	8	200–600	50	32	1600 bit
3	16	600–1000	25	32	800 bit
4	32	1000–2984	62	32	1984 bit
5	64	2984–4968	31	32	992 bit
6	128	4968–8040	24	32	768 bit
Σ					9344 bit

is more close to the resolution of the invariant mass can be chosen. Powers of two have been chosen for all bin widths, as this will make the firmware implementation of the binning easier. A summary of the chosen invariant mass histogram bin ranges and bin widths is shown in Table 9.2.

Until now, only histograms of the invariant mass have been discussed. However, the histogramming concept can easily be extended to other quantities, for example for monitoring the quality of the data used to calculate the invariant mass. Obvious quantities of interest are the kinematic variables that enter the invariant mass calculation. Two histograms of the leading and subleading jet E_T are implemented as well as a two-dimensional histogram showing the $\eta - \phi$ distribution for both jets forming the invariant mass. As a two-dimensional histogram requires more bins and therefore bandwidth compared to a one-dimensional one, it was decided to have a single $\eta - \phi$ histogram for both leading and subleading jet. In addition, the angular distance ΔR between the jets is recorded. Finally, the value of E_T^{miss} is recorded, as high values of E_T^{miss} can give indications on hot calorimeter towers or other detector problems. Similar to the invariant mass histograms, the bin widths and ranges of the histograms have to be defined beforehand. In this case, the minimum and maximum values can be determined from the properties of L1 jets. Since 10 bit are used to represent the 8×8 jet energy, a range from 0 GeV up to 1023 GeV is sufficient to cover all possible values. For ϕ , integer values between 0 and 64 are possible, corresponding to values between 0 and 2π . For η , the possible values range from -39 to 39 , corresponding to coordinates from -3.9 to 3.9 . Table 9.3 shows an overview of the bin widths and bin ranges which were chosen for the additional histograms. In summary, all histogram data sums up to 7047 bins, corresponding to a transmission of 225.504 kB of data. However, since this data has to be transmitted only in

Table 9.3.: Control and data quality histograms, along with their bin widths, ranges and number of bins n_{bins} . The total number of bits required for each histogram is shown. An additional overflow bin is implemented for the $\eta - \phi$ histogram.

histogram	region	bin width	range	n_{bins}	bits per bin	bits total
$E_{\text{T}}^{\text{miss}}$	1	2 GeV	0 GeV to 200 GeV	100	32	3200
$E_{\text{T}}^{\text{miss}}$	2	32 GeV	200 GeV to 808 GeV	19	32	608
leading E_{T}	1	2 GeV	0 GeV to 200 GeV	100	32	3200
leading E_{T}	2	32 GeV	200 GeV to 1022 GeV	27	32	864
subleading E_{T}	1	32 GeV	0 GeV to 200 GeV	100	32	3200
subleading E_{T}	2	32 GeV	200 GeV to 1022 GeV	27	32	864
ΔR	1	0.1	0–10.1	101	32	3232
$\eta - \phi$	1	various	$[-3.9, 3.9] \times [0; 6.3]$	1025	32	32800

intervals of seconds, the resulting data rate is below a MB/s, which is well below the possible IPBus throughput as shown in Figure 4.6, and will therefore not take away bandwidth required for other requests.

9.3. Firmware Implementation

After the definitions of the histograms are fixed, the firmware implementation can be started. This section gives a brief overview of the firmware implementation of the histogramming feature. The discussion in this section is limited to a discussion of the most important features, a more in-depth description of the technical details can be found in Appendix B.1. Figure 9.6 shows a modification of the L1Topo firmware block from Figure 8.3, where an additional block for the histogramming has been introduced outside the RTDP. As the histogramming runs on the same inputs as the algorithm firmware, it is symbolically placed parallel to the algorithm stage. It starts with finding the pair of jets from the JetSort instance. These jets are then used to calculate all desired topological quantities, for example the mass or the angular distance ΔR using the InvariantMassCalc and DeltaRSqrCalc entities introduced in Section 8.3.3. The histogrammed quantities are then fed into the BinMapping entity, which calculates the correct histogram bin number given the information on bin ranges and bin widths described in Section 9.2.1, details on the functionality of the BinMapping entity can be found in Appendix B.1.1. In parallel, the HistCategories entity calculates the value of y^* and assigns the correct invariant mass histogram category to the given pair of jets. Its implementation is described in Appendix B.1.2. Finally, all information on the bin numbers, the assigned category and the topological quantities is fed into the histograms. Due to the particular importance of the histograms, their implementation using the BRAM resources of the FPGA are described in detail in Section 9.3.1. Instead of being connected to an MGT serializer for real-time data output, the expected low-rate histogram data readout occurring only every few seconds allows to use the IPBus protocol for this purpose, similar to the readout of the spy memories of the RTDP. Reading out the histogrammed quantities via IPBus is therefore equivalent to spying on algorithm data. Since the output is not required to be synchronous to the output of the real time algorithms, the latency constraints of the histogramming firmware are more relaxed. 5 BC are required in total starting from the unsorted list of jet TOBs to an invariant mass being filled into a histogram. As explained previously, the

histograms are implemented with two histogram buffers. The EventInfo entity, which runs in parallel to the other blocks of the histogramming, is responsible for determining the currently active of the two histogram buffer. It does so by counting the number of 40 MHz bunch ticks and swapping the buffer connections when the number of ticks reaches a value corresponding to the desired recording time, the counter is reset afterwards. The currently active buffer is then used for writing, while the other buffer is connected to the IPBus readout. When the buffers are swapped, the HistControl IPBus register is set, which indicates to the software that a new histogram is available. In addition, the EventInfo configures a histogram counter which gives a unique number to every read out histogram, and checks whether a BCID veto is applied which is meant to reject data from certain BCIDs. More technical information on the EventInfo entity can be found in Appendix B.1.3.

9.3.1. Histograms

For the histograms themselves, data corresponding to the total number of required bins has to be stored, where each bin contains a 32 bit word. The BRAM resources of Xilinx FPGAs provide the desired functionality, when a memory depth corresponding to the number of bins and a memory width of 32 bit is chosen. All histogram BRAM are implemented using the Xilinx Block Memory Generator IP Core [83], the exact settings and port configurations are listed in Appendix B.1.4. Different configurations are used for the one-dimensional and two-dimensional case, both are introduced in the following. The one-dimensional histograms are configured in the simple dual port mode, where port A is used as a write-only port and port B as a read-only port. Every histogram bin is then represented by a dedicated memory address in the BRAM, where the data word saved at the address stores the event counts. Two of these BRAM instances are wrapped inside

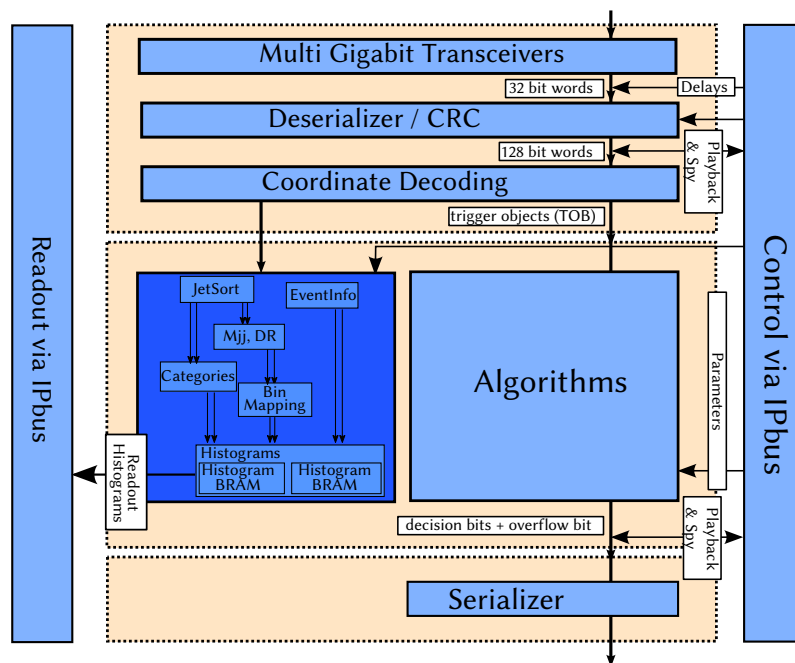


Figure 9.6.: Schematic view of the L1Topo RTDP firmware with an additional block for the histogramming feature.

the HistogramBRAM entity and implement the two histogram buffers as shown in Figure 9.5. Each HistogramBRAM entity represents an IPBus slave, the existence of two histograms is therefore transparent to the IPBus. In order to provide a sufficient number of histogram bins, the bit width of the BRAM memory addresses is defined as

$$w_{\text{addr}} = \lceil \log_2 n_{\text{bins}} \rceil, \quad (9.1)$$


where the ceiling operation $\lceil \cdot \rceil$ rounds its argument to the closest higher integer. In order to synchronously store and readout the histograms, the BRAMs run at four times the 40.08 MHz clock. The timing structure of the subticks is given as follows: in the first subtick of the 160.32 MHz clock, the index of the active histogram provided by the EventInfo is checked to keep or change the currently active buffer. The second subtick sets the enable of the read port of the active buffer to “1” to initialize a read operation at the bin number calculated by the BinMapping entity. In the third subtick, the value read in the second tick is incremented by one. In the fourth and last subtick, the writeEnable of the write port is set to one, and the previously incremented value is written back to the same address the value was initially read from. During this, the inactive buffer is connected to the IPBus output signal. If a change in the IPBus strobe signal indicates an incoming IPBus request, the contents of the inactive buffer are read out. After the read operation has finished, all BRAM contents of the inactive buffer are reset to 0 to avoid any double counting. The two-dimensional histogram utilizes BRAM in the true dual-port mode, where two independent read and write ports are offered. This is required as the coordinates of both leading and subleading jet are filled into the same histogram. Additional logic is required to check whether both jets require writing at the same address, as simultaneous, potentially conflicting write operations have to be avoided. In this case, only one write operation is performed, increasing the memory contents by two instead of one.

9.3.2. Resource Utilization

After the histogramming feature has been designed and all important aspects have been discussed, its firmware can be implemented. The majority of the implementation work was done in all of 2017, as many iterations and bug fixes were necessary. In a standalone implementation the resource utilization is studied to check whether it is feasible to run this firmware alongside the already existing trigger firmware of L1Topo. The resulting utilization is listed in Table 9.4. Most of the logic of the histogramming relies on LUTs and flip-flops. The amount of 27 BRAMs is caused by the required 23 one-dimensional histograms, which each are implemented using a single BRAM block each, and the single two-dimensional histogram, which requires the cascading of four BRAM blocks. As explained in Section 4.2.3, each 36 kbit BRAM consists of two individual 18 kbit BRAMs. Since every histogram buffer can be implemented using a single 18 kbit BRAM, both histogram buffers of the HistogramBRAM are implemented in a single 36 kbit BRAM using the ports independently. The two DSP are required by the calculation of a single invariant mass per event, as described in Section 8.3.3. It is sufficient to implement the histogramming on only one of the four processors, allowing to choose the processor with the lowest logic resource utilization by the rest of the firmware. As presented in Section 8.4, this is given by Topo0 U1. The utilization of this FPGA including the histogramming firmware is shown in Figure 9.7, indicating what parts are required by the histogramming firmware. Compared to the trigger firmware, the histogramming feature requires only small additional amounts of logic. Therefore, the integration of the histogramming firmware into the trigger firmware is possible and has been performed successfully.

Table 9.4.: Absolute and relative numbers of required resources of the histogramming firmware. BRAMs are required for the histograms, DSPs for the invariant mass calculation.

Resource	Amount	Utilization [%]
LUTs	9623	2.22
flip-flops	4294	0.50
BRAMs	27	1.84
DSPs	2	0.06

**Figure 9.7.:** Fraction of resources required by the histogramming on Topo0 U1 (light blue) compared to the total utilization (dark blue). Only a small fraction of the resources is allocated to the histogramming.

9.4. Software Implementation

Besides having an integrated firmware version, an accommodating software implementation is required in order to control the histogramming. Its most important tasks are to initialize IPBus control registers on startup, to perform the read operations and to save the readout data in the appropriate format. Two different implementations of software were used, namely a standalone self-written software package, and an adapted implementation into the existing L1Calo online software. The L1Calo software is responsible for module initialization, control and monitoring of all modules of L1Calo including L1Topo during ATLAS data taking. Since the L1Calo software has many dependencies and runs in a critical production environment, the much simpler standalone version was used for all initial tests and as a proof of concept. In addition, the online software is significantly more complex, meaning that an implementation in this framework requires help from the expert developers. The implementation of the histogramming into the L1Calo software is explained in the following, information on the standalone version is given in Appendix B.2.

9.4.1. Integration to the L1Calo Online Software

The L1Calo online software contains logic to correctly initialize and start-up all L1Calo modules as well as logic for the correct shutdown after a run has finished. In addition, access to debugging information or error messages as well as the access to monitoring histograms is provided during a run. Similar to the standalone software, it relies on the uHAL IPBus software library to communicate with the hardware modules. Using the L1Calo online software has several advantages. As it runs in the TDAQ infrastructure, it provides direct access to the IS to access information on the run status. In addition, access to the OHS provides a convenient possibility to publish histograms and to archive them via the MDA at the end of a run. The logic to publish histogrammed data is already available for monitoring purposes, for example for CMX jet rates or to read the rates for the topological trigger items from L1Topo. Therefore, this already existing logic can be adapted for the histogramming software without large code changes.

In a first step, the L1Calo online software needs to be informed about the names and bin ranges of all histograms, this is done via a json configuration file. In order to address the correct IPBus slaves in the firmware, knowledge on the addresses and properties of IPBus registers are provided in a dedicated address XML file. During initialization, the histogram definitions are read from the configuration file, the requested histograms are created and the IS infrastructure is notified about the existence of these histograms. When a run is ongoing, the software frequently checks the contents of the HistControl register, which indicates whether a new histogram is available for readout. This happens in an interval of 2 s. If a new histogram is available, the read operation is started and all histogram data is read out along with the IPBus error bit and the histogram numbers. In order to allow an association of histograms with the corresponding recording conditions, the software accesses and saves the current recording timestamp. The histogram data is filled into a ROOT histogram [158] and the information on its name, error bit, histogram number and timestamp is saved in the histogram title. When a run is ongoing, the TDAQ IS infrastructure frequently sends a request to publish all statistics, which retrieves the updated histograms and monitoring information, allowing to inspect them already during a run from the online monitoring. In the end, all recorded histograms from a run are archived via the MDA. From there, the histograms can either be downloaded using a web interface or directly accessed on the corresponding file system.

9.4.2. Readout Interval

The initial idea for the histogramming readout interval was to synchronize the start and end time of the readout interval to the start and end times of a lumiblock, as data quality information as well as information on the instantaneous and integrated luminosity and average pile-up are available on a lumiblock level in ATLAS. However, the start and end time information is not directly available on L1Topo, as lumiblock information is formed in and distributed by the CTP, which comes hierarchically after L1Topo in the level-1 trigger. While some basic timing and event information like the BCID signal are transmitted via the TTC interface, no information on the lumiblock is included. Instead, the lumiblock synchronization information is provided to TDAQ software applications via the IS. The IS entry on the lumiblock includes the run number, the lumiblock number and a time stamp, which describes the beginning of the lumiblock with a precision of one BC. Since the IS information is only provided to software applications, this would therefore require to write this information to L1Topo via IPBus. As a consequence, L1Calo would have to subscribe to the CTP IS server to be notified about changes in these variables, which

would also affect other systems of L1Calo besides L1Topo and therefore require interventions by multiple detector experts.

Due to this, an alternative approach is chosen where the histograms are read out in fixed time intervals. The firmware is written in a way that the length of the readout interval is defined in a writable register in the EventInfo entity and can therefore be configured from the software. In the beginning, shorter time intervals of about 1 s were considered. Shorter time intervals are advantageous from a data quality check perspective: If a time interval overlaps with a bad lumiblock with recording conditions not fit for physics, all data collected in this time interval needs to be skipped in an offline analysis. For shorter intervals, this would result in a smaller fraction of data that has to be skipped in case of overlaps. However, shorter intervals also result in a high amount of IPBus transactions, causing additional stress for the system. Therefore, larger recording intervals were chosen to ensure a smooth running of the IPBus transactions. Arbitrarily high time intervals are also not possible, as integer overflows of the counts in histogram bins can occur after a certain time. In the two-dimensional histogram, two entries are filled per BC. Considering a time interval of length t , a total number of $N = t \cdot 2 \cdot 40 \text{ MHz}$ entries are filled into a histogram, leading to a maximum time of

$$t_{\max} = \frac{N}{2 \cdot 40 \text{ MHz}} = \frac{2^{32} - 1}{2 \cdot 40 \text{ MHz}} \approx 53.7 \text{ s}, \quad (9.2)$$

in which no overflows of histogram bins can occur, even if every event is filled into the same bin. Due to this, a recording interval of 50 s was chosen. A closer inspection of the distribution of good and bad quality lumiblocks in an LHC run is shown in Figure 9.8. The figure shows that most of the time many bad lumiblocks or many good lumiblocks are in a row, alternations between those two only rarely occur. Due to this, the additional loss of data due to bad lumiblocks falling into the recording window can be neglected, as this will only concern a small amount of histograms. When determining the integrated luminosity corresponding to such a histogram, it has to be considered that the histogram readout interval will not be identical to a single lumiblock. A method to determine the integrated luminosity corresponding to one or more histograms is presented in Section 10.1.3.




Figure 9.8.: Arrangement of good lumiblocks shown in green and bad lumiblocks shown in red for the LHC run 357962 taken on August 8 and 9 in 2018.

9.5. Histogramming Tests

As a complex, new feature with many components, the histogramming feature needs extensive testing to be ready for deployment. Multiple stages of tests have been performed, starting with

low level tests, for example basic logic simulations using the behavioral simulation as introduced in Section 4.5.4 and module level tests utilizing the test rig of L1Calo as introduced in Section 8.5 starting with a standalone L1Topo module, then successively integrating more and more external components and complexity. Finally, tests can be performed with parts of or the ATLAS experiment itself. A summary of all performed test would go far beyond the scope of this thesis. This section therefore briefly summarizes the standalone tests and focuses on the tests performed with the full ATLAS experiment.

9.5.1. Standalone Tests

The most basic level of tests to check the firmware functionality have been performed using waveform simulations as introduced in Section 4.5.4, starting with unit tests of the functionality of the individual blocks. Some components, for example the jet sorting or the invariant mass calculation, have already been simulated and validated extensively due to their application in trigger algorithms. For this purpose, externally generated test vectors were loaded into the waveform simulation, and the results of the individual events were then checked for correctness. This allowed the identification of bugs in the individual components. For example, a bug in the BinMapping which lead to wrongly assigned bin numbers for the invariant mass was identified this way. Based on these, more complex tests were performed by combining multiple components. To simulate incoming IPBus requests, artificial IPBus requests were issued by manually setting the strobe of the IPBus input signal to 1 after a certain amount of 40 MHz clock ticks have passed. Further examples and details on these tests are given in Appendix B.3.1.

9.5.2. Tests in the Test Rig

For more advanced tests with actual hardware, the test rig of L1Calo as introduced in Section 8.5 is utilized. The test rig allows to playback simple input patterns to a JEM, which in turn results in specific jets being sent to L1Topo. Various tests were performed, for example measurements confirming whether the latency of histogram read requests is compatible with the information given by the IPBus developers in Figure 4.6. Details on these tests are given in Appendix B.3.2. In addition, the test rig allows access to ATLAS TDAQ infrastructure like the IS or the OHS. Therefore, features of the software like the publishing of the histograms to the OHS service can already be tested at this stage. Figure 9.9 shows the resulting histograms from such a test using the OHS browser interface to inspect the most recent histogram created in a run. Several consistency checks can already be made by inspecting these results, for example that the presence of only two entries in the η - ϕ corresponds to the presence of only one entry in the ΔR histogram. In addition, it was confirmed that all required information is correctly included in the histogram title. The main disadvantage of the test rig is the limited phase space coverage of the inputs, as only a limited amount of simple patterns for the JEMs exist. In addition, the number of JEM modules is smaller than in ATLAS, therefore it is not possible to cover the whole η - ϕ space. For more sophisticated inputs, tests in ATLAS are required.

9.5.3. Tests in ATLAS

The final and most important test stage are tests with the modules of the ATLAS detector, including the full L1 trigger system. When no beams are circulating in the LHC and therefore no run is ongoing, subdetector groups are allowed to perform system level tests. Since this test time is limited as it has to be shared between subdetector groups, it is important to gain as much as

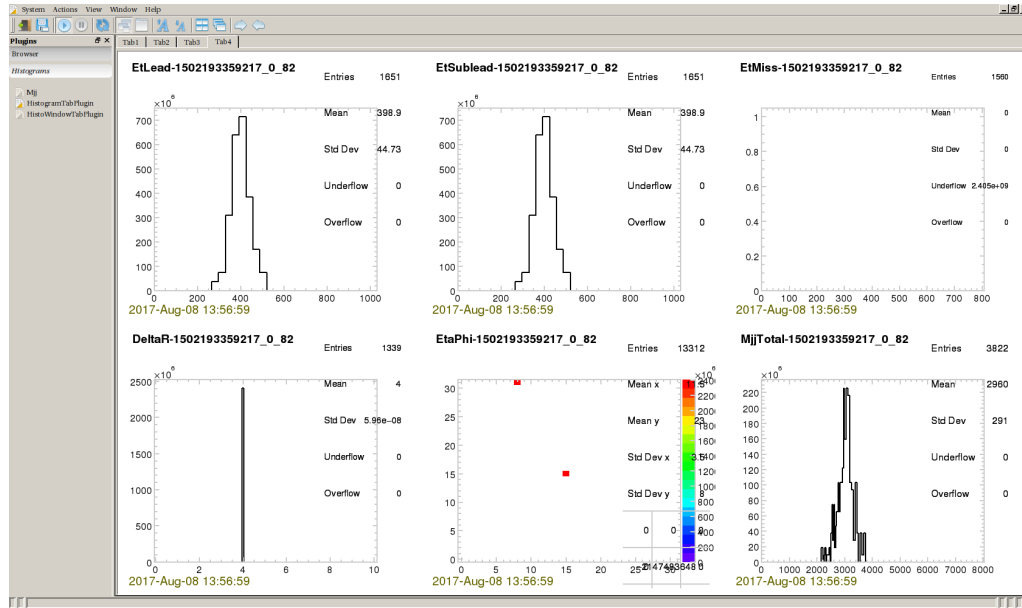


Figure 9.9.: Exemplary overview of histograms in the OHS browser, which allows to directly inspect the most recently published histograms for monitoring. In this case, a simple playback pattern with constant η and ϕ values and varying jet energies was chosen. Note that the displayed $\eta - \phi$ histogram shows the jet bin numbers, which were used in an earlier firmware version.

possible from the individual tests. In the beginning, the main purpose of these tests is to check whether the integration of the histogramming in the existing ATLAS infrastructure is working correctly, namely the publishing of histograms to the OHS and the archiving of the histograms to the MDA. In a preliminary version of the software, it was realized that the histograms were overwritten every time they were published. An additional patch fixed this issues by correctly assigning a unique tag to every publication to the OHS server. Afterwards, it was checked that these histograms appear in the OHS browser. Finally, it was tested whether these histograms were correctly archived to the MDA, and that all tags of a run were archived correctly. Histograms were received successfully both from disk and from the MDA web interface. These histograms were checked for the metadata encoded in the histogram title. The readout timestamp, the IPBus error code and the internal histogram number were found to be correctly encoded. After it was ensured that all infrastructure was set up correctly, the focus of the tests was shifted to investigations of the histogram contents. In a first step, so-called hot tower runs were performed. These special runs were mainly developed for L1Calo and L1Topo algorithm commissioning. By manually changing the noise cuts of a few PPMs, artificial trigger tower signals are created due to electronics noise, which allows feeding inputs to the level-1 trigger algorithms without any external sources. For the purpose of debugging the two-dimensional $\eta - \phi$ map, a special hot tower pattern containing only a single tower was used. The resulting histogram was compared to the one recorded by the official L1Calo monitoring tools of the CMX outputs. A disagreement was observed, which was traced down to an index flip in the bin assignment for the two-dimensional histogram and fixed with a firmware patch. An additional issue was observed when inspecting the ΔR histogram. While only a single jet is present in the event, an entry at a ΔR value of ~ 4.3 is present in the ΔR histogram. This particular value of ΔR was identified to be equal to the ΔR between the coordinates in the $\eta - \phi$ histogram and the coordinates of a *zeroed* jet. A jet containing only zeros




Figure 9.10.: Histogramming test results from a hot tower run with a single jet. The monitoring histogram from the CMX is shown in (a), the histogram from L1Topo in (b). Both show a jet at the identical coordinates.

is mapped to $(\eta, \phi) = (-3.9, \pi/16)$ using the jet TOB η LUT shown in Table A.2. The ΔR value between a zeroed jet and the coordinates from the single visible entry in the $\eta - \phi$ map is then given as

$$\Delta R = \sqrt{\Delta\phi^2 + \Delta\eta^2} = \sqrt{\left(\frac{9\pi}{16} - \frac{\pi}{16}\right)^2 + (0.2 - (-3.9))^2} \approx 4.4, \quad (9.3)$$

corresponding to what is observed in the ΔR histogram. To resolve this issue, a fixed firmware was deployed to veto cases with a zeroed jet for the ΔR and $\eta - \phi$ histograms. Figure 9.10 shows the resulting $\eta - \phi$ histogram after the patches are deployed compared to the CMX monitoring histogram.

Even though the hot tower runs provide an excellent opportunity to validate many aspects of the histogramming, the phase space covered by hot tower events is still limited. The next step is therefore to request test time during ATLAS non-physics runs, for example when the LHC is ramping up again after a technical stop with maintenance work and upgrades on hardware, software and firmware. During ATLAS non-physics runs in September 2017 with reduced number of bunches after a technical stop, namely the runs 336506, 366505 and 336567, the histogramming functionality was activated and tested. In these runs, issues could be discovered which were not visible with discrete inputs as test vectors or hot towers. As an example, Figure 9.11 shows an exemplary E_T^{miss} distribution from run 336506. The E_T^{miss} spectrum observed in this run shows a gap for values from 128 GeV to 200 GeV, which was identified as a bug in the bin mapping and fixed in a new firmware version. After these tests it was concluded that the fixed histogramming firmware is functional and could be deployed for LHC runs in 2018. However, there is one important thing to consider for any further updates. As soon as the histogramming is integrated into a production firmware running in ATLAS, there are only a few possibilities to deploy firmware updates a year, commonly during so-called technical stops with a length of about a week. Even in such a stop, a histogramming fix can only be deployed if the new firmware version also contains




Figure 9.11.: E_T^{miss} spectrum from run 336506. A gap in the spectrum can be observed.

fixes for the trigger firmware. The reason for this is that every new firmware version has to undergo a validation procedure before its triggers can be enabled, significantly limiting the possibilities to deploy any updates. For potential bugs or issues discovered during data taking, it is therefore uncertain whether an updated firmware version containing fixes can be deployed in time.

Towards a Histogramming Data Analysis

In this section, the preparation of an analysis of L1Topo histograms is presented. The most important concept revolves around data quality checks, which are used to select the subset of all data recorded with good and reliable recording conditions. In contrast to offline analyses, where pre-written tools exist to apply data quality criteria, no such tools exist for the analysis of histograms, a new framework to perform these checks for histogram-based data is therefore introduced in Section 10.1. Afterwards, Section 10.2 shows an overview of histogram recorded in 2018 for various LHC demonstrating the functionality of the data quality selection framework. These histograms are used to perform further detailed checks on the features of the histogramming data building on the tests performed in Section 9.5. Only when these are concluded successfully, the final spectrum for the statistical analysis can be derived. As a consequence of these studies, a subtraction procedure is derived in Section 10.3, which handles background events from cosmic rays and calibration effects.

10.1. Data Quality Checks

Various calibrations and data quality checks are performed after the data is recorded during offline event reconstruction or offline analysis, their impact is therefore unknown at the time of the real-time trigger level histogramming. An application of data quality checks is therefore only possible at the level of recorded histograms. As these checks are commonly performed on a lumiblock level in an offline analysis, a method to associate histograms with lumiblocks is required and derived in the following.

10.1.1. Good Run Lists in ATLAS

In ATLAS, the various detector systems will mark the fraction of data which they consider suitable for physics analyses. In case of a defect, for example due to high voltage trips of a subdetector or faulty data transactions in a trigger module, a certain defect is assigned to the corresponding lumiblocks. Data quality status flags are determined from the combined information on defects from all subsystems to provide information on good runs and lumiblocks with recording conditions suitable for a physics analysis. This information is decisive for a data analysis, and is provided in so-called *Good Run Lists (GRLs)* [159]. A GRL contains a nested list of run and lumiblock numbers defined as good for physics and can be queried in an analysis. For an offline analysis of recorded collision data, a lookup in the GRL is straightforward, as the lumiblock numbers are part of the event readout. However, no such lookup can be made for a histogram from L1Topo, making the application of lumiblock level criteria challenging. A different mechanism is therefore

required: the central idea is to associate each histogram with all lumiblocks which overlap with the corresponding histogram recording interval. This requires information on the recording interval, which can be inferred from the histogram timestamp, and on the start and end time of every lumiblock in ATLAS. The information on lumiblock start and end times and other lumiblock-level metadata can be obtained from a dedicated database on ATLAS metadata, which is introduced in the next section.

10.1.2. The Coma Database

In order to quickly access run-level metadata on ATLAS data, the Conditions Metadata for ATLAS (COMA) system has been developed [160]. It is a relational database scheme storing metadata information down to the lumiblock level for most runs of ATLAS.¹ For example, the COMA database contains information on the start and end times of all lumiblocks in a certain run along with the measured instantaneous luminosity. Programmatic access to COMA is possible by using the pyAMI package developed for the access of ATLAS metadata using python [161]. Three tables from the COMA database are used and queried in this analysis. The first and central table of COMA lists all runs in ATLAS identified by a unique run number, their start time, end time and the corresponding run type. An exemplary selection of entries from this table is shown in Table 10.1. For each run number, information on all corresponding lumiblocks can be obtained from two additional tables. The first table contains information on the duration, the start time and the end time of each lumiblock, while the second table contains information on the average instantaneous luminosity \mathcal{L} , the integrated luminosity L and the mean number of interactions $\langle\mu\rangle$. Examples of these tables are shown in Tables 10.2 and 10.3. In order to uniquely identify the lumiblocks, each lumiblock is identified by a lumiblock number. As a lumiblock number is unique only for a given run, the run number is required in addition to uniquely identify each lumiblock from the table with multiple runs.

Table 10.1.: Excerpt from the COMA table containing information on all runs. In the run type field “data18” denotes data recorded in 2018, “900GeV”, “13TeV” the center-of-mass energy in proton-proton collision runs, “cos” denotes runs with cosmic muons and “hi” denotes heavy ion runs.

run number	start time	end time	run type
350013	2018-05-10 18:25:59.0	2018-05-11 10:24:19.0	data18_13TeV
355446	2018-07-12 13:50:54.0	2018-07-12 19:03:49.0	data18_13TeV
358831	2018-08-21 17:04:33.0	2018-08-22 08:01:44.0	data18_cos
360471	2018-09-10 20:04:23.0	2018-09-11 06:59:44.0	data18_13TeV
363437	2018-10-11 05:52:59.0	2018-10-11 05:59:12.0	data18_900GeV
365709	2018-11-12 02:51:08.0	2018-11-12 13:48:39.0	data18_hi

¹Runs corresponding to system specific tests or calibrations runs, for example the hot tower runs with ATLAS as described in Section 9.5, are not listed in the COMA database.

Table 10.2.: Excerpt from the COMA table containing information on all lumiblock and their durations. The exemplary lumiblocks shown all have lengths close to 60 s.

run number	lumiblock number	start time	end time	duration
350013	250	2018-05-10 21:53:47.0	2018-05-10 21:54:47.0	60.12128512
350013	251	2018-05-10 21:54:47.0	2018-05-10 21:55:47.0	60.111664128
350013	252	2018-05-10 21:55:47.0	2018-05-10 21:56:47.0	60.119620096
350013	253	2018-05-10 21:56:47.0	2018-05-10 21:57:47.0	60.117104128
350013	254	2018-05-10 21:57:47.0	2018-05-10 21:58:47.0	60.120696064
350013	255	2018-05-10 21:58:47.0	2018-05-10 21:59:47.0	60.162353152

Table 10.3.: Excerpt from the COMA table containing information on the level of pile-up and instantaneous as well as integrated luminosity for every lumiblock.

run number	lumiblock number	\mathcal{L} [$10^{30} \text{ cm}^{-2} \text{ s}^{-1}$]	$\langle\mu\rangle$	L [μb^{-1}]
350013	250	17388.1523438	48.6236190796	1045398.06477
350013	251	17356.5117188	48.5351371765	1043328.80287
350013	252	17323.875	48.4438781738	1041504.78359
350013	253	17288.28125	48.3443412781	1039321.4041
350013	254	17396.4589844	48.6468467712	1045887.22319
350013	255	17528.0	49.0146827698	1054525.72605

10.1.3. Luminosity Block Association and Luminosity Calculation

In the following, the information stored in the COMA database is associated with the recorded histograms. When a certain histogram timestamp t is provided, the COMA table of all runs is queried to search for the corresponding run number by checking that the timestamp is between the start time and end time information. Using the retrieved run number, all lumiblocks with the corresponding run number are retrieved and their start- and end-times are checked against the timestamp t . As it is not guaranteed that the recording interval coincides with a single lumiblock, it is important to consider the whole recording interval of 50 s to correctly determine all corresponding lumiblocks. Given a histogram recording time stamp t , the corresponding recording time interval (t_0, t_1) is determined as

$$\begin{aligned} t_0 &= t - 50 \text{ s} \\ t_1 &= t \end{aligned} \tag{10.1}$$

Using this, all lumiblocks fulfilling at least one of the three conditions

$$(t_0 \geq \text{start time}) \wedge (t_1 \leq \text{end time}) \tag{10.2a}$$

$$(t_0 \geq \text{start time}) \wedge (t_0 \leq \text{end time}) \tag{10.2b}$$

$$(t_1 \geq \text{start time}) \wedge (t_1 \leq \text{end time}) \tag{10.2c}$$

are selected. Equation (10.2a) selects lumiblocks which completely contain the recording interval, Equation (10.2b) selects additional lumiblocks which started inside the recording interval, but




Figure 10.1.: Schematic view of the overlap interval determination. The start and end time of two lumiblocks are shown in red, the recording intervals of subsequent histograms in blue. The histogram h_i has an overlap time $\Delta t_{h_i, n_{LB}}$ with the first drawn lumiblock and an overlap time of $\Delta t_{h_i, n_{LB}+1}$ with the second shown lumiblock.

possibly ended afterwards, and Equation (10.2c) selects additional lumiblocks which ended inside the recording interval, but possibly started earlier. From this, a list of associated lumiblock numbers is created for every histogram. This list is then used to retrieve information on good or bad lumiblocks from the good run list. Only if all associated lumiblocks are considered as good for physics the histogram is flagged as such.

In addition, information on the mean number of interactions per BC $\langle \mu \rangle$ and the average instantaneous luminosity \mathcal{L} is retrieved for every lumiblock. From these, corresponding values for $\langle \mu \rangle$ and \mathcal{L} can be determined for every histogram. In order to determine the luminosity corresponding to one or multiple histograms, the luminosities of all associated lumiblocks have to be considered. For each lumiblock identified by a number n_{LB} , the overlap of the lumiblock time interval with the recording interval of the histogram h_i is calculated and denoted as $\Delta t_{h_i, n_{LB}}$, as depicted in Figure 10.1. The integrated luminosity of the time overlap interval is then calculated by multiplying the average instantaneous luminosity of the lumiblock with the overlap time. For a histogram, the integrated luminosity is therefore calculated by summing the overlap contributions of all associated lumiblocks as follows

$$L_{h_i} = \sum_{k_{LB}} \Delta t_{h_i, k_{LB}} \mathcal{L}_{k_{LB}}. \quad (10.3)$$

This corresponds to the discrete version of the definition of the integrated luminosity in Equation (5.4) which is used in ATLAS, using lumiblocks as the time unit and assuming a constant instantaneous luminosity within each lumiblock [162].

The integrated luminosity of a sum of histograms is then simply taken as the sum of the individual results of Equation (10.3)

$$L = \sum_i L_{h_i}. \quad (10.4)$$

Similarly, the mean-pile up associated to a histogram $\langle \mu \rangle_{h_i}$ is calculated using the overlap time-weighted average of all lumiblock contributions

$$\langle \mu \rangle_{h_i} = \frac{1}{\sum_{j_{LB}} \Delta t_{h_i, j_{LB}}} \sum_{k_{LB}} \Delta t_{h_i, k_{LB}} \langle \mu \rangle_{k_{LB}}. \quad (10.5)$$

10.1.4. Histogram Filters and Histogram Sums


Using the information in the GRL and the associated values of $\langle\mu\rangle$, it is possible to select only those histograms passing the criteria of one or several so-called *filters*. The first kind of filter performs selections on good run list requirements, requiring that all lumiblocks associated with the histogram are considered good for physics. A second kind of filter selects histogram based on the average pile-up $\langle\mu\rangle$ as calculated in Equation (10.5), allowing selections of histograms in a certain interval $[\langle\mu\rangle_{\min}, \langle\mu\rangle_{\max}]$. Finally, filters on the IPBus error code associated with each histogram are possible, where an error code of 0 indicates a successful transaction. In addition to these lumiblock-level filters, run-level filters allowing the inclusion of only certain run numbers or the selection of runs based on the run type are implemented. Some examples of histograms resulting from various combination of filters are shown in Figure 10.2. Figure 10.2(a) shows a plot of a single total invariant mass histogram. All associated information is shown in the plot, namely the run number, a list of lumiblocks with corresponding $\langle\mu\rangle$ values, the recording time stamp of the histogram and the histogram number. In Figure 10.2(b), the sum of all histograms named “MjjTotal” with a mean pile-up level $\langle\mu\rangle$ of at least 20 are plotted for two cases, once with a good run list requirement and without it. The number of events is reduced by 2% to 5% depending on the invariant mass bin due to the good run list requirement. In Figure 10.2(c), histograms filtered in different ranges of $\langle\mu\rangle$ are shown. The region with higher values contributes more events than the region with lower values of $\langle\mu\rangle$. In rare cases, faulty IPBus transactions indicated by an IPBus error code not equal to zero occur, an example is shown in Figure 10.2(d). The associated lumiblocks are the ones with number three and four, meaning that they are at the very beginning of a run, where the whole TDAQ system is starting up and being configured, IPBus errors can therefore occur. In total, 379 histograms with IPBus transmission errors have been recorded in all runs of 2018. Most of these cases are either in the beginning or at the end of a run, where IPBus errors could occur due to the system being configured or stopped during the IPBus request. In a few cases, IPBus errors are found in the middle of a run. All of those can be associated with so-called TTC restarts, meaning a reconfiguration of TDAQ infrastructure as a consequence of module failure. These IPBus errors are therefore well understood, and in all following cases histograms with a non-zero IPBus error bit are filtered out. In addition, all the associated lumiblocks are not contained in any GRL, correctly indicating that ATLAS was not in a state to record data anyway.

10.2. Recorded Data

In the following, histograms recorded in different kinds of runs are inspected, using the filters described in the previous section. As this is the first time that the data set recorded in 2018 is inspected, a thorough investigation of the features of the recorded histograms is the main focus of this section. These inspections start with simple data from commissioning runs in Section 10.2.1, continue with cosmic muon runs in Section 10.2.2 and concludes on data in physics runs with proton beams in Sections 10.2.3 to 10.2.5.


10.2.1. Beam Splashes

In the very beginning of the LHC recommissioning after an end-of-the-year shutdown, protons are injected into the LHC again, but not yet brought to collision at the interaction points. Therefore, one of the first kind of events that experiments can record are so-called *beam splashes*. By



(a) Exemplary histogram shown together with the most important histogram properties.

(b) Histograms returned using a good run list filter and a $\langle \mu \rangle$ of at least 20.



(c) Application of a $\langle \mu \rangle$ filter.

(d) Histogram with an IPBus error.

Figure 10.2.: Exemplary plots created using histogram filters and information from the COMA database.

purposely closing the collimators in the beam pipe upstream of the detector, the circulating protons create a large spray of particles in the detector. Beam splashes are mostly used to verify that the synchronization between all subsystems works correctly. They also offer a possibility for the first tests to ensure that the histogramming feature works as expected in the corresponding version of the firmware. Histograms were recorded successfully and could be retrieved from the MDA, showing that the firmware is working and the infrastructure running as expected. Figure 10.3 shows two $\eta-\phi$ histograms recorded in a run with beam splashes recorded in April 2018. The histograms allow a distinction between splashes entering the ATLAS detector from either side, depending on whether the jet counts at negative values of η are dominant, see Figure 10.3(a), or the positive values of η are dominant, compare Figure 10.3(b).

10.2.2. Cosmic Data

Another kind of run that is very useful for commissioning and detector calibration are so-called cosmic runs, which refer to special runs using the whole ATLAS detector, but with no beam from the LHC. Instead, cosmic-ray muons, which can penetrate through the material into the ATLAS cavern, hit the detector and cause measurable signals [163]. Most of the cosmic muons




Figure 10.3.: $\eta - \phi$ histograms recorded in events with beam splashes. Splashes entering the ATLAS detector from either side can be distinguished.

come from above and arrive mainly via the two access shafts which were used for the detector installation. The cosmic muons interact with the detector material mostly as minimum-ionizing particles, and can therefore traverse all subdetectors along their flight path. In rare events however, they can leave a larger fraction of energy in the detector, in particular due to energy losses via bremsstrahlung which can also cause calorimeter signals. Figure 10.4 shows various sums of histograms recorded in the cosmic run 358831, namely the leading jet E_T in 10.4(a), the invariant mass in 10.4(b), E_T^{miss} in 10.4(c) and the $\eta - \phi$ map in 10.4(d). The event counts in all histograms are surprisingly high, as rates below a few Hz are expected for events in cosmic data [164]. The E_T spectrum of the leading jet in Figure 10.4(a) shows two peaks, the first one at around 400 GeV, the second one at around 1 TeV. This is in contrast to the expectation of a monotonically falling energy spectrum. As a consequence of this, a peak-like structure can be observed in the invariant mass as well as in the missing transverse energy. As the observed effect cannot be caused by cosmic muons, another explanation has to be found. A hint is given by the resulting $\eta - \phi$ map, which is dominated by events inside $|\eta| \lesssim 1.7$, coinciding with the coverage of the ATLAS tile calorimeter.

