

Searches for Natural Supersymmetry in Hadronic Final States with Heavy Flavor at ATLAS

By Bart Clayton Butler

SEARCHES FOR NATURAL SUPERSYMMETRY IN HADRONIC
FINAL STATES WITH HEAVY FLAVOR AT ATLAS

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF PHYSICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Bart C. Butler

December 2012

© 2012 by Bart Clayton Butler. All Rights Reserved.

Re-distributed by Stanford University under license with the author.



This work is licensed under a Creative Commons Attribution-Noncommercial 3.0 United States License.

<http://creativecommons.org/licenses/by-nc/3.0/us/>

This dissertation is online at: <http://purl.stanford.edu/fk995wq4388>

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Dong Su, Primary Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

David MacFarlane

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Michael Peskin

Approved for the Stanford University Committee on Graduate Studies.

Patricia J. Gumport, Vice Provost Graduate Education

This signature page was generated electronically upon submission of this dissertation in electronic format. An original signed hard copy of the signature page is on file in University Archives.

Abstract

This thesis presents the hadronic-channel supersymmetric searches for direct sbottom and gluino-mediated sbottom and stop production performed on 4.71 fb^{-1} of $\sqrt{s} = 7$ TeV data collected by the ATLAS detector at the Large Hadron Collider. These signatures are characterized by final states with multiple b -tagged jets and missing transverse energy (\cancel{E}_T) and the analysis strategy is chosen accordingly. Particular emphasis is placed on the utilization of the simplified models approach in signal characterization, optimization, and interpretation of results. No significant excess is observed resulting in limits set at 95% confidence level. Relative to the previous versions of the analyses, this iteration represents a several-fold increase in sensitivity to the new physics signatures considered. This is largely due to the use of three b -tag signal regions as well as signal regions based on initial state radiation.

Preface

This work represents a milestone in my continuing journey from knowing nothing to being fully aware of my ignorance. With regard to particle physics, the journey began when, after graduating from college in June, 2006, I relocated to the SLAC National Accelerator Laboratory (Stanford Linear Accelerator Center, as it was known then) to look for rare B -meson decays with the BaBar Collaboration. Having developed an interest in experimental particle physics, and as BaBar was winding down (the last data was taken in April, 2008), I joined the SLAC ATLAS group in January, 2007.

I spent the summer of 2007 at CERN, where I tried to overcome the enormously steep ATLAS learning curve while helping with the final connectivity tests on the newly-inserted pixel detector. April, 2008 marked my semi-permanent relocation to CERN, where I would spend the next two-and-a-half years working on a multitude of disparate performance-oriented projects, including track jets, track-based corrections to calorimeter jets, heavy flavor tagging of track and calorimeter jets, an overhaul of the muon spectrometer alignment code, and refinement of the heavy flavor jet triggers. With both the first (September, 2008) and second (November, 2009) attempts to begin the collision era at the Large Hadron Collider (LHC), this was an exciting time to be at CERN. Thanks to pixel detector and inner detector shifts, I was not only given a front-row seat but was also occasionally on stage throughout this tumultuous and interesting time.

After the first $\sqrt{s} = 7$ TeV data had arrived, it became clear that it was time to change my focus from physics performance to physics analysis. Though some preparation had been done at CERN, the analysis work began in earnest upon my return to SLAC in August, 2010. In collaboration with members of the SLAC theory group,

my experimental colleagues and I developed the optimization strategy which would guide all subsequent ATLAS hadronic-channel heavy flavor supersymmetry (SUSY) searches, a topic which builds on much of my performance work. My interest in efficiency lead to the development of a powerful and flexible analysis framework. This framework has been used for all iterations of the hadronic-channel heavy flavor SUSY searches as well as several other analyses both at SLAC and at other institutions. The first ATLAS conference note based on the new optimization scheme was presented at the European Physical Society conference in July, 2011 [1], and six more have followed as of December, 2012, with three becoming published papers [2, 3, 4, 5, 6, 7].

This thesis is an account of the analyses performed on the full $\sqrt{s} = 7$ TeV 2011 dataset. A large effort was made to focus as much on the “why” of things as the “what”, and not simply to reveal the technical motivations. I hope this document also reflects the wonder we feel as physicists when confronted by the mysteries of nature and our gratitude for the amazing opportunity to learn that the LHC represents.

“If I could remember the names of all these particles, I’d be a botanist.”

– Enrico Fermi

“No man should escape our universities without knowing how little he knows.”

– J. Robert Oppenheimer

“In God we trust; all others must bring data.”

– W. Edwards Deming

Acknowledgements

I am fortunate to be acquainted with, and in some cases also related to, a number of extraordinary people. I am grateful to my mother, Cathy, and my father, Jim, for their unshakable confidence in me and my abilities, and for their unwavering support throughout my life. I would like to thank my brother, Kort, for being my oldest friend and for his continuing mission to make sure I do not take myself too seriously. My fiancée, Ruchita, I would like to thank for her patience as well as for being the wonderful, fun, caring person that she is.

I am indebted to my advisors, Su Dong and Ariel Schwartzman, for their guidance and support, and appreciate above all the flexibility they gave me in letting me pursue my interests. My education and I have benefited enormously from interactions with other current and former members of the SLAC ATLAS group with particular thanks to Charlie Young, Rainer Bartoldus, Per Hansson, Paul Jackson, and Claus Horn. Similarly, my graduate student experience would not have been as scientifically fulfilling or as fun without David Miller, Dan Silverstein, Max Swiatlowski, Josh Cogan, Katie Malone, and Jim Black, my fellow SLAC ATLAS graduate students.

I have also benefited greatly from interactions with non-ATLAS SLAC colleagues. Discussions with Matt Graham and Manuel Franco were instrumental in teaching me analysis and writing this thesis, respectively. Michael Peskin, Tom Rizzo, Joanne Hewett, and, most of all, Jay Wacker have my everlasting gratitude for their earnest attempts to explain supersymmetry to me using short, simple sentences. Their students have also been very helpful, with special thanks to Ahmed Ismail. Michael Peskin and David MacFarlane, my BaBar advisor, were also kind enough to serve on my reading and defense committees, for which I am profoundly grateful.

ATLAS is a huge international collaboration, and I was fortunate to work with some amazingly-talented non-SLAC physicists. The input and assistance of David Lopez, Zach Marshall, Kerstin Perez, Heather Gray, and Shih-Chieh Hsu were crucial to my performance work. The success of the heavy flavor supersymmetry analyses was and continues to be due to the extraordinary competence of the people working on them, in particular Antoine Marzin, Alan Tua, Takashi Yamanaka, Mirjam Fehling-Kaschek, Iacopo Vivarelli, Monica d’Onofrio, and Davide Costanzo.

Contents

Abstract	iv
Preface	v
Acknowledgements	vii
1 Introduction	1
2 The Standard Model	4
2.1 The Strong Sector	4
2.2 The Electroweak Sector	6
2.3 The Higgs Sector and Electroweak Symmetry Breaking	7
2.4 Fermions	10
2.5 Standard Model Particle Content Summary	13
2.6 The Hierarchy Problem	13
3 Supersymmetry and Heavy Flavor	18
3.1 The Minimal Supersymmetric Standard Model	19
3.2 Soft SUSY-breaking in the MSSM	20
3.3 Mass Eigenstates of the MSSM	24
3.4 Naturalness and Heavy Flavor	29
3.5 R-Parity and Dark Matter	32
3.6 Gauge Unification	33

4	Simplified Models	36
4.1	Top-down Approach	36
4.2	Bottom-up/Simplified Models Approach	41
4.3	Direct Sbottom/Stop Production	41
4.4	Gluino-mediated Sbottom/Stop Production	43
4.5	Beyond Supersymmetry	48
5	The Large Hadron Collider	49
5.1	Project Proposal	49
5.2	Machine Design	51
5.3	Operation	53
6	The ATLAS Detector	56
6.1	Coordinate System	58
6.2	Inner Detector	59
6.2.1	Silicon Pixel Tracker	59
6.2.2	Silicon Strip Tracker	62
6.2.3	Transition Radiation Tracker	63
6.3	Calorimetry	63
6.3.1	Liquid Argon Calorimeters	64
6.3.2	Tile Calorimeter	70
6.4	Muon Spectrometer	72
6.4.1	Precision Subsystems	73
6.4.2	Trigger Subsystems	74
6.5	Magnet System	76
6.5.1	Solenoid	76
6.5.2	Toroids	78
6.6	Forward Detectors	78
6.6.1	BCM	79
6.6.2	MBTS	79
6.6.3	LUCID	79
6.6.4	ZDC	80

6.6.5	ALFA	80
6.7	Trigger and Data Acquisition	81
6.7.1	Level 1	81
6.7.2	Level 2	82
6.7.3	Event Filter	83
7	Object Reconstruction and Datasets	84
7.1	Data	84
7.2	Monte Carlo	85
7.3	Tracks and Vertices	87
7.4	Electrons	88
7.5	Muons	89
7.6	Topological Clusters	90
7.7	Jets	91
7.8	Heavy Flavor Tagging	92
7.9	Missing Transverse Energy	99
7.10	Overlap Removal	99
7.11	Triggers	101
7.12	Event Cleaning	101
7.13	Monte Carlo Reweighting	103
8	Search Optimization	105
8.1	Model Kinematics	105
8.1.1	A Representative Example	106
8.1.2	Influence of the Hard-Scatter Process	108
8.1.3	Compressed Scenarios	109
8.2	Optimization Procedure	111
8.3	Variables	113
8.4	Direct Sbottom Signal Regions	116
8.5	Gluino-mediated Sbottom/Stop Signal Regions	122

9	Background Estimation Techniques	126
9.1	Monte Carlo-based Predictions	126
9.2	Monte Carlo Systematic Uncertainties	126
9.2.1	Detector Uncertainties	127
9.2.2	Theoretical/Generator Uncertainties	129
9.3	Transfer Factors	130
9.4	Profile Likelihood	131
9.5	Matrix Method	133
9.6	Jet Smearing	134
10	Direct Sbottom Search	138
10.1	Top Control Regions	138
10.2	Z Control Regions	139
10.3	QCD Control Regions	141
10.4	Results	144
11	Glino-mediated Sbottom/Stop Search	149
11.1	Top Control Regions	150
11.2	QCD Control Regions	151
11.3	Results	155
12	Interpretation of Results	159
12.1	Profile Likelihood Ratio Test and Limits	159
12.2	CL_s	161
12.3	Model-dependent Limits	162
12.4	Model-independent Limits	166
12.5	Discussion and Caveats	167
13	Conclusion and Prospects	173
A	Track Jets	175
A.1	Monte Carlo and Calorimeter Jets	177
A.2	Reconstruction	177

A.3 Optimization	181
A.4 Performance	182
A.5 Pileup	187
A.6 Jet Angular Resolution	188
B PROOF, the Parallel ROOT Facility	194
B.1 SLAC Cluster	195
B.2 Development	196
B.3 Cluster Configuration Optimization	199
C ProofAna, an analysis framework	202
C.1 Design Goals	202
C.2 Class Structure	204
C.3 Event Data Model	206
C.4 Heavy Flavor Supersymmetry Analysis	207
Bibliography	210

List of Tables

2.1	Gauge fields of the Standard Model. The massless electroweak gauge fields W_μ^i and B_μ become the massive W_μ^\pm , Z_μ^0 (weak), and massless A_μ (electromagnetic) fields after electroweak symmetry-breaking via the Higgs mechanism. All gauge fields are spin 1.	13
2.2	Fermionic fields and the Higgs field of the Standard Model. All fermions have spin $\frac{1}{2}$, while the Higgs field has spin 0.	14
3.1	Gauge and mass eigenstates of the supersymmetric particles of the MSSM. A dash (-) indicates that the gauge and mass eigenstates are the same, at least to good approximation with typical choices of the soft SUSY-breaking parameters.	29
7.1	Integrated luminosity passing the Good Runs List as a function of data-taking period in 2011.	85
7.2	Triggers used in the analyses as a function of 2011 data-taking period.	102
8.1	Summary of the optimization-derived event selection for the direct sbottom signal regions.	117
8.2	Summary of the optimization-derived event selection for the gluino-mediated sbottom and stop signal regions.	123
9.1	Selections for “Mercedes” events after the single jet trigger. Jets are ordered by azimuthal distance from the $\vec{\cancel{E}}_T - j_1$ is closest, and j_N is farthest away.	137

10.1	Definition of the control regions adopted for top background estimation, after single lepton trigger and common event cleaning.	139
10.2	For each one-lepton ($t\bar{t}$) control region, the observed event yield is compared to the prediction obtained from the profile likelihood fit and to the expectation from Monte Carlo. The errors on the profile likelihood fit results include all statistical and systematic uncertainties, and thus are partially correlated.	141
10.3	Definition of the control regions adopted for the Z background estimation, after single lepton trigger and common event cleaning.	142
10.4	For each two-lepton (Z and $t\bar{t}$) control region, the observed event yield is compared to the prediction obtained from the profile likelihood fit and to the expectation from Monte Carlo. The errors on the profile likelihood fit results include all statistical and systematic uncertainties considered, and thus are partially correlated.	142
10.5	For each signal region, the observed event yield is compared with the prediction obtained from the profile likelihood fit. The contribution of each SM process to each signal region from Monte Carlo Simulation is also shown. The errors on the profile likelihood fit results include all statistical and systematic uncertainties, and thus are partially correlated.	148
11.1	Definition of the four control regions used to estimate the $t\bar{t}$ background, after the jet + \cancel{E}_T trigger and common event cleaning.	151
11.2	Expected Standard Model event yield and observed data events in the four $t\bar{t}$ control regions. The multi-jet estimate, included in the category “Others” and accounting for $\approx 20\%$ of the total yield, was estimated with the jet-smearing technique. All other background yields were taken from Monte Carlo. The errors include all detector-related and theoretical systematic uncertainties.	151

11.3	Comparison between the results of the profile likelihood fit and the numbers of observed events in the five signal regions. The $t\bar{t}$ event yield predicted by the MC simulation is quoted in parentheses, and all other non-QCD backgrounds are taken from the Monte Carlo in both cases. QCD multi-jet events contribute less than 5% of the total background in all signal regions. The errors include all experimental, theoretical, and statistical uncertainties.	155
12.1	Observed and expected 95% confidence level upper limits on the non-Standard Mode contributions (number of events) to all signal regions. Systematic uncertainties on the SM background estimates are included.	167
A.1	Configurable algorithm parameters for <code>JetTrackZClusterTool</code>	178

List of Figures

2.1	One-loop diagram of the correction to the fermion mass from the Higgs. This correction is proportional to the fermion mass and is logarithmically divergent with the cutoff scale (Equation 2.31)	14
2.2	Fermion (left) and scalar (right) one-loop diagrams which generate opposite-sign quadratically-divergent corrections to the Higgs mass. Cancellation of the quadratic divergences is complete for two scalars and one fermion with couplings obeying the relation $\lambda_{\tilde{f}} = y_f^2$	16
3.1	Example diagrams which can contribute to the $K^0 \leftrightarrow \bar{K}^0$ mixing via loops involving squarks and gluinos. The X symbol represents the flavor-changing interaction. Constraints on kaon mixing can be used to constrain the elements of the SUSY-breaking parameters $m_{\tilde{L}}^2$, $m_{\tilde{e}}^2$ and \mathbf{a}_e	22
3.2	Renormalization group evolution of gauge couplings in the Standard Model (left) and the MSSM (right). In the MSSM the gauge couplings unify at $\approx 10^{16}$ GeV	35
4.1	An example of an mSUGRA-style exclusion plot in terms of m_0 and $m_{1/2}$, the sfermion and gaugino masses [76]	38
4.2	Sample mSUGRA mass spectrum [77]	40
4.3	Direct third-generation squark pair production and decay to two third-generation quarks and two neutralinos.	42

4.4	Mass spectrum of the direct stop production (left) and direct sbottom production (right) simplified models. All other flavors of squarks in each model and the gluino are considered heavy and therefore decoupled.	42
4.5	Dominant leading-order diagrams for third-generation squark pair production at the LHC. All other flavors of squarks and the gluino are assumed heavy and therefore decoupled.	43
4.6	Gluino-mediated production of stop/sbotttom squarks. The final state involves four top or bottom quarks and two neutralinos.	44
4.7	Mass spectra for gluino decays to 4 top quarks (left) or four bottom quarks (right) and 2 neutralinos.	44
4.8	Gluino and sbottom/stop pair production next-to-leading order + next-to-leading log (NLO+NLL) cross sections at the LHC at $\sqrt{s} = 7$ TeV [78, 79, 80] computed using PROSPINO [81].	45
4.9	Dominant leading order diagrams for gluino pair production at the LHC. All squarks are assumed heavy and thus decoupled.	46
5.1	The CERN accelerator complex © 2008 CERN	52
5.2	Peak instantaneous LHC luminosity for each fill for the 2011 $\sqrt{s} = 7$ TeV data-taking period.	54
5.3	Integrated LHC luminosity delivered to and record by ATLAS per day for the 2011 $\sqrt{s} = 7$ TeV data-taking period.	55
6.1	Computer-generated image of the ATLAS detector, with humans for scale. ATLAS Experiment © 2012 CERN.	57
6.2	Diagram of the inner detector and components, ATLAS Experiment © 2012 CERN.	60
6.3	Diagram showing the radial positions of the inner detector components, ATLAS Experiment © 2012 CERN.	61
6.4	Diagram of the calorimeter subsystems, ATLAS Experiment © 2012 CERN.	65
6.5	Comparison between a triangular current pulse in a liquid-argon cell and the output signal after bipolar pulse shaping [25].	66

6.6	The accordion structure of the electromagnetic liquid-argon calorimeters, ATLAS Experiment © 2012 CERN.	67
6.7	Illustration of the segmentation of the electromagnetic barrel (EMB) calorimeter [25].	69
6.8	Diagram of the liquid-argon tubes embedded in absorber matrix within the forward calorimeter (FCAL) [25].	70
6.9	Diagram of a tile calorimeter module consisting of alternating steel absorber and plastic scintillator tile [25].	71
6.10	Diagram of the ATLAS muon systems, ATLAS Experiment © 2012 CERN.	72
6.11	Schematic of cathode strip chamber read-out system [25].	74
6.12	Illustration of two Resistive Plate Chambers (RPCs) [25].	75
6.13	Diagram of the magnet system, including both barrel and end-cap toroids and the inner detector solenoid [25].	77
7.1	Example of parton-level jet reconstruction with the anti- k_t algorithm [143].	93
7.2	MV1 combined tagger performance compared to that of JetFitter and a combination of IP3D and SV1 [153].	95
7.3	Distributions of the signed transverse (left) and longitudinal (right) impact parameter significances in both data and Monte Carlo simulation [151].	97
7.4	Distributions of the secondary vertex invariant mass (top left), secondary vertex p_T fraction (top right), and number of two track vertices (bottom) for the SV1 tagger in both data and Monte Carlo simulation [151].	98
7.5	Distributions of the decay chain invariant mass (top left), decay chain p_T fraction (top right), and number of two track vertices (bottom) for the JetFitter tagger in both data and Monte Carlo simulation [151].	100