10.2.2.1. Tile Laser Calibrations

The most probable explanation for the features observed in cosmic data is related to the bunch group sets introduced in Section 5.4. Every bunch group set contains a category named CalReq foreseen for subdetector calibrations. Events occurring in BCs associated with this category are not used to trigger ATLAS data, as the trigger decision is set in coincidence with the bunch group at the CTP when forming the L1A, vetoing events from these BCIDs. However, this vetoing is not applied for the histogramming analysis on L1Topo which is run before the event reaches the CTP, where the bunch group information is available. Therefore, events from these BC can enter the histograms, even though the reconstructed objects might have no physics origin. A subsystem which is running such calibrations frequently is the Tile calorimeter. The Tile calorimeter reconstructs energy by recording light emitted in scintillating tiles with PMTs,

compare for example Figure 6.6. In order to ensure stable measurements of the energy depositions with small uncertainties, frequent monitoring and adjustment of the PMT gains are required. In addition, the timing behavior of the front-end electronics has to be monitored frequently. For this purpose, a laser calibration system has been developed for the Tile calorimeter [165]. Besides being used in dedicated calibration runs, laser calibrations also occur in standard physics runs during BCIDs in the CalReq bunch group. Two different laser intensities are available and used in order to test both low-gain and high-gain amplification settings of the PMTs [166]. The availability of two different laser intensities corresponds well to the two observed peaks in Figure 10.4(a). In addition, almost all entries in the $\eta - \phi$ map shown in Figure 10.4(d) are inside the coverage of the Tile calorimeter of $|\eta| < 1.7$. Hence, it can be concluded that this background is indeed calibration-induced due to the Tile calorimeter.

In the following, data from physics runs is inspected. If the observed calibration effect is present in physics data as well, a procedure to subtract these peak-like contributions has to be developed.

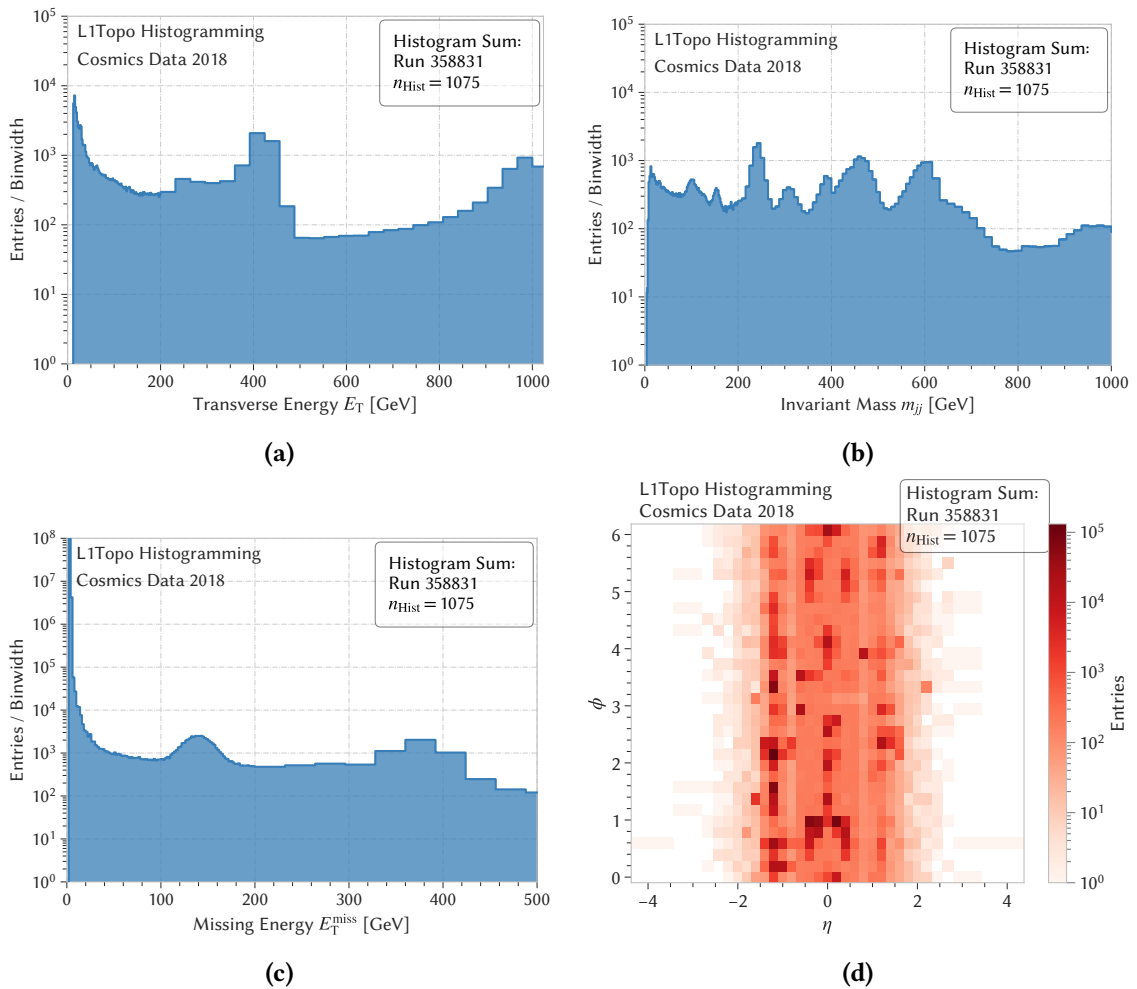


Figure 10.4.: Histograms recorded in cosmic run 358831 showing the energy of the leading jet (a), the total invariant mass (b), the missing transverse energy (c) and the $\eta - \phi$ map in (d).

10.2.3. Physics Data

Physics data corresponds to data collected with colliding proton beams in the LHC. For these runs, good run list filters are always applied. In contrast to the run types presented previously, information on the luminosity and the mean number of interactions per BC $\langle\mu\rangle$ can be extracted from the COMA database. The integrated luminosity and the pile-up distribution as a function of the histogram number are shown for a single run recorded on August 15 and August 16 in 2018 in Figure 10.5.




Figure 10.5.: Luminosity and pile-up as recorded for a run in August 2018 for every recorded histogram. The decay of the instantaneous luminosity over time can be observed.

Histogram sums of all histograms from this run passing the given filters are shown in the following for the total invariant mass in Figure 10.6(a), for the leading jet E_T in Figure 10.6(b), for the angular distance ΔR in Figure 10.6(c) and for the η - ϕ map in Figure 10.6(d). At very low masses $\lesssim 60$ GeV, a turn-on effect can be observed in the invariant mass spectrum. This is caused by the 12 GeV E_T requirement of the level-1 jet algorithm: low invariant masses are often formed by jets with low energies, and at a certain point it occurs frequently that one or both of the required jets fall below the 12 GeV threshold. Above the turn-on region, the invariant mass spectrum falls monotonically. However, deviations from the expected smooth behavior for mass spectra, as for example visible in Figure 9.2, are observed. These non-smooth effects get more prominent for higher invariant masses, in particular above 350 GeV and above 600 GeV. Since the invariant dijet mass is directly related to $\Delta\eta$ and $\Delta\phi$, see Equation (6.9), the limited level-1 jet granularity in both η and ϕ has a limiting effect on the granularity of the invariant mass, leading to a quantization of possible mass values and to fluctuations. This causes the observed, non-smooth oscillation behavior. In particular, when considering forward jet pairs from different forward regions, the minimum invariant mass is given as

$$m_{jj} = \sqrt{2 \cdot 12 \text{ GeV} \cdot 12 \text{ GeV} (\cosh 7.8 - \cos(0))} = 623 \text{ GeV} \quad (10.6)$$

Therefore, the oscillation structure in the invariant mass spectrum starting above 600 GeV is likely

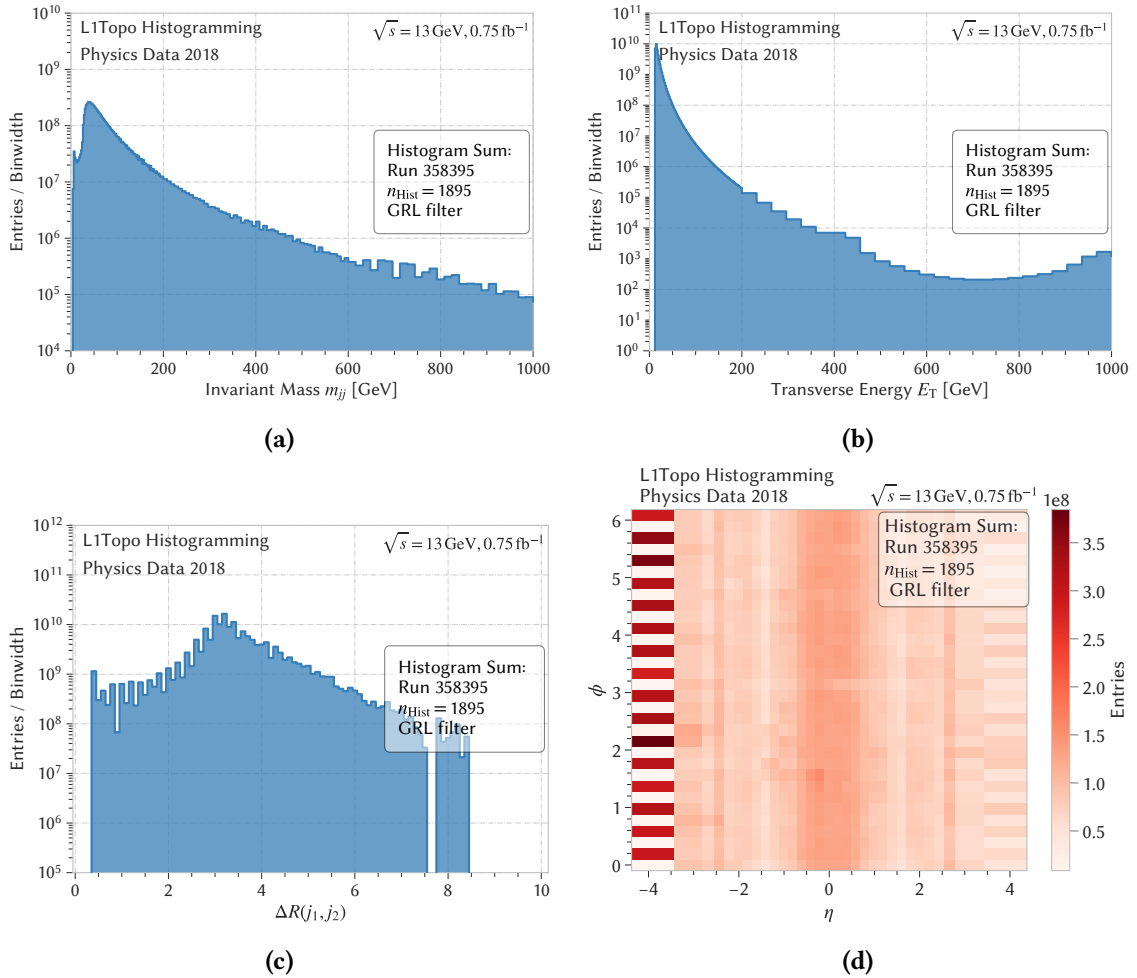


Figure 10.6.: Physics data histograms recorded in 2018: m_{jj} in (a), the leading jet E_T in (b), ΔR in (c) and the η - ϕ map in (d).

caused by low- E_T forward jets pairs. Similarly, the less-pronounced oscillation structure between 350 GeV and 600 GeV is likely caused by the combination of a non-forward jet with a low- E_T forward jet. Since these invariant masses are above the mass range this analysis is targeted at², this effect is not expected to be an immediate show stopper. However, this can potentially limit the range available for a fit, as many common fitting strategies rely on the smoothness of the mass spectrum.

The E_T distribution shows similar features to the one recorded in a cosmic run, compare Figure 10.4(a). In particular, the bump at 400 GeV and rising counts at the higher end of the energy spectrum are not expected in physics data and are similar to what has been observed in cosmic runs. This observation supports the hypothesis of Tile laser calibrations as the source of the peaks. A procedure to subtract this contribution is therefore developed in Section 10.3. The ΔR histogram shows an oscillating behavior due to the discrete possible η and ϕ values, which has not been taken into account when the binning was chosen. No values below $\Delta R = 0.4$ are observed,

²The mass scale of level-1 jets is different from the mass scale in an offline analysis, as will be discussed in Chapter 12. Therefore, a mass of 350 GeV on level-1 corresponds to more than 500 GeV in an offline analysis.




Figure 10.7.: Invariant mass for central jets with $|y^*| < 0.3$. Oscillations at higher masses are present also in this histogram.

as this is the minimum possible distance of jets due to the local maximum criterion of the SLW algorithm. The distribution peaks at $\Delta R \approx \pi$, which corresponds two jets produced back to back in the calorimeter and fits the expectations. Multiple noticeable features are visible in the η - ϕ map. While the spectrum is approximately uniform in ϕ , the different detector and transition regions can be observed in η , with drop-offs in events in the calorimeter transition region at $\eta \approx 1.7$. Due to the large bin width, event counts dominate in the forward region. An oscillating behavior can be observed in the forward region, where neighboring ϕ bins have very different event counts. This can be explained by inspecting the way the jet elements are build in the forward region, compare Figure 7.4. Due to the artificial splitting of forward trigger towers into pairs with equal energies, the corresponding pairs of neighboring jet elements will always contain the same energy depositions. Since the local maximum condition defined in Figure 7.5(b) requires the energy of a local maximum to be larger than the jet element above and larger or equal than the jet element below, a jet will frequently be found in the upper of the two jet elements, explaining the observed structure. However, the asymmetry between both forward regions, where more jets are observed for $\eta \in [-4.9, -3.2]$ than in $[3.2, 4.9]$ cannot be explained by this.

10.2.4. Physics Data: The Invariant Mass Categories

Due to the effects caused by the coarse granularity of forward jets on the invariant mass spectrum, the histogram categories containing only non-forward jets as listed in Table 9.1 will be preferred to the inclusive histogram for an analysis. The invariant mass histogram with the selection $|\eta| \leq 2.0$ for both jets and $y^* < 0.3$ is shown in Figure 10.7. The histogram shows the bumpy structure at 600 GeV, in contrast to the expectations, as these are believed to be caused by forward jets. A possible cause for this would be an error in the assignment of the histogram category of every event as described in Appendix B.1. Therefore, the code determining the histogram categories was inspected, but not a single error was found in simulations of a million randomly

generated events. In order to further debug this behavior, data from hot tower runs is revisited, as introduced in Section 9.5.3. For this purpose, a special pattern is used, where hot towers are generated primarily in the central region. Figure 10.8 shows various invariant mass category histograms and the $\eta - \phi$ map for events where only jets contained in the central region are present. Due to the absence of jets in the outer detector regions, the inclusive sum over all

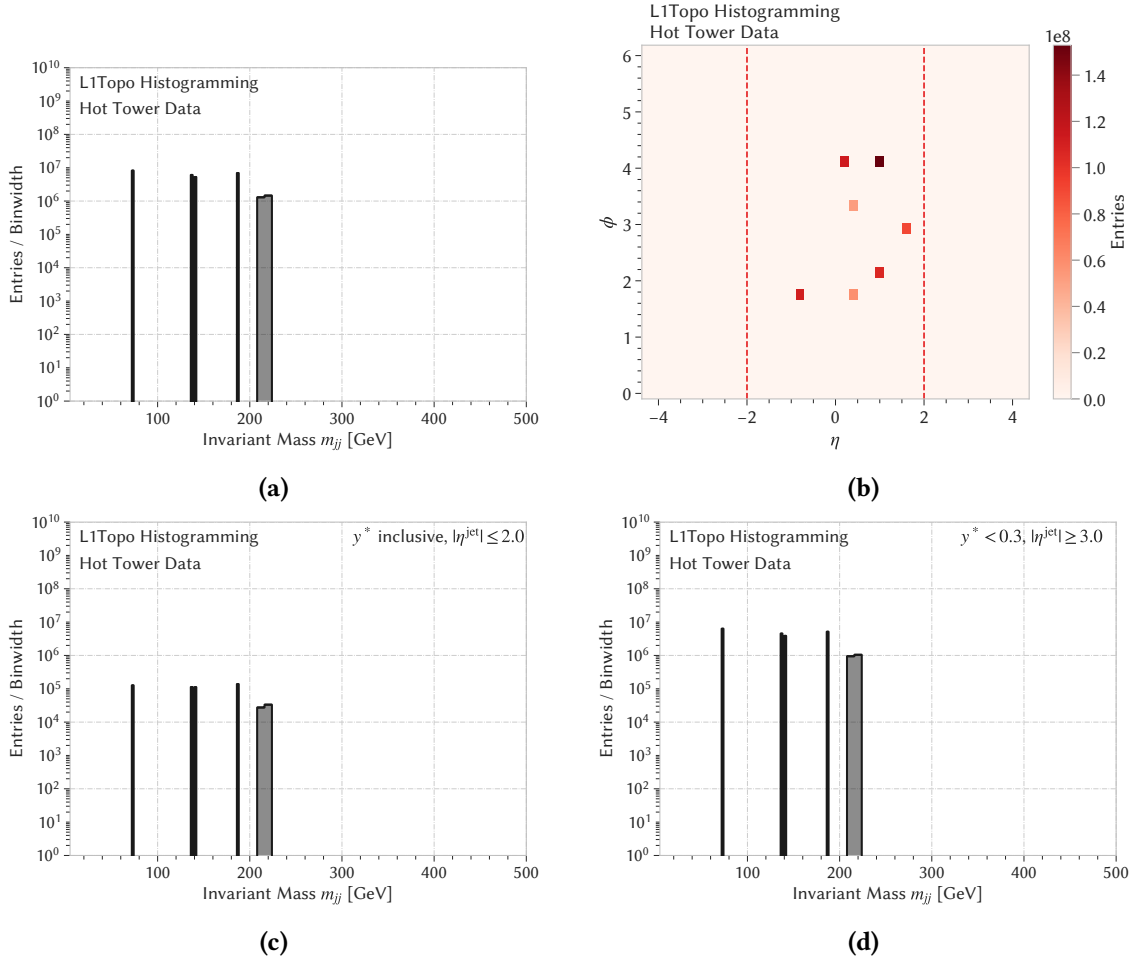


Figure 10.8.: Data collected in hot tower run 352976, showing the total invariant mass (a), the $\eta - \phi$ map (b), the invariant mass of central jets with any y^* value (c) and the invariant mass of forward jets with $y^* < 0.3$ (d).

three central invariant mass categories should correspond to what is visible in the total invariant mass histogram. However, this is clearly not the case, as the entries observed are two orders of magnitude lower than the entries in the total invariant mass histogram. Instead, most of the entries are found in a single histogram category, namely the category $|\eta| > 3.0$ and $y^* < 0.3$. This region should contain only cases in which two jets from the same forward region are present. The observations clearly point at a bug in the category assignment, unfortunately its origins remained unknown for the rest of the data taking.

The bug was found later in the number of clock cycles used inside the different modules of the histogramming. While the invariant mass has a total latency of 6 BC on its path towards the histograms, the histogram category information requires only 5 BC. The category information

which is used is therefore not related to the jets forming the invariant mass, but rather to the jets from the BC directly afterwards. Due to the way the peak finder of the PPM is designed, signals from every trigger tower are assigned to a single BC, compare Section 7.1. Even if the signal spans multiple BCs, no signal will be present in the BC following on the peak, which will therefore likely contain no jets. In firmware, this is signaled by empty, zeroed jets with all information set to zero. However, such empty jets are mapped to $E_T = 0$, $\eta = -3.9$ and $\phi = \pi/16$ as explained in Section 9.5.3. The histogram category assigned to a pair of zeroed jets is given as $|\eta| > 3.0$ and $y^* < 0.3$, corresponding to the category shown in Figure 10.8(d). When this bug was finally discovered, no possibility to deploy a firmware update was available anymore, so this bug is present in all recorded data, making any usage of the histogram categories impossible. However, the total invariant mass histogram is not affected by this bug, and will therefore be used for the analysis in the following, despite its disadvantages.

10.2.5. Physics Data: Luminosity and Pile-Up Profile

Using the information from COMA, it is possible to create summary plots of quantities from COMA, for example the recorded luminosity similar to Figure 5.3(b) or the pile-up profile as shown in Figure 5.2. The time-dependence of the integrated luminosity is shown in Figure 10.9 in a comparison to the luminosity reported by the offline luminosity calculation tool using the same good run list information. Figure 10.10 shows the comparisons of the $\langle\mu\rangle$ profiles. In both distributions slight differences are visible. A possible explanation for this is that the values reported for the luminosity and the number of interactions in COMA correspond directly to the maximum exposure of the detector to the beam, and contain no corrections for dead time, a trigger live time fraction or data quality criteria. Therefore, the luminosity and the mean number of interactions estimated from COMA are expected to be slightly higher than what will be estimated in an offline analysis from the corresponding tools which include this information. In addition, a subtle difference exists between the distributions shown in Figure 10.10 and in Figure 5.2. While the former uses the mean value of μ at the lumiblock level, the latter is calculated from the instantaneous per bunch luminosity, compare Equation (6.11).




Figure 10.9.: COMA luminosity compared to the luminosity reported by the offline luminosity calculation tool.




Figure 10.10.: COMA pile-up profile compared to the pile-up profile reported by the offline LumiCalc tool.

10.3. Subtraction Procedure

From the results of the previous sections, the conclusion that the histogramming is functional can be drawn. However, there are two kind of background effects that have to be considered before a statistical analysis of the total invariant mass spectrum can be started. The first background is caused by muons originating from cosmic rays as described in Section 10.2.2, which also exist as a background for runs with proton beams. For data recorded by ATLAS normally, cosmic-ray backgrounds are suppressed relying on multiple criteria [167]. For jets, the jet cleaning rejects cosmic fake jets based on calorimeter pulse shapes and track-based variables, since most of the cosmic muons will not travel through the interaction point. However, these procedures are only performed for reconstructed events and are therefore not applied for data recorded on L1Topo. The second background is the Tile calibration-induced background, which is present in physics data as was discovered in Section 10.2.3. If such an effect with a peak structure exists inside data, the certainty that a possible discovered peak is caused by an actual new physics signal and not by an unexplained artifact like this is severely reduced. Due to this, a procedure to remove these backgrounds from the spectra in physics data has to be developed. Ideally, the calibration background could be removed by vetoing events from the corresponding BCIDs directly in the firmware. For this it would be required to directly access the current bunch group scheme on L1Topo, as this would allow full control over the rejected BCIDs. However, this is not foreseen in ATLAS, as the bunch group scheme is not available in the level-1 trigger outside the CTP. As an alternative, plans to perform a static veto on certain fixed BCIDs as described in Appendix B.1.3 were developed. A static veto was therefore implemented to veto all events with BCIDs between 3445 and 3564, since the position of the CalReq category is constant at the end of each bunch group set. However, this showed no effect on the resulting spectra, meaning that either the veto was not implemented correctly, the used BCID information is not correct or that the source of the observed peaks is different from what is assumed.

Due to this, the background has to be subtracted directly from the physics histograms, reducing the impact of both cosmic muon events and Tile calibration effects on the invariant mass spectrum. The effect of both backgrounds can be determined from data recorded in cosmic runs by summing up all histograms from cosmic runs in 2018. This corresponds to a total of $n_{\text{hist, cosmic}} = 13164$ histograms, each of those recorded in a 50 s time interval. To normalize this contribution to the expectation in recorded physics data, the possibly different total recording time has to be considered. The cosmic histogram is scaled with a factor $n_{\text{hist, physics}}/n_{\text{hist, cosmic}}$ to let its normalization correspond to the total recording time of the physics data, as each histogram corresponds to the same recording time of 50 s.

In order to illustrate this subtraction, the leading jet E_T histogram is considered, as the effect of the subtraction on this histogram is visible by eye. Figure 10.11 shows an exemplary E_T histogram, where the impact of the peak-like structure is removed by the subtraction. Analogously, the corresponding cosmic invariant mass histogram is subtracted from the total invariant mass histogram. To underline the importance of this correction, the invariant mass spectrum in data is compared to the background contribution and to a possible signal peak in Figure 10.12. A Z' signal with a mass of 350 GeV and a coupling of $g_q = 0.1$ has been superimposed to the spectrum. The ratio of cosmic and physics data is shown as well, reaching values larger than 10^{-4} in a peak-like structure. This effect is in the order of magnitude of possible physics signals, and has therefore to be considered for any analysis. It also can be seen as a boundary which limits the sensitivity of a possible search, since signals which are much weaker than this peak can hardly be found.




Figure 10.11.: Example of subtraction procedure for an E_T histogram from a given run in 2018. The blue curve results from the subtraction of cosmic data (gray) from physics data (green).




Figure 10.12.: Physics data recorded in a run in 2018 with the estimated cosmic muon contribution and a possible Z' signal sample.

The Invariant Mass Spectrum

11.1. The Spectrum in Data

The final invariant mass spectrum for the statistical analysis is derived by summing invariant mass histograms from all physics runs in 2018 using the data quality selections introduced in Chapter 10. For the considered runs, a run type of data18_13TeV is required, to select only data recorded in the correct runs. In addition, a GRL filter using the same GRL as for the offline analysis is applied and only histograms where no IPBus occurred error are considered. All histograms passing these selection criteria are summed. Due to the firmware bug making the invariant mass categories inaccessible, the total invariant mass spectrum is considered, as described in Section 10.2.4. A subtraction procedure as described in Section 10.3 is performed on the spectrum to remove the contribution from cosmic and calibration backgrounds. The resulting invariant mass spectrum is shown in Figure 11.1. The event counts in the spectrum drop over several orders of magnitude, reaching unprecedented, high bin counts of up to $4 \cdot 10^{10}$ in bins of only 2 GeV width. The spectrum is monotonically falling starting at masses of about 50 GeV, at lower masses turn-on effects play an important role. Due to the influx of forward jet-caused oscillation above 350 GeV and above 600 GeV, utilizing these regions for a statistical fitting procedure will be challenging. Therefore, optimizations and studies performed in the following will be focused on the low mass region.

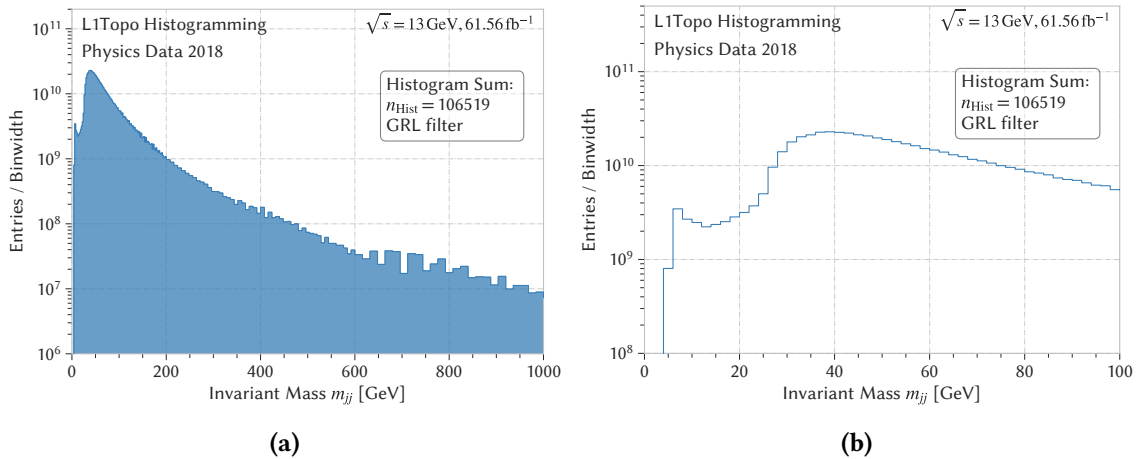


Figure 11.1.: The total invariant mass spectrum using the described filters for the full range in (a) and zoomed into the turn-on region (b) is shown.

11.2. Monte Carlo Samples

For the analysis of the recorded mass spectrum, a detailed understanding of the theoretical descriptions of signal and background processes is required. For this purpose, MC simulation samples are generated and briefly described in the following.

11.2.1. Signal Samples

The signal model used in this thesis has been briefly discussed in Section 2.2.4. The leading order matrix element of the model is evaluated with Madgraph interfaced to Pythia for the parton showering using the NNPDF3.1 PDF sets using the A14 PDF tune [168]. Signal samples are produced by ATLAS for the decay into pairs of quarks and encompass masses above 350 GeV in steps of 100 GeV and couplings g_{SM} between 0.05 and 0.40. To simplify the phase space, the mass of the dark matter particle is set to $m_\chi = 10$ TeV so that the Z' does not decay into dark matter. All official ATLAS samples undergo the GEANT4-based full detector simulation as introduced in Section 6.7. To scan the dependency of the cross section on the coupling and mediator mass, additional samples for lower masses have been generated privately with Madgraph. Figure 11.2 shows the correspondence of the mediator mass, the Standard Model coupling and the production cross section for the generated signal points.




Figure 11.2.: Production cross section for Z' samples as a function of the coupling to quarks g_q in (a) and as a function of the mediator mass in (b).

11.2.2. Background Samples

Since the differential cross section for jet production drops over many magnitudes, the standard QCD samples are sliced in terms of the leading truth jet momentum. This means that a fixed number of events is generated in different ranges of the leading truth jet momentum. To obtain a correct normalization, the different slices have to be scaled with different weights. The generation of slices mainly helps to reduce statistical uncertainties for high- p_T phase space with low cross sections. Table 11.1 shows the definitions of the first six dijet slices, denoted as JZ0,..., JZ5. The first two slices are generated as so-called *minimum bias* events. This generic term refers to an experimental procedure which selects events covering a large part of the overall inelastic pp

cross section, with the goal of accepting “a random sample of everything”. In the context of MC generators, it refers to all inelastic, non-diffractive events. Generating minimum bias for higher p_T slices is resource expensive and inefficient at the same time. Therefore, only the lower slices JZ0 and JZ1 are produced as minimum bias events, all other slices contain a hard p_T jet cut at the generator level. In particular, pile-up is dominated by minimum bias-like events. Therefore, pile-up is simulated in these samples by drawing random events from a sample of minimum bias events and overlaying them with the current interaction [169]. When combining the different slices, every event needs to be multiplied by a weight in order to correctly normalize each sample to the same integrated luminosity L . A weight $\omega_{JZ,i}$ is therefore applied to every event from sample JZi:

$$\omega_{JZ,i} = \frac{\sigma_i \times \epsilon_i}{\sum_k w_{k,i}} \cdot L, \quad (11.1)$$

where σ_i and ϵ_i correspond to the cross section and filter efficiency as given in Table 11.1 for every sample and L is the target luminosity. The term $\sum_k w_{k,i}$ in the denominator corresponds to the sum of MC event weights, which is similar to the number of generated events. Figure 11.3 shows the distribution of the slice-defining leading truth jet momentum weighted to an integrated luminosity of 61.54 fb^{-1} and confirms that the weighting is performed correctly. The invariant mass distribution shows similar features as the one in data, see Figure 11.1. The mass spectrum is clearly dominated by the lowest truth momentum slice JZ0, meaning that the distribution is dominantly formed by soft QCD effects. An additional observation is that the unsmooth effects at masses above 350 GeV and above 600 GeV occur only in the two lowest slices. Using this normalization, the invariant mass spectrum can be created for the simulation samples as well, which is shown in Figure 11.4.

Table 11.1: Pythia QCD MC slice definitions with cross sections, generated events, filter efficiencies and Pythia settings.

Slice	$p_{T,\min}$ [GeV]	$p_{T,\max}$ [GeV]	σ [pb]	Filter Efficiency	Pythia Cut
JZ0	0	20	$7.805 \cdot 10^7$	0.975	minimum bias only
JZ1	20	60	$7.805 \cdot 10^7$	0.0244	minimum bias only
JZ2	60	160	$2.433 \cdot 10^6$	0.00987	15 GeV
JZ3	160	400	$2.645 \cdot 10^4$	0.0117	50 GeV
JZ4	400	800	$2.546 \cdot 10^2$	0.00134	150 GeV
JZ5	800	1300	4.553	0.00145	350 GeV

11.3. Analysis Binning

Before the statistical analysis of the invariant mass spectrum is performed, it is necessary to decide on a suitable binning. A suitable binning is characterized by multiple features. On the one hand, the bin widths should not be too small compared to the mass resolution in order to avoid migration effects between different bins due to resolution effects. On the other hand, the bin widths should be smaller than the expected signal width, so that a possible signal distribution spans over multiple bins. In order to determine the width of signals reconstructed on level-1, multiple $Z' \rightarrow q\bar{q}$ MC samples as described in Section 11.2.1 are used. The available mass points of




Figure 11.3.: Truth jet p_T slices, weighted to the same luminosity. The curve shows a smooth behavior indicating a correct normalization of the different samples.

350 GeV, 450 GeV, 550 GeV, 650 GeV, 750 GeV and 850 GeV are used. For each sample, the invariant mass is calculated from the leading and subleading level-1 jet using a bitwise implementation of the invariant mass algorithm. The resulting histograms use the same binning as implemented on L1Topo. The invariant mass distribution is then fitted with a crystal ball function, and the relative mass resolution $\sigma(m_{\text{initial}})/\langle m_{\text{initial}} \rangle$ is determined for each sample. An example for the fit procedure is shown in Figure 11.5 for the $m_{Z'}$ = 350 GeV sample. Due to the energy scale of level-1 jets, the reconstructed resonance is not centered at 350 GeV but at a lower mass. The resulting dependency of $\sigma(m_{jj})/\langle m_{jj} \rangle$ on $\langle m_{jj} \rangle$ is then fitted with a function of the form $p_0 \oplus \frac{p_1}{\sqrt{m_{jj}}} \oplus \frac{p_2}{m_{jj}}$, which is typically used to describe calorimeter-based resolutions. Figure 11.6 shows the result of the fit. Using the functional form, the value of the mass resolution $\sigma(m_{jj})/\langle m_{jj} \rangle$ can be obtained




Figure 11.4.: Total invariant mass spectrum for the different QCD slices. The slice JZ0 is almost equal to the sum of all slices.





Figure 11.5.: Exemplary fit to determine σ/μ **Figure 11.6.:** Fit of the relative mass resolution for $m_{Z'} = 350$ GeV.

for each mass point. In the following, an iterative procedure to estimate the binning is described, based on the approach described in reference [55] for the offline dijet analysis. At first, an initial lowest bin has to be chosen as the starting value. For this analysis, $m_{\text{initial}} = 48$ GeV was chosen, as the turn-on region at low masses cannot be used for an analysis, compare Figure 11.1(b). In addition, an upper boundary of $m_{jj} < 1000$ GeV is used as a stop criterion. The procedure is then as described in the following, where the value of $\sigma(m_{\text{initial}})/\langle m_{\text{initial}} \rangle$ is obtained from the fit result and the factor f indicates how many bins are desired inside one standard deviation of a resonance. For this analysis $f = 6$ has been chosen, as it was found to be a good compromise between having enough bins while reducing bin-to-bin fluctuations in the low-mass region with very small bin widths.

1. Guess the initial position of the bin center as $m_{\text{center}} = m_{\text{initial}} \left(1 + \frac{f \sigma(m_{\text{initial}})}{2 \langle m_{\text{initial}} \rangle} \right)$
2. Use the derived value of m_{center} to calculate the lower bin edge as $m_{\text{lower}} = m_{\text{center}} \left(1 - \frac{f \sigma(m_{\text{initial}})}{2 \langle m_{\text{initial}} \rangle} \right)$
3. Check, whether m_{initial} and m_{lower} agree within a tolerance of 0.5%. If the disagreement is larger, proceed as follows:
 - If $m_{\text{lower}} > m_{\text{initial}}$, shift m_{center} by -0.25 GeV and repeat step 2.
 - If $m_{\text{lower}} < m_{\text{initial}}$, shift m_{center} by 0.25 GeV and repeat step 2.
4. Once the tolerance criterion is satisfied, set the upper bin edge as $m_{\text{upper}} = m_{\text{center}} \left(1 + \frac{f \sigma(m_{\text{initial}})}{2 \langle m_{\text{initial}} \rangle} \right)$
5. Repeat using m_{upper} as the initial position for the next bin.

After this procedure is finished, it has to be taken into account that no arbitrary bin edges are possible due to the binning of the recorded histograms. Therefore, the resulting edges have to be rounded to the closest possible value of the L1Topo binning. These bins are then used for all following analyses and listed in Table 11.2. As the bin widths are defined as a multiple of the relative signal width, the bin widths increase with increasing masses. This leads to an increase in bin widths compared to the initial bin widths as listed in Table 9.2, in particular in the lower mass region. Widths of 6 GeV to 8 GeV are used for bins below 200 GeV, compared to the initial bin

widths of 2 GeV. The widths between 200 GeV and 440 GeV remain unchanged. Above 440 GeV, widths of 16 GeV are used. This in turn leads to a reduced number of bins in comparison to the original binning. Since this can help to reduce resolution-caused bin-to-bin fluctuations, this could be helpful for the statistical analysis. After the new binning is applied, the spectrum looks as shown in Figure 11.7.

Table 11.2.: Result of the binning procedure. The lowest bin edge is 48 GeV. Bin widths of 6 GeV are used up to masses of 144 GeV, widths of 8 GeV until masses of 440 GeV and 16 GeV until masses of 984 GeV.

m_{jj} range [GeV]	bin width [GeV]
48–144	6
144–440	8
440–984	16




Figure 11.7.: The total invariant mass spectrum using the derived binning is shown for masses up to 600 GeV.

Level-1 Jet Properties

The properties and related uncertainties of reconstructed offline jets are well known, as these are used in a wide range of analyses, compare for example Figure 7.9. In contrast, only few studies exist on the properties of level-1 jets, see for example reference [125], which focus primarily on the rates and efficiencies of level-1 triggers, but contain little to no information on quantities of interest for an analysis, for example the energy scale, the spatial resolution or systematic uncertainties. All of these have an effect on the invariant mass spectrum and need a thorough understanding before the statistical analysis can be performed. Due to this, an analysis of level-1 jet properties is performed in the following, based on comparisons to reconstructed jets as reference objects, using both data recorded with the DAQ system and ATLAS full simulation MC. Properties studied in the following encompass the spatial resolution, the jet energy scale and jet energy resolution as well as reconstruction efficiencies. For all following studies, reference offline jets build by the anti- k_T algorithm with a radius parameter of $R = 0.4$ are used. These jets are calibrated using the procedure shown in Figure 7.7. In addition, only those offline jets marked as good quality by the jet cleaning procedure are considered. In order to compare L1 jets and reference jets, a matching procedure is necessary to establish a one-to-one assignment between the jets. For each L1 jet with $|\eta| < 3.2$, the angular distance ΔR to every reference jet is calculated. If a pair with $\Delta R(\text{L1, reco}) < 0.2$ is found, those two jets are defined as matched. To avoid any ambiguity or double counting in the matching, only those reconstructed offline jets which are isolated are considered. In this case, isolated means that no further reconstructed jets are found in a cone of $\Delta R = 1.0$ around the jet. Cases where two close-by offline jets have only one corresponding L1 jet are avoided by this requirement. In addition, level-1 jets are required to be isolated as well to avoid cases where a large reference jet is reconstructed as multiple jets at level-1. As only a single level-1 η bin exists between 3.2 and 4.9 (-4.9 and -3.2), a ΔR matching as performed in the central region would lead to large matching inefficiencies. Instead, a matching based on $|\Delta\phi|$ is performed, imposing a matching criterion of $|\Delta\phi| < 0.4$. In addition, only those reconstructed jets with $\eta > 2.8$ (or $\eta < -2.8$, respectively) are considered. To justify these matching conditions, the spatial resolution of L1 jets is discussed in the following, which is defined as the distance between a L1 jet and the closest matched reference jet.

12.1. Spatial Resolution

Since the jet coordinates are determined by the properties of the SLW algorithm and the discrete positions of jet elements, a certain displacement of the position with respect to the reference jet is expected. Figure 12.1 shows the distribution of ΔR for central level-1 jets and $\Delta\phi$ for forward jets. From these figures, the previously mentioned choice of the matching parameter can be justified:




Figure 12.1.: Spatial distance ΔR between L1 jets and the closest reconstructed jet for non-forward level-1 jets in (a) and $\Delta\phi$ for forward level-1 jets in (b) estimated in minimum bias MC events. The fraction of jet pairs below the given threshold is denoted as f on both plots.

98 % of all matches in the central region occur with $\Delta R \leq 0.2$ and with $|\Delta\phi| \leq 0.4$ in the forward region. In general, the maximum of the ΔR distribution is not at $\Delta R = 0$, but rather displaced to higher values. A certain displacement is expected due to the L1 granularity, which allows only positions in steps of 0.2 in η in the central region, the actual center of the energy deposition could be anywhere inside that jet element. In addition, the spatial resolution at low values of E_T is worse than at higher energies. Besides the energy-dependent resolution of the calorimeter, a possible cause for this can be the effect of pile-up noise on the level-1 jet reconstruction. For low-energetic jets, pile-up noise can contribute a larger fraction to the jet energy, and therefore possibly displace the position of the local maximum, resulting in a shift of the jet position.

12.2. Jet Energy Scale and Resolution

Other quantities of interest are the energy scale and energy resolution of L1 jets. It is expected that the energy scale of level-1 jets is lower than the energy scale of offline jets due to two reasons, even if the offline jets are considered at the uncalibrated EM scale. The dominant effect is that during event reconstruction, additional calibration procedures are applied to the calorimeter clusters entering the anti- k_T algorithm, see for example reference [130], which are not applied to the trigger towers used in the level-1 trigger. In addition, parts of the jet energy can leak out of the jet energy window of the SLW algorithm. However, this is not considered as a problem for ATLAS L1 jet triggers, as this effect can be compensated by simply adjusting the trigger thresholds to lower values, or by using different trigger thresholds in different calorimeter regions. While it would be possible to perform similar calibrations on trigger towers, a decision against this was taken by ATLAS as no significant improvement on level-1 triggers was observed. Figure 12.2 shows the difference in energy scales in a two-dimensional plot of the energy of L1 jets versus the energy of the matched reconstructed jets. The energy scale of L1 jets is lower than the scale of offline jets, except for a region at small reconstructed jet E_T . This region corresponds to cases where the anti- k_T algorithm has identified two or more jets in the forward region with similar ϕ values, which are merged into a single L1 jet with higher E_T due to the coarse η granularity.




Figure 12.2.: Correlation between the energies of matched pairs of L1 jets and fully calibrated offline jets. The blue dashed line indicates $E_T^{L1} = E_T^{\text{reco}}$. The data used to produce this plot is from a single run in 2018 (run number 358395).

An E_T dependence of the energy scale can be observed, as the relative deviation between the energy scales gets smaller at higher E_T . This can be explained by the fact that bremsstrahlung and $\pi_0 \rightarrow \gamma\gamma$ decays dominate the hadronic shower at higher energies, increasing the electromagnetic component which does not need an additional correction due to non-compensation [137]. Since the response of the ATLAS calorimeter is not uniform in η , the energy scale of a level-1 jet is expected to not only depend on E_T , but also on the position $|\eta|$. The extent of this effect can be quantified by studying the energy response of L1 jets, which is defined as the ratio of energies between level-1 jet and the matched reference jet:

$$\mathcal{R}_{L1} = \frac{E_T^{L1}}{E_T^{\text{reco}}}. \quad (12.1)$$

Studying the distribution of the response is also the starting point for a possible adaption of the JES calibration to level-1 jets, which is therefore discussed in the following.