8.1	Δm (left) and the characteristic momentum p (right) for the direct sbottom model in the $m_{\tilde{b}}$ and $m_{\chi_1^0}$ plane. As long as the neutralino is non-relativistic, the contours have similar shape.	108
8.2	\cancel{E}_T (left) and leading b -jet p_T (right) for the direct sbottom model with $m_{\tilde{b}} = 450$ GeV and $m_{\chi_1^0} = 150$ GeV ($p = 200$ GeV) in the $m_{\tilde{b}}$ and $m_{\chi_1^0}$ plane. Black is the toy model described in Section 8.1.1 with no corrections for the hard-scatter process, while red is the Monte Carlo prediction from HERWIG++. The Monte Carlo distributions are broadened compared to the toy model, mostly at the hard end of the spectrum.	109
8.3	\cancel{E}_T (left) and leading b -jet p_T (right) for the direct sbottom model with $m_{\tilde{b}} = 450$ GeV and $m_{\chi_1^0} = 150$ GeV ($p = 200$ GeV) in the $m_{\tilde{b}}$ and $m_{\chi_1^0}$ plane. Black is the toy model described in Section 8.1.1 with a symmetric application of the HERWIG++ sbottom transverse momentum distribution, while red is the Monte Carlo prediction from HERWIG++. For models with large p , such as this one, the sbottom transverse momentum distribution accounts for all of the difference seen in Figure 8.2.	110
8.4	\cancel{E}_T (left) and leading b -jet p_T (right) for the direct sbottom model with $m_{\tilde{b}} = 750$ GeV and $m_{\chi_1^0} = 700$ GeV ($p \approx 50$ GeV) in the $m_{\tilde{b}}$ and $m_{\chi_1^0}$ plane. Black is the toy model described in Section 8.1.1 with a symmetric application of the HERWIG++ sbottom transverse momentum distribution, while red is the Monte Carlo prediction from HERWIG++. The \cancel{E}_T shows large differences due to the combined effect of small Δm , large $m_{\chi_1^0}$, and the toy model's lack of additional hard-scatter partons. The leading b -jet p_T distribution is much less affected.	111
8.5	Comparison between the significance as calculated using the $\frac{s}{\sqrt{b+\sigma_b^2}}$ approximation, the significance using the Asimov dataset (Z_A), and the significance calculated using toy Monte Carlo [164]. All methods agree well in the small $\frac{s}{b}$ and small $\frac{\sigma_b^2}{b}$ limit.	114

8.6	Most sensitive direct sbottom pair production signal region as a function of the sbottom and neutralino masses.	118
8.7	Signal acceptance, in percent, as a function of mass plane for the four direct sbottom signal regions. Note that the scale is different for each plot.	120
8.8	Relative sensitivity increase in percent of the optimized 4.71 fb^{-1} direct sbottom signal regions over the signal regions from the previous version of the analysis [2]. The most sensitive signal region is used at each point (Figure 8.6). Note that the scale is logarithmic.	121
8.9	Most sensitive gluino-mediated sbottom (left) and stop (right) signal region as a function of the gluino and neutralino masses.	124
8.10	Relative sensitivity increase in percent of the optimized 4.71 fb^{-1} gluino-mediated sbottom (left) and stop (right) signal regions over the signal regions from the previous version of the analysis [3]. The most sensitive signal region is used at each point (Figure 8.9).	125
10.1	The m_T distribution before the upper selection on m_T in CR1L_SR1 (top) and in CR1L_SR3 (bottom). The slashed band includes both detector and theoretical systematic uncertainties. The SM prediction is normalized according to Monte Carlo expectations.	140
10.2	The dilepton invariant mass (m_{ll}) distribution before the upper selection on m_{ll} in CR2L_SR1/CR2L_SR2 (top) and in CR2L_SR3 (bottom). The slashed band includes both detector and theoretical systematic uncertainties. The SM prediction is normalized according to Monte Carlo expectations.	143
10.3	The \cancel{E}_T distribution in the SR1 QCD control region (top) and the leading jet p_T distribution in the SR3 QCD control region (bottom). The slashed band includes both detector and theoretical systematic uncertainties.	145

10.4	Top: m_{CT} distribution in SR1 before the selection on m_{CT} . Bottom: \cancel{E}_T distribution in SR2. The slashed band includes both detector and theoretical systematic uncertainties. All Monte Carlo distributions are normalized to the results of the profile likelihood fit.	146
10.5	Top: Leading jet p_T distribution in SR3a. Bottom: \cancel{E}_T distribution in SR3b. The slashed band includes both detector and theoretical systematic uncertainties. All Monte Carlo distributions are normalized to the results of the profile likelihood fit.	147
11.1	Comparison of the distributions of the effective mass (m_{eff}) between data and Monte Carlo expectation in the two four-jet $t\bar{t}$ control regions CR4-60 (top) and CR4-70 (bottom). The slashed band includes all detector-related and theoretical systematic uncertainties.	152
11.2	Comparison of the distributions of the effective mass (m_{eff}) between data and Monte Carlo expectation in the two six-jet $t\bar{t}$ control regions CR6-70 (top) and CR6-75 (bottom). The slashed band includes all detector-related and theoretical systematic uncertainties.	153
11.3	Comparison of the distributions of the effective mass (m_{eff} , top) and missing transverse momentum (\cancel{E}_T , bottom) between the data and Monte Carlo expectation in QCD multi-jet control region corresponding to three b -tags at 60% b -tagging efficiency. The slashed band includes all detector-related and theoretical systematic uncertainties.	154
11.4	Comparison of the distributions of the effective mass (m_{eff}) between data and Standard Model expectation in the four-jet signal regions SR4-L and SR4-M (top) and SR4-T (bottom). The Monte Carlo distributions have been scaled to the yields from the profile likelihood fit. The slashed band includes all detector-related and theoretical systematic uncertainties.	157

11.5	Comparison of the distributions of the effective mass (m_{eff}) between data and Standard Model expectation in the two six-jet signal regions SR6-L (top) and SR6-T (bottom). The Monte Carlo distributions have been scaled to the yields from the profile likelihood fit. The slashed band includes all detector-related and theoretical systematic uncertainties.	158
12.1	Exclusion contours at 95% confidence level with the CL_s method in the $m_{\tilde{b}}/m_{\chi_1^0}$ plane for direct sbottom pair production, with the sbottom decaying to a bottom quark and neutralino. The theoretical uncertainty on the sbottom cross section is shown as a yellow band. All other systematic uncertainties are included in the nominal limit.	163
12.2	Exclusion contours at 95% confidence level with the CL_s method in the $m_{\tilde{g}}/m_{\chi_1^0}$ plane for off-shell gluino-mediated sbottom production ($pp \rightarrow \tilde{g}\tilde{g}, \tilde{g} \rightarrow b\bar{b} + \chi_1^0$). The theoretical uncertainty on the gluino cross section is shown as a yellow band. All other systematic uncertainties are included in the nominal limit.	164
12.3	Exclusion contours at 95% confidence level with the CL_s method in the $m_{\tilde{g}}/m_{\chi_1^0}$ plane for off-shell gluino-mediated stop production ($pp \rightarrow \tilde{g}\tilde{g}, \tilde{g} \rightarrow t\bar{t} + \chi_1^0$). The theoretical uncertainty on the gluino cross section is shown as a yellow band. All other systematic uncertainties are included in the nominal limit.	165
12.4	Exclusion contours at 95% confidence level with the CL_s method in the $m_{\tilde{g}}/m_{\chi_1^0}$ plane for off-shell gluino-mediated sbottom/stop production with an alternate decay mode ($pp \rightarrow \tilde{g}\tilde{g}, \tilde{g} \rightarrow t\bar{b} + \chi_1^{-1}, \chi_1^{-1} \rightarrow l^- \chi_1^0$ where the lepton from the chargino decay is lost). The theoretical uncertainty on the gluino cross section is shown as a yellow band. All other systematic uncertainties are included in the nominal limit.	166

12.5	Exclusion contours at 95% confidence level with the CL_s method in the $m_{\tilde{g}}/m_{\tilde{b}}$ plane for on-shell gluino-mediated sbottom production ($pp \rightarrow \tilde{g}\tilde{g}, \tilde{g} \rightarrow \tilde{b}\tilde{b}, \tilde{b} \rightarrow b + \chi_1^0$) assuming $m_{\chi_1^0} = 60$ GeV. The theoretical uncertainty on the gluino cross section is shown as a yellow band. All other systematic uncertainties are included in the nominal limit.	168
12.6	Exclusion contours at 95% confidence level with the CL_s method in the $m_{\tilde{g}}/m_{\tilde{t}}$ plane for on-shell gluino-mediated sbottom production ($pp \rightarrow \tilde{g}\tilde{g}, \tilde{g} \rightarrow \tilde{t}\tilde{t}, \tilde{t} \rightarrow t + \chi_1^0$) assuming $m_{\chi_1^0} = 60$ GeV. The theoretical uncertainty on the gluino cross section is shown as a yellow band. All other systematic uncertainties are included in the nominal limit.	169
12.7	Mass spectrum of a pMSSM natural supersymmetric model which evades all limits as of June 2012 [179]. The decay modes of the sbottom in this model are shown in Figure 12.8.	170
12.8	Sbottm decay modes of a pMSSM natural supersymmetric model which evades all limits as of June 2012 [179]. The mass spectrum of the model is shown in Figure 12.7.	171
A.1	3D (left) vs. 2D (right) track jets in a simulated di-jet event with pileup. The minimum track jet p_T cut in both cases is 5 GeV. A line with length 1 cm in the above displays corresponds to a track with p_T of 0.5 GeV. Note that the vertical and horizontal scales are different. 5 jets pass the 5 GeV minimum p_T cut in the 3D case, and each is clearly associated with a single origin interaction. In the 2D case 13 jets pass, and most of these are composite objects with significant contributions from multiple interactions and therefore dubious physics value.	176

A.2	Average number of 2D and 3D anti- k_t track jets vs. number of reconstructed vertices (left) and reconstructed p_T distribution (right) for simulated di-jet events with pileup (2×10^{33} cm $^{-2}$ s $^{-1}$ luminosity). PV refers to 3D track jets matched to the primary vertex of the event (only possible with 3D, where all track jet tracks share a common origin vertex). In the 2D case, the removal of the requirement that all jet constituents (tracks) come from the same interaction allows a large number of “mixed origin” jets to pass the 5 GeV p_T cut in events with many reconstructed vertices.	177
A.3	Diagrams representing the two operational modes of JetTrackZClusterTool. The setting of <i>UseVtxSeeding</i> to <i>True</i> (top) uses the reconstructed vertices of the event to cluster tracks and allows cuts to be applied relative to them. The <i>False</i> setting (bottom) instructs the tool to perform its own clustering based on the high p_T tracks in the event. The default mode is <i>True</i>	180
A.4	Anti- k_t track jet purity vs. reconstructed p_T (left) and reconstructed η (right) for different R simulated di-jet events without pileup ($ \eta < 2.0$). In the right plot, p_T cuts of 5 GeV, 5.5 GeV, and 6 GeV have been applied to $R = 0.4, 0.5$ and 0.6 respectively to ensure a minimum purity level of 75% in bins of reconstructed p_T	182
A.5	Anti- k_t track jet efficiency vs. truth jet p_T for different R simulated di-jet events without pileup ($ \eta < 2.0$). Clockwise from top-left are efficiencies for all jets, isolated jets, findable and isolated jets, and all findable jets. Reconstructed p_T cuts of 5 GeV, 5.5 GeV, and 6 GeV have been applied to $R = 0.4, 0.5$ and 0.6 respectively to ensure a minimum purity level of 75% in bins of reconstructed p_T . All efficiencies are flat with respect to η	183
A.6	Anti- k_t $R = 0.4$ truth jets selected by the findable jet and isolated event criteria vs. truth p_T in simulated di-jets without pileup ($ \eta < 2.0$).	183

A.7	Examples of events with track jet efficiency problems with respect to findable jets. Track jets and tracks are shown in the η/ϕ plane. In the top figure, two adjacent $R = 0.4$ truth jets are merged into a single $R = 0.6$ track jet. In the bottom-left event display, tracks corresponding to a single truth jet are highly separated with respect to the jet axis, resulting in two low-energy $R = 0.4$ track jets, neither of which passes the 5 GeV p_T cut (bottom-right).	185
A.8	Anti- k_t track jet and calorimeter jet efficiencies in simulated di-jet events without pileup (reconstructed jets with $ \eta < 2.0$).	186
A.9	Anti- k_t jet multiplicity per event (top-left), track multiplicity per track jet (top-right), and jet p_T (bottom-left) and η (bottom-right) distributions for simulated di-jet events without pileup. Both calorimeter and track jets are included where appropriate. The additional structure seen for calorimeter jets as a function of η is due to calorimeter structure (gaps) and dead material beyond the tracker.	186
A.10	Anti- k_t track jet and calorimeter jet ϕ resolution ($\phi_{reco} - \phi_{truth}$) in simulated di-jet events without pileup (reconstructed jets with $ \eta < 2.0$) for representative truth p_T bins 15 GeV $< p_T < 20$ GeV (top-left), 20 GeV $< p_T < 25$ GeV (top-right), 25 GeV $< p_T < 30$ GeV (bottom-left), and 35 GeV $< p_T < 40$ GeV (bottom-right). The area under the curves has been normalized to 1 for better direct comparison.	187
A.11	Anti- k_t track jet multiplicity per event (top-right), track multiplicity per track jet (top-right), track jet p_T (middle-left) and η (middle-right) distributions, and number of track jets vs. number of reconstructed vertices (bottom) for simulated di-jet events with pileup (2×10^{33} cm $^{-2}$ s $^{-1}$ luminosity). PV refers to track jets matched to the primary vertex of the event, while MB refers all other jets (minimum-bias/pileup jets).	189

A.12 Anti- k_t calorimeter jet multiplicity per event (top-left), jet p_T (top-right) and η (bottom-left) distributions, and number of jets vs. number of reconstructed vertices (bottom-right) for simulated di-jet events with pileup ($2 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$ luminosity).	190
A.13 Anti- k_t track jet and calorimeter jet ϕ resolution ($\phi_{reco}-\phi_{truth}$) in simulated di-jets (reconstructed $ \eta < 2.0$) with pileup ($2 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$ luminosity) for truth p_T bins $15 \text{ GeV} < p_T < 20 \text{ GeV}$ (top-left), $20 \text{ GeV} < p_T < 25 \text{ GeV}$ (top-right), $25 \text{ GeV} < p_T < 30 \text{ GeV}$ (bottom-left), and $35 \text{ GeV} < p_T < 40 \text{ GeV}$ (bottom-right). The area under the curves has been normalized to 1 for better direct comparison.	191
A.14 Anti- k_t $R = 0.4$ track jet, calorimeter jet, and combination jet resolution ($\phi_{reco}-\phi_{truth}$) in simulated di-jets (reconstructed $ \eta < 0.3$) for truth p_T bins $20 \text{ GeV} < p_T < 30 \text{ GeV}$ (top-left), $40 \text{ GeV} < p_T < 50 \text{ GeV}$ (top-right), $60 \text{ GeV} < p_T < 70 \text{ GeV}$ (bottom-left), and $90 \text{ GeV} < p_T < 100 \text{ GeV}$ (bottom-right).	193
A.15 Anti- k_t $R = 0.4$ track jet, calorimeter jet, and combination jet resolution ($\phi_{reco}-\phi_{truth}$) RMS in simulated di-jets (reconstructed $ \eta < 0.3$) as a function of truth p_T	193

Chapter 1

Introduction

The Standard Model of particle physics is one of the greatest accomplishments of modern science. Painstakingly constructed in fits and starts over the second half of the 20th century, it stands today as the most complete, accurate, and predictive theory of the fundamental interactions ever constructed, at least up to the weak interaction scale. Even the long-sought Higgs boson, the last piece of the Standard Model puzzle, has now likely been found [8, 9].

Despite its success, the Standard Model is a necessarily incomplete and unsatisfying theory, and not only due to its omission of gravity. It provides no viable dark matter candidates, though dark matter accounts for nearly a quarter of the universe's total energy density [10]. While violation of charge-parity symmetry (CP violation) can produce matter and anti-matter asymmetrically, the amount of CP violation in the Standard Model is insufficient to account for the enormous imbalance observed [11]. There is no explanation for the vast difference in mass between the top quark and the other charged fermions and between the charged fermions and the neutrinos. This issue is related to “fine-tuning”, the theoretical concern that Standard Model parameters have seemingly been adjusted to an astounding degree of precision in order to obtain the laws of nature as we know and experience them.

Extensions to the Standard Model known as Beyond the Standard Model (BSM) theories attempt to address these weaknesses. One of most attractive groups of BSM

theories is known as supersymmetry [12, 13, 14, 15, 16, 17, 18, 19, 20], which postulates an additional space-time symmetry between bosons (force carriers) and fermions (matter fields) in order to solve the most serious of the Standard Model’s fine-tuning issues, the gauge hierarchy problem [21]. A semi-miraculous consequence of this new symmetry and another symmetry inserted to address proton stability concerns (R -parity) is that a dark matter candidate often appears in the theory automatically.

Unfortunately, the number of supersymmetric models is enormous, and the dominant experimental signatures often differ considerably. Given that no BSM signal has yet been observed, the most important characteristic of a BSM search is broad sensitivity to possible new physics signatures. One method used to ensure this, while maximizing experimental reach, is known as the simplified models approach [22, 23, 24]. With this prescription, full BSM models are separated into individual final state topologies whose kinematic properties and their dependence on model parameters are thoroughly investigated. The search is then designed leaving the probability of such a decay (cross section and branching ratio) a free parameter. An analysis designed to look for a single topology can constrain any full model which contains a similar final state.

Sensitive new physics searches also require a precision detector, a collider powerful enough to make new heavy particles, and a large amount of recorded data. The ATLAS detector [25], one of the most precise detectors ever built and one of the two general-purpose detectors at the Large Hadron Collider (LHC) [26], fully satisfies the first requirement. While it remains unclear if the LHC is energetic enough to produce new heavy particles, it is unquestionably the most powerful collider ever built, as well as the most luminous (the highest collision rate). The analyses presented here use the ATLAS dataset from the the first full year of high-luminosity LHC data (2011).

The focus of the analyses is hadronic (zero lepton) final states involving multiple b -tagged jets and missing transverse energy (\cancel{E}_T). These signatures arise in supersymmetry primarily due to decay of the third-generation quark supersymmetric partner particles known as the sbottom and the stop, which are some of the most important processes in supersymmetry due to the connection between light third-generation squarks and a “natural” (non-fine-tuned) solution to the gauge hierarchy problem.

These squarks can be produced directly or as quark-squark pairs from gluino decays (the gluino is the superpartner of the gluon). Novel search strategies including radiation-based final state signatures and signal regions requiring three b -tagged jets are utilized (in the latter case, for the first time) to maximize sensitivity to direct sbottom and gluino-mediated sbottom and stop event topologies. The direct stop searches are left to other analyses [27, 28, 29, 30, 31, 32, 33, 34].

Chapter 2 begins with a brief overview of the Standard Model, with a focus on electroweak symmetry-breaking and the gauge hierarchy problem, the primary historical motivation for supersymmetry. In Chapter 3 the Minimal Supersymmetric Standard Model (MSSM) is introduced, and naturalness arguments are followed to illuminate the special place occupied by the third-generation quark superpartners in natural supersymmetry. Chapter 4 contrasts “top-down” with simplified model (“bottom-up”) search strategies, and introduces the specific supersymmetric decay chains which are the focus of subsequent chapters. Chapter 5 gives an overview of the LHC and its history, while Chapter 6 describes the ATLAS detector, its subsystems, and the trigger and data acquisition system. In Chapter 7, the collision and simulated datasets used in the analyses are described as well as the reconstruction and definitions of physics objects and the common event cleaning and reweighting.

Chapter 8 marks a return to simplified models and describes the optimization procedure and results, while Chapter 9 is a technical interlude describing the methods used to produce background estimates as well as the relevant experimental and theoretical uncertainties. Chapters 10 and 11 detail the results of the sbottom pair production and gluino-mediated sbottom and stop searches, respectively, and in Chapter 12 the methods used to interpret the results statistically as well as the interpretation itself are discussed. Finally, Chapter 13 summarizes the most important points from preceding chapters. The near-term prospects for the sbottom and gluino searches in particular and for natural supersymmetry in general are also discussed.

Chapter 2

The Standard Model

The Standard Model (SM) of particle physics is the most accurate theoretical description of Nature ever devised. Developed over the second half of the 20th century with many contributors, it is an amalgam of gauge theories describing the strong (Quantum Chromodynamics) and electroweak (Glashow-Weinberg-Salam theory) interactions, the fermion fields representing the three generations (flavors) of quarks and leptons (matter and anti-matter), and the Higgs field, which both explains electroweak symmetry-breaking and provides the mechanism for generating the fermion masses. The complete Standard Model gauge symmetry group is $SU(3)_C \otimes SU(2)_W \otimes U(1)_Y$.

2.1 The Strong Sector

The theory of strong interactions is referred to as Quantum Chromodynamics (QCD) and is defined by the $SU(3)_C$ gauge group. The eight massless gauge bosons of this group are vector (spin 1) particles known as gluons, and there are three types of strong charge known as color (red, green, and blue). Each gluon carries two color charges, with one type being positive (i.e. red) and the other negative (i.e. anti-green). The only fermions which carry color charge, and thus the only fermions which interact with the gluons, are the quarks. The Lagrangian of the QCD interaction in the Standard Model is given by [35]:

$$\mathcal{L}_{QCD} = i\bar{\psi}(\partial_\mu - ig_s G_\mu^a T^a)\gamma^\mu\psi - \frac{1}{4}G_{\mu\nu}^a G_a^{\mu\nu} \quad (2.1)$$

In this expression, ψ represents a quark field, g_s is the strong coupling constant, G_μ^a are the gluon fields, T^a are the generators of $SU(3)_C$ group, and γ^μ are the Dirac gamma matrices. $G_{\mu\nu}^a$ is analogous to the electromagnetic field strength tensor $F_{\mu\nu}$ and is defined as:

$$G_{\mu\nu}^a \equiv \partial_\mu G_\nu^a - \partial_\nu G_\mu^a - g_s f^{abc} G_\mu^b G_\nu^c \quad (2.2)$$

At a hadron collider such as the LHC, strong interaction scattering dominates the inclusive pp cross section. Two important (and unique within the Standard Model) features of the strong interactions are *asymptotic freedom* and *confinement*.

Asymptotic Freedom Perturbative calculations depend on $\alpha_s \equiv g_s^2/4\pi$ being small. If it is not, the use of Feynman diagrams and other perturbative technology is precluded. However, the strength of the coupling constant evolves with energy according to the renormalization group. The “running” constant implies that a theory with a non-perturbative energy region may also have a perturbative region, depending on the evolution of the coupling constant. This evolution is parameterized by the β function—if the β function is negative the coupling strength decreases with energy. In QCD, the β function is given to lowest non-vanishing order by [36]:

$$\beta(\alpha_s) = -\frac{\alpha_s^2}{2\pi}\left(11 - \frac{4}{3}n_{gen}\right) \quad (2.3)$$

with α_s defined in terms of the strong coupling g_s as:

$$\alpha_s \equiv \frac{g_s^2}{4\pi} \quad (2.4)$$

This result is itself obtained perturbatively (assuming α_s small), so its validity depends on the existence of such a perturbative region. Experimentally, such a region does exist, and for $n_{flav} = 3$ (the Standard Model case) the β function is negative and thus the strong coupling decreases with increasing energy scale—QCD is *asymptotically*

free. Conversely, at low energies α_s becomes large and non-perturbative techniques must be used [37, 38].

Confinement The experimental non-observation of single free quarks indicates that some mechanism of *color confinement* is at work in the strong interactions. There is also considerable theoretical evidence for such confinement though it remains unproven analytically [39]. Confinement states that only color-neutral composite objects, such as mesons and baryons, are directly observable. As two quarks are separated, the force between them does not decrease, and at some point it becomes more energetically favorable to pair produce a quark-antiquark pair than continue the long-distance interaction.

This process is called *hadronization* or *fragmentation* and has tremendous experimental consequences at a hadron collider such as the LHC. Final state quarks and gluons are observable not as single particles but as “jets” of many color-neutral hadrons. For hadronic-channel analyses and analyses involving missing energy such as those discussed in Chapters 11 and 10, these jets of hadrons are of primary importance as are the methods used to reconstruct and calibrate them (Section 7.7).

2.2 The Electroweak Sector

The Glashow-Weinberg-Salam theory [40, 41], or GWS theory, is the Standard Model gauge theory of the unified electroweak interactions. The electroweak gauge symmetry group is $SU(2)_W \otimes U(1)_Y$. The four massless vector gauge bosons W^i ($SU(2)_W$ weak isospin, $i = 1, 2, 3$) and B ($U(1)_Y$ weak hyper charge) become the massive W^+ , W^- , and Z^0 bosons and the massless photon after electroweak symmetry-breaking (Section 2.3). The W^+ , W^- , and Z^0 bosons mediate the weak interactions (familiar to the layperson in the form of radioactivity), and the photon mediates electromagnetism (biology, chemistry, and most of the technology which drives the modern world). The gauge-basis Lagrangian for the electroweak interactions is given by [42]:

$$\mathcal{L}_{EW} = \sum_{\psi} i\bar{\psi}\gamma^{\mu}(\partial_{\mu} + i\frac{g'}{2}YB_{\mu} + i\frac{g}{2}\tau_L^a W_{\mu}^a)\psi - \frac{1}{4}B_{\mu\nu}B^{\mu\nu} - \frac{1}{4}W_{\mu\nu}^a W_a^{\mu\nu} \quad (2.5)$$

In this expression, g and g' are the coupling constants, Y is the weak hypercharge, and τ_L^a are the Pauli matrices which act only on left-handed fermions. $B_{\mu\nu}$ and $W_{\mu\nu}^a$ are constructed in a way analogous to Equation 2.2.

2.3 The Higgs Sector and Electroweak Symmetry Breaking

The unification of the electromagnetic and weak interactions is broken when the weak gauge bosons acquire mass. However, naive addition of a mass term is problematic, because it breaks gauge invariance. This can be seen if we consider an electromagnetic Lagrangian with a photon mass term:

$$\mathcal{L}_{EM} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + \frac{1}{2}m^2 A_{\mu}A^{\mu} \quad (2.6)$$

where $F_{\mu\nu} \equiv \partial_{\nu}A_{\mu} - \partial_{\mu}A_{\nu}$ and the second term is the mass term. If a local $U(1)$ gauge transformation $A_{\mu}(x) \rightarrow A_{\mu}(x) - \partial_{\mu}\eta(x)$ is applied, it is clear that this term is not invariant, and thus such simple mass terms are forbidden by gauge symmetry. The same conclusion applies to the $SU(2)_W$ gauge bosons.

The Standard Model's solution to this problem, the Higgs mechanism (more correctly known as the Englert-Brout-Higgs-Guralnik-Hagen-Kibble mechanism [43, 44, 45]), involves the introduction of a complex scalar $SU(2)_W$ doublet to the theory:

$$\Phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix} \quad (2.7)$$

Using the Gell-Mann-Nishijima relation for the electromagnetic charge [46, 47]:

$$Q = \frac{(\tau_3 + Y)}{2} \quad (2.8)$$

it is clear that this doublet must carry $U(1)_Y$ weak hypercharge $Y_\Phi = 1$ (the sign is arbitrary in the Standard Model). Also added to the Lagrangian is a the scalar potential:

$$V(\Phi) = \mu^2 |\Phi^\dagger \Phi| + \lambda (|\Phi^\dagger \Phi|)^2 \quad (2.9)$$

with $\lambda > 0$. This is the most general scalar potential allowed which is both renormalizable and $SU(2)_W$ invariant [48]. If $\mu^2 < 0$ (a tachyonic mass term) the minimum of the potential is not at zero and the scalar field acquires a *vacuum expectation value* (VEV). In this way, the Lagrangian remains symmetric, but in the ground state of the theory this symmetry is *spontaneously broken*. As the direction of the symmetry-breaking is not constrained by the potential, the VEV can be chosen real and entirely in the direction of the second (electrically neutral) component of Φ :

$$\langle \Phi \rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix} \quad (2.10)$$

with $v = \sqrt{\frac{-\mu^2}{\lambda}}$. The contribution to the Lagrangian from the scalar doublet is given by:

$$\mathcal{L}_{Higgs} = (D^\mu \Phi)^\dagger (D_\mu \Phi) - V(\Phi) \quad (2.11)$$

where

$$D_\mu = \partial_\mu + i \frac{g'}{2} Y B_\mu + i \frac{g}{2} \tau_L^a W_\mu^a \quad (2.12)$$

is the electroweak covariant derivative seen previously in Equation 2.5. Plugging in the VEV from Equation 2.10 into Equation 2.11 (and dropping some terms for clarity), it is apparent that the scalar kinetic term generates terms quadratic in the vector fields:

$$\mathcal{L}_{partial} = \frac{1}{2} \begin{pmatrix} 0 & v \end{pmatrix} \left(\frac{g'}{2} B_\mu + \frac{g}{2} \tau_L^a W_\mu^a \right)^2 \begin{pmatrix} 0 \\ v \end{pmatrix} \quad (2.13)$$

$$\mathcal{L}_{\text{partial}} = \frac{g'^2}{8} B^2 v^2 - \frac{gg'}{4} B_\mu W_3^\mu v^2 + \frac{g^2}{8} W_3^2 v^2 + \frac{g^2}{8} W_1^2 v^2 + \frac{g^2}{8} W_2^2 v^2 \quad (2.14)$$

The last two terms represent the physical W^\pm mass terms, with mass m_W given by:

$$m_W^2 = \frac{1}{4} g^2 v^2 \quad (2.15)$$

The first three terms can be rewritten as:

$$\frac{1}{2} \left(\frac{-g' B^\mu + g W_3^\mu}{\sqrt{g^2 + g'^2}} \right)^2 \frac{1}{4} (g^2 + g'^2) v^2 \quad (2.16)$$

Identifying the quantity in parentheses as the Z^μ :

$$Z^\mu \equiv \frac{-g' B^\mu + g W_3^\mu}{\sqrt{g^2 + g'^2}} \quad (2.17)$$

then its mass m_Z is given by:

$$m_Z^2 = \frac{1}{4} (g^2 + g'^2) v^2 \quad (2.18)$$

The linear combination of B^μ and W_3^μ orthogonal to Z^μ is:

$$A^\mu \equiv \frac{g B^\mu + g' W_3^\mu}{\sqrt{g^2 + g'^2}} \quad (2.19)$$

and remains massless. This is none other than the physical photon. Thus we arrive at the vacuum state where the $SU(2)_W \otimes U(1)_Y$ symmetry of the Lagrangian is broken, but a single unbroken $U(1)$ symmetry (electromagnetism) remains. Three of the four degrees of freedom of the original complex scalar doublet are “eaten” by the W and Z bosons in this gauge and represent the new longitudinal components of their polarizations. The remaining electrically neutral degree of freedom represents the Standard Model Higgs boson.

This scalar degree of freedom can be included in the previous discussion with the

substitution:

$$\langle \Phi \rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix} \rightarrow \Phi = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h \end{pmatrix} \quad (2.20)$$

where instead of using the VEV directly we expand Φ around the VEV and define h as the deviation from the minimum in the neutral (and real) direction. This form can always be achieved via $SU(2)_W$ gauge transformation. Taylor-expanding the scalar potential (Equation 2.9) around the VEV produces a mass term quadratic in h^2 , with mass given by:

$$m_h^2 = 2v^2\lambda \quad (2.21)$$

The preceding discussion can be rewritten entirely in terms of v and m_h , the Higgs boson mass. v can be related directly to the muon lifetime [48], and is about 246 GeV. On July 4th, 2012, both the ATLAS and CMS collaborations announced the discovery of a boson “consistent with” the Standard Model Higgs at around 126 GeV. That is, the boson appears to couple to the W and Z bosons in a manner and strength consistent with the Standard Model Higgs mechanism [8, 9]. Although it is still early, the Higgs mechanism may finally be on as sound footing experimentally as it has been theoretically.

2.4 Fermions

The gauge fields and Higgs field are all bosonic (integer spin) fields, and can be thought of as “force carrier” fields. The other half of the Standard Model menagerie (much more than half, numerically) are the fermionic (half-integer spin) fields, the “matter” fields. All fundamental fermions in the Standard Model are Dirac fermions with spin $\frac{1}{2}$. They are divided into two types, quarks and leptons, and each type has three subtypes or *generations*. While quarks and leptons differ considerably in their gauge couplings, the gauge couplings are the same across quark or lepton generations. The features which distinguish one generation from another are the fermion masses.

Quarks Quarks come in two types: “up”-type which have electric charge $+\frac{2}{3}$ and “down”-type with electric charge $-\frac{1}{3}$. There are three generations, the first of which contain the lightest quarks. These quarks are called up (u) and down (d). The second-generation quarks are called the charm (c) and strange (s) quarks, and the third and by far the heaviest generation is composed of the top (t) and bottom (b) quarks. All quarks carry $SU(3)_C$ color charge and thus couple to the gluons in the strong interactions.

Left-handed up-type and down-type quarks of the same generation also transform under $SU(2)_W$ weak isospin rotations and all quarks carry weak hypercharge—in other words, the quarks interact via the weak force as well.

Leptons There are also three generations of leptons. The charged leptons have electric charge -1 , and in order of increasing mass are the electron, muon, and tau. The left-handed charged leptons form an $SU(2)_W$ weak isospin doublet with the left-handed neutrinos (no electric charge) of the corresponding flavor (i.e, the left-handed electron and left-handed electron neutrino). The leptons do not carry color and therefore do not interact strongly.

In the Standard Model as originally formulated, the left-handed neutrinos were massless and the right-handed neutrinos were thought not to exist. With the discovery of neutrino oscillation the existence of right-handed neutrinos is required assuming that neutrinos are Dirac particles. In its simplest form this would involve introduction of another fermion mass scale hundreds of thousands of times smaller than that of the charged fermions, an idea which is theoretically unpalatable. One can avoid the introduction of a right-handed neutrino by assuming the neutrinos are Majorana fermions (their own anti-particles) [49], though the mass scale issue remains.

The most popular solution is the so-called seesaw mechanism [50, 51, 52], which allows the Dirac mass scale to be similar to that of the charged fermions while suppressing the physical left-handed neutrino masses with a very large right-handed Majorana mass. Whatever the true nature of the neutrinos, the Standard Model approximation of a massless left-handed neutrino is very good for collider applications. Efforts to characterize the nature of the neutrinos and their mass spectrum are

ongoing [53, 54, 55].

Fermion Masses and the Yukawa Couplings Fermion mass terms present a similar problem to the one encountered with the gauge boson mass terms in Section 2.3. In a free theory, simple scalar masses of the form found in the Dirac Lagrangian

$$\mathcal{L}_{Dirac} = i\bar{\psi}\gamma^\mu\partial_\mu\psi - m\bar{\psi}\psi \quad (2.22)$$

are allowed. However, when coupling to the gauge group is added, all terms in a Lagrangian must respect the gauge symmetries. The kinetic term is fixed via promotion of the simple derivative to a covariant derivative. For the mass term, the problem is best seen by expanding it in terms of left-handed and right-handed two component spinors.

$$m\bar{\psi}\psi = m\bar{\psi}_L\psi_R + m\bar{\psi}_R\psi_L \quad (2.23)$$

Left-handed and right-handed fermions transform independently under $SU(2)_W \otimes U(1)_Y$, so these terms cannot be gauge-invariant. Somewhat miraculously, the solution is again provided by the Higgs Mechanism (Section 2.3). Consider the gauge-invariant Yukawa coupling of the Higgs doublet Φ to the down quark:

$$\mathcal{L}_d = -y_d\bar{Q}_L\Phi d_R + h.c. = -\frac{y_d}{\sqrt{2}}\begin{pmatrix} \bar{u}_L & \bar{d}_L \end{pmatrix}\begin{pmatrix} 0 \\ v \end{pmatrix}d_R + h.c. \quad (2.24)$$

Q_L is the $SU(2)_W$ doublet of left-handed up and down quarks, and Φ been replaced by its VEV (Equation 2.10) after the second equality. Multiplying through, these terms reduce to:

$$\mathcal{L}_d = -\frac{y_d v}{\sqrt{2}}\bar{d}_L d_R + h.c. \quad (2.25)$$

The fermion mass can now be identified as:

$$m_d = \frac{y_d v}{\sqrt{2}} \quad (2.26)$$

Field	Symmetry/Charge	Total Number
G_μ^a	$SU(3)_C$ /color	8
W_μ^i	$SU(2)_W$ /weak isospin	3
B_μ	$U(1)_Y$ /weak hypercharge	1

Table 2.1: Gauge fields of the Standard Model. The massless electroweak gauge fields W_μ^i and B_μ become the massive W_μ^\pm , Z_μ^0 (weak), and massless A_μ (electromagnetic) fields after electroweak symmetry-breaking via the Higgs mechanism. All gauge fields are spin 1.

It may seem as if only the labels have changed, but the crucial difference is that the Yukawa interactions are gauge-invariant. The addition of Yukawa couplings to fermions also influences the physics of the Standard Model in fundamental ways, in particular the decay modes of the Higgs boson. With multiple fermion generations, the Yukawa coupling constants become matrices rather than numbers, and do not have to be diagonal. In fact, the Yukawa matrices for the up-type and down-type quarks cannot be simultaneously diagonalized. This is the origin of the difference between the weak interaction eigenstates and quark mass eigenstates and gives rise to the Cabibbo-Kobayashi-Maskawa (CKM) matrix [56, 57] and inter-generational charged weak currents.

In extensions of the Standard Model which account for non-zero neutrino masses, the same mechanism provides the lepton (Dirac) mass terms, with the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix [58, 59] playing the role of the CKM matrix.

2.5 Standard Model Particle Content Summary

The Standard Model gauge fields are summarized in Table 2.1. The fermions and Higgs fields are given in Table 2.2.

2.6 The Hierarchy Problem

Pre-discovery, many reasonable arguments pointed to the Higgs mass lying at or near the electroweak scale [48]. With the discovery of what appears to be a 126 GeV

Name	Field	$SU(3)_C$	$SU(2)_W$	$U(1)_Y$
Left-handed quarks	$Q_L = \begin{pmatrix} u_L \\ d_L \end{pmatrix}$	3	2	$\frac{1}{3}$
Right-handed up quarks	u_R	3	1	$\frac{4}{3}$
Right-handed down quarks	d_R	3	1	$-\frac{2}{3}$
Left-handed leptons	$L_L = \begin{pmatrix} \nu_L \\ e_L \end{pmatrix}$	1	2	-1
Right-handed leptons	e_R	1	1	-2
Higgs	$\Phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix}$	1	2	1

Table 2.2: Fermionic fields and the Higgs field of the Standard Model. All fermions have spin $\frac{1}{2}$, while the Higgs field has spin 0.




Figure 2.1: One-loop diagram of the correction to the fermion mass from the Higgs. This correction is proportional to the fermion mass and is logarithmically divergent with the cutoff scale (Equation 2.31)

Higgs-like boson these arguments have become more academic, but a mystery still surrounds how the Higgs comes to be so light.