12.2.1. Jet Energy Calibration

The most important aspect of the calibration is to make the response uniform in η and thereby independent on the detector region. However, the jet calibration procedure is performed event-by-event in an offline analysis and can therefore not be applied to the histogrammed data retrospectively, due to the lack of event level information. A calibration would have to be performed online in the firmware, where only limited resources, time and information is available. On L1Topo, the available information is restricted to the jet TOBs, i.e. to the four-momentum information. The origin correction, the jet-area based pile-up correction, the η -intercalibration and the global sequential correction rely on information that is not available on L1, these steps can therefore not be applied on level-1. It is therefore unreasonable to expect a calibration with an accuracy

comparable to the offline calibration. As only the jet four-momentum information is required for the absolute MC-based jet energy scale correction, this step of the ATLAS calibration scheme can be adapted to L1 jets. In particular due to limited manpower and the focus on firmware and software developments to implement a functional version of the histogramming, it was not possible to derive and to implement these calibrations for the data taking firmware in 2018. In order to understand the consequences of the lack of a calibration procedure, the following discussion describes how such a calibration could be derived in principle and discusses how this would affect the mass spectrum. Instead of relying on MC descriptions of low-energetic level-1 jets in the soft QCD regime, the chosen approach is data-driven and uses calibrated reconstructed jets in data as reference objects.

An understanding of the response distribution is important in particular for low-energetic L1 jets. Due to this, it is necessary to select a suitable data sample containing sufficient amount of low- p_T jets. Ideally, minimum bias triggers should be used for this purpose to avoid any kind of bias on the leading jet selection. However, this is not feasible, as the prescales on these triggers are far too high to collect sufficiently large samples in the various detector regions and different energy regimes. In order to select a data sample with sufficient jet yield at low E_T values, jets recorded in data selected with primary unprescaled muon triggers are used [170]. The exact trigger requirements are an isolated muon with $p_T \geq 26$ GeV or any reconstructed muon candidate with a p_T of at least 50 GeV. Figure 12.3 shows the response determined from these events in exemplary bins of the reference energy E_T^{reco} and for the possible values of $|\eta^{\text{L1}}|$.




Figure 12.3.: Response depending on E_T^{reco} and on $|\eta|$ of the corresponding L1 jet.

A clear E_T dependency of the response can be observed, with higher responses when considering reconstructed jets with higher E_T values. For low energetic jets, the response can be less than 0.4 in the transition region, while it reaches values of more than 0.7 for higher energies in the barrel and the end-cap. In addition, the response is $|\eta|$ dependent, showing a behavior similar to the offline response in Figure 7.8. For $0 \leq |\eta| \leq 0.8$, the response is approximately constant, afterwards it drops in the transition region between barrel and end-cap detectors at around $|\eta| \approx 1.4$. At higher values, it rises again. Due to the coarse granularity for higher $|\eta|$, effects in the transition




Figure 12.4.: Energy scale for L1 jets in two exemplary bins of E_T^{L1} and $|\eta^{L1}|$, showing low- E_T central jets in (a) and medium- E_T forward jets in (b).

region between end-cap and forward region cannot be directly observed. However, a noticeable drop-off of the response is present in the forward region. The goal of a calibration procedure is to make these distributions more uniform in η , and to bring the response closer to 1. In general, the energy calibration procedure can be described by applying correction factors c_{calib} to the energy E_T^{L1} , which are allowed to depend on both the energy E_T^{L1} and the position of the jet $|\eta^{L1}|$:

$$E_{T,\text{calib}}^{L1} = E_T^{L1} \cdot c_{\text{calib}}(E_T^{L1}, |\eta^{L1}|). \quad (12.2)$$

In order to derive the E_T and $|\eta|$ dependencies of the calibration factors, the properties of the response function are investigated more closely. For this purpose, the response distribution is considered in a two-dimensional binning in E_T^{L1} and $|\eta^{L1}|$. Fits using a Gaussian function are applied to determine the mean μ and width σ of the response curve in every bin. Figure 12.4 shows the distribution of the response for two exemplary cases together with the applied Gaussian fits. Both widths and means of the distribution differ for the various bins reflecting the E_T and $|\eta|$ dependence. Some deviations from a Gaussian shape are visible for jets in the forward region as a result of the coarse granularity in η . To ensure convergence of the fit in all bins, a two-stage fit approach has been developed. In this approach, a first fit is applied over the full range to determine initial values for the mean μ_{initial} and the width σ_{initial} of the Gaussian. The second stage fit uses the guesses from the first stage as its initial parameters. In addition, the window of the second fit is restricted to the interval $\mu_{\text{initial}} \pm n\sigma_{\text{initial}}$, where $n = 2$ for $|\eta| < 3.2$ and $n = 1.3$ otherwise were found to provide good results. The result for μ of the second fit corresponds to the mean response $\langle \mathcal{R} \rangle$ in every bin. To accompany this, the mean energy of all reference jets is calculated in each bin, resulting in a set of points $(\langle E_T^{\text{reco}} \rangle, \langle \mathcal{R} \rangle)$ for each $|\eta^{L1}|$ bin. In order to find a smooth parametrization of the energy response, the dependence of $\langle \mathcal{R} \rangle$ on E_T in a given bin of $|\eta^{L1}|$ is parametrized using an empirical function \mathcal{F} , similar to the one used in the offline calibration in Equation (7.8):

$$\langle \mathcal{R} \rangle(E_T, |\eta^{L1}|) \equiv \mathcal{F}(E_T)|_{|\eta^{L1}|} = \sum_{i=0}^N a_i \cdot (\log E_T)^i. \quad (12.3)$$

The parameter N is a hyperparameter of the model, where $1 \leq N \leq 6$ is allowed. It provides additional flexibility by adding additional higher order terms. The choice of the parameter N is a

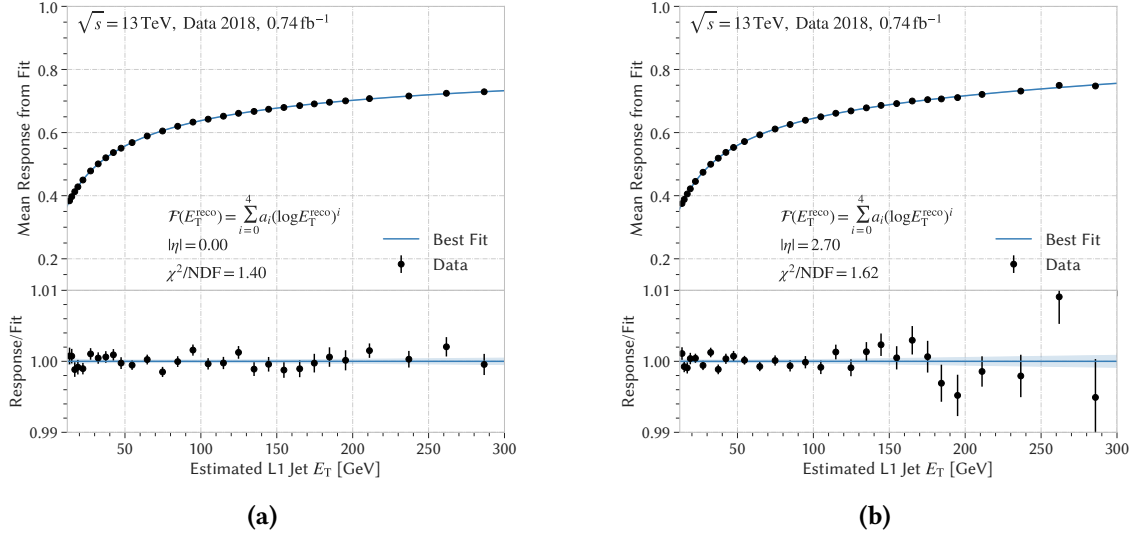


Figure 12.5.: Exemplary fit of the response as a function of the estimated L1 jet E_T for level-1 jets at $\eta = 0.0$ in (a) and jets at $|\eta| = 2.7$ in (b). The ratio of response values to the fit is shown, the shaded region indicates the 68 % confidence band of the fit.

trade-off between having enough degrees of freedom to adequately describe the dependence on E_T and having too much flexibility in the model, which can lead to overfitting. In principle, one could directly choose to parametrize the correction as a function of E_T^{L1} . However, formulating the response directly in terms of E_T^{L1} would require a setup in bins of E_T^{reco} instead of E_T^{L1} , as only in this case $\langle E_T^{L1} \rangle$ could be calculated to form pairs of $(\langle E_T^{L1} \rangle, \langle \mathcal{R} \rangle)$. This would lead to significant deviations from Gaussian shapes for the response, especially at low energies, where a cut-off in the distributions occurs due to the minimum possible value of 12 GeV for E_T^{L1} , see Appendix D.2 for an illustration of this problem. This is circumvented by applying a two-step procedure using the previously defined set of bins in $(E_T^{L1}, |\eta^{L1}|)$. In the first step, the response $\langle \mathcal{R} \rangle_{\text{reco}}(E_T^{\text{reco}}, |\eta^{L1}|)$ as a function of E_T^{reco} is determined from a fit of (12.3) to the points $(\langle E_T^{\text{reco}} \rangle, \langle \mathcal{R} \rangle)$ for every pseudorapidity bin. In the second step the mean energies are transformed by replacing E_T^{reco} with the estimate of the corresponding L1 jet E_T as

$$E_T^{L1, \text{ estimate}} = E_T^{\text{reco}} \times \langle \mathcal{R} \rangle_{\text{reco}}(E_T^{\text{reco}}, |\eta^{L1}|), \quad (12.4)$$

forming pairs of $(\langle E_T^{L1, \text{ estimate}} \rangle, \langle \mathcal{R} \rangle)$. This way, the cut-off and distortion effects are avoided, and the response can now be extracted as $\langle \mathcal{R} \rangle_{L1}(E_T^{L1}, |\eta^{L1}|)$ from the second fit. Two exemplary fits of the response as a function of the estimated E_T^{L1} using $N = 4$ are shown in Figure 12.5. To find an optimal value for the parameter N of the model, a cross-validation procedure is applied. For each value of N , 50 % of all data points are removed from the fit, and the model is fitted to the remaining data points. Afterwards, the prediction quality of each model is evaluated on the dropped data points by calculating the residuals for every dropped point

$$r = \frac{\mathcal{F}(E_T)|_{|\eta^{L1}|} - \langle \mathcal{R} \rangle}{\sigma_{\mathcal{R}}}. \quad (12.5)$$

This procedure is repeated m times. Assuming a Gaussian distribution for the residuals, the

negative log-likelihood for each model on the previously dropped points is calculated as

$$-\log L = -\log \left(\prod_{i=1}^m \exp(-r_i^2) \right) = \sum_{i=1}^m r_i^2, \quad (12.6)$$

and the model with the lowest value of $-\log L$ is taken, as it yields the best predictions for unknown data points. According to this procedure, the model with $N = 4$ yields the best result, providing enough freedom to describe the data while being less susceptible to overfitting than the models with higher values of N . The calibration is then given as the inverse of the response fit in every $|\eta|$ bin

$$c_{\text{calib}}(E_T^{L1}, |\eta|) = \frac{1}{\langle \mathcal{R} \rangle (E_T^{L1}, |\eta|)} = \frac{1}{\mathcal{F}(E_T^{L1})|_{|\eta|}} \quad (12.7)$$

The resulting calibration factors are shown in Figure 12.6 using a binning of 1 GeV in E_T corresponding to the energy granularity of jets at level-1. These factors reflect the structure observed in the response distribution, with higher calibration factors at lower energies and in the problematic detector regions.




Figure 12.6.: L1 jet calibration factors as a function of E_T and $|\eta|$. Higher calibration factors are observed for low transverse energies and in the transition region.

12.2.2. Effect of the Jet Energy Calibration

To estimate the effect caused by the missing calibration on the level-1 mass spectrum, the effect is studied using MC simulation samples for a potential signal with $m_{Z'} = 350$ GeV. As a consequence of the shifted jet energy scale, a shift in the invariant mass scale is expected, meaning that a resonance with a true mass of 350 GeV will not be reconstructed at this mass, but at a lower mass value. The mass scale can be estimated using the fits which were used to determine the binning in Figure 11.6 and is shown in Figure 12.7 for the Z' full simulation samples. For signals in the considered mass range, the energy scale at level-1 is only about 60 % of the true mass. If an excess is found in a level-1 mass spectrum, it is therefore important to consider that the resonance cannot be found at the true mass. If a calibration was applied, this mass scale would be corrected, as is

shown in Figure 12.8. Figure 12.8(a) compares the signal distribution of the uncalibrated with the calibrated case, and Figure 12.8(b) shows the calibrated level-1 signal spectrum compared to the corresponding offline spectrum. An improvement in the mass scale can be observed, moving the spectrum much closer to the reconstructed spectrum and moving its peak close to the true mass of 350 GeV. In addition, an improvement of the relative mass resolution of 32 % can be observed when the calibration is applied. The resonance is now more narrow, making it easier to discover in a resonance search. This allows the conclusion that the lack of a calibration for the signal distribution poses no conceptual problem. However, the potential sensitivity of a search to small signal cross sections is reduced by broader signals.

Another problem could be that combinations of jets from different detector regions could lead to increased occurrences of certain values in the invariant mass background due to the dependence of the response on $|\eta|$. In the worst case, this could lead to distortions of the spectrum which could be interpreted as peaks. Therefore, the effect of the calibration is also studied when applied to minimum bias MC. This is depicted in Figure 12.9. The calibrated spectrum is visibly smoother than the uncalibrated spectrum. Due to the low response at high $|\eta|$, contributions from forward jets are shifted to far higher invariant masses, reducing their impact in the lower mass region which is of main interest for the analysis. A possible conclusion is that many effects of the spectrum are actually caused by the $|\eta|$ dependence of the jet response, and that possible fluctuations could mainly be caused by calibration artifacts, where combinations of objects from different calorimeter regions, in particular badly behaving forward jets, lead to fluctuations. This has to be taken into account in a future analysis. The clear recommendation for any trigger level analysis studies is therefore to derive and study the correct calibration of their objects.

12.2.3. Jet Energy Scale Uncertainty

In the following, it is briefly illustrated how an uncertainty of the jet energy scale of reconstructed jets could be propagated to an uncertainty for level-1 jets. For this, the absolute energy scale




Figure 12.7.: Relative mass scale of Z' signals, showing the ratio of the peak position $\langle m_{jj} \rangle$ on level-1 to the true mass of the mediator $m_{Z'}$. The shown uncertainties are propagated from the uncertainties on the fit parameters, larger uncertainties for high masses result from worse fit performance.

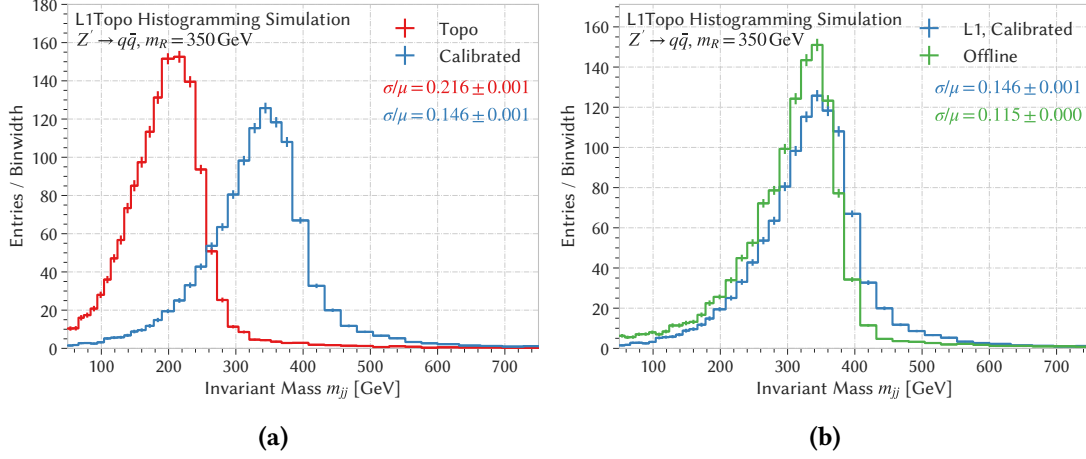


Figure 12.8.: Change of the invariant mass spectrum when a level-1 jet calibration is applied (a) and comparison of the calibrated spectrum to the offline spectrum (b).

uncertainty for reconstructed jets as shown in Figure 7.9 is considered. The calibration factors shown in Figure 12.6 are derived in two modified cases, where the reconstructed jet energies are evaluated at $\pm 1\sigma$ of their uncertainty, leading to modified calibrations given as follows

$$E_{T,\text{corr}}^{L1} = c_{\text{calib}} \cdot \left(1 \pm \frac{\Delta E_T^{\text{reco}}}{E_T^{\text{reco}}} \right) \cdot E_T^{L1}, \quad (12.8)$$

where ΔE_T^{reco} denotes the shift in the reconstructed jet energy caused by the total energy scale uncertainty. The level-1 invariant mass spectrum is then produced multiple times, once using the nominal calibration and then using the shifted calibrations. Figure 12.10 shows the calibrated spectrum along with the spectra with the varied calibrations. Using the uncertainties has a particularly high effect of 20 % in the turn on region, while the effect is at about 5 % at higher masses. These shifted templates could then be used in a statistical analysis to describe effects caused by uncertainties on the level-1 scale. It should be noted that this is only a phenomenological study to consider the general effect that an energy scale correction has on the invariant mass spectrum. Applying only the $\pm 1\sigma$ total uncertainties is a simplification, as correlations between the multiple contributions to the JES have to be taken into account when they are considered in a real analysis [171].

12.3. Jet Energy Resolution

Based on the response fits performed in the discussion of the energy scale, conclusions on the energy resolution can be drawn. For this, the information on the width of the response curves can be taken from the fits as shown in Figure 12.4. By inspecting the relative resolution of the response $\sigma_{\mathcal{R}} / \langle \mathcal{R} \rangle = \sigma_{E_T} / \langle E_T \rangle$, it is possible to obtain the relative E_T resolution [135]. The standard functional form for calorimeter-based resolutions as introduced in Equation (6.10) is used to estimate the contributions of the different terms. Figure 12.11 shows the resolution in different exemplary bins of $|\eta^{L1}|$ as a function of the estimated energy of the L1 jet. It is important to note that these are the resolutions relative to reconstructed jets as a reference objects, which have an inherent resolution themselves, meaning that the overall energy resolution with respect




Figure 12.9.: Comparison of the dijet invariant mass spectrum formed from uncalibrated jets and from calibrated jets.




Figure 12.10.: Comparison of the nominal calibrated spectrum and the spectra resulting from $\pm 1\sigma$ variations of the total JES uncertainty.

to the particle scale is worse than the numbers reported here. The resolution is much worse for low energetic jets, ranging between 15 % and 27 % depending on the calorimeter region. While the behavior is similar for most calorimeter regions at small energies, the resolution in the central region drops faster than in the transition region or in the end-caps. All regions beside the forward region show a similar behavior at small E_T values, the central region of $0 \leq |\eta| < 0.8$ shows the best performance at higher E_T values. It should be noted that the fit does not work well in the forward region ($|\eta| = 3.9$) as well as in the transition region ($|\eta| = 1.4$) due to the non-Gaussian shapes of the response curves.

12.4. Reconstruction Efficiency

Not every particle emerging from a collision will be successfully reconstructed as an object in the ATLAS detector. For example, parts of the energy of the original particle can be deposited outside the detector coverage and parts of the energy depositions can be removed by noise cuts. Even if a jet was successfully reconstructed offline using the anti- k_T algorithm, the effects of the lower energy scale, the worse energy resolution and the fixed window size lead to cases where either no corresponding jet can be found at level-1, or the measured transverse energy at level-1 falls below 12 GeV. Therefore, not every signal event that occurs can be reconstructed successfully, meaning that the reconstruction of a signal from two jets at level-1 has an acceptance A of less than 1.¹ Due to this, the reconstruction efficiency of L1 jets is a quantity of interest for an analysis of level-1 objects and is studied as a function of the jet E_T and η . In this case, the efficiency is defined as ratio of the E_T ($|\eta|$) distribution for reconstructed jets that are matched to a level-1 jet to the distribution for all reconstructed jets. This has the advantage that this property can be determined from data itself, instead of relying on the MC description of low-energetic jets. Figure 12.12(a)

¹Similar to the definition in other dijet resonance analyses [172], in this thesis the acceptance A is defined as the fraction of events passing all kinematic selections, including the 12 GeV jet requirement.




Figure 12.11.: Relative energy resolution in different $|\eta|$ bins, covering various detector regions. The resolution is calculated relative to calibrated offline jets.

shows the efficiency to reconstruct a L1 jet when the anti- k_T offline algorithm has found a jet as a function of the E_T of the offline jet. Two cases are distinguished, reconstructed jets in a central detector region $|\eta| \leq 2.5$, corresponding to the best available granularity at level-1, and an outer region of $|\eta| > 2.5$ covering extended and forward jets at level-1. Similar to this procedure, the efficiency is determined as a function of $|\eta|$, as can be seen in Figure 12.12(b). Again, two selected cases are distinguished, in this case a low E_T region around energy the turn-on region of 30 GeV to 60 GeV, and a region on the plateau with $E_T > 80$ GeV. The reconstruction efficiency as a function of E_T is fitted with a function of the form

$$\epsilon(x) = \frac{\varepsilon}{2} \left[1 + \operatorname{erf} \left(\frac{x - \mu}{\sqrt{2}\sigma} \right) \right], \quad (12.9)$$

which assumes an approximately Gaussian resolution [125]. The parameter ε describes the saturation value of the efficiency curve, μ the efficiency midpoint where an efficiency of 0.5 is reached, and σ the width of the turn-on. The turn-on in the forward region is slightly broader and reaches its efficiency midpoint at higher energies, as is expected due to the worse spatial resolution in that region. A value of close to 100%² efficiency is reached for central jets at about $E_T = 59$ GeV, while this point is only reached at about 66 GeV for jets in the outer regions. A strong $|\eta|$ dependence is observed for low values of the transverse energy, in particular at the calorimeter transition regions $|\eta| \approx 1.5$, at $|\eta| \approx 3.4$ and at $|\eta| > 4.0$. These correspond to the transition region between the Tile extended barrel and the LAr hadronic end-cap, the transition region between the hadronic end-cap and the LAr forward calorimeter, and the forward region at high $|\eta|$. For higher values of E_T , the reconstruction is fully efficient independent of the detector region. In summary, due to the limited reconstruction efficiency of level-1 jets, in particular at low transverse energies, the reconstruction of resonances at low masses is influenced. For example,

²Defined as the 99.7% efficiency point of $\mu + 3\sigma$.

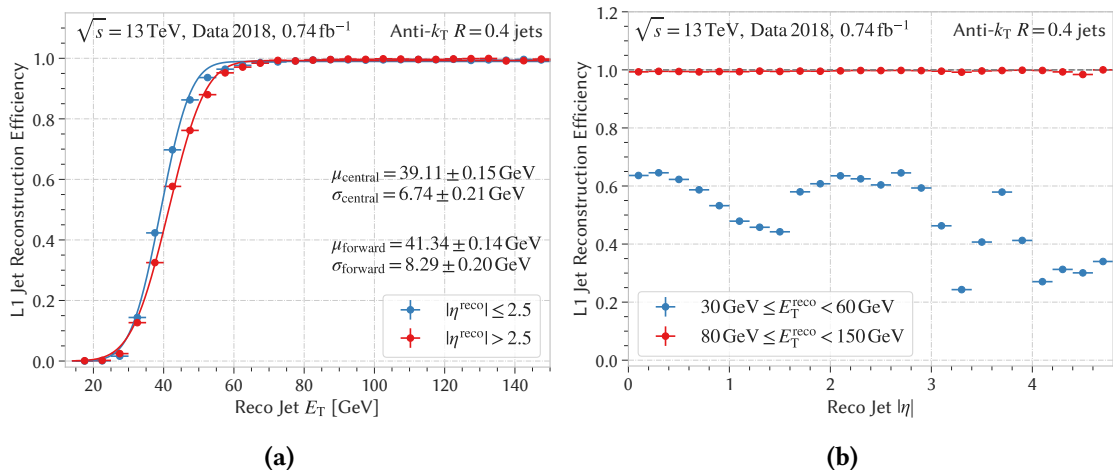


Figure 12.12.: Reconstruction efficiency of L1 jets as a function of the reconstructed jet's E_T (a) and as a function of the reconstructed jet's $|\eta|$ (b).

two reconstructed jets from a $Z \rightarrow qq$ decay which are produced back-to-back could both have about 45 GeV E_T , corresponding to an efficiency of about 70 % per jet, meaning that only half of these events would actually be reconstructed. Observing the Standard Model $Z \rightarrow q\bar{q}$ decays is therefore challenging, as its signal peak is close the turn-on of the mass spectrum and a significant fraction of events is lost due to the reconstruction efficiency. For the $Z' \rightarrow qq$ signal samples of 350 GeV and 450 GeV, this results in acceptances A of 96.7 % and 98.4 %, respectively. This number is expected to be smaller for smaller signal masses. However, the acceptance for most signals of interest is close to 100 % and poses no problem for an analysis.

12.5. Effect of the Jet Cleaning on L1 Jets

For reconstructed offline jets in ATLAS, a standard jet cleaning procedure as described in Section 7.2.2 is used to remove fake-jets. However, it is not possible to apply this procedure to L1 jets on L1Topo, as the information from the inner detector and on energy depositions in the different calorimeter layers required to form the discriminating variables is not available to the modules of the first-level trigger, compare the exact criteria as listed in Appendix A.5. Fake jets, which are normally removed by the jet cleaning, can therefore contribute to the invariant mass spectrum on level-1. Hence, this section is focused on determining the effect of the jet cleaning on L1 jets. For this purpose, a matching as described at the beginning of this chapter in is performed and the jet cleaning is applied on reconstructed offline jets. If a reconstructed jet fails the jet cleaning requirements, the matched level-1 jet is removed as well. This way, the effect of the jet cleaning on level-1 jets can be studied. Figure 12.13 shows the jet cleaning efficiency for reconstructed anti- k_T jets as a function of η and p_T . A notable efficiency drop can be observed in particular at small p_T values for $|\eta| \approx 3.0$. Outside these regions, there are only small inefficiencies observed due to the cleaning procedure. The ratio of dropped level-1 jets due to the jet cleaning to all jets is shown in Figure 12.14. The largest effect of the jet cleaning is again observed at around $|\eta| = 3.0$, where the effect can even be seen for higher values of E_T . These could be related to noise bursts in the HEC, which are one major contributions to fake jets, see Section 7.2.4. However, even in these cases less than 0.5 % of jets are dropped. A slight asymmetry can be observed, with a

smaller efficiency at negative η . In general, only small inefficiencies occur due to the jet cleaning. This is likely caused by the fact that the reconstructed jets rejected by the cleaning procedure tend to have low E_T values. Due to the level-1 energy scale, the corresponding level-1 jets will usually be below the energy threshold of 12 GeV, and will not result in a reconstructed level-1 jet in the first place. As a conclusion, the absence of a jet cleaning procedure is no showstopper for this analysis.




Figure 12.13.: Jet Cleaning efficiency for calibrated anti- k_T $R = 0.4$ jets using the loose criteria.




Figure 12.14.: Jet Cleaning efficiency for L1 jets matched to all calibrated anti- k_T $R = 0.4$ jets.

12.6. Pile-up in Jets

In the following, the effect of pile-up on level-1 jets is studied. Pile-up contributions are expected to affect the invariant mass spectrum in two ways: first by contributing additional noise to the energy of every jet, and second by contributing additional jets to an event. To determine the effect of the first contribution, the fact that two different window sizes for the jet energy window are available is utilized, compare Figure 7.6. For hard-scatter jets, most of the jet energy should be deposited in the 0.4×0.4 -sized center of the jet, while for jets with additional pile-up contributions more energy is expected also in the outer jet elements of the 0.8×0.8 -sized window. The distribution of the ratio $E_T^{4 \times 4} / E_T^{8 \times 8}$ is shown in Figure 12.15 in bins of $\langle \mu \rangle$, where each $\langle \mu \rangle$ bin is normalized to an integral of 1. For lower pile-up levels, the ratio is often close to 1, while lower ratios occur more often for higher pile-up values. In addition, the mean number of jets per event is shown in Figure 12.16 for various bins of $\langle \mu \rangle$. A weak dependence on the number of jets can be observed, as higher numbers of jets are observed more often for higher pile-up levels. These observations confirm the assumptions that were made on the effect that pile-up has on level-1 jets. For offline jets, parts of these pile-up effects are mitigated by the pile-up correction step of the jet energy calibration procedure. However, this correction step used for offline jets relies on information on the mean energy density ρ , the number of vertices N_{PV} and the number of interactions μ , which is not available on L1Topo. An energy density could in principle be calculated using the total energy information provided to L1Topo by the JEP, which could be divided by the total number of jet elements. However, this would only be strictly valid in the




Figure 12.15.: Energy ratio $E_T^{4\times 4}/E_T^{8\times 8}$ in different bins of $\langle\mu\rangle$.




Figure 12.16.: Mean number of jets in different bins of $\langle\mu\rangle$.

central region, where the trigger towers are of equal size, as the area differences are not considered in the energy sum. Another possibility to correct parts of the dependence of the response on pile-up could be from the energy ratio as shown in Figure 12.15. Figure 12.17 shows the response as a function of the energy ratio $E_T^{4\times 4}/E_T^{8\times 8}$. For higher values of the energy ratio, the response distribution is more narrow and takes smaller values on average. When considering lower energy ratios, higher values of the response occur more often, indicating that pile-up contributions to the jet energy play a larger role in this case. This effect is even more pronounced for forward jets, where the response distribution has a long tail. So, an additional correction of the energy scale for a future analysis could be achieved by deriving a pile-up based correction using such an energy ratio information.




Figure 12.17.: Response distribution in different bins of the ratio $E_T^{4\times 4}/E_T^{8\times 8}$.

Monte Carlo Approach

The central challenge of a dijet resonance search is to find a background description of the invariant mass spectrum. Past dijet analyses often applied fits of smooth functions over a wide mass range. However, a similar approach is very challenging for this analysis, due to the quantization effects in the invariant mass spectrum. Due to this, estimations based on MC are investigated in the following. Special challenges for this thesis will be the required amount of statistical precision, the necessity of a detailed detector simulation at the trigger level and the purely phenomenological description of soft QCD effects, which are especially relevant for the lower mass region, is challenging. Due to the latter aspect, background estimations based on MC are rarely used for dijet resonance searches. Therefore, it is uncertain to which extent simulated MC events allow drawing conclusions for data. Since the insights gained from these samples could still be helpful understanding the properties of the recorded data, MC events are considered in the following. In particular, a successful MC-based background estimation would be able to describe the oscillations in the mass spectrum, and would thereby significantly increase the mass range available for the analysis. Both MC full simulation samples produced by ATLAS and privately produced samples using a Delphes fast simulation are inspected.

13.1. Full Simulation

The following studies are based on the QCD full simulation samples introduced in Section 11.2.2. As the GEANT4-based detector simulation is run for these samples, they contain the level-1 jets formed by the simulation of the JEP algorithm, which are input to the L1Topo trigger simulation. These level-1 jet objects are the starting point for a simulation of the histogramming. The goal is to implement a *bitwise* simulation, meaning that all calculations are performed using the same bit-words, bit-lengths and precisions as used in the firmware, as only this allows trustworthy comparisons between recorded histograms and simulated histograms. In particular, the way the invariant mass is calculated in the FPGA firmware needs to be treated with care, compare its firmware implementation as explained in Section 8.3.4. For this purpose, the available floating-point coordinate values for η and ϕ in the simulation are multiplied by 10 and converted to integer values. Starting on these integer values, the LUTs for the $\cosh \Delta\eta$ and $\cos \Delta\phi$ operations are implemented based on the same 7 bit fixed-point representations as in the firmware. The multiplication with the energies is implemented as an integer multiplication and the final rounding is performed with the same bit shift as in the firmware. In addition, the non-restoring square root algorithm of L1Topo is ported to this simulation. Finally, the topological quantities are filled into histograms, using the histogramming definitions given in Table 9.2 and Table 9.3.

13.1.1. Full Simulation Data Comparison

In order to check the correspondence of the simulation and the recorded data, Figure 13.1 shows a comparison of the spectra between the one recorded on L1Topo and the sum of all slices from MC, scaled to the same integrated luminosity.




Figure 13.1.: Comparison of the invariant mass between L1Topo histogramming data and MC for the same luminosity.

A large discrepancy is observed: while the MC spectrum shows fewer events at low invariant masses, the oscillations caused by forward jets are more pronounced than in data and have higher events counts for $m_{jj} \geq 600$ GeV. As discussed in Section 11.2.2, these oscillations are present in the two lowest p_T slices which are produced as minimum bias samples. This means that the level of pile-up strongly affects the mass spectrum, as pile-up can be understood as overlaying additional minimum bias events to the hard interaction. A normalization issue of the MC samples can be excluded, compare Figure 11.3. Therefore, three possible explanations for these discrepancies are considered in the following: either there is a global, inherent problem with these samples, there is an inherent problem with the recorded histograms, or the effect of pile-up on the level-1 energy spectrum is not modeled well in MC. To exclude possible sample issues, the samples are inspected in phase space regions that are more commonly used for analyses. For this purpose, high-energetic reconstructed jets in data recorded with ATLAS are used to check whether a better correspondence between simulation and data can be observed this way. Events are selected using a high- p_T 360 GeV single jet trigger as used in offline dijet analyses, a similar selection is modeled in simulation. The comparison of the resulting leading jet E_T spectra is shown in Figure 13.2. After the turn-on region, where differences are present, a good correspondence between both spectra is observed. The spikes in the full simulation spectrum are caused by low statistics of high- E_T events in the JZ0 slice, which receive large events weights. It is found that the MC samples correspond much better to high- E_T reconstructed jets, indicating that there is no inherent problem with the samples. This allows the conclusion that these effects are caused by level-1 jets, which are only rarely used outside of trigger studies.

To exclude that there is an inherent problem with the histograms recorded on L1Topo, another cross-check is done to verify that the spectrum recorded on L1Topo resembles the corresponding spectrum that is formed from events recorded by ATLAS. For this purpose, a low-energetic jet




Figure 13.2.: Energy spectrum using reconstructed jets in simulation and recorded data.




Figure 13.3.: Comparison of L1 mass histograms from L1Topo, recorded data and simulation, scaled to the same integrated luminosity.

trigger requiring one level-1 jet with $E_T > 25$ GeV is used, in order to reach low invariant masses. The spectrum formed from reconstructed ATLAS data is scaled to the same luminosity as the L1Topo histograms. The comparison of the resulting spectrum to the L1Topo histogram and the simulation histogram is shown in Figure 13.3. This shows a good correspondence between recorded data and L1Topo histograms, even though high statistical fluctuations are visible due to the high trigger prescales of the used low- E_T trigger and the limited data statistics. The distribution from simulation significantly deviates from the other two, indicating that the most probable cause for the discrepancy is a poor description of low-energetic level-1 jets in the full simulation samples. Due to this, the MC distributions of the jet variables entering the invariant mass calculation are investigated, as shown in Figure 13.4. A couple of issues can be observed: in η , the forward region seems to be modeled badly in the simulation. While the shapes are similar in the central region, the simulation contains a much larger fraction of jets in the forward region than the L1Topo data. In ϕ , the full simulation shows an oscillating behavior which is much stronger than the one observed in the data distribution. While the E_T distributions indicate a general normalization issue, additional shape differences can also be observed here, since the ratio of both distributions is not flat. These discrepancies can most likely be explained by an erroneous description of pile-up and its effects on level-1 in the simulation, resulting in pile-up effects being harder in simulation than in data. In particular for low-energetic jets, pile-up can contribute significant parts of the energy. This then results in the oscillating behavior in ϕ , where jets are more likely to be found at certain ϕ positions if neighboring seeds with identical energies exist due to the local maximum comparison operators, see Figure 7.5(b). Since pile-up effects are stronger in the forward region, this would also explain why more jets are found in that region. In general, a full simulation of the level-1 trigger system is challenging, as many complex aspects need to be described, for example the calorimeter pulse shapes and the noise cuts, filter values and the BCID dependent pedestal correction of the PPM. The noise cuts and filter values are carefully determined in dedicated calibration runs on data, they are therefore hard to simulate and will be influenced heavily by an erroneous pile-up description. Due to this, standard ATLAS full simulation MC cannot be used to find an approximate description of the invariant mass spectrum. In order to further study the effects of pile-up, the noise cuts and the coarse granularity, it would be beneficial to vary some underlying parameters. While this is difficult to achieve

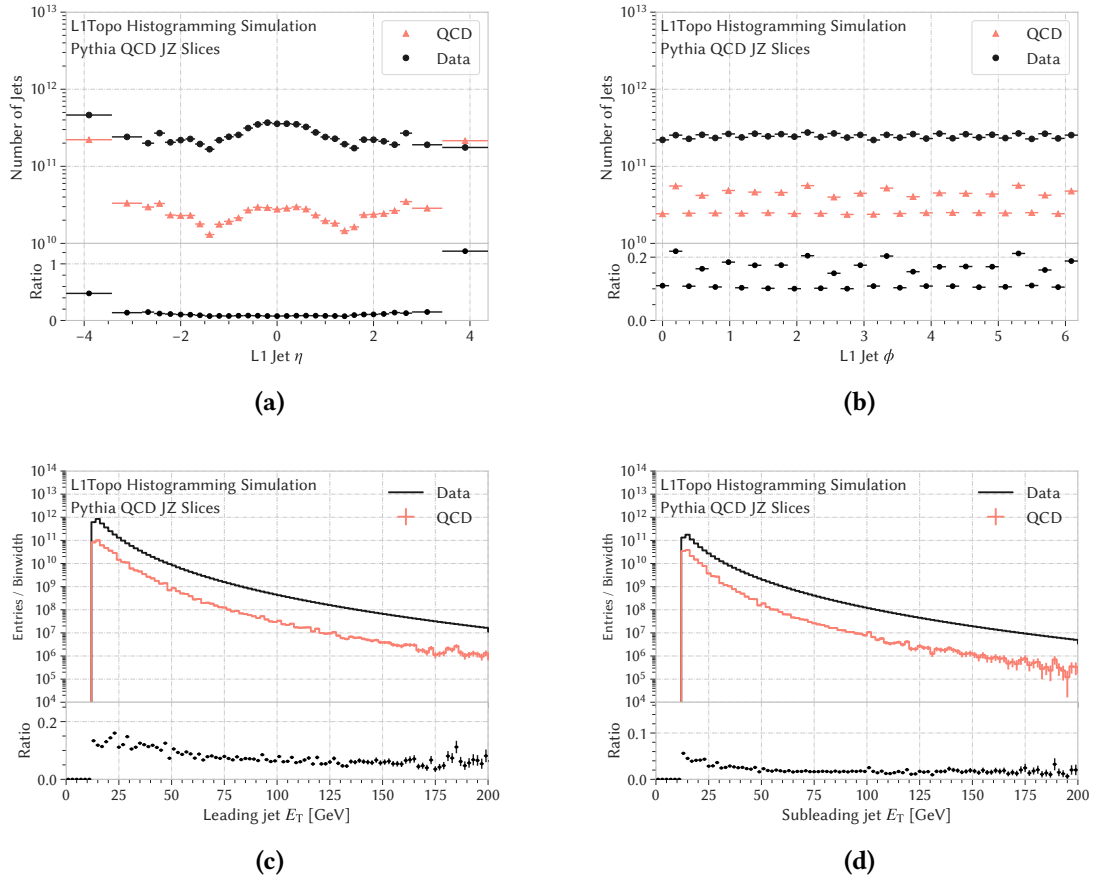


Figure 13.4.: Comparison of kinematic quantities between ATLAS QCD full simulation and histogramming data. The distributions of η (a), ϕ (b), the leading jet E_T (c) and the subleading jet (d) E_T are shown.

for the full simulation samples, as their production is resource intensive and time-consuming, and implementing changes can be very difficult, this is not true for fast simulation approaches. Consequently, an alternative simulation approach is developed and introduced in the following. This approach allows frequent adjustments, therefore various noise cut and pile-up scenarios can be tested and the simulation can be tuned more easily.

13.2. Sliding Window Algorithm Simulation

Based on the idea that many features of the shape of the invariant mass spectrum are driven by the granularity of level-1 jets, the properties of the level-1 jet algorithm and the bitwise mass calculation, a fast-simulation MC-based approach is developed. For this purpose, a simulation of the SLW jet algorithm is implemented, allowing direct and easy control over many settings. Hard QCD events are generated using Madgraph in leading order, these are then interfaced with Pythia for parton showering. A Delphes-based detector simulation is used for these studies, as the fast simulation significantly reduces the time to generate the desired samples. In the configuration of Delphes, the granularity of the calorimeter towers is set to the granularity of the level-1 trigger

towers. Based on these, multiple steps have to be simulated, as described in 7.1, starting with the JEP LUT E_T assignment including noise cuts as well as the correct combination of trigger towers into jet elements. Afterwards, the local maximum finding of the SLW algorithm has to be implemented, and the jet energy has to be calculated in the correct window. Possible adaptations of the energy scale will also be necessary, as Delphes is by default not set up to correspond to the level-1 energy scale at ATLAS. Each of the single steps is discussed in more detail in the following.

13.2.1. Sample Generation

$pp \rightarrow qq'$ events are generated with Madgraph in leading order, additional jets are created in parton showers in Pythia. Since the differential cross section $\frac{\partial\sigma}{\partial p_T}$ drops over several orders of magnitude with higher jet p_T , events are generated in multiple intervals of the leading truth jet p_T in order to generate sufficient statistics in all phase space regions. A summary of the chosen slices¹ and generated events is given in Table 13.1. In order to correctly combine events from the different slices afterwards, the generated numbers of events $n_{\text{events},k}$ and the production cross sections σ_k need to be considered. Therefore, a reweighting of each slice has to be done, where weights w_k are calculated as follows:

$$w_k = \frac{\sigma_k \cdot L}{n_{\text{events},k}}, \quad (13.1)$$

where L is the desired integrated luminosity. The total distribution is then given as

$$h_{m_{jj}} = \sum_{k=1}^4 w_k \cdot h_{m_{jj},k}. \quad (13.2)$$

The effect of this combination is shown in Figure 13.5(a) for the leading jet E_T and in Figure 13.5(b) for the invariant mass for $L = 1 \text{ pb}^{-1}$. These histograms show that different p_T ranges contribute to different parts of the spectrum, where the low-momentum slices dominate the turn-on region and the higher momentum slices take over gradually at higher momenta.


Table 13.1.: p_T boundaries, number of generated events and cross section as reported by Madgraph for every slice.

slice number j	$p_{T\text{min}}$ [GeV]	$p_{T\text{max}}$ [GeV]	n_{events}	σ [pb]
1	20	40	75600000	$6.474 \cdot 10^8$
2	40	60	50600000	$4.234 \cdot 10^7$
3	60	100	58110000	$9.308 \cdot 10^6$
4	100	∞	50600000	$1.256 \cdot 10^6$

13.2.2. Jet Elements and Noise Cuts

The generated samples are fed into a Delphes detector simulation, with a setting for the ATLAS detector, and the calorimeter granularity is set to the granularity of trigger towers shown

¹At the time these samples were generated, the slice definitions of the full simulation samples were unknown to the author, therefore the slice definitions given here differ from the ones from the full simulation.



(a) $E_{T,\text{lead}}$ distribution for the different p_T slices. (b) m_{ij} distribution for the different p_T slices. The high p_T slice only dominates at values above 3 TeV.

Figure 13.5.: Spectrum in different slices, normalized to an integrated luminosity of 1 pb^{-1} .

in Figure 6.10. Based on these towers, a simplified simulation of the PPM is implemented by applying noise cuts on every trigger tower. The noise cuts used in the first level trigger have been introduced in Section 7.1. In the simulation, various scenarios were considered, the 2018 version from Figure 7.3 was found to yield good results and is taken as the baseline configuration. Different sets of noise cuts are implemented for electromagnetic and hadronic layers, energy depositions below these cuts are set to zero. After the noise cuts are applied, the electromagnetic and hadronic component are summed and the trigger towers are combined into jet elements as shown in Figure 7.4, the resulting energies are rounded to integer values. However, it is not possible to simulate all functionality of the PPM. In particular, the BCID dependent pedestal cannot be implemented using such a simplified approach, as no BCID information is available. In addition, it is also important that the energy scale of the towers in Delphes can differ from the energy scale in ATLAS level-1 trigger towers, meaning that it is possible that the scale of the noise cuts is not fully correct.

13.2.3. The Sliding Window Algorithm

In order model the dijet invariant mass distribution of L1Topo, a simulation of the SLW algorithm is implemented, taking into account all the details described in Section 7.1. To implement the algorithm, it is good to reconsider that the JEM firmware implementation is independent on the sizes and coordinates of jet elements, instead it relies entirely on local bin numbers. No distinction between jet elements of different sizes is made, and the correct seed sizes and the correct size of the jet energy window are ensured by relying on incrementing or decrementing jet element bin numbers rather than working with coordinates. For some special cases at the boundary of the calorimeter coverage, a padding layer of jet elements with zero deposited energy is added internally. E_T thresholds of 12 GeV are applied to the jet energy window and to the seed. Figure 13.6 shows the result of the sliding window algorithm simulation for an event with three jets, showing the two-dimensional $\eta - \phi$ map of jet elements over the whole calorimeter. The colors show energy depositions in the individual jet elements, which are also indicated by the numbers written on every jet element on the plot. Blue pluses mark the positions of jets reconstructed using the anti- k_T algorithm implemented in the FastJet package on the original




Figure 13.6.: Example of a sliding window algorithm event. Three jets are identified by the SLW algorithm, corresponding to the jets found by the anti- k_T algorithm. Due to the coarser granularity available, there is a spatial displacement between these jets. The energy depositions in each jet element are given as well as the energy found by the SLW algorithm (red) and the energy found by the anti- k_T algorithm (blue).

towers in Delphes, the blue circles show the region of $R = 0.4$ around the center. Red crosses mark the position of the identified sliding window jets, the yellow dashed rectangles mark the jet seeds and the solid red rectangles the jet energy windows. The blue-colored numbers denote the Delphes jet energies in GeV, the red numbers the energy of the SLW jets. Both central and forward jets have a jet seed of 2×2 jet elements, however the forward jets covers a significantly larger η region. Note that the energy window of the particular forward jet shown in the figure covers only three jet elements in η , as the calorimeter stops at this point. For the algorithm still being well-defined there, an additional padding layer with zero energy is added internally in the algorithm. The identified energies differ due to two reasons: parts of the energy depositions in trigger towers can be removed due to noise cuts. In addition, parts of the energy depositions can fall out of the 4×4 jet energy window, as can for example be seen for the jet with $E_T = 107$ GeV. This implementation of the SLW algorithm is run for all generated events. From the resulting jets, the same bitwise implementation of the L1Topo histogramming as mentioned in Section 13.1 is used to build all histograms. Figure 13.7 shows a comparison of the E_T distribution as well as the m_{jj} distribution between the sliding window simulation histograms and the recorded histograms. From these plots it is evident that the energy scale is not correct and shifted to higher values, resulting in higher event counts.

13.2.4. Energy Scale Correction

In the following, a correction is derived to modify the energy scale of Delphes SLW jets to correspond better to the scale observed at level-1. For this purpose, the procedure to derive




Figure 13.7.: Comparison of the energy spectra (a) and the invariant mass spectra (b) between data and sliding window simulation. Both samples have been normalized to the same luminosity.

the JES correction described in Section 12.2.1 is inverted. Instead of calibrating the energy of ATLAS L1 jets up to the scale of calibrated offline jets, the sliding window jets at the Delphes energy scale are scaled down to the level-1 energy scale observed in ATLAS. To achieve this, two steps are performed: calibration factors are determined from ATLAS data by comparing calibrated reconstructed jets with level-1 jets, these factors are then applied to the SLW jets in Delphes. For the correction derived from ATLAS data, bins are defined in terms of E_T^{L1} and $|\eta^{L1}|$ as in Section 12.2.1, since this allows a parametrization of the response as a function of E_T^{reco} . The same trigger, object selection, and jet matching criteria as explained in Section 12.2 are used for this purpose. Using the same two-step fit procedure, Gaussian fits are applied in each bin, the average response $\langle \mathcal{R} \rangle$ and the standard deviation are estimated from the fit. To derive corrections as a function of E_T^{reco} , the mean value of the energy of all associated reconstructed jets $\langle E_T^{\text{reco}} \rangle$ is calculated for each $(E_T^{L1}, |\eta^{L1}|)$ bin. A smooth parametrization of the response $\langle \mathcal{R} \rangle(E_T^{\text{reco}}, |\eta^{L1}|)$ is then derived using fits of functions as described in Equation (12.3). The obtained response is then applied as a correction in Delphes as follows, given a Delphes SLW jet with E_T^{SLW} and $|\eta^{L1}|$:

$$E_T^{\text{L1,approx}} = \langle \mathcal{R} \rangle(E_T^{\text{SLW}}, |\eta^{L1}|) \cdot E_T^{\text{SLW}}. \quad (13.3)$$

Note that this assumes $E_T^{\text{reco}} \approx E_T^{\text{SLW}}$, meaning that the reconstructed jet energy at the ATLAS scale E_T^{reco} is equivalent to the energy at the Delphes scale E_T^{SLW} . The correspondence of $|\eta^{L1}|$ in Delphes and data is ensured by the way the SLW simulation is set up. In addition to the jet energy scale it is expected that the jet energy resolution differs, as the calorimeter response is determined using a simplified parametrized approach in the fast simulation. Due to this, a residual correction is performed, with the goal of bringing the relative resolution between SLW jets and reconstructed anti- k_T jets in Delphes to the same level as in ATLAS data. From the Gaussian fits performed for the energy scale correction, the relative energy resolution $\sigma_{\mathcal{R}} / \langle \mathcal{R} \rangle$ is determined. This is performed for two cases, in a SLW simulation sample with the applied

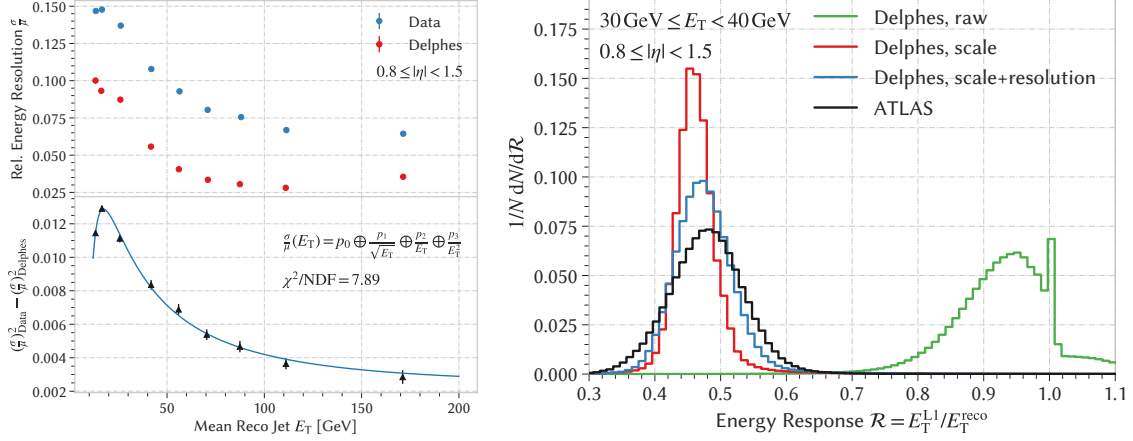


Figure 13.8.: Fit of the quadratic difference in an exemplary bin. **Figure 13.9.:** Energy response distribution at the different correction stages compared to ATLAS data.

energy scale correction (13.3), and in ATLAS data. The residual difference of those two

$$\left(\frac{\sigma}{\mu}\right)_{\text{residual}}^2 \equiv \left(\frac{\sigma_{\mathcal{R}}}{\langle \mathcal{R} \rangle}\right)_{\text{data}}^2 - \left(\frac{\sigma_{\mathcal{R}}}{\langle \mathcal{R} \rangle}\right)_{\text{delphes}}^2 \quad (13.4)$$

is then fitted using a calorimeter-resolution based function as introduced in Equation (6.10) with an additional higher order term proportional to $(E_T^{\text{reco}})^{-2}$. For a given value of E_T^{reco} , the fit result indicates the amount by which the SLW jet energy has to be smeared in order to correct for the existing residual difference. A correction is then applied by drawing random numbers from a Gaussian centered around 0 with the derived resolution as the width:

$$\Delta E \sim \mathcal{N}\left(\mu = 0, \sigma = \left|\left(\frac{\sigma}{\mu}\right)_{\text{residual}}\right| (E_T^{\text{reco}}) \times E_T^{\text{reco}}\right). \quad (13.5)$$

The final E_T for every jet is then calculated using Equations (13.3) and (13.5) as

$$E_T = E_T^{\text{L1,approx}} + \Delta E. \quad (13.6)$$

An example for the relative energy resolutions both in Delphes and in data as well as an exemplary fit of the residual quadratic difference are shown in Figure 13.8. The highest difference between the relative resolutions is observed at around 20 GeV, corresponding to corrections of about 12 % (2.4 GeV) which then decreases for higher energies, reaching a correction of around 4 % (6 GeV) for an energy of 150 GeV. To summarize the various correction steps, the effect of each step on the response distribution is shown in Figure 13.9 compared to the response in data for an exemplary bin in $(E_T, |\eta|)$. Without any corrections, the response in Delphes is close to 1. However, cases where not all energy of the original jet or additional energy contributions are caught lead to responses deviating from 1. The spike at 1 occurs from cases where the position of SLW and anti- k_T jets are very close and the energy deposition is spatially compact, such that both jets contain the same energy. After the response correction is applied, its mean is shifted much closer to the curve from data. However, the curve is more narrow and at a slightly lower mean. This is partly resolved by the energy resolution correction, though some differences remain. This is

caused by the fact that the corrections assumed $E_T^{\text{Delphes,SLW}} = E_T^{\text{Delphes, reco}} = E_T^{\text{ATLAS, reco}}$, which however is only an approximation, as the SLW and anti- k_T algorithm do not always catch the same energy, and the energy scale parametrized inside Delphes is only an approximation of the ATLAS scale.

To summarize these steps, the resulting histograms built from jets corrected using to Equation (13.6) are investigated. The invariant mass spectrum now fits better to the one observed in data, as can be seen in Figure 13.10. The obtained spectrum manages to reproduce the general features of the recorded spectrum, notably for example in the region between 600 GeV to 800 GeV, bringing significant improvements over the full simulation approach. However, deviations of around 5 % for small masses and up to 25 % for higher masses are still visible. In principle, this spectrum could be adjusted further by additional fine-tuning, for example by optimizing the noise cuts. However, the residual resolution correction itself depends on a previous run of the simulation, and therefore depends on the exact choice of noise cuts, which themselves would now depend on the resolution correction in a further optimization. Due to this, there are sets two parameters requiring careful optimizations, a change in either of those requires an optimization of the other one, making these optimizations difficult to perform.




Figure 13.10.: Invariant mass spectrum from Delphes after the energy scale correction has been applied compared to the data spectrum recorded on L1Topo.

13.2.5. Simulation of Pileup

Until now, the utilized samples contain only the particles emerging from the hard scattering itself. In order to be more realistic, additional pile-up events need to be included. In Delphes, pile-up is modeled using a data-overlay approach, similar to the full simulation. Random minimum bias events are drawn from pre-generated samples of events and overlaid to the current interaction according to the required number of interactions μ . In order to estimate the effect that pile-up has on the invariant mass spectrum, Figure 13.11 shows the mean transverse energy depositions of these minimum bias events in the calorimeter for $\mu = 40$ for both trigger towers and jet elements. For trigger towers, this results in average transverse energies of around 0.2 GeV to 1.5 GeV in




Figure 13.11.: Mean energy density of pile-up events as used in Delphes for trigger towers (a) and jet elements (b).

the central and extended region, and between 3 GeV and 5 GeV in the forward region. For jet elements, this corresponds to energies of 0.8 GeV to 2.5 GeV for the non-forward jet elements and additional transverse energies of 6 GeV to 10 GeV for jet elements in the forward regions. These energy depositions alone are comparable in size to the chosen values of the noise cuts. From this it is evident that pile-up will have a strong effect on the invariant mass spectrum, in particular in the forward region. The effect of overlaying additional pile-up with $\mu = 40$ to the spectrum is shown in Figure 13.12. From this it is evident that the effect of pile-up on the spectrum is too large, the effect is even larger than for the full simulation. The jet energy spectrum is shifted to higher values, and the shape of the invariant mass spectrum is heavily distorted, with a rise in event counts above 600 GeV where opposite-sign forward region jet pairs start to contribute. As a result of this, it was concluded that no satisfying description of pile-up exists in simulation with




Figure 13.12.: Comparison of distributions with and without the inclusion of additional pile-up events of $\mu = 40$. The leading jet E_T is shown in (a) and the invariant mass m_{jj} in (b).

the available tools. A possibility to resolve this would be a significant increase in the noise cuts. However, this approach remained unsuccessful and there was little evidence that the required precision could be achieved, in particular due to the cross-dependence on the energy-resolution corrections. Nevertheless, the approach without pile-up still manages to reproduce the invariant mass spectrum with a precision of about 10 % for most bins. While this is not sufficient for a background prediction for an analysis, the spectrum shapes are similar and make it possible to perform phenomenological studies using this spectrum. Two of them are presented in the following, the first one addresses possible effects caused by unsmooth behavior in the energy scale of level-1 jets, and the second one revisits the different invariant mass categories.

13.2.6. Study of Response Unsmoothness

The TLA analysis observed a significant excess, which turned out to be caused by an unsmooth calibration due to localized fluctuations in the jet p_T distribution in the absolute in-situ correction [14]. Therefore, a good understanding of the JES and its associated uncertainties is necessary to exclude such a possible cause for an excess. In this section it is checked whether an unsmoothness in one of the response calibration curves as shown in Figure 12.5 is able to produce an excess in the spectrum, which is comparable in size to the signals of interest. To compare how large such a possible signal is relative to the background, the $m_{Z'}$ = 350 GeV signal sample is considered. For this sample, the ratio of signal to background events is calculated in the central 68 % window, defined as $(\langle m_{jj} \rangle - \sigma_{m_{jj}}, \langle m_{jj} \rangle + \sigma_{m_{jj}})$, with the values obtained from the fits performed in Section 11.3. The working point of $g_q = 0.15$ is chosen, which corresponds to a production cross section of $\sigma = 0.25$ nb. This results in a number of signal events of $n_S = 0.683 \cdot L \cdot \sigma = 0.683 \cdot 61.56 \text{ fb}^{-1} \cdot 0.25 \text{ nb} = 1.05 \cdot 10^7$. By considering the data spectrum in the same mass window, the number of observed background events n_B is obtained as the integral over this window. The relative size of the signal peak is then calculated as $n_S/n_B = 3.06 \cdot 10^{-4}$. In the following, it is checked whether a peak of this relative size could also be caused by a small distortion of the smoothness in the response curves. By introducing an unsmooth jump into the calibration, the invariant mass spectrum is distorted, leading to the presence of deficits and excesses. Figure 13.13(a) shows several jump-like distortions of various heights introduced in the response curves and indicates the mean relative height of every unsmoothness. Figure 13.13(b) shows the effect of each calibration unsmoothness on the mass spectrum, depicted as the ratio of the invariant mass spectrum with an unsmoothness in the response calibration to the spectrum with a smooth response. To estimate the induced height of the peaks, Gaussian fits have been applied in the peak region between 100 GeV and 200 GeV. The induced peak height relative to the nominal mass spectrum is then plotted as a function of the relative height of the unsmoothness in the calibration curve. This is shown in Figure 13.14, where a third order polynomial with the constant term fixed to 1 has been fitted to extrapolate the dependency to the derived values of n_S/n_B . This results in an allowed relative unsmoothness of $\delta\mathcal{F} = 1.86_{-0.39}^{+0.28} \cdot 10^{-3}$ in the response calibration curve. A precision in this magnitude is therefore required in order to ensure that no calibration-induced peak of equal size as a targeted signal is present. This basically tells that the confidence in the origin of a possible peak as a potential discovery is severely reduced, as the jets entering the L1Topo mass calculation are uncalibrated. Some possible unsmoothness in the level-1 energy scale could result in peaks, which are equal in size or larger than the signals that the analysis is interested in. This indicates that uncertainties on the jet energy scale and its calibration should be added as a systematic uncertainty in a fit, in order to avoid the detection of peaks caused purely by calibration artifacts. Of course, the way the unsmoothness was modeled in this

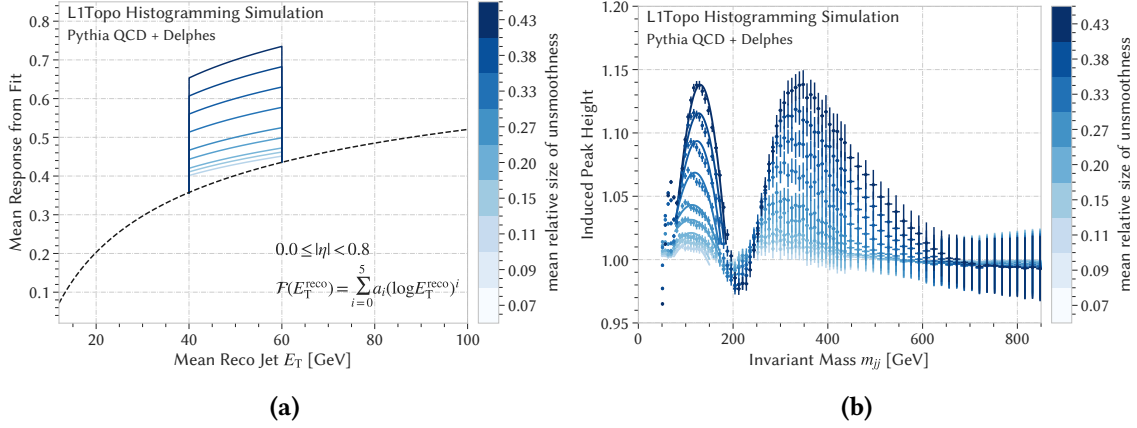


Figure 13.13.: Response calibration jumps (a) and their effect on the invariant mass spectrum (b).




Figure 13.14.: Polynomial fit of the dependence of the induced peak height on the introduced unsmoothness.

study can significantly deviate from what happens in reality, as this is only a phenomenological study. Nevertheless, the drawn conclusions have to be taken into account for this analysis.

13.2.7. The Invariant Mass Categories Revisited

Another advantage of simulation samples is that unlike in the histograms recorded on L1Topo, the invariant mass categories as listed in Table 9.1 are accessible. Figure 13.15(a) shows the total invariant mass histogram in comparison to the different signal region histograms from the SLW simulation, with $y^* < 0.3$ and $y^* < 0.6$ for jets with $|\eta_{1,2}| \leq 2.0$ or $|\eta_{1,2}| \leq 3.0$, respectively. It can be directly observed that these histograms contain far fewer unsmooth oscillations than the total invariant mass histogram. In addition, their bin contents fall much more steeply with rising invariant mass. The more stringent y^* requirement of $y^* < 0.3$ reduces the event count in particular for higher invariant masses, while the $|\eta|$ selections only have an effect at lower invariant masses, for higher masses the curves with only central jets and the one with central and extended jets approach each other. Clearly, the range of these histograms which could be used for a fit is much larger than in the total invariant mass spectrum. In addition to the signal regions, the different contributions from the orthogonal control regions dominated by forward jet effects

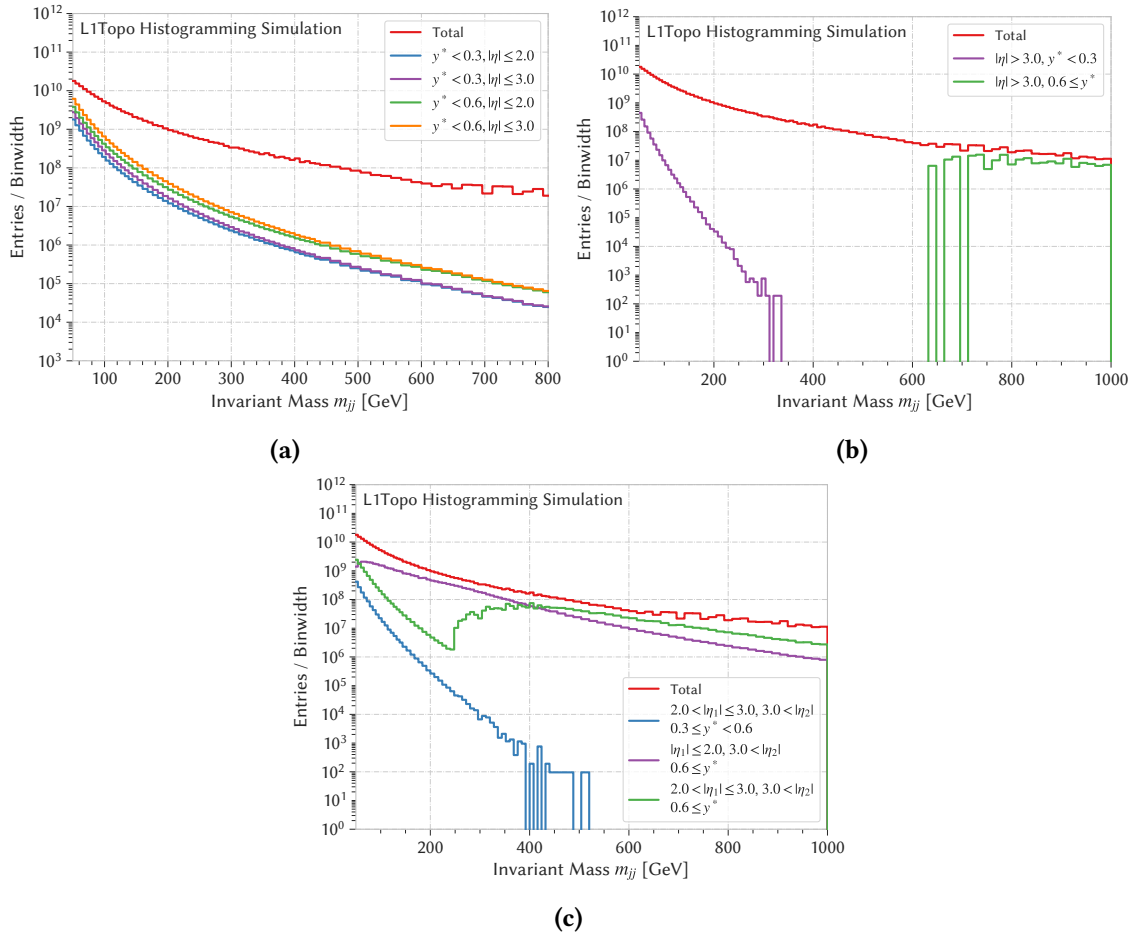


Figure 13.15.: The different invariant mass categories from the SLW simulation. (a) shows the signal regions, (b) the effect of forward jet pairs and (c) the effect of one forward jet.

are studied. Figure 13.15(b) shows the invariant mass distributions which are formed exclusively by forward jet contributions. Two possibilities exist in this case: either both jets come from the same forward region ($y^* < 0.3$), or the jets come from opposing forward regions ($0.6 \leq y^*$). While the pairs from the same region contribute only at small masses with rapidly falling event counts, opposite sign forward region jet pairs cause unsmooth oscillations at above 600 GeV, confirming the assumptions made in Section 10.2.3. Finally, Figure 13.15(c) shows the effect of combining one forward jet with a jet from a different detector region. While the combination of a central and a forward jet shows no striking behavior, this looks different for the combination of a forward jet with an extended jet, where the jets come from opposing sides of the detector ($2.0 \leq |\eta_1| < 3.0$, $3.0 \leq |\eta_2|$, $0.6 \leq y^*$). This contribution falls smoothly down to about 250 GeV, and then suddenly rises and contributes to the oscillations of the mass spectrum observed at around 350 GeV, again confirming the previously made assumption. In summary, the conclusion is that the total invariant mass spectrum is not the ideal choice for an analysis, significant improvements could be expected if one of the signal regions could be used instead.

The Background-Only Fit

Due to the extraordinary features of the recorded invariant mass spectrum and the high, unprecedented available sample size and statistical precision, the central part and main challenge of this analysis is given by the background estimation. As seen in the previous sections, MC descriptions of QCD processes are very challenging, therefore analyses targeting dijet final states usually rely on data-driven background estimations using smooth, empirical fit functions with three to five parameters which provide sufficient flexibility for a reliable background description. For the analysis presented in this thesis, additional challenges come from the properties of level-1 jets, in particular their coarse granularity, the large effects of pile-up and the missing calibration, and the fact that the designated signal region histograms with kinematic selections are not available.

This chapter starts with an introduction into the statistical concepts that are required for the background estimation in Section 14.1. Section 14.2 shows that the direct application of the background-only fit to the recorded data leads to problems with the fit quality. Due to this, a method to parametrize and reduce parts of the pile-up effects in the invariant mass spectrum is introduced in Section 14.3. Section 14.4 describes how the fitting procedure on the pile-up reduced spectrum is set up, and the final fit is performed in Section 14.5.

14.1. Statistics

For the statistical analysis performed in the following, the same approach used by other analyses of a similar kind, for example the offline dijet search [173], the TLA [7] or the dijet+photon search [8], is taken. All of those analyses apply a frequentist framework for a background-only fit, which quantifies whether the recorded invariant mass spectrum is compatible with the Standard Model-only hypothesis, and searches for possible deviations. If no deviations are found, a Bayesian limit setting procedure is applied, which will be introduced in Chapter 15.

Both Bayesian and frequentist approach provide valid, albeit different, interpretations of probability, which usually provide equivalent results [31]. However, they differ fundamentally in the way the results can be interpreted. Depending on the desired interpretation of the answer, the appropriate methods and interpretation can be chosen, as both are mathematically valid [174]. An in-depth discussion of their differences and reasons why these approaches were chosen for dijet analyses is given in reference [175].

14.1.1. The Likelihood

For the following, it is assumed that experimental results are recorded in the form of a histogram of data counts $\mathbf{d} = (d_1, \dots, d_n)$, where n denotes the number of bins. In addition, a background

estimate exists, which is used to create another histogram of counts, possibly depending on some parameters $\boldsymbol{\theta} \mathbf{b}(\boldsymbol{\theta}) = (b_1(\boldsymbol{\theta}), \dots, b_n(\boldsymbol{\theta}))$. The likelihood of the parameter values $\boldsymbol{\theta}$ given the observed data \mathbf{d} is then defined as the probability of the specific observation given the parameter values¹:

$$L(\boldsymbol{\theta} | \mathbf{d}) = P(\mathbf{d} | \boldsymbol{\theta}) \quad (14.1)$$

In the case of histogrammed data with sufficient event counts in every bin, the binned likelihood for the *background-only* hypothesis is given as the product of Poisson distributions for every bin:

$$L(\boldsymbol{\theta} | \mathbf{d}) = \prod_{i=1}^n \text{Pois}(d_i | b_i(\boldsymbol{\theta})) = \prod_{i=1}^n \exp(-b_i(\boldsymbol{\theta})) \frac{b_i(\boldsymbol{\theta})^{d_i}}{d_i!}. \quad (14.2)$$

By maximizing $L(\boldsymbol{\theta} | \mathbf{d})$, or equivalently by minimizing $-\log L(\boldsymbol{\theta} | \mathbf{d})$ with respect to $\boldsymbol{\theta}$, the best-fit values of the unknown parameters $\hat{\boldsymbol{\theta}}$ can be determined. Due to the high number of bins, the minimization of the log-likelihood is chosen, as it provides numerical advantages. The minimization is performed using Minuit [176].

14.1.2. The Background Estimate

In dijet resonance searches, the invariant mass spectrum is generally predicted to be smooth and monotonically falling. Due to this, most past analyses applied fits of empirical functions to model the background. This procedure, which uses a single fit function over the whole mass range is also referred to as the *global fit* [14]. Various functional forms have been found to describe this shape [7, 173, 177], the ones used in this thesis are listed in the following:

$$f_5(z, \boldsymbol{\theta}) = \theta_1 (1 - z)^{\theta_2} z^{\theta_3 + \theta_4 \log z + \theta_5 \log z^2} \quad (14.3a)$$

$$f_{\text{UA2}}(z, \boldsymbol{\theta}) = \frac{\theta_1}{z^{\theta_2}} e^{-\theta_3 z - \theta_4 z^2} \quad (14.3b)$$

$$f_V(z, \boldsymbol{\theta}) = \left(1 - z^{1/3}\right)^{\theta_1} z^{\theta_2 + \theta_3 \log z + \theta_4 \log z^2}. \quad (14.3c)$$

In all three functions, z is defined as $z = m_{jj} / \sqrt{s}$ and $\boldsymbol{\theta}$ denotes the free parameters. To form the background estimate $\mathbf{b}(\boldsymbol{\theta})$, the integral of these functions is then evaluated over every bin. However, increasing statistical precision due to higher integrated luminosities and wide invariant mass ranges covered by recent searches lead to the effect that subtle differences between QCD and the purely empirical approximation with a smooth fit function become more pronounced. As a consequence of this, none of the single functional forms is able to provide a good description of the data of recent searches. To resolve this issue, more recent analyses apply modified strategies for background estimation. The TLA analysis of Run-2 data [7] extends the global fit procedure to a so-called Sliding Window Fit (SWiFt), which is described in the following.

14.1.3. The Sliding Window Fit (SWiFt)

The SWiFt method is built around performing multiple fits, each minimizing a binned negative log likelihood as introduced in Equation (14.2). It relies on the same fit functions as the global fit, but each fit is limited to a window which covers only a subrange of the total range considered for the

¹The likelihood itself is interpreted as a function of the parameter values, and does therefore not necessarily represent a probability, since the integral over the parameters is not fixed to 1.

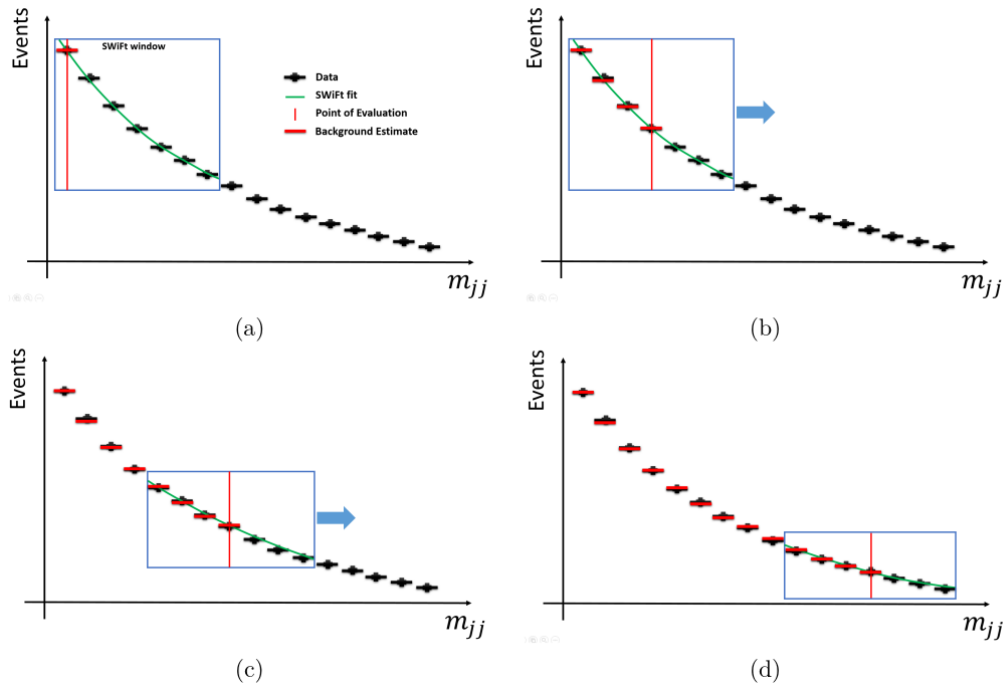


Figure 14.1.: Sketch of the functionality of the SWiFt procedure, from reference [14].

analysis. The SWiFt procedure is depicted in Figure 14.1. A background estimate for the central bin is obtained by evaluating the fitted function at the center of the fit window. Using a fixed window width, the window then slides bin-by-bin over the whole range, creating a background estimate in each window. The fit begins at the lowest window possible that still lies completely within the considered range, as depicted in Figure 14.1(a). In order to avoid edge effects, the first few and last few bins of the histogram are not used as window centers. Instead, the background prediction in the first few bins is determined from evaluating the fit in the first window. As soon as the center of the first window is reached, the point of evaluation is fixed to the window center, as depicted in (b). The fit window then subsequently slides over the whole range (c), until the last possible window is reached. Similar to the first few bins, the last few bins are evaluated from the fit in the last possible window (d). An important hyperparameter of the SWiFt procedure is the window width. Optimizations of this parameter are necessary, since a compromise between two opposing dependencies has to be made. Smaller window widths perform better with respect to localized deviations from a global smooth behavior. However, at the same time this limits the sensitivity to possible signals, as the fit is able to better adapt to local excesses. Larger window sizes, in particular those which also reach to the sideband regions where only little signal is present, are therefore better suited to detect possible signals, however convergence of the fit can be more difficult to achieve than with smaller windows.

14.1.4. Test Statistics

After a background estimate has been found, the level of agreement between the observed data and the background-only hypothesis H_0 can be specified. Two different *test statistics* are used for this purpose. Both are introduced in the following, with brief explanations of how p -values are

defined for these statistics.

14.1.4.1. The χ^2 Statistic

The value of χ^2 is defined as the sum of squared deviations between background and data normalized to the uncertainty of the background prediction:

$$\chi^2 = \sum_{i=1}^n \frac{(d_i - b_i)^2}{\sigma_{b_i}^2}. \quad (14.4)$$

where σ_{b_i} are the uncertainties on the background estimate, for example $\sigma_{b_i}^2 = b_i$ in the case of Poisson errors. The χ^2 allows to quantify the goodness of fit for the background-only model: high values of χ^2 indicate that the deviations between data and background cannot be explained by fluctuations of the background estimate alone, meaning poor quality of the background-only fit. Smaller values indicate a good correspondence between data and background prediction. A disadvantage of the χ^2 statistic is that it does not take bin-to-bin correlations into account, which is a flaw in a search for a localized excess. For the value of χ^2 it does not matter whether there are six adjacent bins with a 2σ excess each, or three bins with a 2σ excess and three bins with a 2σ deficit spread over the whole spectrum. However, the former case is far more interesting in the context of a resonance search, as it can indicate the presence of an excess in the mass spectrum. To cope with these cases, the BumpHunter algorithm is used.

14.1.4.2. The BumpHunter Test

The BumpHunter (BH) algorithm [178] is used to identify the presence of a local excess over a continuous background, as caused by the resonant production of a particle. It scans over all possible windows of adjacent bins $[j, k]$ in the invariant mass spectrum, where $j, k \in \{1, \dots, n\}$ and $j < k$. Initially, window sizes of two are used, but gradually increased up until half of the total number of bins in the mass range is reached. In each of these windows, the Poisson probability t of obtaining a result at least as significant as the observed one if the window is combined into a single bin is calculated. The number of entries when combining the window into a single bin is given as follows

$$d = \sum_{i=j}^k d_i \text{ and } b = \sum_{i=j}^k b_i. \quad (14.5)$$

From this, the probability t is obtained as

$$t = \begin{cases} \gamma(d, b), & \text{if } d \geq b \\ 1 - \gamma(d + 1, b), & \text{if } d < b \end{cases}, \quad (14.6)$$

where the first case describes the probability of finding an excess and the second case the probability of finding a deficit. In this case, γ denotes the incomplete lower gamma function

$$\gamma(d, b) = 1 - e^{-b} \sum_{m=0}^{d-1} \frac{b^m}{m!} \equiv \sum_{m=d}^{\infty} \text{Pois}(m | b). \quad (14.7)$$

By looking at the overall deficit or excess in a certain region, the value of t also accounts for local correlations in the bin-to-bin fluctuations. A test statistic t_0 is then defined as the negative logarithm of the smallest probability t_{\min} obtained for any possible window:

$$t_0 = -\log t_{\min}. \quad (14.8)$$

14.1.4.3. p -values

In order to quantify the probability whether the observed value of a test statistic arose only due to statistical fluctuations of the background, both the χ^2 and the BH test statistic are translated into p -values. A p -value allows to state the probability to obtain a result like the observed data or one that deviates more from the background hypothesis H_0 purely by statistical fluctuations of H_0 . The p -value for a hypothesis H_0 and a test statistic t is defined by the probability to find values of t that are equal or greater than the observed value t_{obs} :

$$p = \int_{t_{\text{obs}}}^{\infty} g(t | H_0) dt, \quad (14.9)$$

where $g(t | H_0)$ describes the probability density function for the test statistic under the hypothesis H_0 . To determine the probability density $g(t | H_0)$, so-called pseudo-experiments are generated. In these, random events are drawn from H_0 by creating Poisson distributed random numbers for every bin i from the background hypothesis b_i . This is repeated several times to create multiple pseudo-spectra, and for every case the test statistic is calculated. For a sufficient number of pseudo-experiments, the obtained distribution of the test statistic approximates the underlying distribution g . The p -value is then defined as the fractions of pseudo-experiments where the test statistic takes a value of at least t_{obs} or higher. This approach has the advantage that no closed-form expression for g is required. A BH p -value of less than 0.01 is required for an excess to be considered as significant. If no significant excess is present, a χ^2 p -value greater than 0.05 is required for a fit to be a valid background estimate in this thesis. An additional advantage of selecting the BH test statistic is that it reports a p -value which already accounts for the *look-elsewhere effect* [14, 178]. The look-elsewhere effect describes the issue when a large number of different hypothesis tests are performed simultaneously, and one or several of them result in small, significant p -values rejecting H_0 simply due to statistical fluctuations and the large parameter space tested.

14.1.5. Uncertainties on the Background Estimate

Due to the finite size of the data set, the best fit values of the parameters θ have associated statistical uncertainties themselves. As a consequence, the background estimate itself has an uncertainty, which is determined from pseudo-experiments. These are generated from the background estimate by drawing the number of events in every bin i from a Poisson distribution $\text{Pois}(\cdot | b_i(\hat{\theta}))$. Each generated pseudo-spectrum is then fitted using the same function and starting parameters that were used for the background estimate itself. For a sufficiently high number of pseudo-experiments, inspecting the distribution of the fit results allows the determination of a central 68.3 % confidence interval for the counts in every bin, which is taken as the uncertainty of the background estimate in the corresponding bin. In dijet analyses, this component is referred to as the *data uncertainty* of the background estimate.

In addition to this uncertainty, the background estimate has also a systematic uncertainty caused by the somewhat arbitrary choice of a fit function for the SWiFt procedure. Since multiple functional forms exist which provide potentially equivalent, but slightly deviating descriptions of the data, there is no clear best choice for the fit. The associated uncertainty is determined by selecting a *nominal* and an *alternate* fit function, which are both fitted to data. In each invariant mass bin, the difference between the fits is calculated in repeated pseudo-experiments, and the root-mean-square of the differences is calculated and taken as the $\pm 1\sigma$ uncertainty.

14.1.6. Potential Discovery and Systematic Uncertainties

In the case of the presence of a significant excess identified by BumpHunter, the corresponding BH window can be excluded from the spectrum. The statistical fit and the determination of the χ^2 p -value is then re-performed on the rest of the spectrum. Afterwards, the BH algorithm is re-run on the new background estimate including the previously excluded window. If the result is not significant anymore, the excess is identified as being caused by bad fit quality only. If however the BH test remains significant in the same mass region even after window exclusion, other possible causes for the observed excess besides new physics need to be considered. An example for such a cause can be unsmoothnesses in the jet calibration, as has been illustrated in Section 13.2.6. A strategy to handle these cases is by formulating possible causes for a fake signal as a systematic uncertainty on the background prediction, which is then applied in addition to the statistical error.

For this purpose, it is possible to run the BH algorithm also in the presence of systematic uncertainties, which is discussed for the JES uncertainties in the following. To include these, it is necessary to obtain the data spectrum in cases where the calibration is not performed at its nominal values, but instead each jet is shifted in the uncertainty of the JES uncertainty component j , for example by $\pm\sigma_j$. By this, a systematically shifted template spectrum $\mathbf{d}_j^{\text{template}}$ is formed. The pseudo-experiments used to determine the p -value are then modified by using the nominal background prediction as the baseline, and adding the difference between nominal and shifted data template, randomly varied using a Gaussian random number $\theta_j \sim \mathcal{N}(\mu = 0, \sigma^2 = 1)$:

$$\mathbf{d}_{\text{PE}} = \mathbf{b}^{\text{nominal}} + \theta_j(\mathbf{d} - \mathbf{d}_j^{\text{template}}). \quad (14.10)$$

In the presence of this additional contribution, unsmoothnesses in the calibration result in unsmoothnesses in the pseudo-experiments due to differences between nominal and shifted data spectrum. This then results in the BH spectrum being shifted to higher values, reducing the significance of the observed p -value. A real signal would not be affected by this procedure: in the presence of a real signal and a smooth calibration the variations should result in $\mathbf{d} \approx \mathbf{d}_j^{\text{template}}$, the BH test statistic in data would therefore be large compared to the BH statistic from pseudo-experiments, leading to a small p -value.

Unfortunately, this procedure cannot be applied for the histograms on L1Topo, as the calculation of $\mathbf{d}_j^{\text{template}}$ requires event-level knowledge of the jets forming the invariant mass. Therefore, it would have been required to implement these systematic template histograms as additional histograms on L1Topo. Due to this, if a peak was discovered in this analysis, the interpretation of its cause needs to be done carefully and this possible explanation has to be taken into account.