Corrections to Fermion Masses from the Higgs The difficulty arises via radiative corrections to the Higgs mass. Following the treatment in References [48] and [60], we consider a theory involving a Higgs-like scalar and a single massive fermion, after electroweak symmetry-breaking, given by the effective Lagrangian:

$$\mathcal{L}_{eff} = i\bar{\psi}\gamma^\mu\partial_\mu\psi - m_f\bar{\psi}\psi + |\partial_\mu\phi_h|^2 - m_h^2|\phi_h|^2 - (y_f\bar{\psi}\psi\phi_h + h.c.) \quad (2.27)$$

The fermion's mass is acquired via electroweak symmetry breaking, with $m_f = \frac{y_f v}{\sqrt{2}}$. The fermion self-energy due to the scalar loop in Figure 2.1 is given by:

$$-i\Sigma_f(p) = \left(\frac{-iy_f}{\sqrt{2}}\right)^2 (i)^2 \int \frac{d^4k}{(2\pi)^4} \frac{(k+m_f)}{[k^2 - m_f^2][(k-p)^2 - m_f^2]} \quad (2.28)$$

The renormalized fermion mass is the bare mass plus a correction:

$$m_f^r = m_f + \delta m_f \quad (2.29)$$

where δm_f is given by:

$$\delta m_f \equiv \Sigma_f(p) \big|_{p=m_f} \quad (2.30)$$

The integral, evaluated with a high energy cutoff Λ , yields:

$$\delta m_f = -\frac{3y_f^2 m_f}{64\pi^2} \log\left(\frac{\Lambda^2}{m_f^2}\right) + \dots \quad (2.31)$$

The ellipsis represents terms which go to zero as $\Lambda \rightarrow \infty$ or are independent of the cutoff. There are two interesting features of this result. The first is that it diverges logarithmically with the cutoff, rather than quadratically. The second feature is that the correction is proportional to m_f . In the limit that the fermion mass vanishes, the correction also vanishes. Setting the fermion mass to zero increases the symmetry of the theory; the corrections to the mass are regulated by the degree to which the symmetry is broken and are said to be *natural*. In the Standard Model, the fermions receive larger corrections from loops with gauge bosons, but all these corrections have these same two features.

Corrections to the Higgs Mass from Fermions Now consider the one-loop correction to the Higgs mass from the fermion loop (Figure 2.2). The correction is:

$$-i\Sigma_h(p) = \left(\frac{-iy_f}{\sqrt{2}}\right)^2 (i)^2 (-1) \int \frac{d^4k}{(2\pi)^4} \frac{\text{Tr}[(k+m_f)((k-p)+m_f)]}{[k^2 - m_f^2][(k-p)^2 - m_f^2]} \quad (2.32)$$

The leading order corrections to the Higgs mass squared as a function of the ultraviolet cutoff Λ are:




Figure 2.2: Fermion (left) and scalar (right) one-loop diagrams which generate opposite-sign quadratically-divergent corrections to the Higgs mass. Cancellation of the quadratic divergences is complete for two scalars and one fermion with couplings obeying the relation $\lambda_{\tilde{f}} = y_f^2$.

$$(\delta m_h^2)_f = -\frac{y_f^2}{8\pi^2} \left[\Lambda^2 + (m_h^2 - 6m_f^2) \log \left(\frac{\Lambda}{m_f} \right) + \dots \right] \quad (2.33)$$

This correction is not proportional to m_h^2 and in fact is *quadratically* divergent with the cutoff Λ . There is no symmetry that protects the Higgs mass from these large corrections, and while top quarks loops are the worst offenders (due to y_t being ≈ 1), gauge boson loops and loops from the quartic Higgs coupling are also problematic.

These large corrections can be dealt with via renormalization. However, with the Higgs at the electroweak scale (126 GeV) the cancellations required between the corrections and the bare mass/counterterms must be precise to the level of ≈ 1 part in 10^{15} . This *fine-tuning* is known as the hierarchy problem (the name arises from the vast difference between the GUT or Planck scales and the electroweak scale).

Corrections to the Higgs Mass from Scalars Contrast the situation with fermionic loops with that of scalar loops (Figure 2.2). Define a new effective Lagrangian which adds a new complex scalar field ϕ_s to \mathcal{L}_{eff} (Equation 2.27):

$$\mathcal{L}'_{eff} = |\partial_\mu \phi_s|^2 - m_s^2 |\phi_s|^2 + \lambda_s \phi_h |\phi_s|^2 + \mathcal{L}_{eff} \quad (2.34)$$

The correction to the Higgs mass is:

$$(\delta m_h^2)_s = \frac{\lambda_s}{16\pi^2} \left[\Lambda^2 - 2m_s^2 \log \left(\frac{\Lambda}{m_s} \right) + \dots \right] \quad (2.35)$$

The scalar contribution is also quadratically divergent, but with the *opposite sign* compared to the fermionic correction. Anticipating the supersymmetric solution, if *two* complex scalars are added to the theory, and if some symmetry enforces the condition:

$$\lambda_s = y_f^2 \tag{2.36}$$

then the quadratic divergences of Equations 2.33 and 2.35 cancel each other. In addition, if the scalar and fermionic masses are near the same, the logarithmic divergences are severely reduced. This allows the bare Higgs mass to lie at the electroweak scale *naturally*, and is one of the most compelling arguments in favor of supersymmetry.

Chapter 3

Supersymmetry and Heavy Flavor

In Section 2.6 it was shown that ultraviolet-divergent corrections to the Higgs mass from fermion loops could be eliminated with the addition of scalar particles with appropriately chosen couplings. While this is a central motivation for supersymmetry, it is ad-hoc, and ignores a glaring weakness in the Higgs mechanism. This weakness is that the Higgs potential itself is an ad-hoc addition to the Standard Model Lagrangian with nothing mandating either the sign or magnitude of μ^2 . A more holistic approach is to incorporate the Standard Model Higgs mechanism into a model where μ^2 arises dynamically, and mandate that it receive no radiative corrections. This requires that some symmetry of the theory forbids a nonzero μ^2 at some sufficiently high energy scale.

There are three ways to arrange such a symmetry [61]. Two of them give rise to “little Higgs” theories [62] and theories with extra space-time dimensions [63, 64]. The third symmetry connects shifts in the Higgs field to a fermion field.

$$\delta\phi = \epsilon \cdot \psi \tag{3.1}$$

The μ^2 mass term can then be forbidden by the same chiral symmetry that forbids fermion mass terms (Section 2.4). Such an additional symmetry has profound consequences for the theory as a whole. The conserved fermionic charge (spin $\frac{1}{2}$) Q_α which generates the symmetry implies a conserved 4-vector charge R_ν given by [61]:

$$\{Q_\alpha, Q_\beta^\dagger\} = 2\gamma_{\alpha\beta}^\nu R_\nu \quad (3.2)$$

For theories with non-trivial scattering, the freedom to add an arbitrary conserved 4-vector charge does not exist [65]. R_ν is therefore constrained to be equal to the existing conserved 4-vector P_ν , the sum of the energy-momentum 4-vectors of everything. This connection implies strongly that “super partners” cannot be added solely for the problematic top quark, but must be added for every particle in the theory. Moreover, the connection to the generator of space-time translations implies a fundamental generalization of space-time symmetry, hence the name “supersymmetry”.

3.1 The Minimal Supersymmetric Standard Model

The *Minimal Supersymmetric Standard Model* (MSSM) is the minimum set of particles necessary for a consistent supersymmetric theory. Each Standard Model particle has a supersymmetric partner with spin differing by $\frac{1}{2}$. As the Q_α operators commute with the generators of the $SU(3)_C \otimes SU(2)_W \otimes U(1)_Y$ gauge group, the new supersymmetric particles have the same internal quantum numbers (with the exception of spin) and gauge interactions as their Standard Model partners.

Sfermions The Standard Model fermions form *chiral supermultiplets* with their supersymmetric spin-0 partners, the sfermions (“s” for “scalar”). For instance, the left-handed quarks (Q) are matched with the left-handed squarks (\tilde{Q}). Note that “handedness” for the sfermions is only a label reflecting the identity of their Standard Model partners (and their weak couplings), as the sfermions themselves are spin 0. Likewise, the leptonic super-partners are known as sleptons.

Gauginos The Standard Model gauge bosons form *gauge supermultiplets* with their spin- $\frac{1}{2}$ super-partners, the gauginos. Thus the gluons have the gluinos (\tilde{g}), the $SU(2)_W$ vector bosons (W_μ^i , $i = 1,2,3$) have the winos (\tilde{W}), and the $U(1)_Y$ vector boson B_μ has the bino (\tilde{B}).

Higgses and Higgsinos One complex scalar Higgs doublet was enough in the Standard Model to account for all electroweak symmetry-breaking. The weak hypercharge Y of the doublet was taken somewhat arbitrarily to be +1 (Section 2.3). In the MSSM, two Higgs doublets with $Y = \pm 1$ must exist in order to cancel a gauge anomaly [60]. These two doublets are denoted by:

$$H_u = \begin{pmatrix} H_u^+ \\ H_u^0 \end{pmatrix}, \quad H_d = \begin{pmatrix} H_d^0 \\ H_d^- \end{pmatrix} \quad (3.3)$$

The subscripts denote which fermions couple to which doublet, with the charged leptons coupling to H_d . The vacuum expectation values (VEVs) acquired by the electrically-neutral components of these doublets after electroweak symmetry-breaking are denoted:

$$v_u = \langle H_u^0 \rangle, \quad v_d = \langle H_d^0 \rangle \quad (3.4)$$

Like in the Standard Model Higgs mechanism, three of the Higgs scalar degrees of freedom after electroweak symmetry-breaking become the longitudinal components of the W^\pm and Z^0 gauge bosons. The other 5 components become two CP-even neutral scalars h^0 and H^0 (in the “decoupling” limit h^0 resembles the Standard Model Higgs boson), one CP-odd neutral scalar A^0 , and two charged scalars H^\pm . Their partners, the spin- $\frac{1}{2}$ Higgsinos, form two $SU(2)_W$ doublets denoted by:

$$\tilde{H}_u = \begin{pmatrix} \tilde{H}_u^+ \\ \tilde{H}_u^0 \end{pmatrix}, \quad \tilde{H}_d = \begin{pmatrix} \tilde{H}_d^0 \\ \tilde{H}_d^- \end{pmatrix} \quad (3.5)$$

3.2 Soft SUSY-breaking in the MSSM

Unbroken supersymmetry, where the supersymmetric partners have the same masses as their SM counterparts, is clearly ruled out by experiment—511 KeV selectrons would certainly have been seen if they existed. However, SUSY can be broken in ways which retain the cancellation of ultraviolet divergences. While specific models of soft SUSY-breaking exist, a common approach is to parameterize the process via all the

renormalizable terms which can be added to the Lagrangian which break SUSY while preserving the cancellation of UV divergences. These are called *soft SUSY-breaking* terms and can be classified as follows [60]:

Gaugino Masses The gauginos in the MSSM are Majorana fermions (they are their own anti-particles), and Majorana mass terms can be added such as:

$$\mathcal{L}_{soft} = -\frac{1}{2}M_1\tilde{B}\tilde{B} - \frac{1}{2}M_2\tilde{W}\tilde{W} - \frac{1}{2}M_3\tilde{g}\tilde{g} + h.c. \quad (3.6)$$

where \tilde{B} is the bino, \tilde{W} are the winos, and \tilde{g} represent the gluinos.

Trilinear Scalar Couplings (A-terms) Gauge-invariant cubic couplings between the left- and right-handed squarks and the Higgs bosons can also be added to the Lagrangian. These are known as *A-terms*.

$$\mathcal{L}_{soft} = -\tilde{u}\mathbf{a}_u\tilde{Q}H_u - \tilde{d}\mathbf{a}_d\tilde{Q}H_d - \tilde{e}\mathbf{a}_e\tilde{L}H_d + h.c. \quad (3.7)$$

In principle, the parameters \mathbf{a} are complex 3 x 3 matrices in generation space. Therefore, they are able to facilitate changes in flavor and mix left- and right-handed squarks. \tilde{Q} and \tilde{L} represent here the left-handed squark and slepton $SU(2)_W$ doublets, respectively, while \tilde{u} , \tilde{d} , and \tilde{e} represent the right-handed sfermion $SU(2)_W$ singlets.

Sfermion Masses The scalar partners of the Standard Model fermions can also receive explicit SUSY-breaking mass terms not generated by the Higgs mechanism. These terms are not proportional to VEVs and do not break $SU(2)_W \otimes U(1)_Y$ gauge symmetry.

$$\mathcal{L}_{soft} = -\tilde{Q}^\dagger\mathbf{M}_Q^2\tilde{Q} - \tilde{L}^\dagger\mathbf{M}_L^2\tilde{L} - \tilde{u}\mathbf{M}_u^2\tilde{u}^\dagger + -\tilde{d}\mathbf{M}_d^2\tilde{d}^\dagger - \tilde{e}\mathbf{M}_e^2\tilde{e}^\dagger \quad (3.8)$$

The matrices \mathbf{M} can also be somewhat arbitrary 3 x 3 matrices in generation space, though they must be Hermitian.




Figure 3.1: Example diagrams which can contribute to the $K^0 \leftrightarrow \bar{K}^0$ mixing via loops involving squarks and gluinos. The X symbol represents the flavor-changing interaction. Constraints on kaon mixing can be used to constrain the elements of the SUSY-breaking parameters $m_{\tilde{L}}^2$, $m_{\tilde{e}}^2$ and a_e .

Higgs Masses and Mixing SUSY-breaking contributions to the Higgs potential and bilinear scalar mixing terms called B terms are also allowed.

$$\mathcal{L}_{soft} = -M_{H_u}|H_u|^2 - M_{H_d}|H_d|^2 - (bH_uH_d + h.c.) \quad (3.9)$$

The SUSY-breaking sector is the primary source of the oft-quoted 100+ free parameters in the MSSM. Fortunately, while there are very few constraints on specific parameters there exist several very strong constraints on the sector as a whole from flavor mixing and measurements of CP violation [66]. However, it should be emphasized that any set of assumptions is far from unique in its ability to suppress experimentally-disallowed flavor-changing or CP-violating processes.

Like the analyses described in this document, the following discussion will center on the squarks, though similar considerations apply to the slepton sector. Consider neutral kaon mixing ($K^0 \leftrightarrow \bar{K}^0$). Several possible contributions to this process involving virtual squarks and gluinos in box diagrams are shown in Figure 3.1.

There are three general approaches to suppressing these amplitudes [67].

1. Make the squarks very heavy.
2. Assume an “alignment” between quark and squark mass matrices.
3. Assume flavor-blind SUSY-breaking (mass degeneracy)

The first option, the “brute force” approach, is simply to make the SUSY particles involved very heavy. As the most stringent limits come from kaon mixing (as opposed to the D and B mesons, which have similar diagrams), it suffices confine this constraint to the first and second generation squarks.

The alignment option assumes that the squark and quark mass matrices have the same or related dynamical origins, and that there is a correspondence between the two which ensures that the squark and quark mass matrices can be diagonalized simultaneously. Of course, unless the squark mass matrices are proportional to the identity, they cannot simultaneously commute with both the up-type and down-type quark mass matrices—this is the origin of the CKM matrix. The diagonalization is imperfect, but can be close enough to evade the limits due to CKM mixing being small and provided the relevant A-terms are small as well.

The final approach, foreshadowed above, is to assume the squark mass matrices are proportional to the identity matrix. This is called *soft supersymmetry-breaking universality*. In fact, it is sufficient to again assume only the first and second generations are degenerate—the neutral flavor-changing processes like those shown in Figure 3.1 are then suppressed via a super-GIM mechanism [67]. The experimental constraints on the third generation are significantly smaller, as the SM contribution to $B^0 \leftrightarrow \bar{B}^0$ mixing is large and has large theoretical uncertainties. Other measurements such as limits on $b \rightarrow s\gamma$ and the electric dipole moment of the neutron and electron further constrain the off-diagonal elements of the squark-mass and A-term matrices and the CP-violating SUSY phases [66].

While it is clear there are many possible ways to evade the experimental bounds on flavor-mixing and CP violation, a consistent narrative emerges. While the first two generations of squarks are constrained to be either degenerate, heavy, or highly aligned with their SM partners with limited mixing via A-terms, the third generation is remarkably free of constraints. Thus, it is free to have smaller elements in the squark mass matrices, large A-terms, or both. These possibilities have large consequences for the mass eigenstates of the third generation (Section 3.3) and thus for naturalness (Section 3.4).

3.3 Mass Eigenstates of the MSSM

The SUSY mass hierarchy after SUSY breaking and electroweak symmetry-breaking is a rich, complicated, and very model-dependent subject. As such, this section will be a predominantly qualitative discussion of the mass eigenstates of the MSSM. It is intended to motivate how certain mass eigenstates could be light enough for naturalness and define the gaugino/higgsino mass eigenstates (neutralinos and charginos) which appear in subsequent chapters.

Sfermions The sfermion mass terms have five potential contributions: the SUSY-breaking masses (Equation 3.8), the A-terms (Equation 3.7), the quartic couplings to the Higgs fields (from so-called *D-terms* in the SUSY potential), and both left-right diagonal and off-diagonal contributions proportional to the Yukawa couplings (arising from the so-called *F-terms* in the SUSY potential). The latter four of these terms arise after electroweak symmetry-breaking and are proportional to a Higgs field VEV. It will be useful in the following discussion to define the ratio of the VEVs of H_u and H_d to be $\tan \beta$:

$$\tan \beta \equiv \frac{v_u}{v_d} \tag{3.10}$$

For the first two generations of fermions, the Yukawa couplings are small are therefore the contributions to the first and second generation sfermion masses which are proportional to the Yukawa couplings can be neglected. The A-terms for the first and second generation are also often assumed to be small to avoid issues with flavor-changing constraints (Section 3.2). In fact, A-terms are often rewritten as proportional to the Yukawa couplings or fermion masses to enforce that only the third generation A-terms can be non-negligible. As a result, the light-flavor (where “light” refers to their fermionic partners) sfermion masses are given by [61]:

$$m_{\tilde{f}}^2 = M_f^2 + \Delta_f \tag{3.11}$$

where Δ_f is a combination of $SU(2)_W \otimes U(1)_Y$ quantum numbers. It is $\mathcal{O}(m_Z^2)$ and thus much smaller than the SUSY-breaking contribution.

For the third generation sfermions, all five terms come into play. Writing the A -terms proportional to the fermion masses, the left-handed/right-handed squark mass matrix is no longer diagonal, and for the sbottoms (\tilde{b}) takes the form:

$$\mathcal{M}_{\tilde{b}}^2 = \begin{pmatrix} M_{\tilde{Q}_3}^2 + \Delta_{\tilde{b}_L} + m_b^2 & m_b(A_b - \mu \tan \beta) \\ m_b(A_b - \mu \tan \beta) & M_{\tilde{d}_3}^2 + \Delta_{\tilde{b}_R} + m_b^2 \end{pmatrix} \quad (3.12)$$

The parameter μ comes from the supersymmetric Higgs potential. The off-diagonal terms mix \tilde{b}_L and \tilde{b}_R to form the mass eigenstates \tilde{b}_1 and \tilde{b}_2 , which are pushed down and up, respectively (this numbering convention, in order of mass, will also apply to the neutralino and chargino mass eigenstates). The structure of the mass matrix for the stau ($\tilde{\tau}$) is identical. The stop (\tilde{t}) mass matrix has a similar structure, with $\tan \beta$ replaced by $\cot \beta$:

$$\mathcal{M}_{\tilde{t}}^2 = \begin{pmatrix} M_{\tilde{Q}_3}^2 + \Delta_{\tilde{t}_L} + m_t^2 & m_t(A_t - \mu \cot \beta) \\ m_t(A_t - \mu \cot \beta) & M_{\tilde{u}_3}^2 + \Delta_{\tilde{t}_R} + m_t^2 \end{pmatrix} \quad (3.13)$$

This mixing is often key to ensuring light third generation squarks and thus naturalness (Section 3.4).

Gluinos The gluinos, being a color octets, cannot mix with any other particle. The gluino mass is simply the SUSY-breaking mass M_3 , neglecting renormalization group evolution. Though constraints on the gluino mass are weak, the strong coupling of the gluino to the squarks mean it should not be heavier than a few TeV [68].

Charginos and Neutralinos The charged gauginos receive a mass contribution from the $SU(2)_W$ SUSY-breaking mass M_2 (Equation 3.6). The charged Higgsinos receive a contribution μ from the Higgs potential. After electroweak symmetry-breaking however, there are cross terms proportional to the VEVs v_u and v_d which mix the charged gauginos and Higgsinos to form mass eigenstates called charginos. The mixing matrix is given by [60]:

$$\mathcal{M}_{\tilde{\chi}^\pm} = \begin{pmatrix} M_2 & gv_u \\ gv_d & \mu \end{pmatrix} \quad (3.14)$$

Diagonalizing this matrix results in two pairs of degenerate chargino mass eigenstates, $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^\pm$.