14.2. Performing the Background Estimate

Using the statistical framework introduced in Section 14.1, the fitting procedure can be applied to data. At first, it will only be applied to a small subset of the full data. While not guaranteed to be completely signal free, possible signal contributions are negligible if the considered data set is small enough, as otherwise the signal contribution would have been observed by one of the other analyses targeting the same mass range, compare Figure 9.4. Figure 14.2 shows a fit to an exemplary histogram recorded on L1Topo, corresponding to a single 50 s recording interval. The fit is restricted to the range up to 350 GeV to avoid the region with large influx from the forward region. The significances which are quoted in the lower half of the fit result plot are calculated




Figure 14.2.: Application of the SWiFt approach to a single histogram recorded on August 15, 2018 in run 358395 corresponding to a pile-up level of $\langle\mu\rangle = 49.1$.

using signed z -values as described in reference [179]. Already using a small subset of the data, the fit procedure fails to produce a reliable background estimate, as indicated by the χ^2 p -value close to zero. This is likely caused by the features of the spectrum, which is affected by contributions from coarse-granularity jets and pile-up effects. Therefore, it can be assumed that using the full recorded data directly in the fitting procedure is destined to fail. These effects are expected to be less severe in the signal region with central jets and $y^* < 0.3$, which is unfortunately not accessible to this analysis. Due to this, another way to mitigate parts of the pile-up effects has to be found. From the recorded data it is possible to obtain information on the pile-up behavior of the invariant mass spectrum by considering the $\langle\mu\rangle$ information associated with each histogram. In the next section, it is discussed how this information can help to improve the background estimation.

14.3. Invariant Mass in Bins of $\langle\mu\rangle$

In order to reduce the influence of pile-up effects on the invariant mass spectrum, the dependence of the spectrum on $\langle\mu\rangle$ is considered. Using the information available in the COMA database, compare Table 10.3, it is possible to associate each histogram with a $\langle\mu\rangle$ value. This can be used to create a two-dimensional histogram of the event counts depending on both the invariant mass and $\langle\mu\rangle$. For this purpose, a set of $m + 1$ bin edges $\mu_0, \mu_1, \dots, \mu_m$ is defined, and n different $\langle\mu\rangle$ filters are processed in parallel as described in Section 10.1.4. The resulting spectra in equidistant $\langle\mu\rangle$ bins with a bin width of five ranging from $\langle\mu\rangle = 0$ to $\langle\mu\rangle = 65$ are shown in Figure 14.3. Multiple observations can be made from these spectra: the event counts differ over several orders of magnitude between the various bins. In addition, the oscillations at 350 GeV and 600 GeV are more pronounced in bins covering higher values of $\langle\mu\rangle$. The large difference in event counts can

be explained by inspecting the pile-up profile shown in Figure 10.10, which shows that more data was recorded at some values of $\langle\mu\rangle$ than at others. The increased oscillation effects are likely caused by increased pile-up noise contributions to jet energies, which have a particularly strong effect on the forward jets. For a better comparison of the different bins, the pile-up profile can be reweighted to a flat spectrum, meaning that the number of events in each $\langle\mu\rangle$ bin is scaled to the same luminosity. The recorded luminosity L_i in a $\langle\mu\rangle$ bin i is calculated as the integral over the pile-up profile $\frac{dL}{d\mu}$, allowing for the determination of weights w_i for every $\langle\mu\rangle$ bin as

$$L_i = \int_{\mu_i}^{\mu_{i+1}} \frac{dL}{d\mu} d\mu \Rightarrow w_i = \frac{1}{L_i} \cdot \frac{L}{m}, \quad (14.11)$$

where $L = 61.56 \text{ fb}^{-1}$ is the total integrated luminosity of the recorded data and m the number of $\langle\mu\rangle$ bins. This way, each bin is scaled to the same luminosity, while keeping the total luminosity unchanged. By this construction, scaling differences caused purely by luminosity differences are removed, and only the dependencies on $\langle\mu\rangle$ itself are left. The resulting spectrum is shown in Figure 14.4. One observation is that the number of events shows a dependency on $\langle\mu\rangle$, with more events observed in bins with higher pile-up levels. In addition, this effect seems to have a strong mass dependence. While the effect seems to be linear and approximating to a constant value at low masses, the effect gets stronger and increasingly nonlinear in the regions above 350 GeV and 600 GeV. It is important to note that for a signal process there should be no such scaling behavior with $\langle\mu\rangle$. To check this, the simulated signal spectrum is inspected in bins of $\langle\mu\rangle$ in Figure 14.5 and in Figure 14.6 with a flat-weighted $\langle\mu\rangle$ profile. Due to the limited number of simulated events, the $\langle\mu\rangle$ bins are chosen larger than those for the data. These figures confirm that no such scaling behavior can be observed in the signal distributions, which is completely independent on $\langle\mu\rangle$ apart from luminosity differences.

In the following, it is explained how the combination of invariant mass and $\langle\mu\rangle$ information allows the quantification of parts of the pile-up effects influencing the mass spectrum in data. In general, two effects caused by pile-up are expected to affect the spectrum. One effect is an increase in energy density in the calorimeter with $\langle\mu\rangle$ due to an increase in soft background activity. The other effect is that sometimes objects from different collisions occurring at the same BC are visible in the detector simultaneously. If objects are selected naively, it can occur that jets from different




Figure 14.3.: Invariant mass spectra in all of 2018 in $\langle\mu\rangle$ bins.

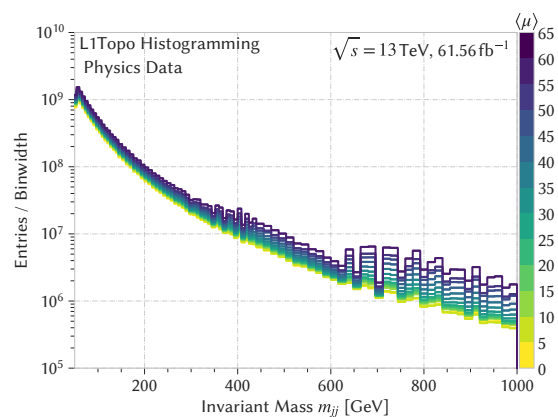


Figure 14.4.: Invariant mass spectra in $\langle\mu\rangle$ bins, each bin is reweighted to the same luminosity.




Figure 14.5.: Invariant mass spectra for a Z' signal with $m_{Z'} = 350$ GeV in $\langle\mu\rangle$ bins.




Figure 14.6.: Z' invariant mass spectra in $\langle\mu\rangle$ bins, normalized to the same luminosity.

collisions are combined to form the invariant mass. In offline analyses, various techniques exist to mitigate these effects. The first effect is mitigated by the pile-up correction step in the jet calibration scheme, which subtracts part of the jet energy. For the second effect, ATLAS analyses often apply a jet vertex tagging [180], which assigns a probability to each jet that it belongs to the primary vertex. The so-called jet vertex fraction is defined as the ratio of the p_T sum of tracks associated to a jet which originate from the primary vertex to the p_T sum of all tracks associated to the jet. Jets which likely originate from the hard scattering take on high values of this ratio, while jets from additional pile-up vertices take on lower values and are rejected by a cut on the jet vertex fraction. Since this cannot be applied on L1Topo, the effect of combining jets from multiple interactions is present in the data. The idea in the following is that this second component can be mitigated by using the $\langle\mu\rangle$ dependency of the spectrum. To illustrate this, Figure 14.7 shows two possible cases of events where two jets are selected as the dijet pair. In one case, jets from the same collision form the dijet pair, while in the other case jets from different, simultaneous interactions form the dijet pair and hence the invariant mass, even though they do not result from the same proton-proton collision. For the intuitive description consider the case of having a well-defined hard scattering at a primary vertex, with some additional collisions occurring at other positions in the detector. For each of the additional vertices, an identical probability of




Figure 14.7.: Simplified depiction of different combination of jets from pile-up events to form dijet pairs. Blue dots indicate vertices of different collisions, occurring in the cylindrical detector at the same time. Lines indicate particle tracks from these vertices, reconstructed jets are indicated by colored cones. Jet pairs can be found combining jets from the same collision (left), or jets from different collisions (right).

contributing an additional jet to the event is assumed on average. As the total rate for each of the two additional jets should scale proportional to μ , the rate of combinations of jets from different collisions should scale proportional to μ^2 . At the same time, the signal contribution is expected to be independent on μ , as was observed in Figure 14.6. In the following, these considerations are used to propose a simple model for the μ dependence in every mass bin. Independent functions of the form

$$N(\mu, \mathbf{p}) = p_0 + p_1(\mu - 1)^2 \quad (14.12)$$

are fitted in every mass bin. The function is designed to contain a hard-scatter term independent of μ described by the parameter p_0 , and a quadratic term scaling with p_1 describing the combinatorial effects, which also allows for a component scaling proportional to μ . As the presence of only one interaction at $\mu = 1$ means that no combination of objects is possible, the quadratic term is designed such that it is equal to zero for $\mu = 1$. For the fits, a non-equidistant μ binning is chosen. From Figure 10.10 it is evident that most data was recorded at high values of μ between 20 and 50, while only little data is recorded for $\mu < 10$. However, the information at small μ is of particular importance to identify the constant contribution independent of μ . Due to this, logarithmic-spaced bins are used for $\mu > 10$, putting more emphasis on the data at smaller values of $\langle\mu\rangle$. All information available below 10 is combined into a single bin due to the limited amount of events available in this region. The luminosity normalization procedure as described in Equation (14.11) is applied also to these bins. Detailed information on the chosen bin edges, centers and the associated luminosities are listed in Appendix D.3. Exemplary results of the fits are shown in Figure 14.8 for four different mass bins. In general, a good quality of fits can be observed, as the distribution of reduced χ^2 values for all fits range between 0.4 and 1.6. In many mass bins, the $\langle\mu\rangle$ bins around 20 show slight overshoots, while the lower $\langle\mu\rangle$ bins show small undershoots, which are however compatible with the fit within the uncertainties. Comparably high uncertainties are visible at low values of $\langle\mu\rangle$, independent on the mass. Due to the weighting procedure described in Equation (14.11), these bins, which contain only a small data set, receive large weights as a result of the small associated luminosities, resulting in large statistical uncertainties. The uncertainties of the fitted parameters are mainly driven by the low luminosities available at low values of $\langle\mu\rangle$. An idea to mitigate this effect was to supply this region with additional data from dedicated low- μ runs at ATLAS. However, it was found that the histograms recorded in these runs are not compatible with the data from normal physics runs, as the PPM noise settings are very different between those two runs [123].

Using the fits allows for an extrapolation of the $\langle\mu\rangle$ dependency by evaluating the fit in every bin for $\langle\mu\rangle \rightarrow 1$. This way, a new invariant mass spectrum is formed, where pile-up dependencies due to jet combinations are reduced. At the same time, a possible signal contribution should be unaltered by this procedure. The uncertainties of the new spectrum are given by the uncertainties of the fits for $\langle\mu\rangle = 1$, which are calculated from the covariance matrix of each fit with the method described in Appendix D.1. The result of the extrapolation compared to the original data spectrum is shown in Figure 14.9, together with their ratio and an indication of the size of the relative uncertainties of the obtained spectrum. The ratio between the spectra is at 0.9 or lower for all bins, and it drops with increasing masses. In addition, an oscillating behavior of the ratio is visible at higher invariant masses, where pile-up effects in forward jets play a significant role. Due to the uncertainties related to the fits, the resulting spectrum has uncertainties which do not correspond to Poisson uncertainties. The relative uncertainties range from 0.2% at low masses up to around 0.8% for higher masses at around 400 GeV. Due to these, it is not a priori clear whether the possible sensitivity of this analysis is comparable to other searches. In the following, the

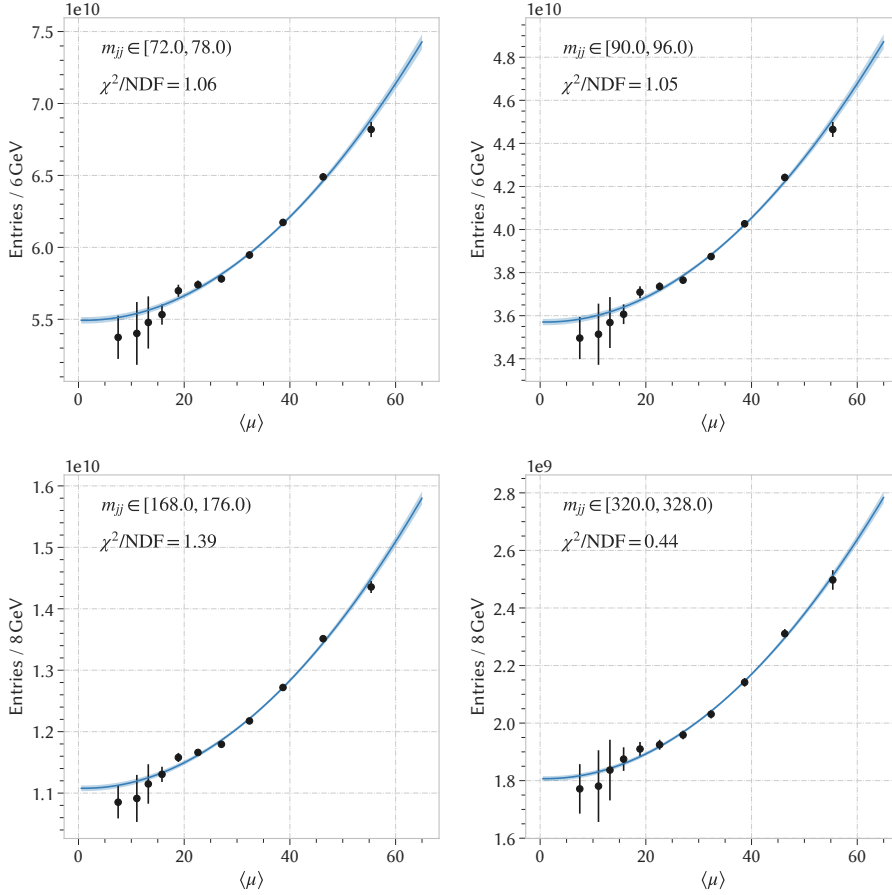


Figure 14.8.: Examples of fits of the $\langle\mu\rangle$ dependency in four invariant mass bins.

background estimation will be applied to the extrapolated spectrum shown in Figure 14.9.

14.4. Setup of the Background-only Fit

The background estimation is run again, this time with the newly derived spectrum. Since the uncertainties of the spectrum do not correspond anymore to Poisson uncertainties, a reweighting procedure is performed and a weighted maximum likelihood fit is applied, compare Appendix D.4. Due to the large number of degrees of freedom in the background parametrization, reliable fit convergence requires reasonable initial parameter values. These are found by first adapting the parameters by hand until the shape of the spectrum looks roughly comparable to the data spectrum. Using these parameters, multiple preliminary fits of the data in the SWiFt window are performed using slight randomizations of the initial parameter guesses. The parameters of the best preliminary fit, which is defined as the one resulting in the smallest χ^2 , are then used as starting parameters for the minimization of the likelihood. In all cases f_5 is chosen as the nominal fit function and either $f_{\text{UA}2}$ or f_{γ} is chosen as the alternate fitting function. Before the fit can be performed, some hyperparameters need to be fixed, namely the fit range and the SWiFt window width. In order to avoid optimizations of hyperparameters on the data itself, it is a common strategy to set up the fitting strategy on a small subset of the data or on high statistics

MC samples. Optimizing the analysis procedure on the full data itself brings the risk of artificially tuning the analysis to make statistical fluctuations look more pronounced, driven by the desire to find an excess in data. By avoiding this, the optimization of the analysis is blind to the actual data itself, and only once the procedure is optimized and fixed the analysis is unblinded and the fit applied to data. For this particular analysis, it is non-trivial to find suitable descriptions of the data spectrum for the optimizations. It is not possible to simply use a small subset of the data, since the invariant mass spectrum was formed from a $\langle\mu\rangle$ -extrapolation procedure and only a small data set is available in the low $\langle\mu\rangle$ region, which is important to fix the constant contribution to the extrapolation fits. As an alternative approach, the MC samples generated with the SLW simulation as introduced in Section 13.2 are used for this purpose. Even though their shape is different from the shape of the extrapolated data spectrum, it is considered as similar enough to tune the fitting procedure, as it still inhibits the same feature of a level-1 invariant mass spectrum. In order to make the results more comparable, the MC spectrum is scaled to the same number of events as the extrapolated spectrum. In addition, the uncertainties are scaled such that the relative uncertainties correspond to those of the extrapolated spectrum. Using MC has the additional advantage that the distribution is guaranteed to be signal-free, avoiding any bias of the optimization due to a possible excess.

14.4.1. Selection of the Fit Range

A fundamental ingredient of the fitting procedure is the mass range used for the fit. Limitations exist for both lower and upper boundary of the fit range. At the lower edge of the spectrum, the level-1 jet threshold of 12 GeV is responsible for a turn-on effect, making this region unusable for a fit. At higher values of the invariant mass, oscillations due to the limited level-1 jet granularity




Figure 14.9.: Extrapolated histogram compared to the original data spectrum. The ratio of this histogram to data is plotted as well as the relative uncertainties on the extrapolation.

at higher $|\eta|$ get increasingly pronounced, limiting the ability to find a suitable background description. For the lower edge, the low mass region of the invariant mass spectrum is fitted in a fine granularity binning with the product of the function f_5 described in Equation (14.3b) and a turn-on sigmoid-like function s :

$$s(m_{jj}; \alpha, \beta) = \frac{1}{1 + \exp\left(-\frac{m_{jj}-\alpha}{\beta}\right)} \quad (14.13)$$

The lower edge of the fit is then defined as the value where the turn-on saturates and its deviation from 1 is less than half of the uncertainties associated with the bin. As the uncertainties reach down to 0.2 %, it was decided to require that the turn-on deviates less than 0.1 % from 1:

$$1 - s(m_{jj}^{\text{lower}}; \alpha, \beta) \stackrel{!}{=} 0.001. \quad (14.14)$$

The result of the fit is shown in Figure 14.10 and results in $m_{jj}^{\text{lower}} = 53.88$ GeV. The next higher bin edge is given as 54 GeV, which is therefore chosen as the lowest bin for the fit. For the higher edge, the goal is to use a range as wide as possible for the fit. If a signal is present, having sidebands without or with only small signal contributions help to constrain the background description. Since signals in the range of interest have widths of $\sigma_{m_{jj}} / \langle m_{jj} \rangle \approx 0.2$, compare Figure 11.6, the distribution can extend across many histogram bins. At the same time, the coarse granularity of forward jets leads to larger fluctuations of the event counts at higher invariant masses, limiting the ability to find a smooth background description. Therefore, the upper bin limit is determined as the largest possible bin where it is still possible to find a reliable background estimate. For this, SWiFt fits of the invariant mass spectrum are performed in a range $[m_{jj}^{\text{lower}}, m_{jj}]$, where the upper boundary is varied between the fits. For each possible upper boundary, the χ^2 p -value of the fit is calculated, the resulting p -values are shown in Figure 14.11. A rapid drop-off can be observed for $m_{jj} > 352$ GeV. Above this value, χ^2 p -values of less than 0.05 are observed, indicating a poor fit which is not able to describe the spectrum. Therefore, the upper limit of the fitting procedure is set to $m_{jj}^{\text{upper}} = 352$ GeV. This result was found to be consistent using various window widths.





Figure 14.10.: Result of a turn-on fit to the low-mass region. The normalization of the turn-on function is scaled up to improve visibility.

Figure 14.11.: χ^2 p -values as a function of the uppermost bin used for the fitting procedure.

14.4.2. Selection of the Window Width

After the fit range is fixed, the SWiFt window width needs to be optimized. For this purpose, the effect of the window width on the fit quality is inspected both without and with the artificial injection of a signal resonance. Figure 14.12 shows two exemplary fits, using window half-widths² of 5 and 14, corresponding to window sizes of 11 and 29. Both choices provide a sufficient background estimate, the χ^2 p -value shows a better correspondence between background-only estimate and data for the smaller window width. At the same time, the BH p -value is larger for the larger window width, giving more confidence that no signal is present. To quantify this effect, Figure 14.13 shows the dependency of the χ^2 p -value on the window half-width. While it is found that smaller windows provide better background estimates, only little dependencies can be seen for window half-widths of eight or larger. To supplement these results, an artificial signal can be injected on top of the spectrum to check how the window width influences the discovery of a signal. For this purpose, generic Gaussian-shaped signals at a mass of 200 GeV are inserted for relative widths of 10 % and 15 %. Signals are inserted with a number of events corresponding to roughly 2.5 times the statistical uncertainties in the region 120 GeV to 280 GeV around the signal peak. Figure 14.14 shows the dependence of the BH p -value on the window half-width for both signal widths, Figure 14.15 shows four exemplary fit results. For the narrow signal width of 10 %, both small and large window half-width are able to find a correct BH interval centered around the signal mass of 200 GeV. However, only the fits using larger window widths are able to discover a significant excess, as indicated by a small BH p -value < 0.01 . Due to the larger sidebands, the larger window width is better able to constrain the background prediction. For the broader signal, all window widths fail to discover the signal and are not able to identify the correct region. The observations confirm the initial claims on the window width, as the sensitivity of this search to signals increases with a higher window width, while at the same time the quality of the background-only fit drops with increasing window width. An additional problem with larger windows is that due to the small available fit range with a limited number of bins, larger widths

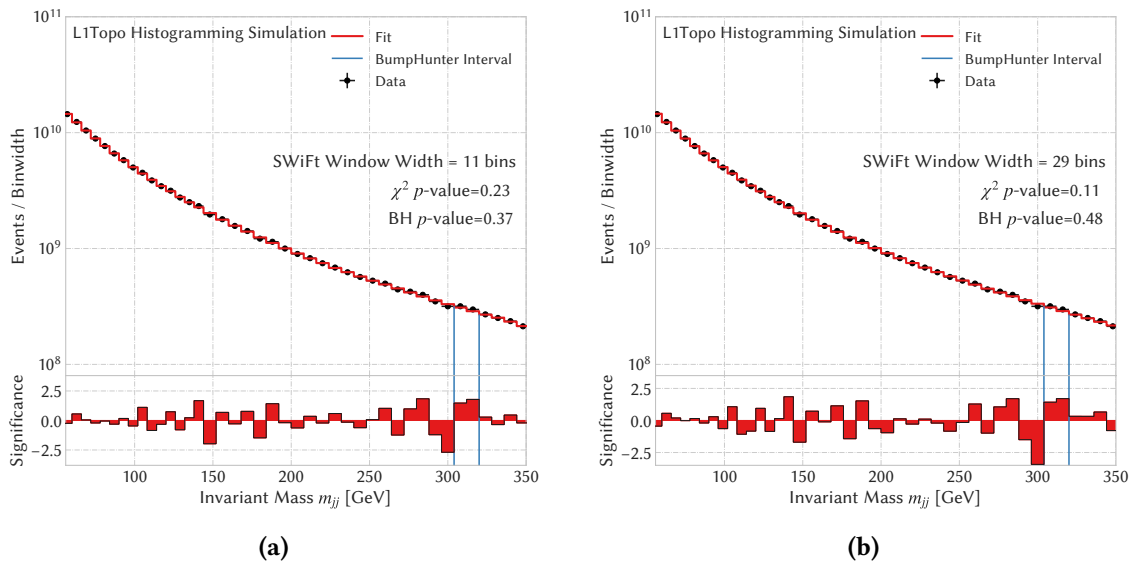


Figure 14.12.: SWiFt fits with window half-widths of 5 bins in (a) and 14 bins in (b).

²The window half-width corresponds to the number of bins in the window on either side of the central bin.




Figure 14.13.: χ^2 p -value as a function of the window half-width for background-only MC.




Figure 14.14.: BH p -value as a function of the window half-width for Gaussian signals.

will at some point range over almost the whole spectrum, diminishing the positive effects of the SWiFit procedure. One observation made during these studies was that smaller run times of the fitting procedure were observed for small window widths, indicating that the fitting procedure converges faster in these cases. Due to this, it was decided to choose a high window half-width of 17 as a baseline, which will be reduced subsequently when the background estimate fails.

14.4.3. Summary of the Fit Strategy

In the following, the fitting procedure is summarized before the fit is applied to the extrapolated spectrum. The chosen strategy largely follows the TLA strategy as described in reference [14], using a fit range of [54 GeV, 352 GeV]. The procedure is then performed as follows using f_5 as the nominal function, f_γ as the alternate function, and a starting window half-width of 17:

1. Perform the SWiFit fitting procedure with the given settings
 - a) If the fit succeeds, go to step 2
 - b) If the fit fails, repeat step 1 and reduce the window half-width by 1
2. Calculate the χ^2 and BH p -values
 - a) If the fit quality is insufficient (χ^2 p -value < 0.05), go to step 3
 - b) If the fit quality is good (χ^2 p -value > 0.05), go to step 4
3. Check whether the BH p -value shows an excess indicated by $p < 0.01$
 - a) If yes, restart the fit and exclude the BH window of the most significant excess
 - b) If no, restart the fit and reduce the window half-width by 1
4. Check whether the BH p -value shows an excess indicated by $p < 0.01$
 - a) If yes and no window has been excluded so far, remove the window of the most significant excess and restart at step 1
 - b) If yes and a BH window has been excluded already, consider the excess significant.
 - c) If no, stop the procedure and consider no excess as significant

If the conclusion of the procedure is that no significant excess is found, exclusion limits will be set using the performed background estimate. If however the conclusion of the procedure is that an excess has been found, its interpretation needs to be done with care, as mentioned in Section 14.1.6. For this, it would be necessary to repeat the fitting procedure including JES systematic uncertainties. Only if the excess still persists with these additional uncertainties, it can be considered as a potential discovery. In the case of this analysis, where the JES uncertainties cannot be applied, an excess at this point can therefore only be seen as a hint for other, supplementary analyses to explore the phase space corresponding to the discovered excess, rather than a potential discovery.

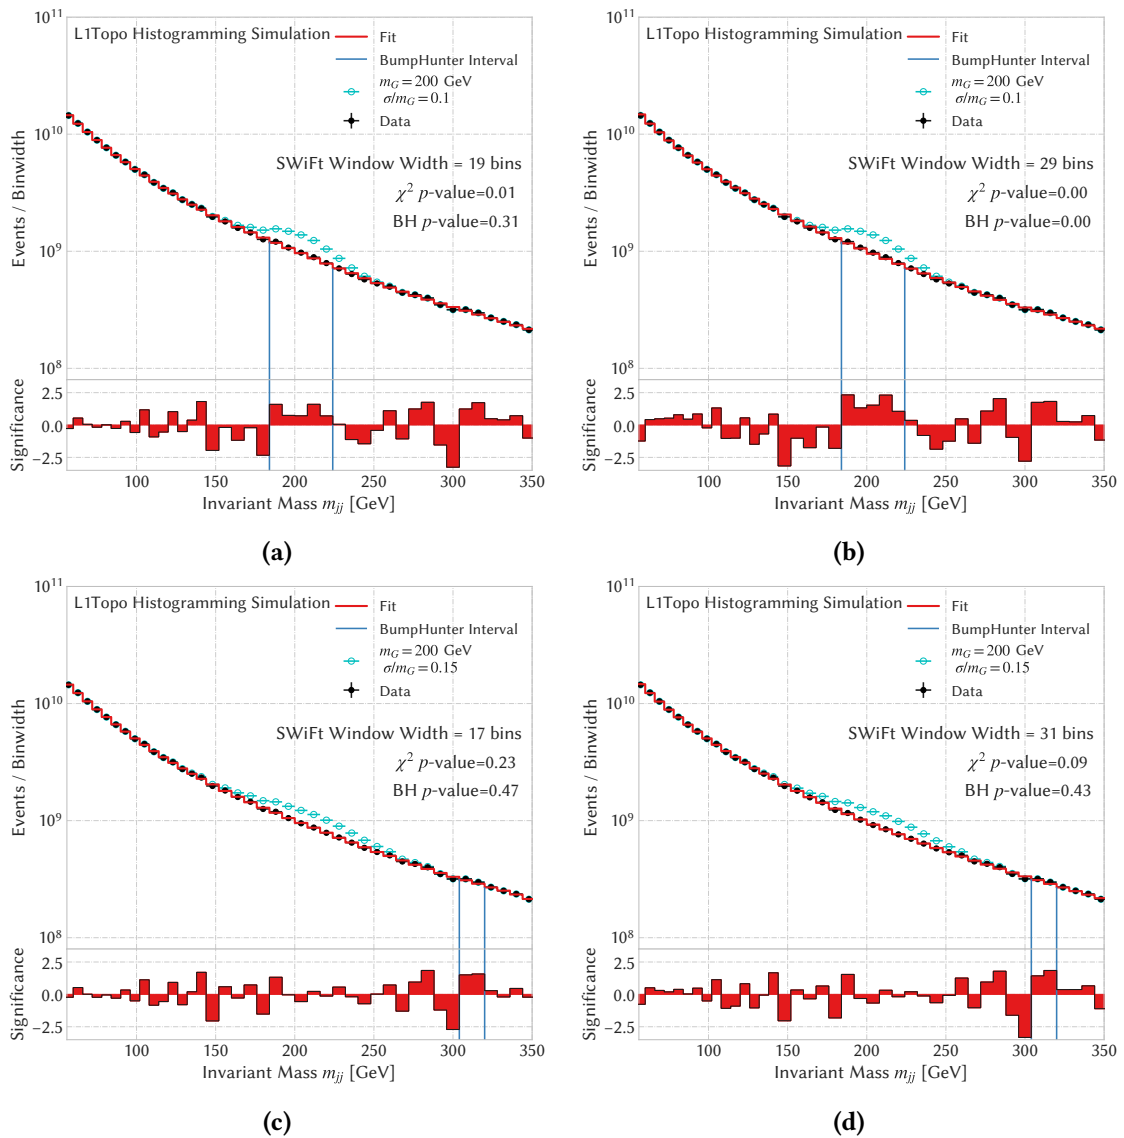


Figure 14.15.: SWiFit fit for different window widths of a spectrum with Gaussian signals at a mass of 200 GeV of widths of 10 % (top row) and 15 % (bottom row). All shown signals are scaled by a factor $1.5 \cdot 10^2$ to be visible.

14.5. Results: Performing the Fit to Data

The fit is applied to the $\langle\mu\rangle$ -extrapolated spectrum as shown in Figure 14.9. After the fitting procedure as described in Section 14.4.3 concluded, a window half-width of 13 was used. Figure 14.16(a) shows the result of the SWiFt fit using the nominal function f_5 and Figure 14.16(b) shows the Poisson p -values for every possible BH interval. The agreement between fit and data is given by a χ^2 of 45.5, corresponding to a p -value of 0.18. The distribution of pseudo-experiments used for the determination of the p -value is shown in Figure 14.17. No significant excess can be observed, with the most significant window ranging from 304 GeV to 344 GeV with a Poisson probability of $t_{\min} = 7.23 \cdot 10^{-3}$ and a BH value of $t_0 = -\log t_{\min} = 4.93$. The corresponding p -value is given as 0.36, and has been determined from pseudo-experiments which are shown in Figure 14.18. A good agreement between fit and data is observed, as the probability of obtaining a result deviating more from the background estimate purely by chance is given by 18 %. Under the background-only hypothesis, the probability of finding an excess which is at least as large as the one observed in data is given as 36 %. Using the nominal and this alternate function, the uncertainties of the background estimate are obtained and shown in Figure 14.19, the background estimate obtained using the alternate function f_y can be found in Appendix D.5. The total fit uncertainty ranges from about 0.2 % to 0.4 %, with larger uncertainties observed at the edges of the fit window. Both components of the uncertainties are shown and contribute to the total error, the uncertainty of the fit due to the data uncertainties is dominant compared to the function choice uncertainty. The higher uncertainties on the edges are expected, since the background estimate in these bins cannot be obtained with these bins as bins centers, compare Figure 14.1. As a summary, it can be stated that a background estimate has been obtained successfully. However, no significant excess has been found. Therefore, the next section focuses on setting exclusion limits.

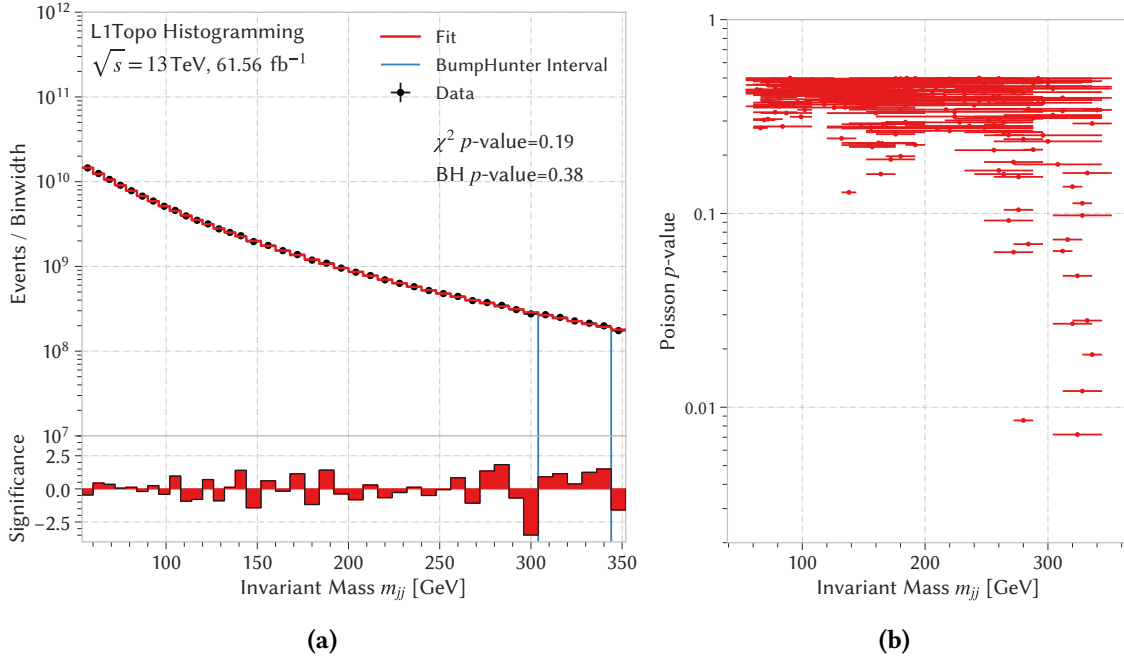


Figure 14.16.: Fit result of the SWiFt fit to the data (a) as well as the distribution of all possible Poisson p -values used for the BH hypothesis test (b).




Figure 14.17.: Distribution of the χ^2 test statistic from 5000 pseudo-experiments (blue) with the value observed in data (red arrow).




Figure 14.18.: Distribution of the BH test statistic from 5000 pseudo-experiments (blue) with the value observed in data (red arrow).

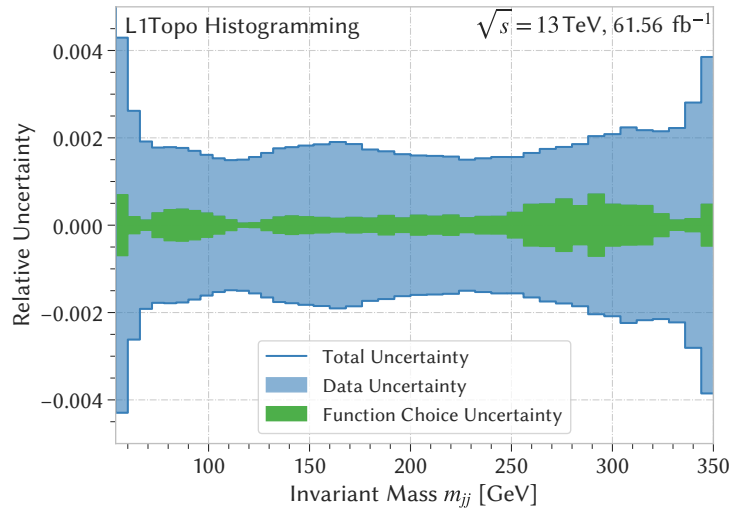


Figure 14.19.: Relative uncertainties of the background estimate, showing both components, namely the data uncertainty (blue) and the function choice uncertainty (green).

Setting Exclusion Limits

The previous chapter showed that no significant deviation is found in the invariant mass spectrum. Therefore, the discussion in the following is focused on setting *exclusion limits*. Exclusion limits indicate which signal cross sections can be excluded by the given data at a certain level of confidence, usually chosen as 95%. This chapter starts with a brief summary of the Bayesian framework for limit settings used in this thesis in Section 15.1. This framework has been developed by ATLAS for dijet searches in Run-2, detailed descriptions can be found in references [55, 175]. The section starts with a short overview of the applied statistical concepts in Sections 15.1.1 and 15.1.2 and continues with a brief discussion of sources of systematic uncertainties and their inclusion into a limit setting procedure in Section 15.1.3. Section 15.1.4 introduces the likelihood used for the limit setting and describes the implementation in the Bayesian framework using so-called Markov Chain Monte Carlo methods. After this, the limit setting procedure is applied in Section 15.2 for two different cases: model-independent limits using generic Gaussian signal shapes are set in Section 15.2.1, and limits for a Z' dark matter mediator model are set in Section 15.2.2. The latter allows for a direct comparison with the results of other searches, which is performed in Section 15.3.

15.1. Bayesian Limits Framework

15.1.1. Theory of Bayesian Limit Setting

In the Bayesian framework, limits are calculated using *Bayes theorem*. In its most generic formulation, Bayes theorem describes how the conditional probabilities of two events A and B are related and can be transformed into each other [181]. Let $P(A|B)$ denote the conditional probability of event A occurring given that B is true. Then, Bayes theorem states the following relation between $P(A|B)$ and $P(B|A)$:

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}, \quad (15.1)$$

where $P(A)$ and $P(B)$ denote the probabilities of observing the events A and B , respectively. In physics, this is applied to the scenario where A represents a hypothesis, which can possibly depend on several parameters, and B represents the observed data. In this scenario, $P(A|B)$ is the probability assigned to the hypothesis after inspecting the data, and is called *posterior probability*. This probability is the quantity of interest of the statistical analysis, as it allows a direct statement of the probability for a particular hypothesis after data has been observed. Analogously, $P(B|A)$ is the probability of the data given that the theory is true. When regarded as a function of A for fixed data B , this is equivalent to the likelihood $\mathcal{L}(A|B)$, which can be directly formulated as soon

as the data is available. $P(A)$ is the probability assigned to the hypothesis before the experiment is done, the so-called *prior probability*, and $P(B)$ is chosen such that the correct normalization of the probability is ensured. Often, the hypothesis depends on multiple parameters of which some are parameters of interest, for example a signal scaling parameter μ , and some additional parameters \mathbf{v} , which are not of immediate interest, but have to be included in the analysis. The latter are commonly referred to as *nuisance parameters*. The prior probability for the hypothesis parameters is then denoted as $\pi(\mu, \mathbf{v})$ and reflects possible prior knowledge about their values. There is no strict rule on how priors should be chosen, the actual choice of priors is discussed briefly in Section 15.1.2. While the signal scaling is an intrinsic property of the theory behind the hypothesis itself, nuisance parameters often describe systematic uncertainties of the measurement process, for example due to the limited knowledge of the detector resolution. Therefore, the parameter of interest μ is often statistically independent of all nuisance parameters. In addition, the nuisance parameters are usually constructed in a way to be independent of each other, as this allows to factorize the prior as $\pi(\mu, \mathbf{v}) = \pi(\mu) \prod_i \pi(v_i)$. Using these identifications, Bayes theorem allows to express the posterior probability of the hypothesis parameters $P(\mu, \mathbf{v} | \text{data})$ in dependence on the priors and the likelihood $\mathcal{L}(\mu, \mathbf{v} | \text{data})$ as follows:

$$P(\mu, \mathbf{v} | \text{data}) = \frac{\mathcal{L}(\mu, \mathbf{v} | \text{data}) \pi(\mu, \mathbf{v})}{\int d\mu' d\mathbf{v}' \mathcal{L}(\mu', \mathbf{v}' | \text{data}) \pi(\mu', \mathbf{v}')}, \quad (15.2)$$

where the denominator acts as a normalization constant. For the limit setting itself, only the posterior distribution of the signal scaling parameter μ itself is of interest. In order to remove the dependency on the nuisance parameters, the *marginal* posterior distribution $P(\mu | \text{data})$ is considered, where the integral over the nuisance parameters is performed:

$$P(\mu | \text{data}) = \int d\mathbf{v} P(\mu, \mathbf{v} | \text{data}) \propto \int d\mathbf{v} \mathcal{L}(\mu, \mathbf{v} | \text{data}) \pi(\mu) \prod_i \pi(v_i), \quad (15.3)$$

Equation (15.3) is the version of Bayes theorem which is used for the limit setting procedure, the integral is also-called *marginalization*. By checking the 95 % quantile of the posterior probability distribution for the parameter μ , the upper limit μ_{upper} can be set from the marginal posterior as

$$0.95 \stackrel{!}{=} \int_{-\infty}^{\mu_{\text{upper}}} d\mu P(\mu | \text{data}). \quad (15.4)$$

Due to the chosen Bayesian approach, Equation (15.4) allows the direct interpretation that the probability of μ being smaller than μ_{upper} is given by 95 %. In addition to μ_{upper} , the *observed* upper limit, so-called *expected* upper limits are also reported frequently. The expected upper limits indicate the expectations on an upper limit purely due to statistical fluctuations from a background-only hypothesis. They are obtained by setting all nuisance parameters to their maximum likelihood values and the signal strength to 0. From the corresponding spectrum pseudo-experiments are drawn, and the upper limits are calculated using the same marginalization procedure as for the observed limit. By inspecting the quantiles of the resulting limit distribution, a central value and bands corresponding to Gaussian 1 and 2σ intervals are determined.

15.1.2. Choice of the Priors

In general, there is no a priori rule which describes how a prior should be chosen. Different approaches exist, either choosing an informative prior, which incorporates knowledge from

previous analyses, for example by taking the posterior distribution from a previous iteration of the analysis, or choosing an uninformative prior, which is constructed to contain as little information as possible. In this case, an uninformative, flat prior is chosen for the signal scaling parameter μ , consistent with previous dijet analyses. As no less than 0 signal events are expected, the prior is cut off below 0. In order to make the prior normalizable, another cut-off is defined at a sufficiently large, but finite number of signal events μ_{\max} .

$$\pi(\mu) = \begin{cases} 1/\mu_{\max}, & \text{if } \mu \text{ in } [0, \mu_{\max}] \\ 0, & \text{otherwise,} \end{cases} \quad (15.5)$$

where μ_{\max} is defined as the value of μ where the likelihood is 10^5 times smaller than its maximum value for fixed nuisance parameters. For the nuisance parameters describing systematic uncertainties, Gaussian priors are commonly assigned, with a standard deviation σ_{v_i} equal to the uncertainty associated to the nuisance parameter v_i . The systematic uncertainties are discussed in more detail in the next section.

15.1.3. Treatment of Systematics

In the limit setting, different kinds of systematic uncertainties are possible. There can be systematic uncertainties on the obtained background estimate, and systematic uncertainties on the MC signal distribution. For dijet resonance searches, the majority of systematics uncertainties are given by components of the JES uncertainty, both on data and signal. However, an application of these is not possible for the data recorded in this thesis, as has been explained before. Due to this, no extensive analysis of systematic uncertainties is performed in the following. As this analysis is expected to be limited by uncertainties on the background estimation due to the uncertainties of the $\langle\mu\rangle$ extrapolation method, this is not seen as a showstopper. Since other systematic uncertainties could still be relevant in a future analysis of new data, where one of the designated signal regions is available, the general treatment of systematics in the limit setting is explained in the following. A few systematic uncertainties that are easy to implement are applied in the limit setting as an illustration, namely the data uncertainty, the function choice uncertainty and the luminosity uncertainty.

Uncertainties: Background Estimate

In order to include the uncertainty on the background estimate due to the finite size of the data set, the derived uncertainties from Section 14.1.5 are defined as the $\pm 1\sigma$ uncertainty in the background estimate. Let b_i denote the nominal background estimate in bin i , and σ_{b_i} the uncertainty. The assumption is now that the fit result varies linearly, i.e. the influence of the systematic is given as $b_i + v\sigma_{b_i}$, with $v \sim \mathcal{N}(0, 1)$.

To include the fit function choice uncertainty in the Bayesian setting, the relative difference between nominal and alternate fit result is calculated for every bin, and a nuisance parameter $v_{\text{choice}} \sim \mathcal{N}(0, 1)$ is added which ranges from 0 (only nominal fit) to 1 (only alternate fit):

$$b = f_{\text{nominal}} \left(1 + v_{\text{choice}} \frac{f_{\text{alt}} - f_{\text{nominal}}}{f_{\text{nominal}}} \right)$$

This way, the background estimate can be varied between the two obtained background estimates, but not beyond that range, with a slight builtin preference for the nominal function.

In addition to these, JES uncertainties influence the shape of the data spectrum itself, which in turn affects the background estimate. These uncertainties would be incorporated by considering the different possible sources of JES uncertainties, and creating variations of the invariant mass spectrum for each nuisance parameter in small steps, for example in the range between $\pm 0.5\sigma$ and $\pm 3\sigma$. The effect of each nuisance parameter is assumed to be linear between these. However, an application in data recorded from L1Topo would require an explicit implementation of every variation for every nuisance parameter as an additional histogram, making an implementation of this approach challenging in reality.

Other uncertainties could be considered for this analysis, for example due to the angular resolution of level-1 jets. Jet positions are only known at the coarse level of jet elements, however the true center of the jet could be anywhere inside a given jet element. Incorporating such an uncertainty could be done by shifting jets in η and ϕ according to their uncertainty in their spatial resolution to obtain a systematically shifted spectrum. The difference between the nominal and the shifted spectrum would then be a measure for the uncertainty, and could accommodate for effects caused by the coarse jet granularity. Again, this would require the firmware implementation of additional histograms. Furthermore, this would require internal data formats allowing for granularities finer than 0.1 in η , which is not feasible due to limited FPGA resources, as the jet data format is shared with the RTDP algorithms, whose resource usage would grow significantly.

Uncertainties: Signal

As the integrated luminosity of the recorded data set enters the limit setting as a parameter of the total number of signal events $n = \sigma \times A \times \text{BR} \times L$, an uncertainty on the luminosity will have an effect on the number of signal events. Two kinds of luminosity uncertainties are considered in the following. Firstly, the luminosity measurement performed by ATLAS has an assigned uncertainty [106]. For Run-2, the total luminosity uncertainty is given as $\sigma_{L,1} = 1.7\%$. The dominant contribution to this uncertainty comes from nonlinearities in the calibration transfer, which describes the extrapolation of the absolute luminosity calibration of LUCID obtained in Van der Meer scans with special low- μ conditions to actual physics running conditions at high μ . In addition, the luminosity obtained in this thesis is slightly higher than the official one, due to the way the luminosity obtained from the COMA database is calculated. The deviation between the luminosity obtained for histograms directly from COMA and for recorded data using that same GRL is taken from Figure 10.9. This results in a deviation of $\sigma_{L,2} = 1 - 58.5/61.5 = 5.0\%$. Due to this, the luminosity 1σ uncertainty is estimated as the combination of both contributions as $\sigma_L = \sqrt{\sigma_{L,1}^2 + \sigma_{L,2}^2} = 5.3\%$.

In addition, the JES uncertainties would also have an effect on the signal distribution. For the signal, which is obtained from MC, it would be easier to apply such an uncertainty as the four-vectors of the level-1 jets are available, allowing to implement the corresponding templates for each variation. It is important to note that this would not only change the signal normalization, but also the shape of the signal distribution itself. For Gaussian signals, which are not obtained from MC, the distributions would be simply shifted in mass and resolution to mimic the effect of JES uncertainties.

Possible additional sources of signal uncertainties are given by the choice of a PDF set for the signal generation as well as a particular MC generator, which are relevant for the produced signal. These could be obtained by checking the variation of the signal distribution for different PDF sets and different MC generators, producing a nominal and an alternate spectrum for each of those. However, they are expected to be negligible compared to JES uncertainties.

15.1.4. Implementation of the Limit Setting

Since the marginalization (15.3) to obtain the posterior distribution for the signal scaling parameter $P(\mu | \text{data})$ can be a highly multidimensional integral, it is often very difficult to perform. In particular in cases of complex likelihoods, where it is not possible to find an analytic expression for the posterior probability, numerical optimization tools have to be applied. In the dijet Bayesian framework, the Bayesian Analysis Toolkit (BAT) is utilized for this purpose [182]. BAT is an analysis package based on Bayes theorem, which determines the posterior distribution and performs all desired marginalizations. For this, it has to be supplied with a user-defined likelihood and prior distributions for all parameters. The main challenge is therefore a suitable implementation of the likelihood, which is given as follows using the background estimate obtained in the background-only fit:

$$L(\mu, \boldsymbol{\nu}) = \prod_{i=1}^n \text{Pois} \left(d_i \mid \mu s_i + b_i + \sum_{\text{background NPs } j} v_j (b_{ji} - b_i) + \sum_{\text{signal NPs } j} v_j (s_{ji} - s_i) \right), \quad (15.6)$$

where b_{ji}, s_{ji} describe the $\pm 1\sigma$ shifts in bin i associated with the nuisance parameter v_j and uncertainty σ_{ji} . When calculating the likelihood, it is important to apply the different systematic components in the correct order, as some of them only change the normalization, but some also change the shape of the distribution. Due to this, the signal distribution is first scaled to unity, so that μ corresponds directly to the number of signal events. All template-based systematics, for example the different JES components on the signal, are also scaled by the initial normalization of the signal template. For fixed values of $(\mu, \boldsymbol{\nu})$, each template systematic contribution is then added, as these can alter the shape of the distribution. Afterwards, the signal normalization is applied and finally systematics affecting the normalization are applied, in this case only the luminosity uncertainty on the signal. For the background estimate, no scale-only uncertainties are present, the different contributions are therefore applied one-by-one on top of the background estimate. Afterwards, signal and background are added together to obtain the expected number of events $d_i(\mu, \boldsymbol{\nu})$, allowing to calculate the value of the likelihood for the given set of parameters $(\mu, \boldsymbol{\nu})$. The formulation of the likelihood now allows the determination of the posterior and its marginalized distributions. BAT uses the Markov Chain Monte Carlo (MCMC) technique for this purpose. MCMC is used to construct and sample from arbitrary complex posterior distributions, in particular in cases where it is either impossible or computationally challenging to sample directly from the posterior [183]. It is based on a random walk in the parameter space by sequentially drawing values of $(\mu, \boldsymbol{\nu})$ from approximate distributions. Each draw is then corrected to make the following draw correspond better to the target distribution until convergence is achieved. Samples drawn after this procedure converges can be used to sample directly from the posterior distribution. Details on the functionality of MCMC is given in Appendix D.6.

15.2. Results: Setting Limits on Data

In the following, limits are derived for two cases: generic, model-independent limits are derived in Section 15.2.1, and limits for a Z' dark matter mediator model are set in Section 15.2.2. Since only events that pass the analysis reconstruction can be considered in the limit setting, and since only dijet final states are considered, limits are set on the signal production cross section σ , multiplied with the acceptance A and the branching ration BR.

15.2.1. Generic Limits

A simple possibility to set limits is to model hypothetical signals as generic shape-based resonances in the invariant mass spectrum, which are not based on any particular model. Gaussian-shaped signals are often used for this purpose. Their advantage comes from the fact that they allow recasting of limits for new theoretical models that have not been explicitly checked by the analysis [184, 185]. In this analysis, the Gaussian resonances are produced with their peak at the signal mass m_G that would be measured in an experiment corresponding to the level-1 mass scale. Signals with their resolution σ_G/m_G on the mass resolution curve are considered in the following, compare Figure 11.6. The widths σ_G of these signals are assumed to be dominated by detector resolution effects, meaning that the intrinsic width of the resonance is negligible in comparison. In addition, several cases of fixed width signals are considered: very narrow signals below the level-1 mass resolution with $\sigma_G/m_G = 0.15$ to check whether an improvement can be obtained in this case and signals at larger fixed widths of $\sigma_G/m_G = 0.2$ and $\sigma_G/m_G = 0.28$. Gaussian signals are produced at mass points in steps of 25 GeV to 50 GeV, depending on the resolution. To ensure that each signal peak is contained in the fit range, only signals where $m_G - \sigma_G \geq m_{jj}^{\text{lower}}$ and $m_G + \sigma_G \leq m_{jj}^{\text{upper}}$ are considered. In addition, the tails of the distribution are trimmed and only the 95 % central region is retained, in order to avoid effects of statistical fluctuations in the tails. The resulting contribution from the Gaussian signal is then added to the likelihood of the limit setting procedure, resulting in an upper limit on the number of events $\mu_{G,\text{upper}}$. By dividing this number by the integrated luminosity of the data sample, the upper limit can be interpreted as the product of the total cross section, the acceptance and the branching ratio $\sigma \times A \times \text{BR}$. Exemplary limits obtained for fixed widths of 0.15 and 0.2 can be seen in Figure 15.1(a) and Figure 15.1(b). The black points show the observed limits obtained from data, while the dashed line indicates the central values of the expected limits. Central 68 % and 95 % bands for the expected limit are shown in green and yellow, respectively. For both curves, it can be observed that the exclusion limits tend to improve with higher masses. However, the improvements get smaller or even disappear at the upper edge of the spectrum. In general, observed and expected limits agree well with each other. Figure 15.2 shows the observed limits for signal points on the resolution curve in comparison




Figure 15.1.: Expected and observed 95 % upper limits on $\sigma \times A \times \text{BR}$ for Gaussian signals of relative width 0.15 (a) and relative width 0.22 (b).

to the various curves at fixed resolutions. The general observation is that the limits tend to get stronger for narrower signals, where the fit is better able to distinguish a signal contribution from localized bin-to-bin fluctuations. However, at the same time the fluctuations of observed limits seem to get larger with smaller signal widths, as these are more likely to be influenced by localized statistical fluctuations in a single bin. Obtained upper limits on $\sigma \times A \times \text{BR}$ range from 8 nb to 15 nb at low masses of 100 GeV to 150 GeV, from 2.5 nb to 5 nb at masses around 200 GeV and from 1.5 nb to 3.5 nb at higher signal masses around 300 GeV. The observed curve at the level-1 mass resolution of the detector can be seen as a boundary of the region accessible to this analysis. Curves below this resolution indicate how the limits could improve from an improvement in the mass resolution of a signal.




Figure 15.2.: Comparison of 95 % upper limits on $\sigma \times A \times \text{BR}$ for various Gaussian signal widths. The curve in black corresponds to signals on the fractional signal mass resolution on level-1.

15.2.2. Model-dependent Limits

In addition to these generic limits, limits can also be interpreted in the context of a specific model. In this section, upper limits on a simplified dark matter mediator Z' model are derived. Exclusion limits are calculated for two of the mass points described in Section 11.2.1 at 350 GeV and 450 GeV, as these are the lowest mass points for which officially produced MC samples are available. While lower masses would also certainly be interesting for this analysis, mass samples were not available, and it was decided against running the time and resource intensive production. For each mass point, the posterior distribution $P(\mu | \text{data})$ is derived using the methods described in Section 15.1.4. Figure 15.3 shows the marginalized posterior probability for the 450 GeV sample along with the determined 95 % quantile. The 95 % upper limit on the number of signal is derived as $1.94 \cdot 10^8$. This corresponds to a value of $S/B = 0.002$ in a $\pm 1\sigma$ window around the signal peak position, which is comparable to the uncertainties of the $\langle \mu \rangle$ extrapolation, indicating that the sensitivity is limited by this uncertainty component. In order to cross-check the fitting procedure, the marginalized posterior distributions of the nuisance parameters can be checked. Two exemplary plots comparing the posterior and the prior distribution are shown in Figure 15.4 for the function




Figure 15.3.: Marginalized posterior for the signal scaling parameter for a Z' model with a mass of 450 GeV. In addition, the obtained 95 % exclusion limit is shown.

choice uncertainty and the luminosity uncertainty. The posterior for the function choice shows a preference of the fit towards the nominal fit model f_5 , while the luminosity uncertainty poses no constraints on the fit result. This is expected, as the luminosity uncertainty is only applied to the signal distribution, where the parameter μ can compensate for changes in the luminosity. In order to convert the obtained limit μ_{upper} into a cross section times acceptance times branching ratio, it is divided by the corresponding integrated luminosity of 61.54 fb^{-1} . Figure 15.5 shows the resulting limits for the two Z' samples. Note that in this case, all limits are given at the true value of the resonance mass, not at the level-1 mass scale. In addition to the observed and expected limit curves, a theoretical cross section curve of the Z' model with $g_q = 0.75$ is shown. For this curve, each theoretical cross section is divided by the obtained experimental acceptance at the same mass, an exponential function is used to extrapolate between these points. The observed curve intersects with the theoretical cross section curve at around 410 GeV, which can be interpreted as a 95 % lower bound on the mass of a potential Z' resonance for this particular model in the considered range. Observed excluded cross sections are given as 5.07 nb for a true signal mass of 350 GeV, corresponding to a level-1 mass of about 202 GeV, and as 3.14 nb for a signal mass of 450 GeV (279 GeV at level-1).

15.3. Results and Comparison to other Searches

In order to interpret the obtained limits, a comparison with the results of other searches is performed. For this, the result for $\sigma_{\text{limit}} \times A \times \text{BR}$ is translated into a limit on the coupling strength to quarks g_q . In a first step, a value of the cross section σ_{limit} is calculated. For this, the obtained result is divided by the signal acceptances mentioned in Section 12.4. As the utilized model assumes only decays of the Z' into pairs of quarks, $\text{BR}(Z' \rightarrow q\bar{q})$ is considered equal to one. Based on this, the values of (g_q, σ_q) obtained for the mass points shown in Figure 11.2 are considered. A reference point $(g_{\text{ref}}, \sigma_{\text{ref}})$ is selected, and the fact that the leading order cross section scales proportional to the squared coupling is utilized to extrapolate the coupling to a cross section of

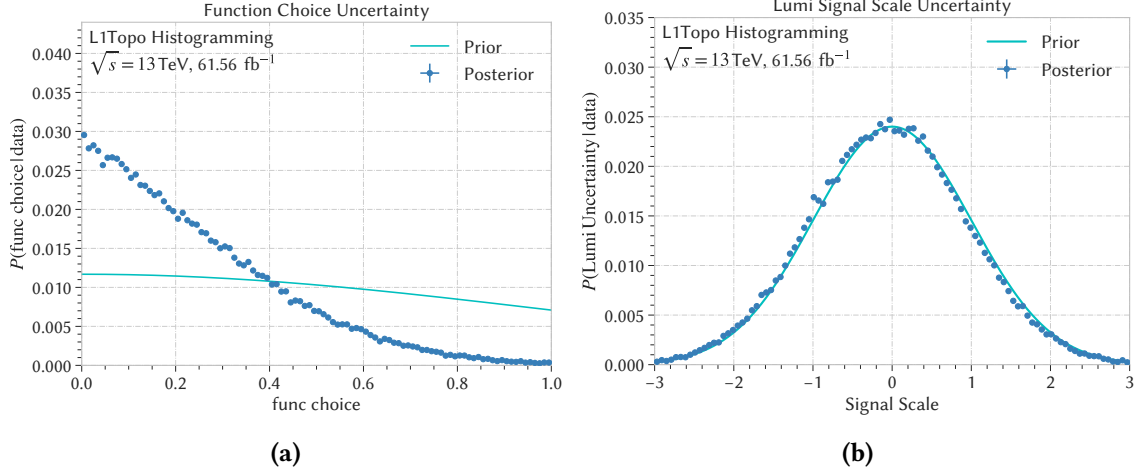


Figure 15.4.: Comparison of prior and posterior for the function choice uncertainty (a) and the luminosity uncertainty (b).

σ_{limit} :

$$\frac{\sigma_{\text{limit}}}{\sigma_{\text{ref}}} = \left(\frac{g_{\text{limit}}}{g_{\text{ref}}} \right)^2 \Rightarrow g_{\text{limit}} = \sqrt{\frac{\sigma_{\text{limit}}}{\sigma_{\text{ref}}}} \cdot g_{\text{ref}}. \quad (15.7)$$

This results in excluded couplings of $g_q = 0.66$ for a mass of 350 GeV and $g_q = 0.81$ for a mass of 450 GeV at 95 % credibility. To put this result into a context, Figure 15.6 shows the obtained limits in comparison to a wide range of limits from different ATLAS analyses using Run-2 data. The limits from this analysis are weaker than the results from other recent analyses, and no new phase-space region can be excluded using the data recorded in this thesis. This result was expected due to several reasons, which are discussed in the following. The primary effect is that the designated signal regions, for example the histogram category $y^* < 0.3$ and $|\eta_{1,2}| \leq 2.0$, are not accessible due to a bug in the histogramming firmware. Hence, jets from the end-cap and forward region, which have worse granularities than more central jets, enter the invariant mass calculation. These cause oscillations in the invariant mass spectrum, since the coarse granularities cause some values of the invariant mass to occur more often than other values. In addition, the non-uniform response of jets leads to a deformation of the invariant mass spectrum. Due to the low response in the forward region, these jets contribute already at smaller masses as they would in a calibrated spectrum. Furthermore, the level-1 spectrum is strongly affected by pile-up. Jets receive additional contributions to their energy from pile-up noise. Besides this, combinations from pile-up jets from different interactions in the same BC also affect the invariant mass spectrum. This effect was partially resolved with a $\langle \mu \rangle$ extrapolation in every mass bin. However, this procedure introduced increased uncertainties, reducing the statistical precision of the available data set. This limits the sensitivity available to this analysis, which is dominated by the uncertainty on the background estimate due to the increased data uncertainty. Manually lowering the errors coincided with convergence problems in the fit, which are attributed to non-smoothness effects caused by the level-1 jet properties. Another problem is that the mass range usable for the fit was limited to the region below 352 GeV due to the effect of forward jets on the spectrum. To fully explore the strength of such a search, a wider mass range is beneficial, as having sidebands with little to no signal allows for larger SWiFt windows, increasing the significance to potential signals.




Figure 15.5.: Upper limits on $\sigma \times A \times BR$ obtained for two Z' samples. Observed limits (solid line), expected limits (dashed line) and the 68 %, 95 % uncertainties of the expected limits (green, yellow) are shown. In addition, a theoretical cross section curve is shown (dashed blue).




Figure 15.6.: ATLAS limit summary plot of Run-2 results, superimposed with the L1Topo limit results from this thesis (blue). The original figure is taken from reference [186]. No improvements of the limits can be achieved using the data from this thesis.

Conclusion and Future Plans

16.1. Summary and Conclusion

In this thesis, data corresponding to an unprecedented amount of statistical precision has been recorded using a novel, trigger-level histogramming approach utilizing the Level-1 Topological Processor at the ATLAS experiment. By creating histograms of dijet invariant masses directly on the FPGAs of the level-1 trigger, every single event occurring at ATLAS in 2018 could be utilized for a dijet resonance search. Unlike conventional searches, where the statistical precision in the low mass region $\mathcal{O}(100 \text{ GeV})$ is limited due to the bandwidth capacity limitations of the TDAQ system, the full amount of statistics is available for this analysis. For this purpose, new firmware for the histogramming has been written and added to L1Topo, and new software for its control and readout has been written and tested extensively. In order to avoid using histograms that were recorded while detector problems occurred, which could render the histograms unfit for an analysis, a framework for the application of data quality criteria and filtering based on various run conditions has been developed. Since level-1 jets are objects that have not been used before for an analysis, an extensive study of their properties, for example their spatial resolution and energy scale, was performed. In this context, a possible jet energy scale calibration for level-1 jets was derived and the consequences of a missing calibration were discussed. Afterwards, two different MC-based approaches have been studied to find a background description for the invariant mass shape. Using the MC improved the understanding of the invariant mass spectrum and how the different considered aspects as the limited granularity or pile-up influence the spectrum. However, these approaches were not able to provide a sufficiently accurate description, therefore a search for localized excesses was performed using data-driven phenomenological functions in the SWiFt procedure. To improve the fit quality, a method to reduce parts of the pile-up contributions affecting the shape of the invariant mass spectrum has been developed. Applying a background-only fit to the pile-up modified spectrum, no deviations to the Standard Model prediction were found, meaning that there is no evidence for the presence of dijet resonances in the data. Exclusion limits at 95 % credibility level were set using a Bayesian approach. No improvement of the limits could be made compared to other searches at ATLAS. However, despite the exceptional challenges faced during the development of this analysis, be it implementing an analysis in the FPGA firmware in a production critical environment, or performing an analysis based on objects with previously only sparsely studied properties, an analysis has been performed successfully. Therefore, the results in this thesis can be seen as a successful proof-of-concept for a first-level trigger analysis.

The limits from other searches could not be improved mainly due to a faulty event selection in the firmware, rendering the designated signal regions unusable. As a consequence of this, the

total invariant mass histogram had to be used for the analysis, which however limited the fit range to a region at smaller masses where coarse-granularity forward jet contributions are less dominant. In addition, contributions from pile-up jets had to be mitigated using an extrapolation procedure, which however resulted in reduced event counts and increased uncertainties, limiting the sensitivity of the analysis.

Due to the limited time and manpower available, not all steps of the statistical analysis were performed as extensive as they could have. For example, it was not possible to implement a jet energy scale calibration for level-1 jets into the firmware of L1Topo in the given time constraints, which should be considered as mandatory for a future analysis. In addition, the vetoing of certain BCID is recommended in order to veto events from BCID in the Calreq bunch group at the end of each bunch group scheme. Otherwise, effects from dedicated calibration runs can enter the invariant mass spectrum, and lead to peak-like structures. In this thesis, the effects from the Tile laser calibrations were below the achievable sensitivity, and did therefore not affect the analysis. However, these effects are important to consider in a future analysis with improved sensitivity.

16.2. Possible Improvements and Outlook for Run-3

In a possible future iteration of the histogramming, several improvements could be gained compared to the analysis performed in this thesis. A new version with an improved firmware would mean that the designated signal regions, for example the region requiring $y^* < 0.3$ and $|\eta_{1,2}| < 2.0$, would be available for the analysis. In order to investigate the improvement that this region provides, the spectrum of this signal region obtained from the sliding window algorithm simulation in Section 13.2.7 is considered for a luminosity of 61.56 fb^{-1} . The same background-only fit and limit setting procedure as used for the statistical analysis in this thesis is performed, with the results being shown in Figure 16.1(a) for the background-only fit and in Figure 16.1(b) for the limit setting procedure. For this fit, a higher mass range is available, ranging up to 600 GeV due to the absence of coarse-granularity forward jets. The fit converges with a χ^2 p -value of 0.31 and a BH p -value of 0.43. The resulting exclusion limits are again transformed into limits

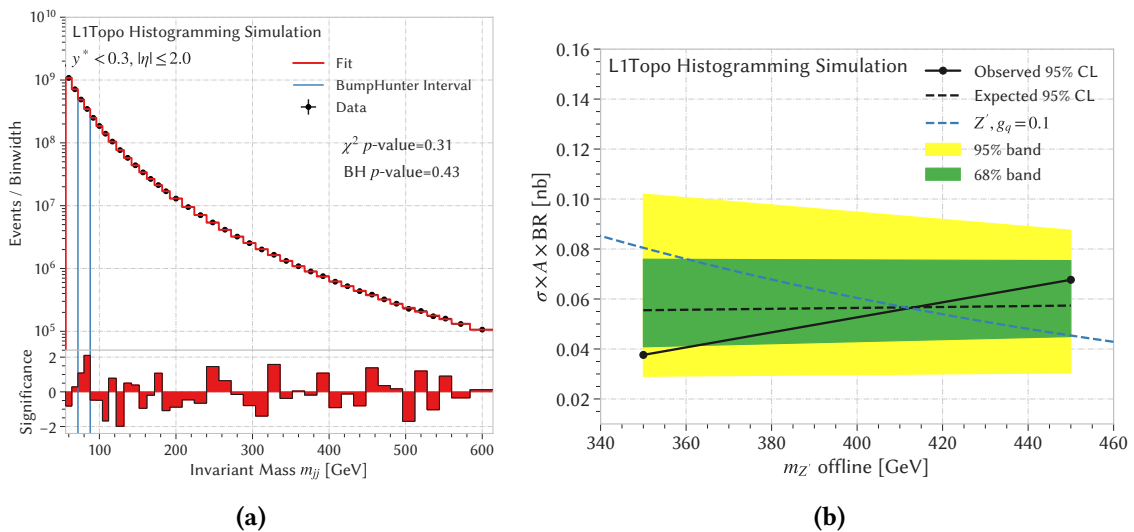


Figure 16.1.: Fit result in the signal region in simulation (a) and Z' exclusion limits (b).




Figure 16.2.: Limits summary plot, modified from [186]. The hypothetical contribution from the L1Topo simulation signal region has been superimposed to the other results (dark cyan).

on the coupling, which are given as 0.16 and 0.25 at masses of 350 GeV and 450 GeV respectively. They are shown in comparison to other limits in Figure 16.2. A significant improvement can be observed by using this histogram for the analysis. Considering only statistical uncertainties, these limits are comparable to the limits obtained from the resolved dijet+photon analyses, as shown in yellow and red. This shows that the potential sensitivity of this approach is comparable to other searches, and that a second iteration of the histogramming could be applied for Run-3.

Additional improvements to the histogramming analysis sensitivity can be expected due to a major upgrade of the ATLAS level-1 trigger for Run-3 of the LHC [139]. In order to maintain trigger efficiencies at low thresholds for expected luminosities above a level of $2 \cdot 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ along with corresponding high pile-up values, new trigger modules are introduced to L1Calo and supplied with increased granularity data. Among these modules is the Jet Feature Extractor (jFEX), which will use improved algorithms on higher-granularity calorimeter information to identify both small and large- R jet TOBs, hadronic tau TOBs and energy sums. In addition, a new L1Topo module is introduced to perform multiplicities and topological algorithms with the increased granularities and precisions. Both modules use up to date technology based on Xilinx UltraScale+ FPGAs[187], providing the capabilities to run improved, finer-granularity trigger algorithms. The jFEX will apply a Sliding Window algorithm to identify jets, and is therefore of particular interest for this analysis. In contrast to the JEP, the jFEX system runs on trigger towers with a granularity of 0.1×0.1 , improving by a factor of two on the granularity of the jet element-based Run-2 jets in both η and ϕ in the central region. In addition, the jFEX will provide a significantly improved granularity in the forward region. Instead of having a rectangular shape, jets are built with $\Delta R = 0.4$, approximating a more round shape similar to the anti- k_T algorithm. Two example for jets built on the jFEX are shown in Figure 16.3 together with the available trigger

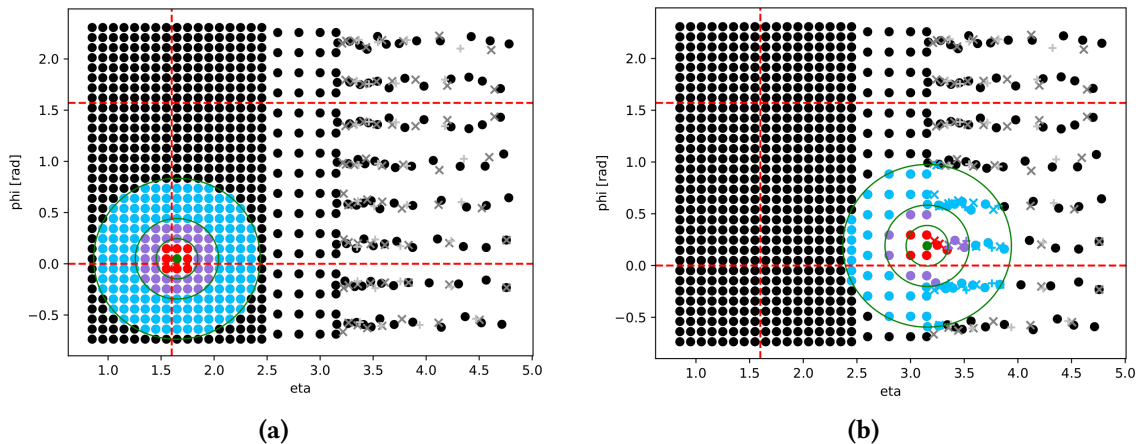


Figure 16.3.: Example for a seed (red), small- R jet (purple) and large- R jet (cyan) as identified by the jFEX in the central region (a) and in the forward region (b). The black dots indicate the center of trigger towers and thereby all possible jet positions. From [127].

towers, which correspond to all possible jet positions. As a result of the improved granularity in η and ϕ , the precision of the LUT of the cosh and cos operations will improve, leading to a smoother mass spectrum due to reduced quantization effects. In addition, an increase in energy resolution is expected, as energy deposition are available at a finer granularity. Furthermore, studies on a pile-up subtraction procedure for the energies of towers entering the jet algorithm are currently performed and likely to be implemented. A jFEX-level energy correction, which aims at equalizing the jet response in η , is also currently under development. Both of those are expected to further improve the recorded histogramming mass spectrum and are expected to improve the possible sensitivity of the analysis.

Besides this, the set-up and validation of the SWiFt fitting procedure can be performed more extensively in a future iteration of the analysis. For example, the choice of functions was mainly oriented at the choices made by similar analyses. In a full study, the choice of function is commonly performed using Wilk's test [188], which can determine whether an alternative hypothesis provides a significantly better background estimate than the null hypothesis based on the likelihood ratio of the two hypotheses. To further validate the background estimate, more *signal injection* tests could be performed to check whether the SWiFt procedure manages to extract the correct signal distribution and to assess the expected sensitivity of the fit. In addition, so-called *spurious signal* tests should be performed, which are used to estimate the fit bias, meaning the difference between inserted and extracted signal yield.

Especially at the unprecedented statistics explored in this thesis, it is not clear whether the used fitting methods can continue to provide the desired results. Various functionless methods to model smooth backgrounds are currently explored, as analyses in ATLAS experience problems with background descriptions when the statistical uncertainties get smaller [189]. This analysis could also benefit from an improved background description procedure. If it was possible to find a better MC description of the mass spectrum, a combination of MC and data-driven procedures could be possible. Since the invariant mass spectra seem to be dominated by effects caused by the jet energy scale and position quantization, the background estimate obtained from smooth functions could be weighted by factors coming from MC template histograms which contain information on the bin-to-bin fluctuations and quantization effects.

Experiment and Object Reconstruction

A.1. FPGA Resources

Table A.1.: Resources of the FPGA device XC7VX690TFFG1927--3 reported using Vivado 2016.4.

BRAMs	DSPs	FlipFlops	Gb Transceivers	LUT elements
1470	3600	866400	80	433200

A.2. Object Reconstruction

Additional objects are reconstructed at level-1 and used in the topological algorithms of L1Topo.

A.2.1. Energy Sums

The JEP algorithm for energy sums, for example the missing transverse energy, runs on the same jet elements as the jet algorithm. In order to obtain the components of the transverse energy, each jet element's E_T is multiplied by the appropriate geometrical constants:

$$\begin{aligned} E_x &= E_T \cdot \cos \phi \\ E_y &= E_T \cdot \sin \phi \end{aligned} \tag{A.1}$$

The values of $\cos \phi$ and $\sin \phi$ are obtained from LUTs, since the coordinates of the jet elements are fixed. Afterwards, the individual components are summed on each JEM. The results are then provided to the CMX, where the global values of $\sum E_T$, E_x and E_y are calculated from the results of all JEMs. The missing transverse energy is then determined from the components as

$$(E_T^{\text{miss}})^2 = E_x^2 + E_y^2. \tag{A.2}$$

A.2.2. Electron/Photon and Tau

The goal of the electron/photon and τ algorithms is to identify narrow, high E_T showers in the electromagnetic calorimeter. A depiction of the relevant quantities used in the CP algorithms is shown in Figure A.1. Both algorithms work similarly and are based on 0.1×0.1 trigger towers as inputs. In the ECAL, all four possible 2×1 E_T sums in the center of the 4×4 trigger tower window are performed. For an electron/photon object, at least one of the four possible 2×1 E_T sums is




Figure A.1.: Trigger towers used in the electron/photon/tau algorithms are shown as well as variables used in the electromagnetic isolation criteria. Modeled after a figure in [112].

required to surpass a predefined threshold. In addition, the three isolation energies, namely the electromagnetic isolation ring, the hadronic isolation ring and the hadronic core veto are required to be smaller than some adjustable, but static thresholds. For the tau algorithm, the energy in the hadronic core is added to each of the four $2 \times 1 E_T$ sums and at least one of these sums is required to pass a threshold. Hadronic and electromagnetic isolation are also required, while the hadronic veto is not used as showers in the hadronic calorimeter are expected for hadronically decaying taus.

A.2.3. Muons

The muon trigger algorithms are similar in RPC and TGC. Muons originating from the interaction region are identified by searching patterns of coincident hits in the different trigger stations within a *road*. A road tracks the path of a muon from the interaction point through the detector. The width of the road is related to the muon momentum, higher p_T values correspond to smaller road widths. Examples of high- p_T and low- p_T roads for trigger algorithms are shown in Figure A.2. Six p_T triggers are operated in parallel in this way, requiring the coincident hits to be contained inside the corresponding road widths.

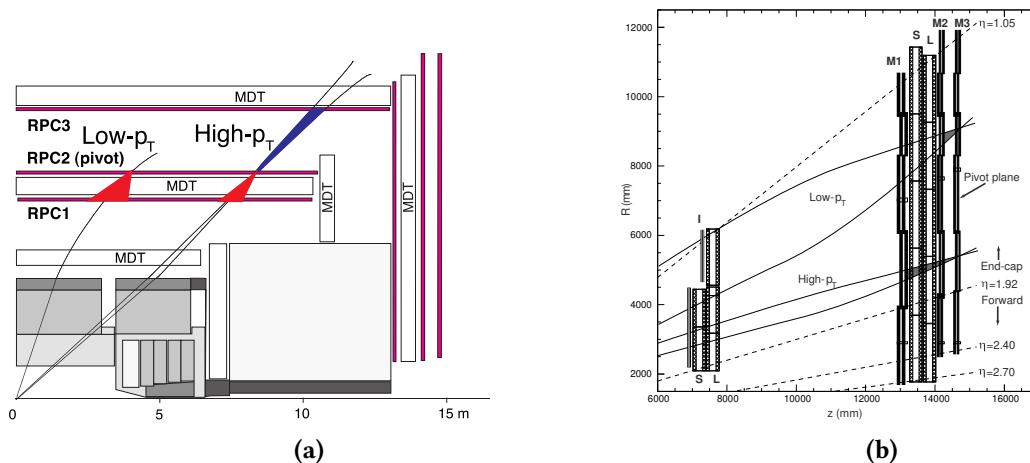


Figure A.2.: The scheme of the Level-1 Muon barrel (a) and end-cap (b) trigger algorithms are shown along with examples for both high- p_T and low- p_T roads, from [89].

A.3. L1Topo Firmware: Data Types

The LUT used for the decoding of level-1 jet positions is shown in Table A.2.

Table A.2.: Jet η LUT. $\eta_{\text{global bin}}$ denotes the global bin number, η_{generic} the integer representation of that coordinate returned by the LUT and η the corresponding float coordinate.

$\eta_{\text{global bin}}$	η_{generic}	η (float)
0	-39	-3.9
1	-29	-2.9 (-2.95 rounded)
2	-27	-2.7 (-2.65 rounded)
3	-25	-2.5 (-2.45 rounded)
4	-22	-2.2
5	-20	-2.0
6	-18	-1.8
7	-16	-1.6
8	-14	-1.4
9	-12	-1.2
10	-10	-1.0
11	-8	-0.8
12	-6	-0.6
13	-4	-0.4
14	-2	-0.2
15	0	0.0
16	2	0.2
17	4	0.4
18	6	0.6
19	8	0.8
20	10	1.0
21	12	1.2
22	14	1.4
23	16	1.6
24	18	1.8
25	20	2.0
26	22	2.2
27	25	2.5 (2.45 rounded)
28	27	2.7 (2.65 rounded)
29	29	2.9 (2.95 rounded)
30	39	3.9
31	39	3.9 (dummy)

A.4. Firmware Algorithms

A.4.1. Invariant Mass

Several algorithms exist which utilize this invariant mass calculation, for example the Invariant-MassInclusive1 and InvariantMassInclusive2 algorithms. Both calculate squared invariant masses of all possible pairs of inputs of one or two lists, respectively. The exact criteria is that any pair of TOBs fulfills

- $(E_{T,1} > E_{Tmin,1} \wedge E_{T,2} > E_{Tmin,2}) \vee (E_{T,1} > E_{Tmin,2} \wedge E_{T,2} > E_{Tmin,1})$
- $M_{min}^2 \leq M_{12}^2 \leq M_{max}^2$,

where the values of $E_{Tmin,1}$, $E_{Tmin,2}$, M_{min}^2 and M_{max}^2 are adjustable via IPBus.

A.4.2. DeltaRSqr

The DeltaRSqr algorithm calculates the angular distance ΔR^2 , compare Equation (6.8), between all possible combinations of its inputs. For each pair of TOBs, the topological quantities $\Delta\eta_{12} = |\eta_1 - \eta_2|$ and $\Delta\phi_{12}$ are calculated via the dedicated entities, where $\Delta\phi_{12}$ is calculated as

$$\Delta\phi = \begin{cases} 64 - |\phi_1 - \phi_2|, & \text{if } |\phi_1 - \phi_2| > 32 \\ |\phi_1 - \phi_2|, & \text{else} \end{cases} \quad (\text{A.3})$$

From these, the quantity $\Delta R^2 = \Delta\phi^2 + \Delta\eta^2$ is calculated using the DeltaRSqrCalc entity. All required operations involve only the addition or multiplication of integers, and are therefore straightforward to implement. Additional cuts on the E_T of the TOBs can be required. In summary, the algorithm fires if at least one TOB pair fulfills all the following criteria:

- $(E_{T,1} > E_{T,1,min} \wedge E_{T,2} > E_{T,2,min}) \vee (E_{T,1} > E_{T,2,min} \wedge E_{T,2} > E_{T,1,min})$
- $\Delta R_{min}^2 \leq \Delta R_{12}^2 \leq \Delta R_{max}^2$.

where $E_{T,1,min}$, $E_{T,2,min}$, ΔR_{min}^2 and ΔR_{max}^2 are parameters whose values are configurable via IPBus. To adjust for the representation of ΔR as an integer, the IPBus parameters have to be declared using the same representation. For example, a trigger condition of $0.2 \leq \Delta R < 1.5$ would require setting $\Delta R_{min}^2 = (10 \cdot 0.2)^2 = 4$ and $\Delta R_{max}^2 = (10 \cdot 1.5)^2 = 225$.

A.4.3. The SimpleCone algorithm

By calculating the ΔR between objects, it is possible to identify and combine close-by objects already on the first trigger level. The SimpleCone algorithm is designed for this purpose, its functionality is depicted in Figure A.3. It runs on an abbreviated list of ten objects above a certain E_T threshold. For every object of the list, an additional E_T cut can be applied. The distance ΔR is then calculated for all pairs of objects. For each TOB, the energies of all TOBs which are contained in a circle with radius ΔR_{max} around the TOB are summed up. If a configurable E_T threshold is passed, the trigger fires.

In order to parallelize the required addition operations as good as possible, i.e. to minimize the amount of subsequent summation operations, the sum is calculated with a sum tree, as depicted in Figure A.4. Compared to an iterative addition of each energy to the sum of all previous energies, this approach has the advantage that more sums can be performed in parallel, reducing the depth of the summation and lowering the required latency.

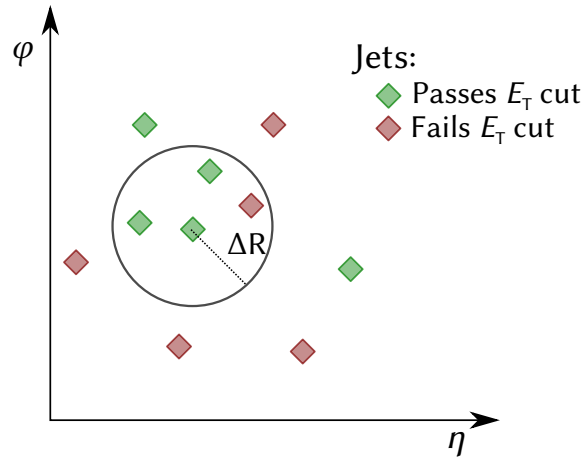


Figure A.3.: Depiction of the SimpleCone algorithm. The algorithm fires if the sum of energies of the green TOBs inside the ΔR circle is larger than a predefined threshold.




Figure A.4.: Depiction of an energy sum tree. Compared to a sequential sum, which would require three subsequent additions, only two stages of additions are required, as the first two sums can be performed in parallel.

A.5. Jet Cleaning

The following criteria correspond to the loose working point. A jet is identified as *BadLoose* and rejected if it fulfills at least one of the following criteria:

- $f_{\text{HEC}} > 0.5$ and $|f_Q^{\text{HEC}}| > 0.5$ and $\langle Q \rangle > 0.8$,
where f_{HEC} is the fraction of energy deposited in the HEC, f_Q^{HEC} is the deposited energy fraction with poor signal quality and $\langle Q \rangle$ describes the signal quality of the pulse shape.
- $|E_{\text{neg}}| > 60 \text{ GeV}$,
where E_{neg} is the energy sum of all cells with negative energy.
- $f_{\text{EM}} > 0.95$ and $f_Q^{\text{LAr}} > 0.8$ and $\langle Q \rangle > 0.8$ and $|\eta| < 2.8$,
where f_{EM} is the jet energy fraction deposited in the electromagnetic calorimeter and f_Q^{LAr} the fraction of all energy deposited in cells with poor signal quality in the LAr calorimeter.
- $f_{\text{max}} > 0.99$ and $|\eta| < 2$,
where f_{max} describes the maximum jet energy fraction deposited in any calorimeter layer.
- $f_{\text{EM}} < 0.05$ and $f_{\text{ch}} < 0.05$ and $|\eta| < 2$,
where f_{ch} is the ratio of the p_{T} sum of tracks associated to the jet and the total jet p_{T} .
- $f_{\text{EM}} < 0.05$ and $|\eta| \geq 2$

L1Topo Histogramming

B.1. Firmware Implementation

This section discusses supplementary implementation details for the histogramming firmware. Figure B.1 shows a depiction of the top module of the histogramming firmware. Some of the newly introduced components are explained in detail in the following, the BinMapping entity in Appendix B.1.1, the HistCategories entity in Appendix B.1.2, the EventInfo in Appendix B.1.3 and the BRAM configuration in Appendix B.1.4.