The situation is similar for the neutral gauginos and higgsinos, except that both M_1 and M_2 are involved (the bino and the neutral wino, respectively). The four neutralino mass eigenstates arise from diagonalizing the symmetric mixing matrix [60]:

$$\mathcal{M}_{\tilde{\chi}^0} = \frac{1}{\sqrt{2}} \begin{pmatrix} \sqrt{2}M_1 & 0 & -g'v_d & g'v_u \\ 0 & \sqrt{2}M_2 & gv_d & -gv_u \\ -g'v_d & gv_d & 0 & -\sqrt{2}\mu \\ g'v_u & -gv_u & -\sqrt{2}\mu & 0 \end{pmatrix} \quad (3.15)$$

The four mass eigenstates are denoted $\tilde{\chi}_0^i$ with $i = 1, 2, 3, 4$ in order of increasing mass. The lightest neutralino is often taken to be the lightest supersymmetric particle (LSP) for cosmological reasons (Section 3.5).

Renormalization Group Evolution and the Third Generation Thus far, this chapter has mostly been concerned with parameters at the input scale. In order to ensure the absence of large logarithms in the calculations of observables at the weak scale, the couplings and soft SUSY-breaking parameters must be evolved down to the weak scale via the renormalization group. A third mechanism for generating relatively light third-generation squark masses is provided by this renormalization group evolution.

First, it is important to note that, as a result of the *supersymmetric non-renormalization theorem* [60], the β functions (Section 2.1) which control the evolution of supersymmetric parameters such as the Yukawa couplings and μ are proportional to the parameters themselves and combinations of dimensionless couplings. This ensures that they will not receive large radiative corrections from unknown heavy particles. This theorem also applies to the gauge couplings (Section 3.6).

Parameters which break SUSY however have no such protection. Consider the A-term mixing matrices \mathbf{a} . If they are approximately diagonal and only nonzero for the third generation:

$$\mathbf{a}_u \approx \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & a_t \end{pmatrix} \quad \mathbf{a}_d \approx \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & a_b \end{pmatrix} \quad (3.16)$$

then the renormalization group equations for a_t and a_b are given at one-loop by [60]:

$$\begin{aligned} \frac{1}{16\pi^2} \frac{d}{d \log Q} a_t &= \frac{1}{16\pi^2} \beta_{a_t} = a_t \left[18y_t^* y_t + y_b^* y_b - \frac{16}{3} g_3^2 - 3g_2^2 - \frac{13}{15} g_1^2 \right] \\ &+ 2a_b y_b^* y_t + y_t \left[\frac{32}{3} g_3^2 M_3 + 6g_2^2 M_2 + \frac{26}{15} g_1^2 M_1 \right] \end{aligned} \quad (3.17)$$

$$\begin{aligned} \frac{1}{16\pi^2} \beta_{a_b} &= a_b \left[18y_b^* y_b + y_t^* y_t + y_\tau^* y_\tau - \frac{16}{3} g_3^2 - 3g_2^2 - \frac{7}{15} g_1^2 \right] \\ &+ 2a_t y_t^* y_b + 2a_\tau y_\tau^* y_b + y_b \left[\frac{32}{3} g_3^2 M_3 + 6g_2^2 M_2 + \frac{14}{15} g_1^2 M_1 \right] \end{aligned} \quad (3.18)$$

Not only are these β functions not proportional to the respective parameters, there is a piece of each which involves only the gauge couplings, the Yukawa couplings, and the gaugino masses. That is, even if the \mathbf{a} matrices are zero at the input scale, renormalization group evolution will *generate* A-terms at the electroweak scale for the third-generation.

There is also the matter of the soft SUSY masses, which are also unprotected from large corrections. In the approximation of diagonal squark mass matrices with the first two generations degenerate, the one-loop β functions for the first two generations are [60]:

$$\frac{1}{16\pi^2} \beta_{M_{\tilde{Q}_{1,2}}^2} = -\frac{32}{3} g_3^2 |M_3|^2 - 6g_2^2 |M_2|^2 - \frac{2}{15} g_1^2 |M_1|^2 + \frac{1}{5} g_1^2 S \quad (3.19)$$

$$\frac{1}{16\pi^2}\beta_{M_{\bar{u}_{1,2}}^2} = -\frac{32}{3}g_3^2|M_3|^2 - \frac{32}{15}g_1^2|M_1|^2 + \frac{4}{5}g_1^2S \quad (3.20)$$

$$\frac{1}{16\pi^2}\beta_{M_{\bar{d}_{1,2}}^2} = -\frac{32}{3}g_3^2|M_3|^2 - \frac{8}{15}g_1^2|M_1|^2 + \frac{2}{5}g_1^2S \quad (3.21)$$

where S is defined as:

$$S \equiv M_{H_u}^2 - M_{H_d}^2 + \text{Tr}[\mathbf{M}_{\bar{Q}}^2 - \mathbf{M}_{\bar{L}}^2 - 2\mathbf{M}_{\bar{u}}^2 + \mathbf{M}_{\bar{d}}^2 + \mathbf{M}_{\bar{e}}^2] \quad (3.22)$$

and is often relatively small. The feature to note is that the dominant terms (those involving the gaugino masses) are all negative. Thus, the SUSY-breaking masses are *heavier* at the weak scale than they are at the input scale. For the third generation, there are additional contributions from the Yukawa couplings and A-term parameters. Define $X_{t,b}$ as:

$$X_t \equiv 2|y_t|^2(M_{H_u}^2 + M_{Q_3}^2 + M_{\bar{u}}^2) + 2|a_t|^2 \quad (3.23)$$

$$X_b \equiv 2|y_b|^2(M_{H_d}^2 + M_{Q_3}^2 + M_{\bar{d}}^2) + 2|a_b|^2 \quad (3.24)$$

Note that these quantities are positive, and can be quite large. The β functions for the third generation squarks are then [60]:

$$\frac{1}{16\pi^2}\beta_{M_{Q_3}^2} = X_t + X_b + \frac{1}{16\pi^2}\beta_{M_{Q_{1,2}}^2} \quad (3.25)$$

$$\frac{1}{16\pi^2}\beta_{M_{\bar{u}_3}^2} = 2X_t + \frac{1}{16\pi^2}\beta_{M_{\bar{u}_{1,2}}^2} \quad (3.26)$$

$$\frac{1}{16\pi^2}\beta_{M_{\bar{d}_3}^2} = 2X_b + \frac{1}{16\pi^2}\beta_{M_{\bar{d}_{1,2}}^2} \quad (3.27)$$

The positive contributions from the $X_{t,b}$ cancel some of the negative contributions from the gaugino masses, which makes the third generation evolve more slowly than

Name	Spin	Gauge Eigenstates	Mass Eigenstates
Squarks	0	$\tilde{u}_L \tilde{u}_R \tilde{d}_L \tilde{d}_R$	-
		$\tilde{c}_L \tilde{c}_R \tilde{s}_L \tilde{s}_R$	-
		$\tilde{t}_L \tilde{t}_R \tilde{b}_L \tilde{b}_R$	$\tilde{t}_1 \tilde{t}_2 \tilde{b}_1 \tilde{b}_2$
Sleptons	0	$\tilde{e}_L \tilde{e}_R \tilde{\nu}_e$	-
		$\tilde{\mu}_L \tilde{\mu}_R \tilde{\nu}_\mu$	-
		$\tilde{\tau}_L \tilde{\tau}_R \tilde{\nu}_\tau$	$\tilde{\tau}_1 \tilde{\tau}_2 \tilde{\nu}_\tau$
Gluinos	$\frac{1}{2}$	\tilde{g}^a	-
Charginos	$\frac{1}{2}$	$\tilde{W}^\pm \tilde{H}_u^+ \tilde{H}_d^-$	$\tilde{\chi}_1^\pm \tilde{\chi}_2^\pm$
Neutralinos	$\frac{1}{2}$	$\tilde{B}^0 \tilde{W}^0 \tilde{H}_u^0 \tilde{H}_d^0$	$\tilde{\chi}_1^0 \tilde{\chi}_2^0 \tilde{\chi}_3^0 \tilde{\chi}_4^0$

Table 3.1: Gauge and mass eigenstates of the supersymmetric particles of the MSSM. A dash (-) indicates that the gauge and mass eigenstates are the same, at least to good approximation with typical choices of the soft SUSY-breaking parameters.

the first two and results in lighter third generation squarks at the weak scale, assuming degenerate SUSY-breaking masses at the input scale. If the input scale third generation masses are lighter already (plausible, given how “special” the third generation is in the SM), then renormalization group effects serve to intensify the difference at the weak scale. It should be noted that this same effect pushes μ^2 negative for the Higgs field H_u and leads to the acquisition of a VEV.

MSSM Gauge and Mass Eigenstate Summary The gauge and mass eigenstates of the supersymmetric particles of MSSM are summarized in Table 3.1. Antiparticles for the squarks and sleptons are implied.

3.4 Naturalness and Heavy Flavor

The possibility of relatively light third-generation squarks has been emphasized throughout this brief introduction to supersymmetry as being important for naturalness. The reason for this is that while softly-broken supersymmetry preserves cancellations of ultra-violet divergences, the remaining logarithmic corrections can still be large, and

if they are parts of the model can still be significantly fine-tuned. However, all of these logarithmic corrections are not created equal, and the most important involve the third-generation squarks.

The supersymmetric Higgs potential, after setting $H_u^\pm = 0$, is given by [60]:

$$V = (\mu^2 + M_{H_u}^2)|H_u^0|^2 + (\mu^2 + M_{H_d}^2)|H_d^0|^2 - (bH_u^0H_d^0 + h.c.) + \frac{1}{8}(g^2 + g'^2)(|H_u^0|^2 - |H_d^0|^2)^2 \quad (3.28)$$

One will note that, in contrast to the SM version (Equation 2.9), the quartic coupling is not a new parameter λ but instead is related to the electroweak gauge couplings. As H_u and H_d have opposite weak hypercharges, b , $v_u = \langle H_u^0 \rangle$ and $v_d = \langle H_d^0 \rangle$ can all be real and positive without loss of generality. Minimizing this potential with respect to H_u^0 and H_d^0 and eliminating the VEVs and gauge couplings in favor of $\tan \beta$ and the Z mass m_Z , respectively, yields:

$$M_{H_u}^2 + \mu^2 - b \cot \beta - \frac{m_Z^2}{2} \cos(2\beta) = 0 \quad (3.29)$$

$$M_{H_d}^2 + \mu^2 - b \tan \beta + \frac{m_Z^2}{2} \cos(2\beta) = 0 \quad (3.30)$$

Eliminating b and solving for m_Z^2 in terms of the soft masses M_{H_u} and M_{H_d} and the parameter μ^2 yields:

$$m_Z^2 = 2 \frac{M_{H_d}^2 - M_{H_u}^2 \tan^2 \beta}{\tan^2 \beta - 1} - 2\mu^2 \quad (3.31)$$

m_Z is known to be small (≈ 90 GeV). In order for this to be natural and not dependent on large cancellations, the terms on the right should be relatively small as well. As applied to μ^2 , which sets the mass scale of the Higgsinos, this motivates a light neutralino and chargino mass spectrum. As for the soft mass terms, they receive radiative corrections. The largest of these comes from loops involving top squarks due to the large top Yukawa coupling y_t . At one loop, this correction has the form [68]:

$$\Sigma_u = \frac{3y_t^2}{16\pi^2} \times m_{\tilde{t}}^2 \times \text{small log} \quad (3.32)$$

where $m_{\tilde{t}}$ is the geometric mean of masses of the two top squark mass eigenstates. The quadratic dependence on the top squark masses implies that in order for Equation 3.31 to be natural the top squarks must be relatively light.

If the Higgs-like particle discovered is assumed to be the lightest supersymmetric Higgs boson h^0 , this requirement can cause significant tension within the MSSM, as similar radiative corrections are relied upon to push m_{h^0} up from the tree-level result:

$$m_{h^0}^2 \leq m_Z^2 \cos^2(2\beta) \quad (3.33)$$

The positive radiative contribution from incomplete top quark/squark loop cancellation at one loop is given by [69]:

$$\delta m_{h^0}^2 = \frac{3g^2}{8\pi^2} \frac{m_t^4}{m_W^2} \ln\left(\frac{m_{\tilde{t}}^2}{m_t^2}\right) \quad (3.34)$$

and must be large enough for m_{h^0} to reach 126 GeV. Slightly non-minimal formulations of SUSY can relax this tension considerably [70].

While the bottom squarks are not typically involved in large radiative corrections due to the much smaller bottom Yukawa coupling y_b , there are more indirect reasons to expect at least one sbottom to be relatively light as well. The left-handed stop and sbottom form an $SU(2)_W$ doublet, and the SUSY-breaking masses must respect gauge symmetry (Equation 3.8). Therefore at the input scale, before electroweak symmetry-breaking, the left-handed stop and sbottoms are degenerate with mass $M_{\tilde{Q}_3}$.

Electroweak symmetry-breaking generates additional terms as shown in Equations 3.12 and 3.13, but both the Δ_f terms, $\mathcal{O}(m_Z^2)$, and the fermion mass squared terms are significantly smaller than $M_{\tilde{Q}_3}^2$. Thus, some of the left-handed degeneracy should survive in approximate form, though large stop mixing can weaken the relationship considerably. Significant sbottom mixing, perhaps due to large $\tan \beta$, can help compensate.

3.5 R-Parity and Dark Matter

Throughout this chapter SUSY has been described primarily as a means of providing a natural solution to the hierarchy problem. However, SUSY provides other benefits “automatically” which are in some ways just as important. One of these features is a viable cold dark matter candidate.

There are several renormalizable couplings possible in the MSSM which violate either lepton number or baryon number. Individually, they are highly constrained. Together, they lead to very fast ($\approx 10^{-2}$ seconds) proton decay, which is clearly in conflict with reality. A common way to prevent this is to postulate an exactly conserved discrete symmetry, called *R-parity*, which enforces conservation of lepton and baryon numbers.

$$R_p = (-1)^{3(B-L)+2S} \quad (3.35)$$

B is baryon number, L is lepton number, and S is spin. All Standard model particles have R-parity of +1, and all SUSY particles have R-parity of -1. Exact conservation has profound experimental consequences—SUSY particles are always produced in pairs, and the lightest supersymmetric particle (LSP) is absolutely stable (as is the lightest baryon and lepton, in agreement with the observed stability of the electron and proton). This also means that each SUSY final state will include at least two invisible, weakly interacting particles carrying off momentum. This manifests itself experimentally as *missing energy* and is one of the primary experimental handles in searches for SUSY.

The LSP is typically the lightest neutralino, a particle with only weak interactions. A stable, massive, weakly-interacting, electrically-neutral particle happens to be exactly the right kind of particle needed to explain “cold” dark matter and the corresponding relic density [71]. This “coincidence” between the weak interaction scale and the dark matter relic density is called the *WIMP* (Weakly Interacting Massive Particle) *miracle*. It means that SUSY with R-parity, though motivated originally to solve the hierarchy problem (Section 2.6), also provides a prime dark matter candidate—something the Standard Model conspicuously lacks.

3.6 Gauge Unification

Dark matter is not the only seemingly-unrelated feature obtained “for free” with minimal SUSY. Grand Unified Theories, or GUTs, attempt to explain the observed quantization of electric charge (exact multiples of $\frac{1}{3}$ the fundamental electron charge) despite the fact that the $U(1)_Y$ group in the Standard Model by itself would seem to have no such restriction. A GUT is also seen as the next step in unifying the strong, weak, and electromagnetic interactions with gravity in a Theory of Everything. GUTs are typically formulated as a higher symmetry of Nature, such as $SU(5)$. It is unbroken at the GUT scale ($\approx 10^{16}$ GeV) and then spontaneously broken to become $SU(3)_C \otimes SU(2)_W \otimes U(1)_Y$ at the weak scale.

A feature of such theories is a common coupling strength for the strong and electroweak interactions at the GUT scale. Following the treatment in [61], for $SU(5)$ at the GUT scale the couplings obey the relation:

$$g_1 = g_2 = g_3 = g_{unified} \quad (3.36)$$

where

$$g_1 = \sqrt{\frac{5}{3}}g', \quad g_2 = g, \quad g_3 = g_s \quad (3.37)$$

Defining a “fine structure” constant for each of these couplings:

$$\alpha_i \equiv \frac{g_i^2}{4\pi} \quad (3.38)$$

and using the one-loop renormalization group (RG) equation for gauge couplings:

$$\frac{d\alpha_i}{d \log Q} = \beta(\alpha_i) = -\frac{b_i}{2\pi}\alpha_i^2 \quad (3.39)$$

yields the equation governing the evolution of α_i from scale M to scale Q :

$$\alpha_i^{-1}(Q) = -\frac{b_i}{2\pi} \log \frac{Q}{M} + \alpha_i^{-1}(M) \quad (3.40)$$

Assuming the α_i are equal at scale Q allows elimination of $\alpha_i^{-1}(Q)$ and $\log \frac{Q}{M}$,

which yields the relation:

$$\alpha_3^{-1} = (1 + B)\alpha_2^{-1} - B\alpha_1^{-1} \quad (3.41)$$

where

$$B \equiv \frac{b_3 - b_2}{b_2 - b_1} \quad (3.42)$$

This relation is valid at any scale within the one-loop approximation. From precision electroweak data at $Q = m_Z$, the values of α_i^{-1} are known to be [72]:

$$\alpha_1^{-1} = 59.00 \pm 0.02, \quad \alpha_2^{-1} = 29.57 \pm 0.02, \quad \alpha_3^{-1} = 8.50 \pm 0.14 \quad (3.43)$$

with the constant B at one-loop then given by:

$$B = 0.716 \quad (3.44)$$

Whether the couplings unify then becomes a question of whether the b_i coefficients of the theory, when combined into B using Equation 3.42, are consistent with the measured value of B in Equation 3.44. For the Standard Model, B is given by [61]:

$$B_{SM} = \frac{110 + 5n_h}{220 - 2n_h} \quad (3.45)$$

For n_h (the number of Higgs doublets) = 1, B_{SM} is 0.527, which is clearly inconsistent with gauge unification. In contrast, for the MSSM B is given by:

$$B_{MSSM} = \frac{30 + 5n_h}{60 - 2n_h} \quad (3.46)$$

For two Higgs doublets, this reduces to $B_{MSSM} = \frac{5}{7} = 0.714$, almost exactly in agreement with Equation 3.44. The unification of gauge couplings in the MSSM and the non-unification in the SM are illustrated graphically in Figure 3.2.




Figure 3.2: Renormalization group evolution of gauge couplings in the Standard Model (left) and the MSSM (right). In the MSSM the gauge couplings unify at $\approx 10^{16}$ GeV

Chapter 4

Simplified Models

In Chapter 3, naturalness arguments were used to motivate light third-generation squarks in a general way. However, this feature itself is not sufficient to constrain the high-scale supersymmetry parameters significantly, nor do all SUSY models containing light third-generation squarks share the same final state kinematics. Moreover, there are often multiple production modes and decay paths contributing to the same experimental final state (i.e. n jets and missing transverse energy), with contributions highly dependent on choice of model parameters. In this chapter, the traditional top-down approach to new physics searches will be contrasted with the bottom-up or *simplified models* approach [22, 23, 24] which seeks to address some of these limitations.

4.1 Top-down Approach

In traditional searches conducted at the Tevatron and with the first LHC data, the first step was to choose a reference model. In many cases the reference model of choice was minimal supergravity (mSUGRA) [73, 74, 75]. Via assumptions such as sfermion and gaugino mass universality at the SUSY-breaking scale, mSUGRA models reduce the MSSM parameter space to four parameters and a sign. This set of parameters is significantly more tractable than the 100+ free parameters in the MSSM. The next step involves generating simulated events over a representative grid

in this reduced parameter space, and then optimizing selections to maximize search sensitivity. Sensitivity is quantified statistically using the significance of the expected deviation from background hypothesis.