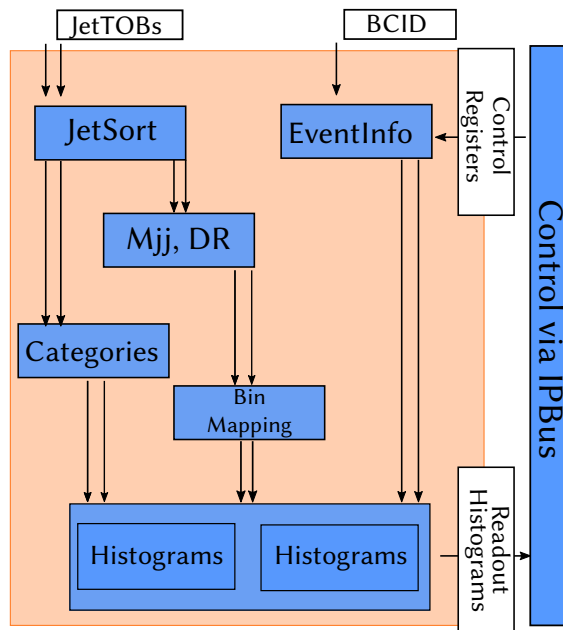


Figure B.1. Schematic depiction of the histogramming top-level module showing the most important entities.

B.1.1. BinMapping

In the bin mapping entity, the value of each topological quantity is translated into the corresponding bin number, taking the varying bin widths into account. In the following, the procedure is explained for the invariant mass histogram as listed in Table 9.2, however it is also valid for all other one-dimensional histograms. For each region with a number k and a corresponding range

$\left[m_{\text{jj, lower}}^{(k)}, m_{\text{jj, upper}}^{(k)} \right)$ and a bin width $b^{(k)}$, the bin number of the lowest bin of the region $n_{\text{lower}}^{(k)}$ is determined as

$$n_{\text{lower}}^{(k)} = \sum_{i=1}^{k-1} \frac{m_{\text{jj, upper}}^{(i)} - m_{\text{jj, upper}}^{(i)}}{b^{(i)}}. \quad (\text{B.1})$$

If $m_{\text{jj}} \in \left[m_{\text{jj, lower}}^{(k)}, m_{\text{jj, upper}}^{(k)} \right)$, then the bin number is calculated as

$$n_{\text{bin}} = n_{\text{lower}}^{(k)} + \left\lfloor \frac{m_{\text{jj}} - m_{\text{jj, lower}}^{(k)}}{b^{(k)}} \right\rfloor, \quad (\text{B.2})$$

where the operation $\left\lfloor \cdot \right\rfloor$ denotes an integer division, which rounds the result to the nearest smaller integer. In firmware, this is implemented utilizing the fact the all bin widths are powers of two, therefore the integer division by a factor of 2^n are performed by simply dropping the last n bits. As an example, an invariant mass of $m_{\text{jj}} = 673$ GeV would fall into the m_{jj} range 3, with

$$n_{\text{base}}^{(3)} = \frac{200 - 0}{2} + \frac{600 - 200}{8} = 150 \quad (\text{B.3})$$

and $b^{(k)} = 16$ GeV, resulting in a bin number of

$$n_{\text{bin}} = 150 + \left\lfloor \frac{673 - 600}{16} \right\rfloor = 154. \quad (\text{B.4})$$

If an invariant mass value larger than the highest defined range is observed, an overflow bin is filled. For the two-dimensional histogram, the used logic is different. In this case, the global bin numbers n_ϕ and n_η of the jet TOBs as discussed in Section 8.3.1 are used instead of the global coordinates. Both of them are represented using 5 bit of information. A unique bin number for the two-dimensional histogram is then calculated as

$$n_{\text{bin}} = n_\phi \cdot 2^5 + n_\eta. \quad (\text{B.5})$$

B.1.2. The HistCategories Entity

Inside this entity, the value of y^* is calculated as

$$2 \cdot y^* = |\eta_1 - \eta_2|. \quad (\text{B.6})$$

By multiplying both sides of the equation by two, the necessity of performing an integer division in the firmware with possible rounding precision loss can be avoided. To compensate for this, the cuts on y^* as shown in Table 9.1 are multiplied by two as well when applied in the firmware. The information on y^* , η_1 and η_2 is then used to determine an index ranging from 0 to 17 which enumerates the histogram categories. This is used later to address the correct invariant mass category histogram. In parallel, all histogrammed quantities are fed into the BinMapping entity. This entity determines the correct bin number for each topological quantity and the given histogram binning. The bin number information is then forwarded to the histogram block. Additional information is provided by the EventInfo block. This block determines which of the two histograms is currently active and which buffer is connected to the IPBus readout, configures a histogram counter which gives a unique number to every read out histogram, and checks whether a BCID veto is applied which is meant to reject data from certain BCIDs. In addition,

an event veto is calculated for cases where one of the two jets has an energy overflow. A total latency of 6 BC is required for the whole histogramming top module. Some blocks introduced in Figure B.1, like the JetSort or invariant mass algorithms, have already been introduced before. The remaining blocks are discussed in more detail in the following.

B.1.3. EventInfo

In the EventInfo entity, all firmware IPBus registers that are used to control the functionality of the histogramming are defined. The most important registers are listed in Table B.1.

Table B.1.: List of registers to configure the histogramming using IPBus requests.

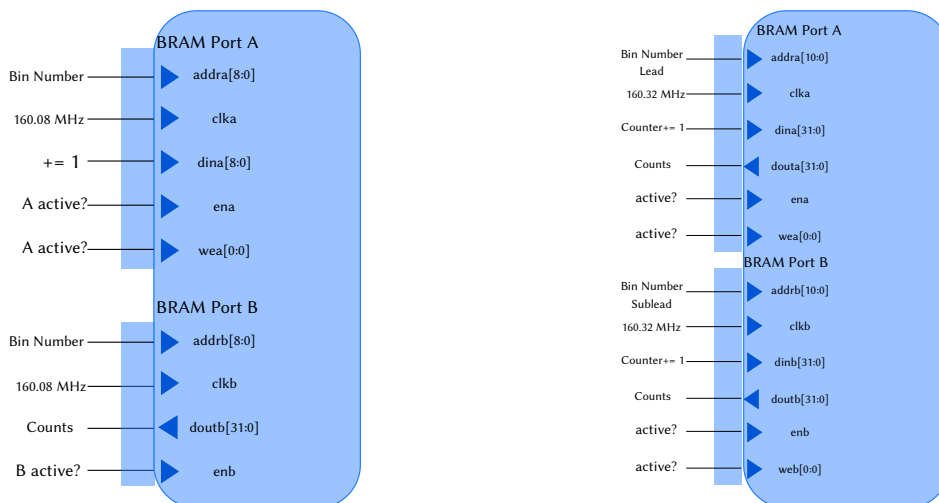
name	range	description	Read/Write
Histograms	65536	base address for all histograms	r
HistStatus	1	Information on IPBus error during histogram read	r
HistogramNumber	1	counts histograms, unique ID of histogram	r
HistControl	1	needs to be set to '1' if histogram was read	r/w
BCIDControl	1	set bit 0 to '1' activate BCIDs vetoing	r/w
NumLHCTurns	32	Number of LHC turns after which buffers are swapped	r/w
BCIDVeto	3654	Indicate accept or veto for each BCID	r/w

To read out the histogram data, IPBus registers have to be defined for the histograms themselves. A base address is defined for the very first histogram bin of the total invariant mass histogram, the subsequent addresses correspond to the following bins of the following histograms. The HistStatus register is filled in the firmware with the IPBus error bit of the histogram read transaction, which can read out along with the histogram to allow vetoing cases in which an IPBus error occurred. Every histogram is assigned a unique HistogramNumber, which is a circular 32 bit counter which increases by one every time the histogram buffers are swapped. This register is read only and useful for debugging purposes. The HistControl register is both readable and writable. Inside the firmware, the value is set to 1 as soon as a new histogram is available for readout. This allows the software to check if a new histogram is available and to reset this register after performing a successful IPBus read operation. The BCIDControl field can be used to activate the BCID veto functionality, which vetoes events from certain BCIDs. A list of accepted and rejected BCID is indicated by the BCIDVeto bits.

In addition, the EventInfo entity determines the histogram buffer which is active for writing. Depending on the index of the active histogram buffer, the data is filled in the corresponding histogram. The value of NumLHCTurns determines for how many LHC turns the buffers are used before being swapped. Per default, it is set to an integer multiple of 11246, the number of LHC turns in a second. Inside the firmware, a counter maximum is defined by multiplying this value by the number of BCID inside a single LHC turn, which is given by 3564. A counter in the EventInfo counts the number of ticks, which is reset to 0 as soon as the maximum is reached. In this case, the value of HistControl is set to 1 to indicate to the software that a new histogram is available. In the meantime, the contents of the other buffer are connected to the IPBus and can be accessed using a read request at any time. When the active buffer is swapped, multiplexers change the connections of both buffers.

B.1.4. Block RAM configuration

The first bin of a histogram has a dedicated memory address which is incremented by one for each bin, reading out all bins can therefore be done using an incremental read operation. Figure B.2(a) shows the BRAM configuration for a histogram with 294 bins with an address width of $w_{\text{addr}} = 9$ bit. This is the configuration used for the invariant mass, the other one-dimensional histograms differ only by the address widths. Figure B.2(b) shows the configuration for a 2D histogram with 1025 bins. The shown names of the ports of the BRAM are the ones used by Xilinx, compare Figure 4.5.



(a) Configuration of Xilinx BRAM for a one-dimensional histogram.

(b) Configuration of Xilinx BRAM for a 2D histogram. It is implemented as true dual port RAM with 1025 bins.

Figure B.2.: BRAMs configurations used in the histogramming firmware.

B.2. Standalone Software

The standalone is based on the uHAL library, which has been introduced in Section 4.3. Using the uHAL interface to perform IPBus read and write transactions, small test suites are written to perform readout operations of histogram data via the command line. The data received from this software was saved in a binary file format, which was afterwards converted into histograms with the correct binning. The standalone software makes it possible to quickly test the histogram features without many external dependencies, allowing for flexible tests already very early in the development cycle. In addition, the build times for the software are much shorter than for the official software.

However, the standalone approach has several disadvantages if used beyond the scope of standalone tests. Since an official L1Calo software already exists, it is not desired to run an additional software in parallel, sending possibly conflicting IPBus requests. In addition, accessing IS data like the ATLAS run status and whether L1Topo is included in the run from the IS is possible only with the official software. Furthermore, infrastructure to publish and archive histograms already exist via the MDA, Corresponding implementations exist in the L1Calo software for monitoring

histograms. If the tests using this software show to be promising, an implementation in the official L1Calo software is necessary.

B.3. Additional Histogramming Tests

Some supplementary information on histogramming tests is given in the following.

B.3.1. Waveform Simulation

Figure B.3 shows an exemplary test result of the histogram bin mapping functionality.

In this case, an invariant mass of 65 GeV was wrongly mapped to bin number 65¹, as a bin width of 1 GeV was erroneously assumed in the firmware for this invariant mass region. Similar tests like these were performed for all individual blocks.

Some simulations are more complex, in particular when the IPBus is required to be simulated. Since no real IPBus requests can be issued in the simulation, another solution is required. To simulate incoming IPBus requests, artificial IPBus requests were issued by manually setting the strobe of the IPBus input signal to 1 after a certain amount of 40 MHz clock ticks have passed. Figure B.4 shows a waveform window of these simulations. However, this procedure only allowed testing of individual parts of the firmware code. The IPBus read data is written to a file and inspected, its contents are compared to the expectations.




Figure B.3.: Simulation of the bin mapping functionality, showing the squared invariant mass, the invariant mass and the resulting bin address. Note that all values are given as hexadecimal integers.

B.3.2. Latency Measurement in the Test Rig

Using the standalone software implementation, first functionality tests were performed in the test rig. By performing overnight tests, in which histograms were read periodically after a certain time interval, the stability of IPBus read operations was tested. In one of these tests, the latency of issuing an IPBus block read request in the software until the response from L1Topo arrived was measured. While the IPBus latency has been measured by its developers, compare Figure 4.6, the number of words readout for the histogramming is significantly higher than what was considered in the measurements. This test was therefore useful to estimate what latency can be expected. Figure B.5 shows the result of measurements of the time required for the complete IPBus read command to finish. Requests to read all histogramming data had a latency between 20 ms and 80 ms in all but two cases, matching the expectations and showing no concerning performance drops. The varying time can be explained by the presence of additional IPBus requests from the control software. In this case, the IPBus ControlHub is responsible to schedule the incoming requests from all different sources.

¹65₁₀ = 41₁₆

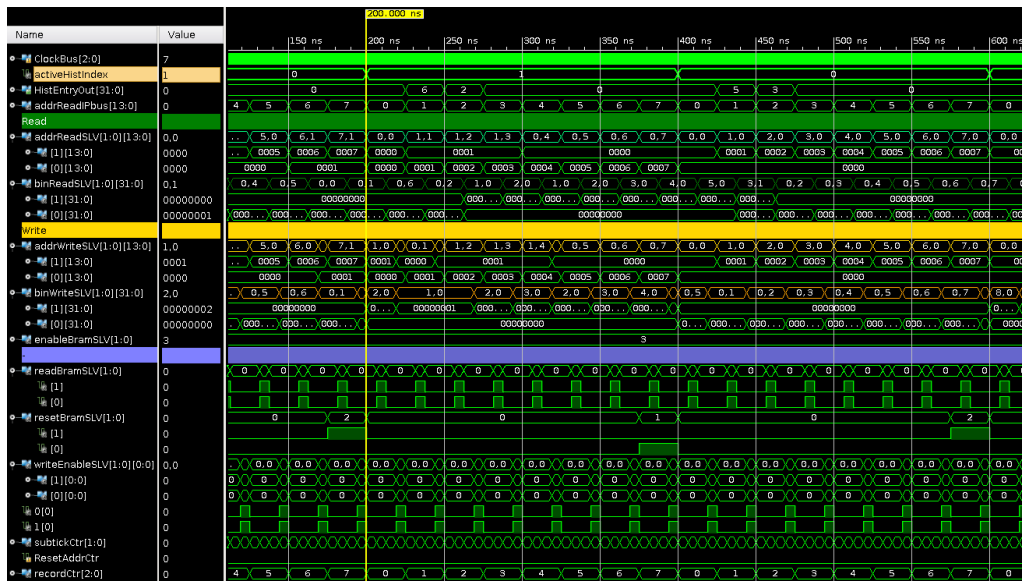


Figure B.4.: Simulation of the histogramming feature including IPBus requests.





Figure B.5.: Overview of the time delay a histogram read request takes to finish.

Precision of the Invariant Mass Algorithm

In the following the precision of the invariant mass algorithm is studied. Due to the fixed-point precision used in the firmware, some deviations from a floating-point calculation of the invariant mass are expected. For this, the relative deviation

$$\delta x = \frac{|x_{\text{round}} - x_{\text{true}}|}{|x_{\text{true}}|}, \quad (\text{C.1})$$

is calculated for all possible cases, where $x = \cos \Delta\phi, \cosh \Delta\eta$. For the \cos operation, the points with $\cos \Delta\phi \equiv 0$ have been skipped, as Equation (C.1) is not defined in this special case. Box diagrams for all possible values of the corresponding LUTs are shown in Figure C.1(a) for $\cosh \Delta\eta$ and in Figure C.1(b) for $\cos \Delta\phi$. The box marks the 50 % central region, the line inside the box the median. The whiskers indicate the rest of the distribution besides some points that are considered as outliers when they are outside $1.5 \times$ the interquartile range. For the \cosh operation, more than 75 % of rounding error are below 0.05 %. In a few cases, outliers with rounding errors of 0.1 % to 0.35 % are observed. All these cases belong to the same category fulfilling $\cosh \Delta\eta \approx 1$, where it is difficult to precisely represent small deviations from 1 using fixed point numbers. For the \cos operation the uncertainties are larger, with the central 50 % region ranging from 0.1 % to 0.45 %. However, all of these are negligible with respect to the spatial resolution of level-1 jets. Figure C.2 shows how these propagate into an uncertainty of the invariant mass itself due to the angular



(a) Box diagram of the relative error of $\cosh \Delta\eta$ due to rounding using the fixed point representation.

(b) Box diagram of the relative error of $\cos \Delta\phi$ due to rounding using the fixed point representation.

Figure C.1.: L1Topo LUT rounding errors.




Figure C.2.: Box plot of the invariant mass errors due to rounding. The whiskers were chosen to indicate the 5 and 95 percentiles.

difference term $\cosh \Delta\eta - \cos \Delta\phi$. All cases where the error is larger than 0.05 are caused by events where $\cosh \Delta\eta \approx \cos \Delta\phi = 1$. These correspond to cases where $\Delta\eta \approx 0$ and $\Delta\phi \approx 0$. These cases are clearly unphysical, since a minimal separation of $\Delta\eta = 0.4$, $\Delta\phi = 0.4$ or $\Delta\eta = 0.2$ and $\Delta\phi = 0.2$ is required between two jets due to the local maximum requirement. Some other cases with small values for the angular distances result in uncertainties between 1% and 5%. As the angular difference is close to 0 in these cases, very large jet energies are required for them to result in relevant invariant masses. For example, a dijet pair with $E_{T,1} = 150$ GeV and $E_{T,2} = 100$ GeV with $\Delta\eta = 0.4$ and $\Delta\phi = 0$ would still result in an invariant mass below the turn-on region:

$$\sqrt{2 \cdot 150 \text{ GeV} \cdot 100 \text{ GeV} \cdot (\cosh(0.4) - \cos(0))} \approx 49 \text{ GeV} \quad (\text{C.2})$$

Such cases are very rare and completely negligible compared to low- E_T jets with more realistic angular separation. Due to this, the effect of the invariant mass precision is of no relevance for this thesis, as other effects are much more dominant.

Statistical Analysis

D.1. Fit Uncertainty Calculation

For some fits, for example the response fits and the $\langle \mu \rangle$ extrapolation, uncertainty bands were drawn to show the $\pm 1\sigma$ uncertainty of the fit function. The method to determine these bands is taken from [181] and is explained briefly in the following. Given a model $\mathcal{F}(x; \mathbf{a})$ depending on some parameters \mathbf{a} with best fit values $\hat{\mathbf{a}}$, the covariance matrix of the fit model is calculated using the propagation of uncertainties. By this, it can be expressed using the Jacobian J of the fit function with respect to its parameters \mathbf{a} and the covariance matrix of the fit parameters $\text{Cov}_{\hat{\mathbf{a}}}$:

$$\text{Cov}_{\mathcal{F}} = J \cdot \text{Cov}_{\hat{\mathbf{a}}} \cdot J^T. \quad (\text{D.1})$$

The estimate for the standard deviation of $\sigma_{\mathcal{F}}$ is then calculated as

$$\sigma_{\mathcal{F}} = \sqrt{\text{Cov}_{\mathcal{F}}} = \sqrt{J \cdot \text{Cov}_{\hat{\mathbf{a}}} \cdot J^T} \quad (\text{D.2})$$

If the true standard deviation is unknown, the estimator for the mean value of normal distributed data is not normal distributed. Instead, the estimator follows the Student's t distribution, which has to be considered in this case. Due to this, a $(1 - \alpha)$ confidence interval is calculated as follows:

$$\mathcal{F}(x; \hat{\mathbf{a}}) - t_{N-n, 1-\frac{\alpha}{2}} \cdot \sigma_{\mathcal{F}} \leq \mathcal{F}(x; \mathbf{a}_{\text{true}}) \leq \mathcal{F}(x; \hat{\mathbf{a}}) + t_{N-n, 1-\frac{\alpha}{2}} \cdot \sigma_{\mathcal{F}} \quad (\text{D.3})$$

where $t_{N-n, 1-\frac{\alpha}{2}}$ is the $1 - \frac{\alpha}{2}$ quantile of the Student's t distribution with $N - n$ degrees of freedom, where N corresponds to the number of data points used for the fit and n to the number of parameters. Using Equation (D.3), the boundaries of the uncertainty bands are defined. For all bands shown in this thesis, $1 - \alpha = 0.683$ is chosen.




Figure D.1.: Comparison of the distribution of level-1 energies in a bin of reconstructed jet E_T (left) and the distribution of reconstructed jet E_T in a bin of level-1 jet E_T (right).

D.2. Estimated Level-1 E_T

If the binning of the response is chosen in bins of E_T^{reco} to directly extract the dependence on E_{TL1} , the shapes of the responses deviates from a Gaussian shape in the low-energetic bins. Figure D.1 shows a comparison of the response distribution using bins in E_T^{reco} and bins in E_T^{L1} . The distribution in a E_T^{reco} bin shows a cut-off due to the minimum E_T^{L1} of level-1 jets, a Gaussian fit is therefore not possible. Due to this, the two-step procedure mentioned in the main text is applied. Figure D.2 shows two exemplary response fits of the first step, where the response is given as a function of E_T^{reco} .

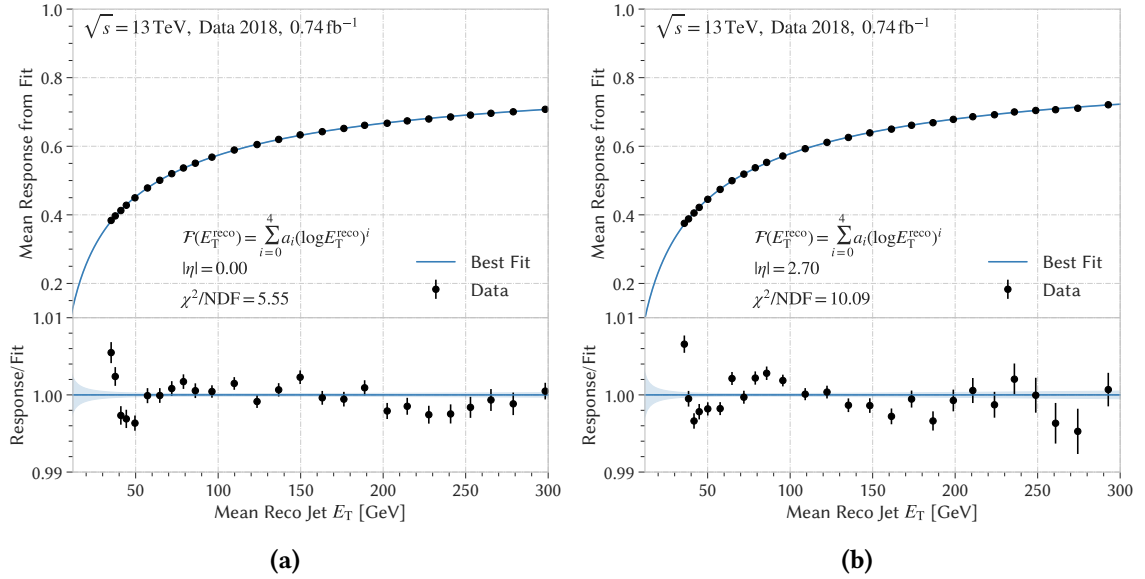


Figure D.2.: Fits of the response as a function of the mean reconstructed jet energy.

D.3. Definition of the Bins

The $\langle\mu\rangle$ bins are listed in Table D.1, including lower and upper bin edges as well as the integrated luminosity associated with every bin. Everything below $\langle\mu\rangle = 10$ is combined into a single bin, afterwards ten logarithmically spaced bins are chosen between $\langle\mu\rangle = 10$ and $\langle\mu\rangle = 60$, meaning that every bin edge is a constant multiple of the previous bin edge, i.e. $\langle\mu\rangle_{i+1} = c \cdot \langle\mu\rangle_i$ for $i = 1, \dots, 10$. As the scaling behavior inside every bin is expected to be proportional to μ^2 , the chosen bin centers are modified to reduce the bias which would be introduced when choosing the central value of the bin. Let $E(\mu)$ denote the expectation value of μ in every bin, which is calculated for a quadratic probability density function $f(\mu) \propto \mu^2$ as:

$$E(\mu)_i = \int_{\mu_i}^{\mu_{i+1}} d\mu f(\mu)\mu = \frac{3}{\mu_{i+1}^3 - \mu_i^3} \int_{\mu_i}^{\mu_{i+1}} d\mu \mu^2 \cdot \mu = \frac{3}{4} \frac{\mu_{i+1}^4 - \mu_i^4}{\mu_{i+1}^3 - \mu_i^3} \quad (\text{D.4})$$

The bin centers are chosen to correspond to this expectation value for every bin. The presented binning was found to provide smaller uncertainties of the fit parameters for $\langle\mu\rangle \rightarrow 1$, and was therefore chosen in favor of an equidistant binning.

Table D.1.: Definition of the $\langle\mu\rangle$ bins used for the pile-up extrapolation procedure, listing bin edges, bin centers and the recorded luminosity associated with every bin.

Bin	$\langle\mu\rangle_{\text{low}}$	$\langle\mu\rangle_{\text{high}}$	Center	Luminosity [fb^{-1}]
1	0.00	10.00	7.50	0.07
2	10.00	11.96	11.04	0.04
3	11.96	14.31	13.21	0.07
4	14.31	17.12	15.80	0.60
5	17.12	20.48	18.90	2.18
6	20.48	24.49	22.61	5.13
7	24.49	29.30	27.04	8.09
8	29.30	35.05	32.35	10.51
9	35.05	41.93	38.69	13.81
10	41.93	50.16	46.29	16.42
11	50.16	60.00	55.37	4.70

D.4. Reweighting Procedure

In the case of histograms with errors deviating from simple Poisson uncertainties, the log-likelihood method is modified, in order to let every bin better correspond to its statistical power and to let statistical tests that require Poisson distributed bin contents still be well-defined. In order to apply the methods assuming Poisson distributed uncertainties, a reweighting procedure is applied. An effective number of events is calculated for every bin as

$$n_{\text{eff},i} = \left(\frac{d_i}{\sigma_{d_i}} \right)^2$$

with corresponding weights for every bin

$$w_{\text{eff},i} = \frac{d_i}{n_{\text{eff},i}} = \frac{\sigma_{d_i}^2}{d_i}$$

Essentially, this transforms the distribution to let every bin reflect the effective number of events contributing to this bin, indicating its statistical power. The preliminary fits are applied to the nominal spectrum, obtaining good starting parameters for the background estimate b_i for a certain bin:

$$b_i = \int_{m_i}^{m_{i+1}} dm f\left(\frac{m}{\sqrt{s}}; \theta\right), \quad (\text{D.5})$$

Using these definitions, a weighted likelihood fit is applied, where the contribution of every bin are modified using the corresponding weight

$$L = \prod_{i=m_1}^{m_2} \text{Pois}(n_{\text{eff},i}, \frac{b_i}{w_{\text{eff},i}}),$$

where (m_1, m_2) are the borders of the current SWiFt fit window. Note that if this method was applied for Poisson uncertainties $\sigma_{d_i} = \sqrt{d_i}$, the original spectrum would remain unchanged [175]. In the case where the deviation from Poisson uncertainties are caused by adding events with different weights, for example in the case of trigger prescales or MC event weights, this method has been applied and shown to provide good results [190].

D.5. Background Estimate

Figure D.3 shows the background estimate and distribution of Poisson probabilities as obtained from the alternate fitting function f_y , which is used to calculate the uncertainties on the background estimate.

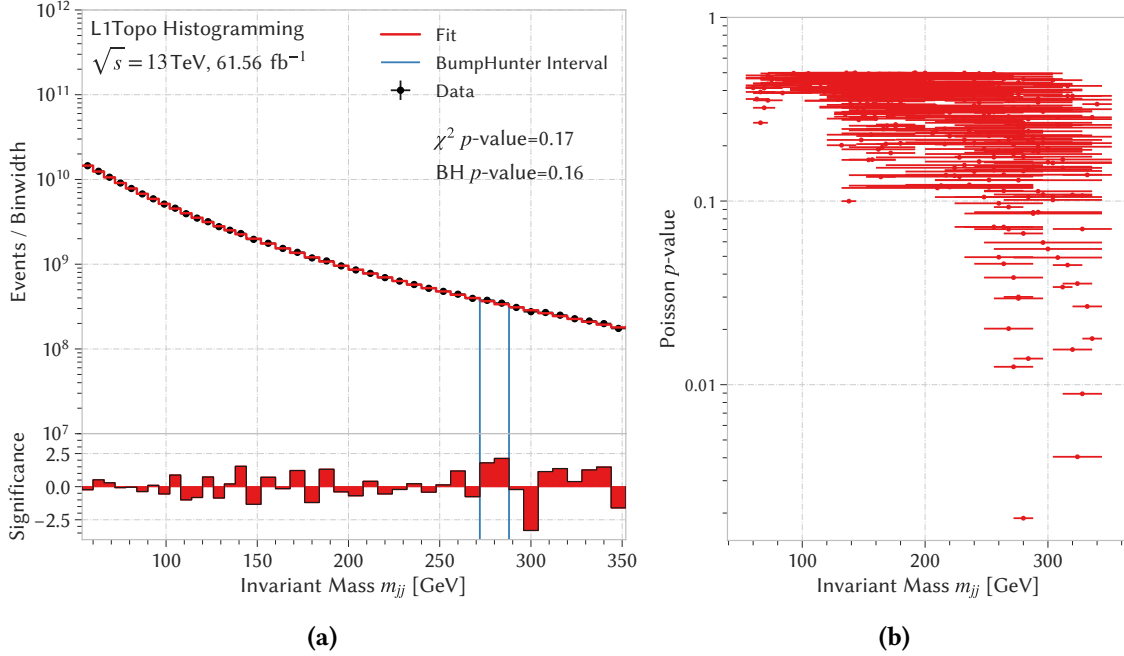


Figure D.3.: Fit result (a) as well as the distribution of all possible Poisson intervals (b).

D.6. Markov Chain Monte Carlo

MCMC is a method to construct and sample from arbitrary complex posterior distributions, in particular in cases where it is either impossible or computationally challenging to sample directly from the posterior [183]. It is based on a random walk in the parameter space by sequentially drawing values of (μ, ν) from approximate distributions. Each draw is then corrected to make the following draw correspond better to the target distribution. Due to this, the sequence of drawn random variables $\theta^0, \theta^1, \dots$ fulfills the property that θ^k only depends on the previous value θ^{k-1} , which is the defining condition of a *Markov Chain*. Since the approximate distributions are improved in each step, the goal is that they converge to the target posterior probability distribution after a sufficient number of steps have been performed. For this, the properties of Markov chains are valuable, as they ensure convergence towards a stationary distribution under a few conditions, compare for example Section 11.2 of reference [183] for a proof and [191] for details on how convergence is achieved in BAT. Before the drawn samples are used for the analysis itself, a so-called burn-in period is performed to achieve convergence, samples obtained in this period are dropped. Several independent chains are started at different points in parameter space and required to converge to the same distribution. Only after this convergence is achieved, drawn samples from the posterior are used for the statistical analysis. A commonly used algorithm to generate the Markov chains is the Metropolis algorithm [192], which is explained briefly in the

following. Let $\theta = (\mu, \nu)$ denote all relevant parameters. The algorithm requires the choice of a starting distribution $P_0(\theta)$, from which the initial samples are drawn, and a jumping distribution $J_t(\theta^* | \theta^{t-1})$, which proposes new samples based on the previous one. It then proceeds as follows:

1. Draw a starting point θ^0 with $P(\theta^0 | \text{data}) > 0$ from P_0 .
2. For $t = 1, 2, \dots$
 - a) Sample a proposal for a new position θ^* from $J_t(\theta^* | \theta^{t-1})$
 - b) Calculate the ratio

$$r = \frac{P(\theta^* | \text{data})}{P(\theta^{t-1} | \text{data})} \quad (\text{D.6})$$

- c) Set the new value as

$$\theta^t = \begin{cases} \theta^* & \text{with probability } \min(r, 1) \\ \theta^{t-1} & \text{otherwise} \end{cases}$$

As points which improve the probability density are more likely to be accepted, it is ensured that more points in regions of high posterior density are sampled, and the distribution of draws converges towards the target distribution. The functionality of the algorithm is depicted in Figure D.4 with examples on obtained marginalized posterior distributions.




Figure D.4.: Depiction of the Metropolis algorithm for a two-dimensional parameter space. The chain starts at the point in the lower left. Steps with dashed red arrows are rejected, steps with filled blue arrow are accepted. The number at each point denotes how many iterations have been performed before a new draw has been accepted. The distributions on the top and the right of the plot indicate the obtained marginalized distributions, the yellow curves the corresponding 68 % central region. From reference [193].

Bibliography

- [1] S. Artz, *Dijet Analysis on the Level-1 Topological Processor in the ATLAS Experiment*, PhD thesis: Johannes Gutenberg-Universität Mainz, 2017, unpublished thesis (cit. on pp. [iii](#), [99](#)).
- [2] *Standard Model Summary Plots Spring 2020*, tech. rep. ATL-PHYS-PUB-2020-010, CERN, 2020, URL: <https://cds.cern.ch/record/2718937> (cit. on pp. [1](#), [57](#)).
- [3] Particle Data Group, *Review of Particle Physics, Progress of Theoretical and Experimental Physics* **2020** (2020), 083C01, ISSN: 2050-3911, DOI: [10.1093/ptep/ptaa104](https://doi.org/10.1093/ptep/ptaa104), eprint: <https://academic.oup.com/ptep/article-pdf/2020/8/083C01/33653179/ptaa104.pdf> (cit. on p. [1](#)).
- [4] ATLAS Collaboration, *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, *Phys. Lett.* **B716** (2012) [1](#), DOI: [10.1016/j.physletb.2012.08.020](https://doi.org/10.1016/j.physletb.2012.08.020), arXiv: [1207.7214](https://arxiv.org/abs/1207.7214) [[hep-ex](#)] (cit. on pp. [1](#), [6](#)).
- [5] CMS Collaboration, *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, *Phys. Lett.* **B716** (2012) [30](#), DOI: [10.1016/j.physletb.2012.08.021](https://doi.org/10.1016/j.physletb.2012.08.021), arXiv: [1207.7235](https://arxiv.org/abs/1207.7235) [[hep-ex](#)] (cit. on pp. [1](#), [6](#)).
- [6] E. Simioni, *Upgrade of the ATLAS Level-1 Trigger with event topology information*, *Journal of Physics: Conference Series* **664** (2015) [082052](#) (2015), DOI: [doi:10.1088/1742-6596/664/8/082052](https://doi.org/10.1088/1742-6596/664/8/082052) (cit. on pp. [1](#), [60](#), [77](#), [78](#)).
- [7] ATLAS Collaboration, *Search for low-mass dijet resonances using trigger-level jets with the ATLAS detector in pp collisions at $\sqrt{s}=13$ TeV*, (2018), arXiv: [1804.03496](https://arxiv.org/abs/1804.03496) [[hep-ex](#)] (cit. on pp. [2](#), [26](#), [98](#), [163](#), [164](#)).
- [8] ATLAS Collaboration, *Search for low-mass resonances decaying into two jets and produced in association with a photon using pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, Submitted to: *Phys. Lett.* (2019), arXiv: [1901.10917](https://arxiv.org/abs/1901.10917) [[hep-ex](#)] (cit. on pp. [2](#), [98](#), [163](#)).
- [9] S. Corrodi, *Einführung in die Teilchenphysik: Masterclass*, Masterclass, 2016, URL: <https://www.uzh.ch/cmsssl/physik/dam/jcr:84bc6466-e86a-49e7-9b4b-aa8762b56f47/MC2016-SMIntro.pdf> (cit. on p. [3](#)).
- [10] J. Woithe, G. J. Wiener, and F. F. V. der Veken, *Let's have a coffee with the Standard Model of particle physics!* *Physics Education* **52** (2017) [034001](#), DOI: [10.1088/1361-6552/aa5b25](https://doi.org/10.1088/1361-6552/aa5b25) (cit. on p. [3](#)).
- [11] C. P. Burgess and G. D. Moore, *The standard model: A primer*, Cambridge University Press, 2006, ISBN: 9780521860369 (cit. on pp. [5](#), [6](#)).
- [12] C. Patrignani et al., *Review of Particle Physics*, *Chin. Phys.* **C40** (2016) [100001](#), DOI: [10.1088/1674-1137/40/10/100001](https://doi.org/10.1088/1674-1137/40/10/100001) (cit. on pp. [5](#), [14](#), [53](#)).
- [13] M. Thomson, *Modern Particle Physics*, Cambridge University Press, 2013, DOI: [10.1017/CB09781139525367](https://doi.org/10.1017/CB09781139525367) (cit. on p. [5](#)).

- [14] H. M. zu Theenhausen, *Search for Light Dijet Resonances using Trigger Jets with the ATLAS Experiment at the Large Hadron Collider*, PhD thesis: Universität Heidelberg, 2018 (cit. on pp. 5, 9, 19, 160, 164, 165, 167, 177).
- [15] M. Chala, F. Kahlhoefer, M. McCullough, G. Nardini, and K. Schmidt-Hoberg, *Constraining Dark Sectors with Monojets and Dijets*, **JHEP** **07** (2015) 089, DOI: [10.1007/JHEP07\(2015\)089](https://doi.org/10.1007/JHEP07(2015)089), arXiv: [1503.05916](https://arxiv.org/abs/1503.05916) [hep-ph] (cit. on pp. 5, 15–17).
- [16] D. Abercrombie et al., *Dark Matter Benchmark Models for Early LHC Run-2 Searches: Report of the ATLAS/CMS Dark Matter Forum*, (2015), arXiv: [1507.00966](https://arxiv.org/abs/1507.00966) [hep-ex] (cit. on pp. 5, 15).
- [17] H. An, X. Ji, and L.-T. Wang, *Light Dark Matter and Z' Dark Force at Colliders*, **JHEP** **07** (2012) 182, DOI: [10.1007/JHEP07\(2012\)182](https://doi.org/10.1007/JHEP07(2012)182), arXiv: [1202.2894](https://arxiv.org/abs/1202.2894) [hep-ph] (cit. on pp. 5, 14, 15).
- [18] G. Bertone, D. Hooper, and J. Silk, *Particle dark matter: evidence, candidates and constraints*, **Physics Reports** **405** (2005) 279, ISSN: 0370-1573, DOI: [10.1016/j.physrep.2004.08.031](https://doi.org/10.1016/j.physrep.2004.08.031) (cit. on p. 5).
- [19] P. W. Higgs, *Broken Symmetries and the Masses of Gauge Bosons*, **Phys. Rev. Lett.** **13** (1964) 508, [160(1964)], DOI: [10.1103/PhysRevLett.13.508](https://doi.org/10.1103/PhysRevLett.13.508) (cit. on p. 6).
- [20] F. Englert and R. Brout, *Broken Symmetry and the Mass of Gauge Vector Mesons*, **Phys. Rev. Lett.** **13** (9 1964) 321, DOI: [10.1103/PhysRevLett.13.321](https://doi.org/10.1103/PhysRevLett.13.321) (cit. on p. 6).
- [21] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble, *Global Conservation Laws and Massless Particles*, **Phys. Rev. Lett.** **13** (20 1964) 585, DOI: [10.1103/PhysRevLett.13.585](https://doi.org/10.1103/PhysRevLett.13.585) (cit. on p. 6).
- [22] D. Galbraith and C. Burgard, *Example: Standard model of physics*, 2016, URL: <http://www.texample.net/tikz/examples/model-physics/> (cit. on p. 7).
- [23] S. L. Glashow, *Partial Symmetries of Weak Interactions*, **Nucl. Phys.** **22** (1961) 579, DOI: [10.1016/0029-5582\(61\)90469-2](https://doi.org/10.1016/0029-5582(61)90469-2) (cit. on p. 6).
- [24] S. Weinberg, *A Model of Leptons*, **Phys. Rev. Lett.** **19** (1967) 1264, DOI: [10.1103/PhysRevLett.19.1264](https://doi.org/10.1103/PhysRevLett.19.1264) (cit. on p. 6).
- [25] A. Salam, *Weak and Electromagnetic Interactions*, Conf. Proc. **C680519** (1968) 367 (cit. on p. 6).
- [26] G. Arnison et al., *Experimental observation of isolated large transverse energy electrons with associated missing energy at $\sqrt{s}=540$ GeV*, **Physics Letters B** **122** (1983) 103, ISSN: 0370-2693, DOI: [10.1016/0370-2693\(83\)91177-2](https://doi.org/10.1016/0370-2693(83)91177-2) (cit. on p. 8).
- [27] M. Banner et al., *Observation of single isolated electrons of high transverse momentum in events with missing transverse energy at the CERN pp collider*, **Physics Letters B** **122** (1983) 476, ISSN: 0370-2693, DOI: [10.1016/0370-2693\(83\)91605-2](https://doi.org/10.1016/0370-2693(83)91605-2) (cit. on p. 8).
- [28] G. Arnison et al., *Experimental observation of lepton pairs of invariant mass around 95 GeV/c² at the CERN SPS collider*, **Physics Letters B** **126** (1983) 398, ISSN: 0370-2693, DOI: [10.1016/0370-2693\(83\)90188-0](https://doi.org/10.1016/0370-2693(83)90188-0) (cit. on p. 8).
- [29] P. Bagnaia et al., *Evidence for $Z^0 \rightarrow e^+e^-$ at the CERN pp collider*, **Physics Letters B** **129** (1983) 130, ISSN: 0370-2693, DOI: [10.1016/0370-2693\(83\)90744-X](https://doi.org/10.1016/0370-2693(83)90744-X) (cit. on p. 8).