This procedure was reasonable for the early data. mSUGRA was a well-motivated, complete supersymmetric model and seemed like a leading contender to be the version of SUSY chosen by Nature. mSUGRA also conveniently predicts several dramatic, hard-to-miss final-state signatures. However, choosing a single SUSY model, any model, is inherently restrictive. After two years of LHC data with no sign of SUSY, we can no longer afford to limit ourselves to convenient models only.

High Scale vs. Weak Scale Parameters One difficulty with the top-down approach is that the model parameters, such as the gaugino and squark masses, are specified at high scale rather than at the weak scale. The weak-scale parameters can always be generated via renormalization group evolution, but this evolution is seldom trivial, and as a result the weak-scale consequences of high-scale parameter choices are often unclear or under-appreciated. An example of this is the so-called *compressed scenarios*, the name given to the final-state signatures which arise from a SUSY mass spectrum with large weak-scale degeneracy, in particular, between the neutralino and gluino/squark masses. Due to the restricted phase space, these models often result in soft final-state signatures which require novel search strategies. mSUGRA itself is not generic enough to produce compressed spectra, and the fact that analyses designed to find mSUGRA tended to be insensitive to these scenarios was not well-appreciated during the early days of the LHC.

One of the reasons such weaknesses went unnoticed was the practice of interpreting results exclusively in terms of high-scale parameters as shown in Figure 4.1. Such a plot cannot reveal lack of sensitivity to certain regions of phase space if such regions are not accessible in the m_0 (sfermions) and $m_{1/2}$ (gauginos) plane. Moreover, results presented in this way make reinterpretation of results in alternative models difficult—the high scale limits must be reinterpreted in terms of weak-scale parameters, and then the weak-scale parameters translated back to the high-scale parameters of another model.




Figure 4.1: An example of an mSUGRA-style exclusion plot in terms of m_0 and $m_{1/2}$, the sfermion and gaugino masses [76]

Single vs. Multiple Decay Chains New physics searches are naturally organized around particular *channels*. These are final states which share similar particle content and kinematic features. An example would be final states with one lepton (electron or muon) and additional jets, defined as clusters of final state particles (Section 7.7). However, very rarely is there a one-to-one correspondence between final states and decay chains in a given model. A sample mSUGRA mass spectrum, with possible decays and branching ratios, is shown in Figure 4.2. A final state with no electrons or muons and two or more b -jets and missing transverse energy would have contributions from stop pair production and subsequent decay to the lightest neutralino ($\tilde{t} \rightarrow t + \chi_1^0, t \rightarrow b + W, W \rightarrow \tau\nu$) and sbottom decay to the lightest neutralino ($\tilde{b} \rightarrow b + \chi_1^0$). Also contributing, however, would be decays via the second-lightest neutralino χ_2^0 , which would add additional jets to the final state, in some cases additional b -jets ($\chi_2^0 \rightarrow \chi_1^0 + h, h \rightarrow b\bar{b}$), as well as decays involving the charginos if the mass difference between the lightest neutralino and lightest charginos is small (thus rendering the leptons from such a decay unobservable). Gluino production can also contribute via decays to squark-quark pairs ($\tilde{g} \rightarrow \tilde{q}q$).

Thus, an analysis optimized for a particular model in a given channel is tailored, to at least some degree, to the particular mix of masses, branching ratios, and production modes of that model. All of these separate decay chains will, in general, have different kinematic features. It follows that such an analysis is suboptimal for any individual decay chain, which limits its sensitivity to alternate models with similar final state signatures.

Cross Sections and Branching Ratios Cross sections and branching ratios have another negative effect on the generic sensitivity of top-down searches in that they are key to the classification of a region of phase space as “excluded”. In terms of the reference model, this is true. In other models with similar final states, it may not be excluded if the cross section or branching ratio is suppressed relative to the reference model. If the same reference model is used in the next iteration of the analysis, the excluded region is often ignored in the optimization, and sensitivity to this region of phase space can be lost.




Figure 4.2: Sample mSUGRA mass spectrum [77]

4.2 Bottom-up/Simplified Models Approach

The simplified model approach can be thought of as a change of basis in model space to weak-scale parameters combined with individual consideration of each production mode and decay chain. Thus, sbottom pair production and decay to two bottom quarks and two neutralinos would be one simplified model, while gluino-mediated sbottom production with the same sbottom decay mode would be another. The weak-scale masses are taken to be the parameters of the models. In the direct sbottom production case, the simplified model is specified by two parameters, the sbottom mass $m_{\tilde{b}}$ and the neutralino mass $m_{\chi_1^0}$ (the sbottom lifetime, a third potential parameter, is assumed to be small enough to be irrelevant). Varying the model parameters then gives a complete survey of the final state configurations produced by this decay chain and exposes the weaknesses of a given analysis in an obvious way. It should be noted that simplified models are not full SUSY models—only the parameters relevant to the single decay chain are included. Limits expressed in terms of the weak scale parameters need to be adjusted for branching ratio and, in some cases, for cross section when applied to full models.

4.3 Direct Sbottom/Stop Production

As previously mentioned, the direct squark production model has two parameters: the squark mass and the mass of the neutralino. The entire event topology is shown diagrammatically in Figure 4.3. The mass spectrum is shown in Figure 4.4.

There can be no other parameters, neglecting possible production-decay spin correlations, because energy and momentum conservation completely determine the kinematics in the center-of-mass frame of the squark. The decay matrix element is therefore a constant.

The overall event topology is also dependent on the angle and momentum of the squarks coming out of the hard-scatter. These are determined by the squark mass and the proton parton density functions (PDFs). As the squark is a colored particle, the dominant leading-order (LO) production diagrams involve the strong interactions




Figure 4.3: Direct third-generation squark pair production and decay to two third-generation quarks and two neutralinos.




Figure 4.4: Mass spectrum of the direct stop production (left) and direct sbottom production (right) simplified models. All other flavors of squarks in each model and the gluino are considered heavy and therefore decoupled.




Figure 4.5: Dominant leading-order diagrams for third-generation squark pair production at the LHC. All other flavors of squarks and the gluino are assumed heavy and therefore decoupled.

(Figure 4.5).

The analysis described in Chapter 10 is targeted at the sbottom case. It is also sensitive to stop decay to a bottom quark and a chargino (an alternative decay chain) provided the mass difference between the chargino and neutralino is small (they are higgsino-like) and the lepton from the chargino decay is lost. ATLAS searches for other stop decay modes can be found in References [27, 28, 29, 30, 31, 32, 33, 34].

4.4 Gluino-mediated Sbottom/Stop Production

The second type of simplified model considered is production of gluinos which subsequently decay to a top-antitop or bottom-antibottom pair and a neutralino via off-shell stops or sbottoms (Figure 4.6).

This decay proceeds via an off-shell stop or sbottom. The term “off-shell” indicates that the squark is heavier than the gluino (Figure 4.7). The stop and sbottom decays are considered separate models.




Figure 4.6: Gluino-mediated production of stop/sbottom squarks. The final state involves four top or bottom quarks and two neutralinos.




Figure 4.7: Mass spectra for gluino decays to 4 top quarks (left) or four bottom quarks (right) and 2 neutralinos.




Figure 4.8: Gluino and sbottom/stop pair production next-to-leading order + next-to-leading log (NLO+NLL) cross sections at the LHC at $\sqrt{s} = 7$ TeV [78, 79, 80] computed using PROSPINO [81].

As noted in Section 3.3, the constraints on the gluino mass are weak. It could easily be too heavy to produce at current energies, or it could be relatively light. However, if it is light, the production cross section at the LHC are large: typically fifty times higher than direct third-generation squark pair production for the same mass (Figure 4.8).

In addition, the final state of gluino-mediated squark production has four quarks. In the third-generation case, there will be four b -jets. This is a dramatic signature with very little Standard Model background, and thus is the preferred channel for discovery *if* the gluino happens to be light. The dominant LO strong production diagrams are shown in Figure 4.9.

The gluino in this model decays directly to a three-body final state. In such




Figure 4.9: Dominant leading order diagrams for gluino pair production at the LHC. All squarks are assumed heavy and thus decoupled.

decays, energy and momentum conservation alone do not fully determine the final state kinematics. Two degrees of freedom remain, and therefore the decay matrix element can have nontrivial dependence on the decay kinematics and spins [82]. The matrix element also depends to some extent on model-specific parameters, such as the elements neutralino mixing matrix (Equation 3.15). This is also true for direct sbottom production and decay, but there it affects only the value of the constant matrix element and not the kinematic distributions. Fortunately, for SUSY such effects tend to be small [83], in particular for a heavy intermediate squark. The decay is assumed to proceed according to phase space with the matrix element assumed to have negligible kinematic and spin dependence. The model is then completely parameterized by the gluino and neutralino masses.

Gluino-mediated models with on-shell squarks are fundamentally different in that their kinematics depend on three parameters: the gluino mass, the squark mass, and the neutralino mass. However, in the case of the sbottom squark, from the perspective of the final state there are two main categories of signatures which arise

from on-shell gluino decays. In one broad category, the gluino-sbottom and sbottom-neutralino mass differences are both large enough to produce energetic b -jets and the final states qualitatively resemble those of the off-shell decays. In the other category, either the gluino and squark or squark and neutralino are almost degenerate and the final state has only two detectable b -jets instead of four. This final state resembles that of the direct sbottom production. Therefore, two analyses targeting direct sbottom production and gluino production and decay via an off-shell sbottom should have some sensitivity to on-shell gluino decays as well. Note, however, that the production cross sections will differ in the on-shell case from those shown in Figure 4.8 due to additional diagrams involving intermediate light squarks. Fortunately, the optimization procedure is largely independent of production cross section/branch ratio (Chapter 8).

For gluino decays to a top-antitop and a neutralino, the large top mass virtually guarantees that the b -jets in the final state will be energetic enough to be detectable regardless of the choice of stop mass. The differences in available phase space between models is still observable, but the difference between on-shell and off-shell decays is reduced.

An additional simplified model with the gluino decaying to a top quark, bottom quark, and chargino is also considered for interpretation of results.

$$pp \rightarrow \tilde{g}\tilde{g}, \tilde{g} \rightarrow t\bar{b} + \chi_1^-, \chi_1^- \rightarrow l^- + \chi_1^0 \quad (4.1)$$

The chargino is almost degenerate with the neutralino, and therefore the leptons are too soft to be reconstructed. The chargino final state can be the dominant decay mode of the stop or sbottom in some models, including the case of an on-shell stop where $m_{\tilde{t}} - m_{\chi_1^0} < m_t$.

In summary, the gluino decay model via an off-shell squark is used for optimization and the primary interpretation because of its relative simplicity (two parameters instead of three) and the fact that it is expected to encompass the on-shell case to a reasonable degree. However, one should note that the direct squark production analyses should also cover the on-shell case and possibly be more sensitive.

4.5 Beyond Supersymmetry

The simplified models described in Section 4.2 make minimal SUSY-specific assumptions. Therefore, analyses targeted at these models should have some sensitivity to other models with similar colored particles and decay chains. One such model is Universal Extra Dimensions [64]. The analogue of the gluino decay in SUSY is the decay of the first-level Kaluza Klein (KK) excitation of the gluon to quark-antiquark and the first-level KK excitation of the weak hypercharge boson B [82]:

$$pp \rightarrow g^1 g^1, g^1 \rightarrow q\bar{q} + B^1 \quad (4.2)$$

Here g^1 is a color octet like the gluino, but has spin 1. The KK quarks are spin $\frac{1}{2}$ rather than spin 0 like the squarks, and the B is spin 1. If spin-correlations are subdominant then the simplified model analyses should also be sensitive to this topology. No interpretation in this context is attempted, but the fact that such an interpretation could be done on a ‘‘SUSY’’ analysis simply by rescaling cross sections illustrates well the power of the simplified models approach.

Chapter 5

The Large Hadron Collider

The Large Hadron Collider (LHC) [26] is the world's largest and most energetic particle collider, as well as the largest machine ever built by humankind. At its heart is a proton-proton (pp) synchrotron housed within a 27-kilometer subterranean ring situated approximately 100 m below the border between France and Switzerland. While the interaction points (IPs) experience temperatures up to 500,000 times hotter than those at the center of the sun, the superconducting magnets which bend the proton beams are cooled to 1.9 K, a temperature lower than that of the Cosmic Microwave Background. Meanwhile, the beam pipe vacuum is maintained at a pressure of 10^{-13} atmospheres, lower than that of interplanetary space. These individual engineering feats together form a precision instrument capable of colliding protons at rates upward of millions per second and open up exploration of the tera-electronvolt (TeV) scale, an energy range in elementary particle physics never before explored by colliders.

5.1 Project Proposal

The concept of re-using the Large Electron Positron (LEP) collider tunnel for a subsequent hadron collider was discussed as early as 1977, only two years after the idea of LEP itself. In fact, the insistence that the LEP tunnel be of large enough circumference so as to not compromise the energy of an eventual hadron collider was one of the considerations during the discussions leading up to LEP approval in 1981

[84]. When the 40 TeV Superconducting Super Collider (SSC) was approved by the United States Congress in 1987, serious doubts regarding the existence of sufficient international support for two large colliders. The large energy disadvantage the LHC would have relative to the SSC almost derailed the project entirely. Lyn Evans, LHC project leader, has stated that it was only "the resilience and conviction of Carlo Rubbia" that the kept LHC project alive.

The cancellation of the SSC in October, 1993 resulted in renewed interest in the LHC from both sides of the Atlantic. This did not, however, signal the end of the the project's financial woes. In December, 1993, a proposal to build the machine over a 10-year period by reducing CERN's other experimental commitments to the bare minimum (with the exception of LEP) was generally well-received. However, it was blocked by Germany and the United Kingdom, who were uncomfortable with the required increase in the CERN budget. During this period, the CERN Council voting rules were amended at Germany and the UK's instigation from a simple to double majority where the second major criterion was weighted towards large financial contributors (member states with large GNP).

In June, 1994, the LHC proposal was made to the Council again but, despite a large simple majority in favor, approval was blocked. Germany and the UK demanded significant additional contributions from France and Switzerland, the CERN host states, and that financial planning proceed under pessimistic budget expectations (1% annual reduction in real budget). The host-state impasse was resolved when France and Switzerland offered to make extra voluntary contributions. The overall budget constraint forced the LHC to be proposed as a two-stage machine, where only two-thirds of the dipole magnets would be installed in the first stage. This schedule was to be revised in 1997 after non-member state contributions were known.

Non-member state financial contributions were forthcoming, with Japan announcing a financial contribution in 1995; India, Russia, and Canada in 1996; and the United States in 1997. However in June, 1996, Germany announced that the financial cost of reunification necessitated a unilateral reduction of its CERN contribution by between 8% and 9%. The UK soon followed suit. The eventual solution to this crisis came in the 1996 decision to allow CERN to take on debt to finance a one-stage LHC project,

with the eventual completion date adjustable based on a future cost estimate revision. This revision came in 2001 when the estimated LHC price tag was increased by 18%.

5.2 Machine Design

In order to reduce costs and exploit existing facilities, the LHC is actually the final stage in a string of daisy-chained accelerators at the CERN complex, some of which have been in operation (with periodic upgrades) for more than 50 years. The CERN accelerator complex with approximate completion dates for each of its components is shown schematically in Figure 5.1 (not to scale).

CERN operates a rich non-LHC particle physics program, and not all components of the CERN accelerator complex are part of the LHC accelerator chain. Those that are, and the corresponding proton energies, are: the Linac 2 (50 MeV), where the protons are initially accelerated from rest; the Proton Synchrotron Booster (PSB or "Booster", 1.4 GeV); the Proton Synchrotron (PS, 25 GeV); the Super Proton Synchrotron (SPS, 450 GeV); and, finally, the LHC (0.45 - 7 TeV). The LHC both accelerates the injected proton bunches and collides them at several experimental interaction points (IPs) around the 26.6 km ring. It does this with an ensemble of 9,300 superconducting magnets of various types, 16 radio-frequency (RF) cavities, and a supporting infrastructure which includes the liquid-helium cryogenics system necessary for magnet operation.

Out of the 9,300 LHC magnets, the largest and most important are the 1,232 dipole magnets which bend the beam along its approximately circular trajectory and the 858 quadrupole magnets which are the primary magnets responsible for beam focusing. The other magnets in the 9,300 total are smaller correction magnets (sextupoles, octupoles, decapoles, etc.) which are also employed in beam shaping and are often contained inside the cryogenic volume of the larger dipoles and quadrupoles. The dipole magnets represent the largest technical challenge, as the peak dipole magnetic field, to first order, controls the maximum energy of the accelerator. For the LHC, the dipole magnets were designed to provide a peak field of 8.3 T, which is sufficient for 7 TeV proton beams (the design energy of the LHC).

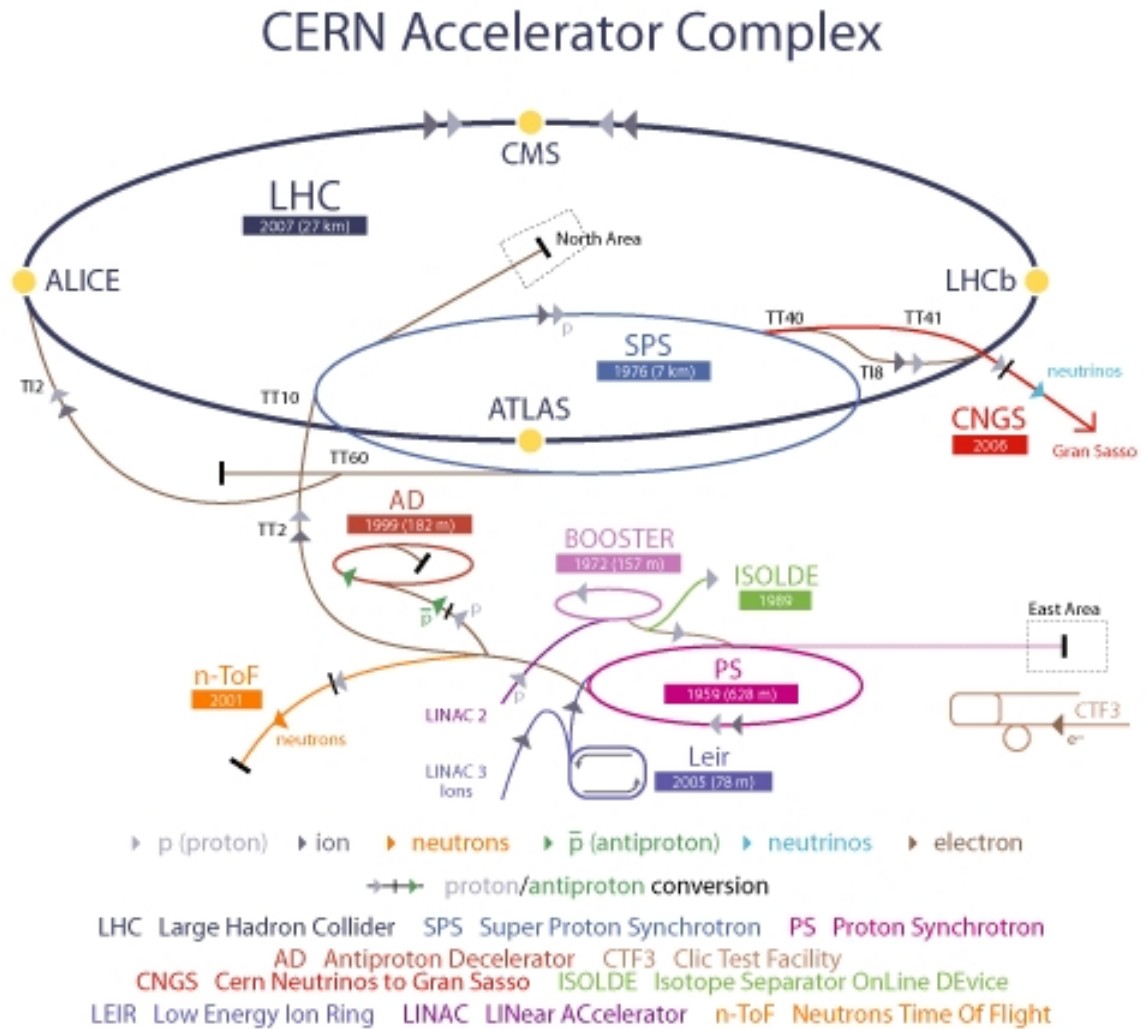


Figure 5.1: The CERN accelerator complex © 2008 CERN

5.3 Operation

Proton beams at the injection energy (450 GeV) were circulated in the LHC for the first time on September 10th, 2008, to much fanfare. On September 19th a faulty dipole magnet interconnect caused a quench in approximately 100 dipole magnets [85]. The abrupt temperature rise in the affected magnets and associated loss of liquid-helium cooling resulted in mechanical damage to the affected sectors. Moreover, it was determined that additional magnet interconnects may have been affected. This has limited the pp beam energy to 4 TeV as of the end of 2012. One of the primary machine goals for the multi-year shutdown beginning in 2013 is to ensure that the design beam energy of 7 TeV can be reached when collisions resume.

Most of 2009 was spent understanding the 2008 incident and conducting related repairs. On November 20th, 2009, 450 GeV proton beams were circulated in the LHC for the first time since 2008, and on November 23rd the first pp collisions were achieved in four LHC detectors. Short, low-luminosity runs at center-of-mass energies (\sqrt{s}) of 900 GeV and 2.36 TeV followed. On March 30th, 2010, beams were collided for the first time at $\sqrt{s} = 7$ TeV (3.5 TeV beam energy) where the LHC would remain for the next two years.

The start of 7 TeV physics runs marked the true beginning of the LHC research program. However, the collision rate, or luminosity, is almost as important as the energy for physics analysis. The luminosity depends both on the number of protons circulating in the collider and their configuration. The LHC design luminosity is 10^{34} $\text{cm}^{-2} \text{s}^{-1}$ [26]. While the luminosity was steadily increased during 2010, the maximum peak luminosity achieved was 2×10^{32} $\text{cm}^{-2} \text{s}^{-1}$, a factor of 50 lower than design. As a result, the 2010 dataset consisted of only 40 pb^{-1} of data.

The high-luminosity $\sqrt{s} = 7$ TeV run began on March 13th, 2011. The highest peak luminosity achieved during this period was 3.65×10^{33} $\text{cm}^{-2} \text{s}^{-1}$, only a factor of 3 lower than design. The total integrated luminosity delivered by the LHC to ATLAS during 2011 was 5.61 fb^{-1} , of which 5.25 fb^{-1} was recorded. This dataset is used for the searches described in this document, and is described in further detail in Section 7.1. Peak luminosity per fill and total integrated luminosity as a function of time for




Figure 5.2: Peak instantaneous LHC luminosity for each fill for the 2011 $\sqrt{s} = 7$ TeV data-taking period.

the $\sqrt{s} = 7$ TeV 2011 data period are shown in Figures 5.2 and 5.3.




Figure 5.3: Integrated LHC luminosity delivered to and record by ATLAS per day for the 2011 $\sqrt{s} = 7$ TeV data-taking period.

Chapter 6

The ATLAS Detector

Several experiments are tasked with recording and analyzing the collisions produced by the LHC. The detectors for these experiments are placed at different points along the 27 km LHC ring, called interaction points (IPs). The two all-purpose detectors are ATLAS [25] and CMS [86], and are located at IP1 and IP5, respectively. A specialized detector designed specifically for heavy ion collisions is called ALICE [87] and is located at IP2, while another detector built specifically for b -hadron physics (LHCb [88]) is situated at IP8.

The analyses presented in this document use data recorded with the ATLAS (A Toroidal LHC ApparatuS) detector, shown in Figure 6.1. ATLAS was designed with two main goals: facilitate the precision measurement of physics objects over a large solid angle (necessary for the successful prosecution of the LHC physics program) and survive the extraordinarily harsh environment near the interaction point. ATLAS follows a traditional 4π detector construction, with a barrel section and two end-cap sections with gaps for the beam line.

Each barrel and end-cap section is further divided into tracker, calorimeter, and muon spectrometer sections, in order of increasing distance from the interaction point. The inner detector with its 2 T solenoidal magnetic field is a multi-component tracker for charged particles and is described in Section 6.2. The calorimeter, located with the superconducting solenoid in a cryostat outside the inner detector, is designed to measure the energy of interacting particles and is described in Section 6.3. The




Figure 6.1: Computer-generated image of the ATLAS detector, with humans for scale. ATLAS Experiment © 2012 CERN.

ATLAS muon spectrometer and toroidal magnetic field, representing the majority of ATLAS' 44 m by 25 m volume, sits outside the calorimeter. A description of the muon spectrometer is given in Section 6.4, and the magnet system is described in Section 6.5. Specialized sub-detectors designed for luminosity and other measurements in the far forward regions are discussed in Section 6.6. A dedicated trigger and data acquisition (TDAQ) system is used to cope with the enormous LHC luminosity and select those events most interesting for physics. This system is described in Section 6.7.

6.1 Coordinate System

The coordinate system used in ATLAS is a right-handed Cartesian system centered at the nominal interaction point. The z axis lies along the beam line. In the plane transverse to the beam, the positive x direction points towards the center of the LHC ring and the positive y direction points upward. The direction of positive z (counterclockwise on the LHC ring when viewed from above) is denoted Side A of the detector, while Side C is the side in the negative- z direction.

Due to the cylindrical shape of the detector, a more common set of directional coordinates, in particular for physics objects, is the azimuthal angle ϕ and the *pseudorapidity* η , and for absolute locations in the detector the transverse radius r , ϕ , and z . The azimuthal angle is measured from the positive x axis as usual. The pseudorapidity is defined in terms of the polar angle from the positive z axis (θ) as:

$$\eta \equiv -\ln \tan \left(\frac{\theta}{2} \right) \quad (6.1)$$

The advantage of η over z as a coordinate is that in the massless limit, η approximates the rapidity y :

$$y \equiv \frac{1}{2} \ln \left(\frac{E + p_z}{E - p_z} \right) \quad (6.2)$$

The difference in rapidity between two particles is invariant under Lorentz boosts along the beam line. The rapidity is used over the pseudorapidity in operations

involving possibly massive objects, such as jet-finding (Section 7.7). The distance between two objects in the ϕ/η plane, denoted ΔR , is also a commonly used quantity, in particular for physics object overlap removal (Section 7.10).

$$\Delta R \equiv \sqrt{(\Delta\phi)^2 + (\Delta\eta)^2} \quad (6.3)$$

6.2 Inner Detector

The inner detector, or central tracking system, is charged with reconstructing the paths of charged particles via high-resolution position measurements known as hits in the three tracking sub-detectors: the silicon pixel tracker (Pixel), the silicon strip tracker (SCT), and the transition radiation tracker (TRT). Combined with a 2 T solenoidal magnetic field aligned with the beam line, these tracks allow the direction and momentum of charged particles to be determined with high accuracy (transverse momentum and impact parameter resolution of $0.05\% \text{ GeV}^{-1} \times p_T \oplus 1\%$ and $10 \mu\text{m}$, respectively [25]). At design luminosity, the inner detector will be required to measure $\mathcal{O}(1,000)$ charged particle tracks above the minimum p_T threshold of 0.5 GeV every 25 ns [89]. The track acceptance/fiducial volume is defined by $|\eta| < 2.5$, and the reconstruction efficiency varies from 78% at the minimum track p_T to above 85% for tracks above 10 GeV [90]. Diagrams of the inner detector components and their relative positions are shown in Figures 6.2 and 6.3.

6.2.1 Silicon Pixel Tracker

The silicon pixel tracker, or pixel detector, is the innermost part of the tracking system and consists of 1,744 pixel modules, the basic unit of data acquisition for the pixel detector. Each module contains 47,232 pixels of size $50 \mu\text{m}$ by $400 \mu\text{m}$, and is read out by sixteen radiation-hard front-end chips bonded to the back of the module for a grand total of 80.4 million pixel read-out channels. In the barrel section, the pixel modules are oriented parallel to the beam line and have resolutions of $10 \mu\text{m}$ in r - ϕ and $115 \mu\text{m}$ in z [89]. In the end caps, the modules are arranged radially relative




Figure 6.2: Diagram of the inner detector and components, ATLAS Experiment
 © 2012 CERN.




Figure 6.3: Diagram showing the radial positions of the inner detector components, ATLAS Experiment © 2012 CERN.

to the beam line, resulting a resolutions of $10\ \mu\text{m}$ in $r\text{-}\phi$ and $115\ \mu\text{m}$ in r .

The barrel section and each end cap is divided into three layers (in the end caps these are referred to as disks). The innermost barrel layer is situated 50.5 mm from the beam line and is known as the b -layer or vertexing layer due to its key role in the reconstruction of both primary and secondary vertices. Both primary and secondary vertices are used in the identification of jets derived from heavy-flavor quarks. The outermost layer is located 122.5 mm from the beam line.

As the ATLAS detector closest to the interaction region, the pixel detector receives the highest radiation dose—approximately 158 kGy/year at the b -layer and 25.4 kGy/year at the outermost pixel barrel layer at design luminosity [25]. This will eventually damage the b -layer to the such an extent that it will be unable to perform its physics role. The pixel detector was originally scheduled for upgrade after $300\ \text{fb}^{-1}$ of integrated luminosity with an Insertable b -Layer (IBL) placed 31 mm from the beam line between the current vertexing layer and a new, narrower beam pipe [91, 92]. This upgrade schedule has since been advanced to the long shutdown beginning in 2013, giving ATLAS four functioning pixel layers for the first 14 TeV data [93].

6.2.2 Silicon Strip Tracker

Outside the pixel detector sits a silicon strip detector known as the Semi-Conductor Tracker (SCT). Like the pixel detector, it is divided into barrel ($|\eta| < 1.1$) and end-cap components ($1.1 < |\eta| < 2.5$). There are 4,088 SCT modules, each composed of two 64 mm silicon strip sensors attached to one another with a 40-mrad angular offset to provide two-dimensional hit localization. In the barrel these are arranged with an average strip pitch of $80\ \mu\text{m}$ in four layers parallel to the beam line and to the solenoid magnetic field. In the end caps, the modules are arranged radially into eighteen disks, nine in each end cap. The radial space occupied by the SCT barrel is between 275 mm and 560 mm from the beam line. The hit resolution in the transverse direction ($r\text{-}\phi$) is $17\ \mu\text{m}$, while the longitudinal hit resolution (z in the barrel, r in the end caps) is $580\ \mu\text{m}$ [25, 89]. Like the pixel detector, the strips are read out with radiation-hard front-end chips for a total of 6.3 million read-out channels [94]. The

radiation dose received by the innermost barrel layer of the SCT is estimated to be 7,600 Gy/year [25], which is significantly less than pixel detector radiation flux.

6.2.3 Transition Radiation Tracker

Outside the silicon detectors sits the transition radiation tracker, or TRT, which provides charged track measurements for $\eta < 2.0$. The TRT consists of straw tubes (proportional drift tubes) 4 mm in diameter. These are arranged longitudinally in up to 76 layers in the barrel section, and radially in the end caps in 160 planes. The 52,544 barrel tubes are 1,440 mm in length and they populate two active regions divided at $\eta = 0$ by a plastic support. Each half is read out separately. They are further arranged in three cylindrical groupings and divided into 32 sectors in ϕ [95]. Each end cap consists of 122,880 370 mm-long tubes arranged in eighteen “wheels”. In total, the TRT has 350,848 read-out channels.

While the r - ϕ resolution of TRT hits is $130 \mu\text{m}$ [25], the relative lack of precision as compared to the silicon trackers is compensated by the large number of TRT hits per track (typically 30) [89]. In addition, the barrel drift tubes are embedded in a matrix of polypropylene-polyethylene fibers with a diameter of $19 \mu\text{m}$ which aid in electron identification via transition X-rays. These X-rays are generated as charged particles cross materials with different dielectric constants. These photons are measured by a radiation-hard discriminator in the front-end electronics [96]. In the end caps, the transition radiation is generated by foil interleaved between the straws.

6.3 Calorimetry

The ATLAS physics program and high-luminosity LHC environment require high-resolution and high-granularity energy measurements of electromagnetic and hadronic physics objects out to $|\eta| < 4.9$. To achieve this, two detector technologies are utilized. For the electromagnetic calorimeters and the hadronic end-cap calorimeters, a liquid-argon (LAr) active medium is used due to its radiation hardness and fast response. These calorimeters require cryogenic cooling to 80 K using liquid nitrogen,

and cryostats to house them. They are discussed in Section 6.3.1. The barrel section hadronic calorimeter, called the tile calorimeter, utilizes plastic scintillator tiles separated by steel absorbers (Section 6.3.2). This technology was motivated by robustness, high energy resolution due to its opacity to hadrons (9.7 nuclear interaction lengths at $\eta = 0$), and lower cost as compared to liquid argon. The tile calorimeter is cooled with water.

All of these systems are *non-compensating, sampling* calorimeters with extensive transverse and longitudinal segmentation for high-resolution measurements of electromagnetic and hadronic showers. The term “non-compensating” refers to the fact that some energy from the particle interaction with the calorimeter is lost and therefore not present in the signal from the active medium. The term “sampling” indicates the design choice of alternating layers of absorber and active medium with the read-out signal proportional to the energy loss in each layer of active medium. The total energy of the incident particle is inferred from the size and distribution of the sampled energy deposits.

A baseline calibration was obtained via electron test beam measurements for the electromagnetic calorimeters [97, 98, 99, 100] and cosmic and test beam muons for the hadronic calorimeters [101, 102, 103, 100]. The $Z \rightarrow ee$ invariant mass from collision data is then used to refine the baseline test beam/cosmic muon calibration. A diagram of the ATLAS calorimeter system is shown in Figure 6.4.

6.3.1 Liquid Argon Calorimeters

The liquid-argon (LAr) calorimeters have considerable diversity in design and choice of absorber. These differences are related to the individual calorimeters’ differing locations and tasks. However, the common choice of liquid argon as the active medium means all signals originate in the same way. An *ionization gap* is formed by applying a bias voltage across a region filled with liquid argon. A charged particle entering this gap ionizes the liquid argon and the charge is collected by the read-out. The charged particles which generate the ionization can be direct products of the hard-scatter interaction, the result of conversion (in the case of photons), or from showers




Figure 6.4: Diagram of the calorimeter subsystems, ATLAS Experiment © 2012 CERN.




Figure 6.5: Comparison between a triangular current pulse in a liquid-argon cell and the output signal after bipolar pulse shaping [25].

originating in the absorbers.

The ionization signal is then shaped by electronics which induce a bipolar pulse of integral zero (Figure 6.5). This shaping is designed to cancel the effect of a constant energy injection into the calorimeter and therefore corrects for the high-luminosity LHC environment and multiple interactions in an average sense. All liquid-argon calorimeter signals are shaped in this way though each detector's pulse peaking time and integration length is optimized independently [104].

Electromagnetic Calorimeters The electromagnetic barrel (EMB) and end-cap (EMEC) calorimeters feature a novel accordion structure (Figure 6.6) which minimizes gaps in ϕ while also ensuring an approximately-constant liquid-argon ionization gap and sampling fraction within regions of the detector. The absorbers for these calorimeters are lead plates laminated with stainless steel. The ionization gap




Figure 6.6: The accordion structure of the electromagnetic liquid-argon calorimeters, ATLAS Experiment © 2012 CERN.

is formed by these plates in combination with spacers made of three copper layers and honeycombed polyimide sheets. The spacers contain the active medium with the two outer copper layers providing the high-voltage bias and the inner layer reading out the signal via capacitive coupling [25].

The EMB covers $|\eta| < 1.475$. For the $|\eta| < 0.8$ region, the lead absorber plate thickness is 1.53 mm while beyond the thickness is reduced to 1.13 mm to maintain the high- η sampling fraction. Both the transverse and longitudinal granularity of the EMB is a function of the detector cell radial layer as shown in Figure 6.7. The primary EMB layer, Layer 2, has a granularity of 0.025 in both $\Delta\phi$ and $\Delta\eta$, while Layers 1 and 3 have cells of size 0.1×0.0031 and 0.025×0.05 in $\Delta\phi \times \Delta\eta$, respectively. These differences in granularity aid in shower shape determination. The EMB is housed inside a cryostat which it shares with the inner detector's solenoid magnet (Section

6.5.1) to reduce the amount of dead material in front of the calorimeter.

The EMEC, hadronic end-cap calorimeter, and forward calorimeter are housed in two additional cryostats at each end of ATLAS. The EMEC covers the region defined by $1.375 < |\eta| < 3.2$ and also features the accordion geometry (Figure 6.6), though its orientation differs from that of the EMB. The cell granularity also tends to be larger. Like the EMB, the EMEC uses lead plates as absorbers with thicknesses of 1.7 mm for the $|\eta| < 2.5$ region and 2.2 mm for the $|\eta| > 2.5$ region. The design energy resolution of the EM calorimeters is $10\% \text{ GeV}^{\frac{1}{2}}/\sqrt{E} \oplus 0.7\%$ [105] and they provide a minimum depth of 22 radiation lengths. A pre-sampler inside the first layer is present for the region $|\eta| < 1.8$. It is used to estimate energy losses from dead material encountered before the calorimeter. The EMB contains 109,568 read-out channels while the EMEC has 63,744.

Hadronic and Forward LAr Calorimeters Behind the EMEC (as seen from the interaction region) sits the hadronic end cap ($1.5 < |\eta| < 3.2$), or HEC, which also uses liquid argon as an active medium due to the intense radiation environment at high $|\eta|$. Instead of the accordion geometry, a more traditional parallel-plate geometry is used. The absorbers are made from copper and provide a depth of 10 nuclear interaction lengths. The design energy resolution of the HEC is $50\% \text{ GeV}^{\frac{1}{2}}/\sqrt{E} \oplus 3\%$ [105], and it has 5,632 read-out channels [25].

Covering the high pseudorapidity region $3.1 < |\eta| < 4.9$ is the forward calorimeter, or FCAL. It is recessed slightly relative to the face of the EMEC to reduce neutron reflection. While it also uses liquid argon, it has a distinct geometry designed to help deal with the extremely intense radiation at high $|\eta|$ (Figure 6.8). Tubes of liquid argon with a central copper electrode are embedded within an absorber matrix. The innermost of the three detector layers (the electromagnetic layer) uses copper as the absorber, while the outer two layers (the hadronic layers) use tungsten. The FCAL as a whole provides a depth of approximately 10 nuclear interaction lengths and accounts for 3,524 read-out channels [100]. The FCAL design energy resolution is $100\% \text{ GeV}^{\frac{1}{2}}/\sqrt{E} \oplus 10\%$ [105].




Figure 6.7: Illustration of the segmentation of the electromagnetic barrel (EMB) calorimeter [25].




Figure 6.8: Diagram of the liquid-argon tubes embedded in absorber matrix within the forward calorimeter (FCAL) [25].