- [30] R. Ellis, W. Stirling, and B. Webber, *QCD and Collider Physics*, Cambridge Monographs on Particle Physics, Nuclear Physics and Cosmology, Cambridge University Press, 2003, ISBN: 9780521545891 (cit. on pp. 9, 19, 25).
- [31] M. Tanabashi et al., *Review of Particle Physics*, *Phys. Rev. D* **98** (3 2018) 030001, DOI: [10.1103/PhysRevD.98.030001](https://doi.org/10.1103/PhysRevD.98.030001) (cit. on pp. 10–12, 163).
- [32] L. Canetti, M. Drewes, and M. Shaposhnikov, *Matter and antimatter in the universe*, *New Journal of Physics* **14** (2012) 095012, DOI: [10.1088/1367-2630/14/9/095012](https://doi.org/10.1088/1367-2630/14/9/095012) (cit. on p. 11).
- [33] K. Hirata et al., *Observation of a Neutrino Burst from the Supernova SN 1987a*, *Phys. Rev. Lett.* **58** (1987) 1490, [727(1987)], DOI: [10.1103/PhysRevLett.58.1490](https://doi.org/10.1103/PhysRevLett.58.1490) (cit. on p. 11).
- [34] Q. R. Ahmad et al., *Direct evidence for neutrino flavor transformation from neutral current interactions in the Sudbury Neutrino Observatory*, *Phys. Rev. Lett.* **89** (2002) 011301, DOI: [10.1103/PhysRevLett.89.011301](https://doi.org/10.1103/PhysRevLett.89.011301), arXiv: [nuc1-ex/0204008](https://arxiv.org/abs/nuc1-ex/0204008) [nuc1-ex] (cit. on p. 11).
- [35] T. Mannel, *Theory and Phenomenology of CP Violation*, *Nuclear Physics B - Proceedings Supplements* **167** (2007) 170, Proceedings of the 7th International Conference on Hyperons, Charm and Beauty Hadrons, ISSN: 0920-5632, DOI: [10.1016/j.nuclphysbps.2006.12.083](https://doi.org/10.1016/j.nuclphysbps.2006.12.083) (cit. on p. 11).
- [36] C. A. Baker et al., *Improved Experimental Limit on the Electric Dipole Moment of the Neutron*, *Phys. Rev. Lett.* **97** (13 2006) 131801, DOI: [10.1103/PhysRevLett.97.131801](https://doi.org/10.1103/PhysRevLett.97.131801) (cit. on p. 12).
- [37] W. D. Boer, *Grand unified theories and supersymmetry in particle physics and cosmology*, *Progress in Particle and Nuclear Physics* **33** (1994) 201, ISSN: 0146-6410, DOI: [10.1016/0146-6410\(94\)90045-0](https://doi.org/10.1016/0146-6410(94)90045-0) (cit. on p. 12).
- [38] G. Bertone, D. Hooper, and J. Silk, *Particle dark matter: evidence, candidates and constraints*, *Physics Reports* **405** (2005) 279, ISSN: 0370-1573, DOI: [10.1016/j.physrep.2004.08.031](https://doi.org/10.1016/j.physrep.2004.08.031) (cit. on pp. 12, 13).
- [39] K. G. Begeman, A. H. Broeils, and R. H. Sanders, *Extended rotation curves of spiral galaxies: Dark haloes and modified dynamics*, *Mon. Not. Roy. Astron. Soc.* **249** (1991) 523 (cit. on p. 13).
- [40] X-ray: NASA/CXC/CfA/M.Markevitch et al.; Optical: NASA/STScI; Magellan/U.Arizona/D.Clowe et al.; Lensing Map: NASA/STScI; ESO WFI; Magellan/U.Arizona/D.Clowe et al., *1E 0657-56: NASA Finds Direct Proof of Dark Matter*, <http://chandra.harvard.edu/photo/2006/1e0657/index.html> (cit. on p. 14).
- [41] D. Clowe, A. Gonzalez, and M. Markevitch, *Weak-Lensing Mass Reconstruction of the Interacting Cluster 1E 0657-558: Direct Evidence for the Existence of Dark Matter*, *The Astrophysical Journal* **604** (2004) 596, DOI: [10.1086/381970](https://doi.org/10.1086/381970) (cit. on p. 12).
- [42] W. Hu and S. Dodelson, *Cosmic Microwave Background Anisotropies*, *Ann. Rev. Astron. Astrophys.* **40** (2002) 171, DOI: [10.1146/annurev.astro.40.060401.093926](https://doi.org/10.1146/annurev.astro.40.060401.093926), arXiv: [astro-ph/0110414](https://arxiv.org/abs/astro-ph/0110414) [astro-ph] (cit. on p. 13).
- [43] N. Aghanim et al., *Planck 2018 results. VI. Cosmological parameters*, (2018), arXiv: [1807.06209](https://arxiv.org/abs/1807.06209) [astro-ph.CO] (cit. on p. 13).

- [44] P. A. R. Ade et al., *Planck 2015 results. XIII. Cosmological parameters*, *Astron. Astrophys.* **594** (2016) A13, DOI: [10.1051/0004-6361/201525830](https://doi.org/10.1051/0004-6361/201525830), arXiv: [1502.01589](https://arxiv.org/abs/1502.01589) [[astro-ph.CO](#)] (cit. on p. 15).
- [45] D. S. Akerib et al., *First Results from the LUX Dark Matter Experiment at the Sanford Underground Research Facility*, *Phys. Rev. Lett.* **112** (9 2014) 091303, DOI: [10.1103/PhysRevLett.112.091303](https://doi.org/10.1103/PhysRevLett.112.091303) (cit. on p. 17).
- [46] CMS Collaboration, *Search for dark matter, extra dimensions, and unparticles in monojet events in proton–proton collisions at $\sqrt{s}=8$ TeV*, *The European Physical Journal C* **75** (2015) 235, ISSN: 1434-6052, DOI: [10.1140/epjc/s10052-015-3451-4](https://doi.org/10.1140/epjc/s10052-015-3451-4) (cit. on p. 17).
- [47] UA2 Collaboration, *A search for new intermediate vector bosons and excited quarks decaying to two-jets at the CERN $\bar{p}p$ collider*, *Nuclear Physics B* **400** (1993) 3, ISSN: 0550-3213, DOI: [10.1016/0550-3213\(93\)90395-6](https://doi.org/10.1016/0550-3213(93)90395-6) (cit. on p. 17).
- [48] ATLAS Collaboration, *Search for new phenomena in the dijet mass distribution using pp collision data at $\sqrt{s}=8$ TeV with the ATLAS detector*, *Phys. Rev. D* **91** (5 2015) 052007, DOI: [10.1103/PhysRevD.91.052007](https://doi.org/10.1103/PhysRevD.91.052007) (cit. on p. 17).
- [49] ATLAS Collaboration, *Search for a dijet resonance produced in association with a leptonically decaying W or Z boson with the ATLAS detector at $\sqrt{s}=8$ TeV*, ATLAS-COM-CONF-2013-078 (2013) (cit. on p. 17).
- [50] CDF Collaboration, *Search for new particles decaying into dijets in proton-antiproton collisions at $\sqrt{s}=1.96$ TeV*, *Phys. Rev. D* **79** (11 2009) 112002, DOI: [10.1103/PhysRevD.79.112002](https://doi.org/10.1103/PhysRevD.79.112002) (cit. on p. 17).
- [51] CMS Collaboration, *Search for resonances and quantum black holes using dijet mass spectra in proton-proton collisions at $\sqrt{s}=8$ TeV*, *Phys. Rev. D* **91** (5 2015) 052009, DOI: [10.1103/PhysRevD.91.052009](https://doi.org/10.1103/PhysRevD.91.052009) (cit. on p. 17).
- [52] S. Höche, *Introduction to parton-shower event generators*, *Proceedings, Theoretical Advanced Study Institute in Elementary Particle Physics: Journeys Through the Precision Frontier: Amplitudes for Colliders (TASI 2014): Boulder, Colorado, June 2-27, 2014*, 2015 235, DOI: [10.1142/9789814678766_0005](https://doi.org/10.1142/9789814678766_0005), arXiv: [1411.4085](https://arxiv.org/abs/1411.4085) [[hep-ph](#)] (cit. on pp. 19, 21).
- [53] J. M. Campbell, J. W. Huston, and W. J. Stirling, *Hard Interactions of Quarks and Gluons: A Primer for LHC Physics*, *Rept. Prog. Phys.* **70** (2007) 89, DOI: [10.1088/0034-4885/70/1/R02](https://doi.org/10.1088/0034-4885/70/1/R02), arXiv: [hep-ph/0611148](https://arxiv.org/abs/hep-ph/0611148) [[hep-ph](#)] (cit. on pp. 19, 23).
- [54] R. M. Harris and K. Kousouris, *Searches for Dijet Resonances at Hadron Colliders*, *Int. J. Mod. Phys. A* **26** (2011) 5005, DOI: [10.1142/S0217751X11054905](https://doi.org/10.1142/S0217751X11054905), arXiv: [1110.5302](https://arxiv.org/abs/1110.5302) [[hep-ex](#)] (cit. on pp. 19, 24–26).
- [55] L. A. Beresford, *Searches for dijet resonances using $\sqrt{s}=13$ TeV proton-proton collision data recorded by the ATLAS detector at the Large Hadron Collider*, PhD thesis: University of Oxford, 2017 (cit. on pp. 20, 133, 181).
- [56] F. D. Aaron et al., *Combined Measurement and QCD Analysis of the Inclusive $e^\pm p$ Scattering Cross Sections at HERA*, *JHEP* **01** (2010) 109, DOI: [10.1007/JHEP01\(2010\)109](https://doi.org/10.1007/JHEP01(2010)109), arXiv: [0911.0884](https://arxiv.org/abs/0911.0884) [[hep-ex](#)] (cit. on p. 21).

- [57] M. Arneodo et al., *Measurement of the proton and deuteron structure functions, $F_2(p)$ and $F_2(d)$, and of the ratio σ_L / σ_T* , *Nucl. Phys.* **B483** (1997) 3, DOI: [10.1016/S0550-3213\(96\)00538-X](https://doi.org/10.1016/S0550-3213(96)00538-X), arXiv: [hep-ph/9610231](https://arxiv.org/abs/hep-ph/9610231) [hep-ph] (cit. on p. 21).
- [58] ATLAS Collaboration, *Measurement of the high-mass Drell–Yan differential cross-section in pp collisions at $\sqrt{s}=7$ TeV with the ATLAS detector*, *Phys. Lett.* **B725** (2013) 223, DOI: [10.1016/j.physletb.2013.07.049](https://doi.org/10.1016/j.physletb.2013.07.049), arXiv: [1305.4192](https://arxiv.org/abs/1305.4192) [hep-ex] (cit. on p. 21).
- [59] S. Chatrchyan et al., *Measurement of the Muon Charge Asymmetry in Inclusive $pp \rightarrow W + X$ Production at $\sqrt{s} = 7$ TeV and an Improved Determination of Light Parton Distribution Functions*, *Phys. Rev.* **D90** (2014) 032004, DOI: [10.1103/PhysRevD.90.032004](https://doi.org/10.1103/PhysRevD.90.032004), arXiv: [1312.6283](https://arxiv.org/abs/1312.6283) [hep-ex] (cit. on p. 21).
- [60] R. D. Ball et al., *Parton distributions from high-precision collider data*, *Eur. Phys. J.* **C77** (2017) 663, DOI: [10.1140/epjc/s10052-017-5199-5](https://doi.org/10.1140/epjc/s10052-017-5199-5), arXiv: [1706.00428](https://arxiv.org/abs/1706.00428) [hep-ph] (cit. on p. 22).
- [61] ATLAS Collaboration, *Measurement of inclusive jet and dijet cross-sections in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Journal of High Energy Physics* **2018** (2018) 195, ISSN: 1029-8479, DOI: [10.1007/JHEP05\(2018\)195](https://doi.org/10.1007/JHEP05(2018)195) (cit. on pp. 25, 27).
- [62] M. Zinser, *Search for new heavy charged bosons and measurement of high-mass Drell-Yan production in proton-proton collisions*, 2016, URL: <https://cds.cern.ch/record/2239722> (cit. on p. 26).
- [63] J. Alwall et al., *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*, *Journal of High Energy Physics* **2014** (2014) 79, ISSN: 1029-8479, DOI: [10.1007/JHEP07\(2014\)079](https://doi.org/10.1007/JHEP07(2014)079) (cit. on p. 28).
- [64] T. Sjöstrand et al., *An introduction to PYTHIA 8.2*, *Computer Physics Communications* **191** (2015) 159, ISSN: 0010-4655, DOI: [10.1016/j.cpc.2015.01.024](https://doi.org/10.1016/j.cpc.2015.01.024) (cit. on p. 28).
- [65] B. Andersson, G. Gustafson, G. Ingelman, and T. Sjöstrand, *Parton fragmentation and string dynamics*, *Physics Reports* **97** (1983) 31, ISSN: 0370-1573, DOI: [10.1016/0370-1573\(83\)90080-7](https://doi.org/10.1016/0370-1573(83)90080-7) (cit. on p. 28).
- [66] *ATLAS Public Event Displays from Run 2 Exotics Group Analysis*, Accessed: 2019-09-06, URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/ExoticsDisplayRun2> (cit. on p. 29).
- [67] R. Bajaj and S. Fahmy, *Mapping for maximum performance on FPGA DSP blocks*, *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* **35** (2015) 1, DOI: [10.1109/TCAD.2015.2474363](https://doi.org/10.1109/TCAD.2015.2474363) (cit. on p. 32).
- [68] Xilinx, *7 Series FPGAs Data Sheet: Overview*, DS180, v2.6, Xilinx, 2018 (cit. on p. 31).
- [69] N. Nottbeck, *Implementation of Deep Neural Networks on FPGAs for Sub-Microsecond Real-Time Triggering at the ATLAS Detector*, Masterarbeit: Johannes Gutenberg-Universität Mainz, 2019 (cit. on p. 31).
- [70] S. Maldaner, *Algorithmenentwicklung für den topologischen Trigger am ATLAS Detektor*, Diplomarbeit: Johannes Gutenberg-Universität Mainz, 2013 (cit. on pp. 31, 85, 86).
- [71] Xilinx, *7 Series FPGAs Configurable Logic Block: User Guide*, UG474, v1.8, Xilinx, 2016 (cit. on p. 32).

- [72] Xilinx, *UltraScale Architecture Libraries Guide*, UG974, v2018.3, Xilinx, 2018 (cit. on pp. 32, 33).
- [73] Xilinx, *UltraScale Architecture Configurable Logic Block: User Guide*, UG574, v1.5, Xilinx, 2017 (cit. on p. 33).
- [74] Xilinx, *7 Series DSP48E1 Slice: User Guide*, UG479, v1.10, Xilinx, 2018 (cit. on pp. 33, 34).
- [75] Xilinx, *7 Series FPGAs Memory Resources: User Guide*, UG473, v1.14, Xilinx, 2019 (cit. on pp. 34, 35).
- [76] Xilinx, *7 Series FPGAs GTX/GTH Transceivers: User Guide*, UG476, v1.12.1, Xilinx, 2018 (cit. on p. 34).
- [77] W. W. Peterson and D. T. Brown, *Cyclic Codes for Error Detection*, *Proceedings of the IRE* **49** (1961) 228, ISSN: 0096-8390, DOI: [10.1109/JRPROC.1961.287814](https://doi.org/10.1109/JRPROC.1961.287814) (cit. on p. 35).
- [78] *AdvancedTCA® Standard Overview*, <https://www.picmg.org/openstandards/advancedtca/>, Accessed: 2019-06-21 (cit. on pp. 35, 77).
- [79] C. G. Larrea et al., *IPbus: a flexible Ethernet-based control system for xTCA hardware*, *Journal of Instrumentation* **10** (2015) C02019, URL: <http://stacks.iop.org/1748-0221/10/i=02/a=C02019> (cit. on pp. 35, 36).
- [80] *IEEE Standard VHDL Language Reference Manual, IEEE Std 1076-2008 (Revision of IEEE Std 1076-2002) (2009) c1*, DOI: [10.1109/IEEESTD.2009.4772740](https://doi.org/10.1109/IEEESTD.2009.4772740) (cit. on p. 36).
- [81] *IEEE Standard for SystemVerilog—Unified Hardware Design, Specification, and Verification Language, IEEE Std 1800-2017 (Revision of IEEE Std 1800-2012) (2018) 1*, DOI: [10.1109/IEEESTD.2018.8299595](https://doi.org/10.1109/IEEESTD.2018.8299595) (cit. on p. 36).
- [82] Xilinx, *Vivado Design Suite User Guide: Design Flows Overview*, UG892, v2018.3, Xilinx, 2018 (cit. on p. 37).
- [83] Xilinx, *Block Memory Generator v8.3: LogiCORE IP Product Guide*, PG058, Xilinx, 2017 (cit. on pp. 39, 104).
- [84] Xilinx, *Vivado Design Suite Tutorial: Logic Simulation*, UG937, v2018.3, Xilinx, 2018 (cit. on p. 40).
- [85] L. Evans and P. Bryant, *LHC Machine*, *Journal of Instrumentation* **3** (2008) S08001 (cit. on p. 41).
- [86] *CERN who we are: Our Member States*, Accessed: 2019-10-16, URL: <https://home.cern/about/who-we-are/our-governance/member-states> (cit. on p. 41).
- [87] E. Mobs, *The CERN accelerator complex. Complexe des accélérateurs du CERN*, (2016), General Photo, URL: <https://cds.cern.ch/record/2197559> (cit. on p. 42).
- [88] H. Damerou, *Creation and storage of long and flat bunches in the LHC*, PhD thesis: Darmstadt, Tech. U., 2005, URL: <http://weblib.cern.ch/abstract?CERN-THESIS-2005-048> (cit. on p. 41).
- [89] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, *JINST* **3** (2008) S08003, DOI: [10.1088/1748-0221/3/08/S08003](https://doi.org/10.1088/1748-0221/3/08/S08003) (cit. on pp. 42, 47, 51, 52, 54, 55, 196).

- [90] CMS Collaboration, *The CMS experiment at the CERN LHC*, Journal of Instrumentation **3** (2008) S08004, URL: <http://stacks.iop.org/1748-0221/3/i=08/a=S08004> (cit. on p. 42).
- [91] *LHC Experiments*, Accessed: 2018-12-23, URL: <https://home.cern/science/experiments> (cit. on p. 42).
- [92] LHCb Collaboration, *The LHCb Detector at the LHC*, Journal of Instrumentation **3** (2008) S08005, URL: <http://stacks.iop.org/1748-0221/3/i=08/a=S08005> (cit. on p. 42).
- [93] ALICE Collaboration, *The ALICE experiment at the CERN LHC*, Journal of Instrumentation **3** (2008) S08002, URL: <http://stacks.iop.org/1748-0221/3/i=08/a=S08002> (cit. on p. 42).
- [94] W. Herr and B. Muratori, *Concept of luminosity, Intermediate accelerator physics. Proceedings, CERN Accelerator School, Zeuthen, Germany, September 15-26, 2003*, 2003 361, URL: <https://cds.cern.ch/record/941318/files/p361.pdf> (cit. on p. 43).
- [95] F. Bordry, *LHC Status and Plans*, Rencontres de Moriond, 2018 (cit. on p. 43).
- [96] P. Collier, *Status and Plans for the CERN accelerator Complex*, 29th International Symposium on Lepton Photon Interactions at High Energies, 2019, URL: https://indico.cern.ch/event/688643/contributions/3410189/attachments/1889928/3116533/190805-CERN_Status_and_Plans.pdf (cit. on p. 43).
- [97] *ATLAS Luminosity Public Results*, Accessed: 2019-08-01, URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2> (cit. on pp. 44, 45).
- [98] ATLAS Collaboration, *Luminosity determination in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector at the LHC*, Eur. Phys. J. **C76** (2016) 653, DOI: [10.1140/epjc/s10052-016-4466-1](https://doi.org/10.1140/epjc/s10052-016-4466-1), arXiv: 1608.03953 [hep-ex] (cit. on pp. 44, 56).
- [99] M. Ferro-Luzzi, W. Herr, and T. Pieloni, *LHC bunch filling schemes for commissioning and initial luminosity optimization*, 2008, URL: <http://cds.cern.ch/record/1114612> (cit. on p. 45).
- [100] R. Bailey and P. Collier, *Standard Filling Schemes for Various LHC Operation Modes*, 2003, URL: <https://cds.cern.ch/record/691782> (cit. on p. 45).
- [101] A. Hoecker and J. Stelzer, *ATLAS Run Query Tool*, 2019, URL: <https://atlas-runquery.cern.ch/> (cit. on p. 46).
- [102] J. Pequenaó, *Computer generated image of the whole ATLAS detector*, 2008, URL: <http://cds.cern.ch/record/1095924> (cit. on p. 48).
- [103] K. Potamianos, *The upgraded Pixel detector and the commissioning of the Inner Detector tracking of the ATLAS experiment for Run-2 at the Large Hadron Collider*, PoS **EPS-HEP2015** (2015) 261, arXiv: 1608.07850 [physics.ins-det] (cit. on p. 50).
- [104] G. Mullier, *The upgraded Pixel Detector of the ATLAS experiment for Run-2 at the Large Hadron Collider*, Journal of Instrumentation **11** (2016) C02061 (cit. on p. 50).
- [105] H. Kolanoski and N. Wermes, *Teilchendetektoren: Grundlagen und Anwendungen*, 1. Auflage, Springer, 2016, DOI: [10.1007/978-3-662-45350-6](https://doi.org/10.1007/978-3-662-45350-6) (cit. on p. 53).

- [106] ATLAS Collaboration, *Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC*, 2019, URL: <https://cds.cern.ch/record/2677054> (cit. on pp. 55, 184).
- [107] O. Viazlo and A. L. Collaboration, *ATLAS LUCID detector upgrade for LHC Run 2*, 2015, URL: <https://cds.cern.ch/record/2062038> (cit. on p. 55).
- [108] S. Maettig, *Luminosity Measurements with the ATLAS Detector*, PhD thesis: Hamburg U., 2012 (cit. on p. 55).
- [109] ATLAS Collaboration, *The ATLAS Data Acquisition and High Level Trigger system*, *Journal of Instrumentation* **11** (2016) P06008, DOI: [10.1088/1748-0221/11/06/p06008](https://doi.org/10.1088/1748-0221/11/06/p06008) (cit. on p. 56).
- [110] *ATLAS Approved Plots DAQ*, Accessed: 2019-06-28, URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/ApprovedPlotsDAQ> (cit. on p. 58).
- [111] S. Kolos, G. Boutsoukis, and R. Hauser, *High-Performance Scalable Information Service for the ATLAS Experiment*, *Journal of Physics: Conference Series* **396** (2012) 012026, DOI: [10.1088/1742-6596/396/1/012026](https://doi.org/10.1088/1742-6596/396/1/012026) (cit. on p. 57).
- [112] R. Achenbach et al., *The ATLAS Level-1 Calorimeter Trigger*, *Journal of Instrumentation* **3** (2008) P03001, URL: <http://stacks.iop.org/1748-0221/3/i=03/a=P03001> (cit. on pp. 59, 68, 196).
- [113] S. Artz et al., *The ATLAS Level-1 Muon Topological Trigger Information for Run 2 of the LHC*, *Journal of Instrumentation* **10** (2015) C02027 (cit. on p. 60).
- [114] S. Artz, *Physics performances with the new ATLAS Level-1 Topological trigger in Run 2*, 2016 (cit. on p. 60).
- [115] S. Artz et al., *Upgrade of the ATLAS Central Trigger for LHC Run-2*, 2014, URL: <https://cds.cern.ch/record/1969488> (cit. on p. 60).
- [116] R. Pöttgen, *Search for Dark Matter in events with a highly energetic jet and missing transverse momentum in proton-proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS Detector*, PhD thesis: Universität Mainz, 2014 (cit. on p. 60).
- [117] ATLAS Collaboration, *Performance of the ATLAS Trigger System in 2015*, *Eur. Phys. J. C* **77** (2017) 317, DOI: [10.1140/epjc/s10052-017-4852-3](https://doi.org/10.1140/epjc/s10052-017-4852-3), arXiv: [1611.09661](https://arxiv.org/abs/1611.09661) [hep-ex] (cit. on p. 61).
- [118] A. Ruiz-Martinez and the ATLAS Collaboration, *The Run-2 ATLAS Trigger System*, 2016, URL: <https://cds.cern.ch/record/2133909> (cit. on p. 61).
- [119] ATLAS Collaboration, *The ATLAS Simulation Infrastructure*, *The European Physical Journal C* **70** (2010) 823, ISSN: 1434-6052, DOI: [10.1140/epjc/s10052-010-1429-9](https://doi.org/10.1140/epjc/s10052-010-1429-9) (cit. on pp. 61, 62).
- [120] S. Agostinelli et al., *Geant4—a simulation toolkit*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **506** (2003) 250, ISSN: 0168-9002, DOI: [10.1016/S0168-9002\(03\)01368-8](https://doi.org/10.1016/S0168-9002(03)01368-8) (cit. on p. 61).
- [121] The DELPHES 3 collaboration, *DELPHES 3: a modular framework for fast simulation of a generic collider experiment*, *Journal of High Energy Physics* **2014** (2014) 57, ISSN: 1029-8479, DOI: [10.1007/JHEP02\(2014\)057](https://doi.org/10.1007/JHEP02(2014)057) (cit. on p. 62).

- [122] M. Cacciari, G. P. Salam, and G. Soyez, *FastJet user manual*, *The European Physical Journal C* **72** (2012) 1896, ISSN: 1434-6052, DOI: [10.1140/epjc/s10052-012-1896-2](https://doi.org/10.1140/epjc/s10052-012-1896-2) (cit. on p. 62).
- [123] ATLAS Collaboration, *Performance of the upgraded PreProcessor of the ATLAS Level-1 Calorimeter Trigger*, (2020), arXiv: [2005.04179](https://arxiv.org/abs/2005.04179) [[physics.ins-det](#)] (cit. on pp. 64–66, 172).
- [124] J. Jongmanns, *The Upgrade of the PreProcessor of the ATLAS Level-1 Calorimeter Trigger for the LHC Run-2*, PhD thesis: Universität Heidelberg, 2017 (cit. on p. 64).
- [125] A. C. Daniells, *Pile-Up Suppression in the ATLAS Level 1 Calorimeter Trigger and Searches for Higgs Boson Pair Production*, PhD thesis: Birmingham U., 2016, URL: http://www.ep.ph.bham.ac.uk/publications/thesis/acd_thesis.pdf (cit. on pp. 65, 135, 145).
- [126] C. Kahra and A. Watson, Private Communication, 2015–2019 (cit. on p. 67).
- [127] M. Weirich, Private Communication, 2020 (cit. on pp. 68, 194).
- [128] ATLAS Collaboration, *Performance of the ATLAS Track Reconstruction Algorithms in Dense Environments in LHC Run 2*, *Eur. Phys. J.* **C77** (2017) 673, DOI: [10.1140/epjc/s10052-017-5225-7](https://doi.org/10.1140/epjc/s10052-017-5225-7), arXiv: [1704.07983](https://arxiv.org/abs/1704.07983) [[hep-ex](#)] (cit. on p. 67).
- [129] ATLAS Collaboration, *Reconstruction of primary vertices at the ATLAS experiment in Run 1 proton–proton collisions at the LHC*, *Eur. Phys. J.* **C77** (2017) 332, DOI: [10.1140/epjc/s10052-017-4887-5](https://doi.org/10.1140/epjc/s10052-017-4887-5), arXiv: [1611.10235](https://arxiv.org/abs/1611.10235) [[physics.ins-det](#)] (cit. on p. 67).
- [130] ATLAS Collaboration, *Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1*, *Eur. Phys. J.* **C77** (2017) 490, DOI: [10.1140/epjc/s10052-017-5004-5](https://doi.org/10.1140/epjc/s10052-017-5004-5), arXiv: [1603.02934](https://arxiv.org/abs/1603.02934) [[hep-ex](#)] (cit. on pp. 68, 136).
- [131] M. Cacciari, G. P. Salam, and G. Soyez, *The anti- k_t jet clustering algorithm*, *JHEP* **04** (2008) 063, DOI: [10.1088/1126-6708/2008/04/063](https://doi.org/10.1088/1126-6708/2008/04/063), arXiv: [0802.1189](https://arxiv.org/abs/0802.1189) [[hep-ph](#)] (cit. on p. 69).
- [132] ATLAS Collaboration, *Jet energy measurement with the ATLAS detector in proton-proton collisions at $\sqrt{s} = 7$ TeV*, *Eur. Phys. J.* **C73** (2013) 2304, DOI: [10.1140/epjc/s10052-013-2304-2](https://doi.org/10.1140/epjc/s10052-013-2304-2), arXiv: [1112.6426](https://arxiv.org/abs/1112.6426) [[hep-ex](#)] (cit. on p. 69).
- [133] ATLAS Collaboration, *Jet energy scale measurements and their systematic uncertainties in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Rev.* **D96** (2017) 072002, DOI: [10.1103/PhysRevD.96.072002](https://doi.org/10.1103/PhysRevD.96.072002), arXiv: [1703.09665](https://arxiv.org/abs/1703.09665) [[hep-ex](#)] (cit. on pp. 69, 71, 72).
- [134] M. Cacciari and G. P. Salam, *Pileup subtraction using jet areas*, *Physics Letters B* **659** (2008) 119, ISSN: 0370-2693, DOI: <https://doi.org/10.1016/j.physletb.2007.09.077> (cit. on p. 70).
- [135] ATLAS Collaboration, *Determination of jet calibration and energy resolution in proton-proton collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector*, (2019), arXiv: [1910.04482](https://arxiv.org/abs/1910.04482) [[hep-ex](#)] (cit. on pp. 71, 143).

- [136] *JetEtMiss Public Plots: Jet energy scale and uncertainties in 2015-2017 data and simulation*, Accessed: 2020-02-20, URL: <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PLOTS/JETM-2018-006/> (cit. on p. 72).
- [137] C. Doglioni, *Measurement of the inclusive jet cross section with the ATLAS detector at the Large Hadron Collider*, PhD thesis: Oxford U., 2011, URL: <https://ora.ox.ac.uk/objects/uuid:1d2872e7-9fc4-4dd4-90bd-0b4e96fbb327> (cit. on pp. 73, 137).
- [138] 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*, 2010 (cit. on p. 73).
- [139] ATLAS Collaboration, *Technical Design Report for the Phase-I Upgrade of the ATLAS TDAQ System*, CERN-LHCC-2013-018 (2013), Final version presented to December 2013 LHCC. (cit. on pp. 77, 193).
- [140] U. Schäfer, *L1Topo: Run-2 \rightarrow Phase-1 (\rightarrow Phase-2 L0Topo)*, L1Calo Joint Meeting, 2016, URL: <https://indico.cern.ch/event/463317/contributions/1968407/attachments/1225015/1792861/Uli-Topo2.pdf> (cit. on p. 78).
- [141] *Avago MiniPOD 12x14G Receiver Module AFBR-824VxyZ*, <https://www.broadcom.com/products/fiber-optic-modules-components/networking/embedded-optical-modules/minipod/afbr-824vxyz>, Accessed: 2020-07-21 (cit. on p. 77).
- [142] S. Artz, Private Communication, 2017 (cit. on p. 79).
- [143] C. Kahra, Private Communication, 2019 (cit. on p. 80).
- [144] M. Palka, *L1Calo Topo Phase I readout specification*, 2019, URL: <https://cds.cern.ch/record/2654187> (cit. on p. 81).
- [145] S. Artz et al., *L1Topo Firmware Documentation*, ATLAS Internal NOTE, 2017, URL: <https://gitlab.cern.ch/sartz/L1TopoFirmwareDocumentation> (cit. on p. 83).
- [146] W. T. Fedorko and P. P. Plucinski, *CMX Base Function FPGA Firmware functionality*, 2016, URL: https://web.pa.msu.edu/hep/atlas/l1calo/cmz/specification/5_firmware_review/CMX_BaseFunction_Firmware_V1p1_after_review.pdf (cit. on p. 83).
- [147] J. F. Damp, *Level-1 Topological Processor (L1Topo) ATLAS TDAQ Phase-I Upgrade Firmware Algorithms*, 2019, URL: <https://cds.cern.ch/record/2654103> (cit. on p. 87).
- [148] R. Simoniello, *The ATLAS Level-1 Topological Processor: from design to routine usage in Run-2*, 2018, URL: <https://cds.cern.ch/record/2649959> (cit. on p. 91).
- [149] *ATLAS L1Calo Trigger Public Results*, Accessed: 2019-09-20, URL: https://twiki.cern.ch/twiki/bin/view/AtlasPublic/L1CaloTriggerPublicResults#L1Calo_Performance_plots_2018_Ru (cit. on p. 92).
- [150] ATLAS Collaboration, *The ATLAS Tau Trigger in Run 2*, (2017) (cit. on pp. 92, 93).
- [151] *ATLAS Trigger Operation Public Results: Summary Plots*, Accessed: 2019-03-12, URL: <https://twiki.cern.ch/twiki/pub/AtlasPublic/TriggerOperationPublicResults/l1topo.326945.short.pdf> (cit. on p. 94).

- [152] *ATLAS Trigger Operation Public Results: Summary Plots*, Accessed: 2019-03-12, URL: https://twiki.cern.ch/twiki/pub/AtlasPublic/TriggerOperationPublicResults/invM_mumu.pdf (cit. on p. 94).
- [153] C. W. Kalderon, N. Sherafati, A. Martyniuk, and D. R. Zaripovas, *Jet Trigger Public Plots – October 2017*, 2017, URL: <https://cds.cern.ch/record/2288217> (cit. on p. 94).
- [154] mcviz, *mcviz*, <https://github.com/mcviz/mcviz>, 2015 (cit. on p. 95).
- [155] ATLAS Collaboration, *Search for new phenomena in dijet events using 37 fb^{-1} of pp collision data collected at $\sqrt{s}=13\text{ TeV}$ with the ATLAS detector*, *Phys. Rev. D* **96** (2017) 052004, DOI: 10.1103/PhysRevD.96.052004, arXiv: 1703.09127 [hep-ex] (cit. on p. 97).
- [156] B. A. Dobrescu and F. Yu, *Coupling-Mass Mapping of Dijet Peak Searches*, *Phys. Rev. D* **88** (2013) 035021, [Erratum: *Phys. Rev. D* **90**, no. 7, 079901 (2014)], DOI: 10.1103/PhysRevD.88.035021, 10.1103/PhysRevD.90.079901, arXiv: 1306.2629 [hep-ph] (cit. on p. 98).
- [157] *ATLAS Exotics Public Results: Dark Matter Coupling Summary Plots 2017*, Accessed: 2019-02-12, URL: https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/CombinedSummaryPlots/EXOTICS/ATLAS_DarkMatterCoupling_Summary/ATLAS_DarkMatterCoupling_Summary_201703.png (cit. on p. 99).
- [158] R. Brun and F. Rademakers, *ROOT: An object oriented data analysis framework*, *Nucl. Instrum. Meth. A* **389** (1997) 81, ed. by M. Werlen and D. Perret-Gallix, DOI: 10.1016/S0168-9002(97)00048-X (cit. on p. 107).
- [159] M. Baak et al., *Data Quality Status Flags and Good Run Lists for Physics Analysis in ATLAS*, 2009, URL: <https://cds.cern.ch/record/1168026> (cit. on p. 113).
- [160] E. J. Gallas et al., *Conditions and configuration metadata for the ATLAS experiment*, *Journal of Physics: Conference Series* **396** (2012) 052033, DOI: 10.1088/1742-6596/396/5/052033 (cit. on p. 114).
- [161] J. Fulachier, J. Odier, and F. L. and, *ATLAS Metadata Interface (AMI), a generic metadata framework*, *Journal of Physics: Conference Series* **898** (2017) 062001, DOI: 10.1088/1742-6596/898/6/062001 (cit. on p. 114).
- [162] S. Mättig, *The Online Luminosity Calculator of ATLAS*, *Journal of Physics: Conference Series* **331** (2011) 022035, DOI: 10.1088/1742-6596/331/2/022035 (cit. on p. 116).
- [163] ATLAS Collaboration, *Studies of the performance of the ATLAS detector using cosmic-ray muons*, *The European Physical Journal C* **71** (2011) 1593, ISSN: 1434-6052, DOI: 10.1140/epjc/s10052-011-1593-6 (cit. on p. 118).
- [164] ATLAS Collaboration, *Beam-induced and cosmic-ray backgrounds observed in the ATLAS detector during the LHC 2012 proton-proton running period*, *JINST* **11** (2016) P05013, DOI: 10.1088/1748-0221/11/05/P05013, arXiv: 1603.09202 [hep-ex] (cit. on p. 119).
- [165] J. Abdallah et al., *The Laser calibration of the ATLAS Tile Calorimeter during the LHC run 1*, *Journal of Instrumentation* **11** (2016) T10005, DOI: 10.1088/1748-0221/11/10/t10005 (cit. on p. 120).
- [166] F. Scuri, *Performance of the ATLAS Tile LaserII Calibration System*, tech. rep. ATL-TILECAL-PROC-2015-020, CERN, 2015, DOI: 10.1109/NSSMIC.2015.7581768 (cit. on p. 120).

- [167] ATLAS collaboration, *Selection of jets produced in 13 TeV proton-proton collisions with the ATLAS detector*, tech. rep. ATLAS-CONF-2015-029, CERN, 2015, URL: <https://cds.cern.ch/record/2037702> (cit. on p. 126).
- [168] ATLAS Collaboration, *ATLAS Pythia 8 tunes to 7 TeV datas*, tech. rep. ATL-PHYS-PUB-2014-021, CERN, 2014, URL: <https://cds.cern.ch/record/1966419> (cit. on p. 130).
- [169] Z. Marshall, *Re-re-defining the Standard QCD Di-Jet Samples: Beginning to like event weights*, 2015, URL: <https://cds.cern.ch/record/2016630> (cit. on p. 131).
- [170] ATLAS Collaboration, *Performance of the ATLAS muon triggers in Run 2*, (2020), arXiv: 2004.13447 [hep-ex] (cit. on p. 138).
- [171] ATLAS Collaboration, *A method for the construction of strongly reduced representations of ATLAS experimental uncertainties and the application thereof to the jet energy scale*, tech. rep. ATL-PHYS-PUB-2015-014, CERN, 2015, URL: <https://cds.cern.ch/record/2037436> (cit. on p. 143).
- [172] ATLAS Collaboration, *A search for new physics in dijet mass and angular distributions in pp collisions at $\sqrt{s}=7$ TeV measured with the ATLAS detector*, *New Journal of Physics* **13** (2011) 053044, DOI: 10.1088/1367-2630/13/5/053044 (cit. on p. 144).
- [173] ATLAS Collaboration, *Search for new phenomena in dijet events using 37 fb^{-1} of pp collision data collected at $\sqrt{s}=13$ TeV with the ATLAS detector*, *Phys. Rev.* **D96** (2017) 052004, DOI: 10.1103/PhysRevD.96.052004, arXiv: 1703.09127 [hep-ex] (cit. on pp. 163, 164).
- [174] R. D. Cousins, *Why isn't every physicist a Bayesian?* *American Journal of Physics* **63** (1995) 398, DOI: 10.1119/1.17901, eprint: <https://doi.org/10.1119/1.17901> (cit. on p. 163).
- [175] K. Pachal, *Search for new physics in the dijet invariant mass spectrum at 8 TeV*, 2015, URL: <https://cds.cern.ch/record/2063032> (cit. on pp. 163, 181, 212).
- [176] F. James and M. Roos, *Minuit - a system for function minimization and analysis of the parameter errors and correlations*, *Computer Physics Communications* **10** (1975) 343, ISSN: 0010-4655, DOI: 10.1016/0010-4655(75)90039-9 (cit. on p. 164).
- [177] CDF Collaboration, *Search for new particles decaying into dijets in proton-antiproton collisions at $\sqrt{s}=1.96$ TeV*, *Phys. Rev. D* **79** (11 2009) 112002, DOI: 10.1103/PhysRevD.79.112002 (cit. on p. 164).
- [178] G. Choudalakis, *On hypothesis testing, trials factor, hypertests and the BumpHunter*, *PHYSTAT 2011*, 2011, arXiv: 1101.0390 [physics.data-an] (cit. on pp. 166, 167).
- [179] G. Choudalakis and D. Casadei, *Plotting the differences between data and expectation*, *The European Physical Journal Plus* **127** (2012) 25, ISSN: 2190-5444, DOI: 10.1140/epjp/i2012-12025-y (cit. on p. 169).
- [180] ATLAS Collaboration, *Tagging and suppression of pileup jets with the ATLAS detector*, tech. rep. ATLAS-CONF-2014-018, CERN, 2014, URL: <https://cds.cern.ch/record/1700870> (cit. on p. 171).
- [181] O. Behnke, K. Kröniger, G. Schott, and T. Schörner-Sadenius, *Data Analysis in High Energy Physics*, John Wiley & Sons, Ltd, 2013, ISBN: 9783527653416, DOI: 10.1002/9783527653416 (cit. on pp. 181, 209).

- [182] A. Caldwell, D. Kollár, and K. Kröninger, *BAT – The Bayesian analysis toolkit*, **Computer Physics Communications** **180** (2009) 2197, ISSN: 0010-4655, DOI: [10.1016/j.cpc.2009.06.026](https://doi.org/10.1016/j.cpc.2009.06.026) (cit. on p. 185).
- [183] A. Gelman et al., *Bayesian Data Analysis, Third Edition*, Chapman & Hall/CRC Texts in Statistical Science, Taylor & Francis, 2013, ISBN: 9781439840955, URL: <http://www.stat.columbia.edu/~gelman/book/> (cit. on pp. 185, 213).
- [184] J. Dandoy, *Search for New Phenomena in Dijet Topologies from Proton-Proton Collisions at $\sqrt{s} = 13$ TeV*, 2016, URL: <https://cds.cern.ch/record/2224312> (cit. on p. 186).
- [185] ATLAS Collaboration, *Search for new phenomena in the dijet mass distribution using pp collision data at $\sqrt{s} = 8$ TeV with the ATLAS detector*, **Phys. Rev. D** **91** (5 2015) 052007, DOI: [10.1103/PhysRevD.91.052007](https://doi.org/10.1103/PhysRevD.91.052007) (cit. on p. 186).
- [186] ATLAS Collaboration, *Dark matter summary plot: Update May 2020*, tech. rep. ATL-PHYS-PUB-2020-015, CERN, 2020, URL: <https://cds.cern.ch/record/2719323> (cit. on pp. 190, 193).
- [187] Xilinx, *UltraScale Architecture and Product Data Sheet: Overview*, DS890, v3.6, Xilinx, 2018 (cit. on p. 193).
- [188] S. S. Wilks, *The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses*, **Ann. Math. Statist.** **9** (1938) 60, DOI: [10.1214/aoms/1177732360](https://doi.org/10.1214/aoms/1177732360) (cit. on p. 194).
- [189] N. Berger, *Recommendations for the Modeling of Smooth Backgrounds*, tech. rep. ATL-COM-PHYS-2020-418, CERN, 2020, URL: <https://cds.cern.ch/record/2719221> (cit. on p. 194).
- [190] G. Bohm and G. Zech, *Statistics of weighted Poisson events and its applications*, **Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment** **748** (2014) 1, ISSN: 0168-9002, DOI: [10.1016/j.nima.2014.02.021](https://doi.org/10.1016/j.nima.2014.02.021) (cit. on p. 212).
- [191] F. Beaujean, A. Caldwell, D. Kollár, K. Kröninger, and S. Pashapour, *Bayesian Analysis Toolkit in Searches*, (2011) 209, DOI: [10.5170/CERN-2011-006.209](https://doi.org/10.5170/CERN-2011-006.209) (cit. on p. 213).
- [192] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, *Equation of state calculations by fast computing machines*, **The journal of chemical physics** **21** (1953) 1087 (cit. on p. 213).
- [193] B. A. T. Developers, *Bayesian Analysis Toolkit Manual*, 2020, URL: <https://github.com/bat/bat/tree/manual/doc/manual/figures> (cit. on p. 214).

Acknowledgements

Personenbezogene Daten sind aus Datenschutzgründen aus der elektronischen Version gelöscht.

Curriculum Vitae

Personenbezogene Daten sind aus Datenschutzgründen aus der elektronischen Version gelöscht.