6.3.2 Tile Calorimeter

The tile hadronic calorimeter consists of a central barrel section covering the region $|\eta| < 1.0$ and two extended barrel sections for coverage of the region $0.8 < |\eta| < 1.7$. These sections consist of 64 modular wedges, one of which is shown in Figure 6.9. Each wedge is serviced independently by off-detector equipment. The wedges use 3 mm-thick plastic scintillator tile as the active medium surrounded by approximately 5 mm of steel absorber. The azimuthal and radial tile lengths are between 200 mm and 400 mm and between 97 mm and 187 mm, respectively. Signals are transmitted from the scintillating tiles to photomultiplier tubes (PMTs) via wavelength-shifting fibers. The fiber mapping arranges the tile calorimeter into three logical longitudinal layers, the innermost two spanning $\Delta\eta = 0.1$ with the outermost spanning $\Delta\eta = 0.2$. The three layers are 1.5, 4.1, and 1.8 nuclear interaction lengths thick in the central barrel, and 1.5, 2.6, and 3.3 interaction lengths in the extended barrel [25]. The tile calorimeter has 4,672 read-out channels and a test beam energy resolution for isolated pions of $56\% \text{ GeV}^{\frac{1}{2}}/\sqrt{E} \oplus 5.5\%$ at $\eta = 0.35$.




Figure 6.9: Diagram of a tile calorimeter module consisting of alternating steel absorber and plastic scintillator tile [25].




Figure 6.10: Diagram of the ATLAS muon systems, ATLAS Experiment © 2012 CERN.

6.4 Muon Spectrometer

The muon spectrometer system sits outside the calorimeters and accounts for most of the ATLAS detector's volume and its imposing profile. It consists of four subsystems, two of which are intended for precision measurements and two for triggering. The magnetic field for the muon system is provided by the barrel and end-cap air-core toroid magnets (Section 6.5). The layout of these components is shown in Figure 6.10.

Cosmic ray data from 2008 and 2009 were used to measure the efficiency and performance of the muon system. Including the gap near $\eta = 0$, the integrated reconstruction efficiency is 94%, whereas outside this gap region it rises to 97% [106].

In the central region ($|\eta| < 1.1$), the relative momentum resolution is given by:

$$\frac{\sigma_{p_T}}{p_T} = \frac{(0.29 \pm 0.03) \text{ GeV}}{p_T} \oplus (0.042 \pm 0.002) \oplus (4.1 \pm 0.4) \times 10^{-4} \text{ GeV}^{-1} \times p_T \quad (6.4)$$

The term inversely proportional to p_T arises from the energy loss of the muon in the detector, the constant term is due to multiple scattering, and the term that is linear in p_T is due to the difficulty in measuring the curvature of nearly-straight (high momentum) tracks.

6.4.1 Precision Subsystems

The two precision muon trackers are the Monitored Drift Tube chambers (MDTs) used in the central region ($|\eta| < 2.7$ for the two outermost layers, $|\eta| < 2.0$ for the innermost layer), and the Cathode Strip Chambers (CSCs) used in the forward end-cap regions ($2.0 < |\eta| < 2.7$).

Monitored Drift Tubes There are 1,150 MDTs arranged in three layers or stations. Each MDT is composed of two multi-layers of drift tubes attached to either side of the chamber support structure. Each multi-layer, in turn, is composed of between three and eight layers of drift tubes depending on the size and location of the MDT. The drift tubes themselves are made of an aluminum-manganese alloy, which acts as a cathode and is 29.97 mm in diameter [107]. The anode is a 50 μm gold-plated tungsten-rhenium wire which is suspended at the center of each tube. Each tube is pressured to 3 bar with a mixture of 93% argon and 3% carbon dioxide. This mixture was selected for its aging properties and small likelihood of forming deposits within the tube. A small amount (up to 300 ppm) of water vapor is also added to improve high-voltage stability. An optical alignment system allows an MDT chamber to measure the centroid of a track segment with a resolution of about 30 μm . There are approximately 354,000 MDT read-out channels.




Figure 6.11: Schematic of cathode strip chamber read-out system [25].

Cathode Strip Chambers The CSCs are used in the end-cap regions of the muon spectrometer due to the higher charged-particle flux and track density. They are multi-wire proportional chambers. A plane of wire anodes are arranged radially from the beam line (the central wire is radial, the other wires in a given CSC are arranged parallel to it). To either side of the wire plane are cathode strip planes separated from the anode plane by 2.5 mm. One set of cathode strips is arranged in the direction perpendicular to the anode wires to provide a spatial resolution of $60 \mu\text{m}$ in the bending plane (r - z), while the strips on the opposite side are arranged parallel for a spatial resolution of 5 mm. The read-out is attached to the cathode strips (Figure 6.11). This geometry can robustly distinguish between multiple tracks at the the same time and is resistant to changes in temperature and pressure [25]. There are 32 CSCs in the ATLAS detector, corresponding to 31,000 read-out channels.

6.4.2 Trigger Subsystems

The muon trackers used for triggering are designed for fast response and are also used to augment the measurements of the precision detectors in ϕ . The Resistive Plate Chambers (RPCs) are situated in the very central region ($|\eta| < 1.05$), while the Thin Gap Chambers (TGCs) cover the more forward region defined by $1.05 < |\eta| < 2.4$.

Resistive Plate Chambers The RPCs are arranged in three concentric layers, or stations, in the barrel region of the ATLAS detector. The spacing is such that

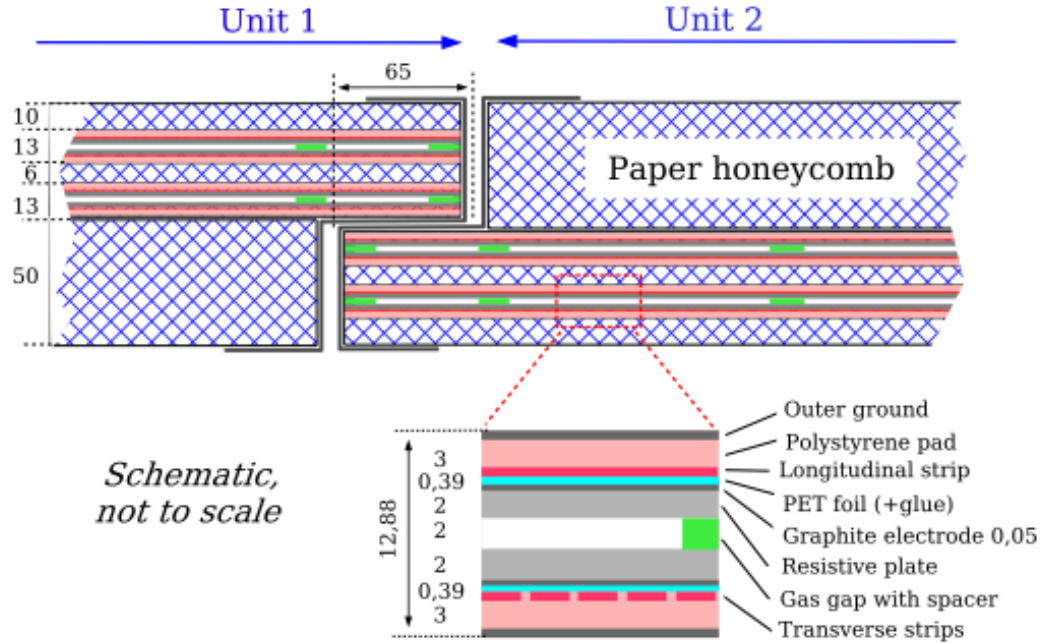


Figure 6.12: Illustration of two Resistive Plate Chambers (RPCs) [25].

the minimum muon p_T for the trigger is 6 GeV. Each RPC consists of two partially-overlapping RPC units, and each RPC unit contains two orthogonal (for two-dimensional read-out) detection layers. Two RPC units are shown in Figure 6.12. Each detection is composed of two phenolic-melaminic plastic laminate resistive plates coated with graphite on the outside and separated by a gas-filled gap. The gas is a mixture of 94.7% tetrafluoroethane, 5.0% isobutane, and 0.3% sulfur hexafluoride [25]. The gas gap and plates are biased with high voltage which creates an electric field of ≈ 4.9 kV/mm. An ionizing muon track creates a charge avalanche whose signal is read out via capacitively-coupled copper strips. RPCs are capable of operating at rates of ≈ 1 kHz/cm² with a time resolution of 1.5 ns [108].

Thin Gap Chambers Muons are triggered in the end caps via the 3,588 TGCs, which form an inner wheel with coverage of the $1.05 < |\eta| < 2.4$ region and an outer wheel covering the $1.05 < |\eta| < 1.92$ region. The inner wheel is composed of doublet TGC units (two TGCs), while the outer wheel has both doublet and triplet

units (three TGCs). The TGCs themselves are multi-wire proportional chambers, like the CSCs, but with anode plane/cathode plane separation between 1.4 and 1.8 mm. This small distance coupled with a high electric field provides a time resolution of 4 ns, well within the 25 ns bunch-crossing window, and a maximum operating rate over 20 kHz/cm². This high electric field necessitates a quenching gas mixture of 55% carbon dioxide and 45% n-pentane. The anode wires are gold-plated tungsten, oriented parallel to the beam line, while the cathode planes are fiberglass sheets with an inner surface coated with graphite. The outer surface is covered with segmented copper cladding which serves as the interface for the 318,000 TGC read-out channels [25]. The threat of fire from potential interaction between the flammable n-pentane and the high voltage motivated coating the fiberglass sheets with flame-retardant epoxy binding and enclosing each TGC in an airtight envelope continuously flushed with carbon dioxide.

6.5 Magnet System

The ATLAS magnet system provides the magnetic fields necessary to measure the momentum of tracks in the inner detector (Section 6.2) using the central solenoid, and in the muon spectrometer (Section 6.4) using the barrel and end-cap toroids. All magnets are superconducting and composed of aluminum-stabilized niobium-titanium/copper conductor cable [109]. The four systems have a stored energy of 1.6 GJ and provide a field of at least 50 mT over a 12,000 m³ volume [25]. They are shown in Figure 6.13.

6.5.1 Solenoid

The solenoid sits just outside the inner detector and shares a cryostat with the liquid-argon barrel electromagnetic calorimeter (Section 6.3.1) to reduce dead material in front of the calorimeter. The cold volumes, however, are separate due to the different cryogenic temperatures required (80 K for liquid argon vs. 4.5 K for the solenoid). The total depth of the solenoid is ≈ 0.66 radiation lengths [25]. The solenoidal




Figure 6.13: Diagram of the magnet system, including both barrel and end-cap toroids and the inner detector solenoid [25].

magnetic field of 2 T is created with 7.73 kA flowing through a single length of 30 mm by 4.25 mm superconductor cable wound around an aluminum support structure 1,154 times [109]. The solenoid can be charged and discharged in ≈ 30 minutes, and in the case of a quench the energy can be safely discharged into the cryostat cold mass [110].

6.5.2 Toroids

The barrel and end-cap air-core toroid magnets have the responsibility of generating the magnetic field within the muon system. The eight barrel toroids are enclosed in stainless-steel racetrack-shaped vessels with two coils of 57 mm by 12 mm conductor cable wound 120 times. They extend 25.3 m axially and radially between 4.7 m and 10.05 m from the beam line. A coil current of 20.5 kA corresponds to a magnetic field between 0.2 and 2.5 T within the barrel toroid volume [109, 111].

The end-cap toroids extend 5 m along the z axis (beam line) at both ends of the detector and fill the volume between 0.825 m and 5.25 m in radius. Each end-cap toroid is composed of eight pairs of coils made of 41 mm by 12 mm conductor cable. 20.5 kA of nominal current produces a magnetic field between 0.2 and 3.5 T within the end-cap toroid volume. The eight pairs of coils share a gear-shaped aluminum vacuum vessel, with a coil pair placed in each of the gear teeth. These teeth then fit in the gaps left by the barrel toroids.

6.6 Forward Detectors

The forward detectors are tasked with measuring luminosity and beam conditions, and are often located at the edges or outside the main detector volume. The information collected by these detectors is crucial for both ATLAS and the running of the collider, and in many cases it is forwarded directly to the LHC control room as well as recorded by ATLAS.

6.6.1 BCM

The Beam Conditions Monitor (BCM) provide ATLAS and the LHC operators measurements of the ATLAS collision rate as well as notification should a beam incident occur. Two BCM units are placed at ± 1.84 m in z along the beam line (and within the detector volume). Each BCM unit has four detectors, each with two back-to-back sensors. The detectors are a radial distance of 55 mm from the beam line in a “cross” pattern. Relative to the interaction point, the detectors are at $|\eta| = 4.2$. The sensors have an acceptance of 0.5 cm^2 and are made from radiation-hard polycrystalline Chemical-Vapor-Deposition (pCVD) diamond. This material can withstand 10^{15} pions/cm², and features a signal response of 1 ns which can be read out at over 100 MHz [112]. Timing coincidence between the BCM units on either side of the interaction region is used to discriminate between a collision event (particles arrive at each BCM unit almost simultaneously) or a beam incident (a several ns difference between the two BCM unit measurements).

6.6.2 MBTS

The Minimum-Bias Trigger Scintillators (MBTS) are plastic scintillating tiles connected to photomultiplier tubes (PMTs). They are located 3.56 m from the interaction point on the inner wall of the calorimeter cryostats and to either side of the interaction region. Each MBTS is divided into two regions in pseudorapidity ($2.09 < |\eta| < 2.82$ and $2.82 < |\eta| < 3.84$), and eight regions in ϕ . In the early data, this system was used to trigger on minimum-bias events by operating the MBTS as simple counters. One or two hits on each side or a coincidence of hits on both sides of the detector were required to fire the trigger. As expected, it is no longer usable due to radiation damage.

6.6.3 LUCID

LUCID (LUminosity measurement using a Cerenkov Integrating Detector) is an even more egregious example of acronym abuse than “ATLAS”. Out of all the forward detectors located inside the main detector volume, LUCID sits farthest away from the

interaction region. It is designed to measure the instantaneous ATLAS luminosity via measurement of inelastic pp scattering. The two LUCID modules are located at ± 17 m from the interaction point in z and 100 mm in r , and cover a pseudorapidity range of $5.61 < |\eta| < 5.93$. They each consist of 20 1.5 m by 15 mm polished aluminum tubes surrounding the beam pipe and contained within an aluminum vessel pressurized with perfluorobutane at 1.3 bar. At the far side of each tube is an appropriately-sized photomultiplier tube (PMT). This configuration results in Cerenkov thresholds of 10 MeV for electrons and 2.8 GeV for pions. The Cerenkov photons are reflected down the polished tubes and captured by the PMTs. Multiple particles in a single tube are distinguishable via measurement of the pulse threshold. On its own, it is able to measure the instantaneous ATLAS luminosity at $\approx 25\%$ precision.

6.6.4 ZDC

The closest forward detector which sits outside the main detector volume is the Zero-Degree Calorimeter (ZDC), located 140 m to each side of the interaction point. The ZDC modules sit in the gap created by the forking of the beam pipe. They consist of one electromagnetic module and three hadronic modules, all constructed from tungsten plates placed perpendicular to the beam. Incident particles shower in the tungsten plates and produce Cerenkov photons, which are then guided to photomultiplier tubes using quartz rods. The rods penetrate the tungsten plates and lie parallel to the beam, and the quartz strips lie perpendicular to and between the plates. The differences between the electromagnetic module (29 radiation lengths) and hadronic modules (1.14 nuclear interaction lengths for each) are essentially limited to spatial resolution and ordering relative to the incident particles. While the ZDC is primarily used to detect high- η neutrons from heavy ion collisions, it can also be used for minimum-bias triggering and measurement of beam-gas and beam-halo effects.

6.6.5 ALFA

The Absolute Luminosity for ATLAS detector (ALFA) is the final and most distant of the ATLAS forward detectors. The ALFA modules measure luminosity via the

relationship given by the optical theorem between forward elastic scattering amplitude and total cross section. As a result, they must be positioned very close to the beam (1 mm) and very far from the interaction region (four ALFA modules are placed at ± 240 m and ± 244 m from the interaction point). The proximity requirement makes ALFA the only mobile ATLAS detector (the modules can be retracted when not in use). The ALFA modules are “Roman pots”, a term referring to the fact that the detectors penetrate the beam vacuum vessel while contained within their own secondary vacuum vessel, or pot. This configuration helps prevent contamination of the beam vacuum. The detector inside each pot is a scintillating fiber tracker with aluminized 0.5 mm fibers arranged in 10 double-sided, cross-hatched planes. The spatial resolution is about $36 \mu\text{m}$ when multiple scattering is taken into account [25].

6.7 Trigger and Data Acquisition

The ATLAS trigger and data acquisition system (TDAQ) [89] is used to identify and read out interesting events, as well as to configure the detector. ATLAS uses a three-level trigger system. The Level 1 (L1) trigger can accept events at 75 kHz, upgradable to 100 kHz, and the Level 2 (L2) trigger accepts events at a maximum rate of 3.5 kHz. The Event Filter (EF) can accept events at 200 Hz. The L2 and EF together form the High-Level Trigger (HLT) which is software. The L1 is a hardware-based trigger configured via firmware.

6.7.1 Level 1

Ignoring the now-inoperative MBTS and other specialized triggers, all Level 1 triggers arise in either the calorimeter or muon subsystems (L1Calo and L1Muon), and have a target accept or reject window of $2.5 \mu\text{s}$. The L1Muon trigger signals are built from at least three hits in the RPC or TGC stations with up to 6 p_T thresholds allowed in firmware. The L1Calo triggers have significantly more variety. They are built from $\approx 7,000$ 0.1 by 0.1 (in η and ϕ) calorimeter trigger towers. At large η , the trigger tower size is also larger. These towers are used to trigger on electrons, photons, taus, jets,

and compound event-level variables such as missing transverse energy (\cancel{E}_T) and the scalar sum of the transverse energy ($\sum E_T$). Isolation requirements are implemented by requiring angular separation of trigger towers. They are used in the photon, electron, and tau triggers.

L1Calo and L1Muon decisions are passed to the Central Trigger Processor (CTP), which handles the final Level 1 trigger decision. During this time, the detector data is stored in on-detector buffers. The CTP has a trigger menu with up to 256 distinct signatures, most of which combine information from multiple calorimeter towers and/or Level 1 muons. The gluino-mediated analysis discussed in Chapter 11, for instance, relies on a Level 1 trigger requiring both large \cancel{E}_T and a high energy jet.

With an affirmative CTP decision, the detector buffers are read out into the data acquisition (DAQ) system's Read-Out System (ROS) using 1,574 Read-Out Links (ROLs), each of which is connected to its own Read-Out Buffer (ROB). The positions of the trigger objects of interest in the detector are fed to both the DAQ system and the L2 trigger system as Regions Of Interest (ROIs).

It should be noted that the integration time for both the calorimeters and the muon system is longer than the nominal 25 μs latency (design) between bunch crossings. This leads to what is known as "out-of-time" pileup, defined as event contamination from collisions other than the one that fired the trigger.

6.7.2 Level 2

The Level 2 trigger is the first stage of the software HLT, and the first trigger stage with access to inner detector information (hits). Much more precise calorimeter information than available at Level 1, such as the energy deposited in individual cells, is also available. The ROIs from the L1 trigger are used to request relevant detector data from the DAQ system. This data can then be used to construct objects such as tracks and showers to be used in L2 decisions. The L2 can do this partial event reconstruction at the 75 kHz L1-accept rate via massive parallel processing (each event takes about 40 ms in CPU time to reconstruct). However, the tight constraints imposed by this rate restrict the L2 algorithms' focus to speed and consistency rather

than efficiency and resolution. The L2-accept rate is 3.5 kHz. Events failing the L2 trigger are removed from the ROS.

6.7.3 Event Filter

If the event passes the L2-accept, the detector data is passed from the ROS to the DAQ system's Sub-Farm Inputs (SFIs), which construct a full event object from detector inputs (no ROIs) and pass it along to the Event Filter network. This network applies standard ATLAS reconstruction tools and algorithms to make a final trigger decision. The latency of the event filter is around 4s because the event reconstruction is not parallelized as the ROI reconstruction is at L2. However, the Event Filter is able to trigger events at 200 Hz by using multiple SFIs and event filter network nodes to process events simultaneously. Affirmative EF decisions transfer the full event and decision to the Sub-Farm Outputs (SFOs) which subsequently send the data to tape at a maximum rate of 400 Hz. Each event is ≈ 1.6 MB in size [25].

Chapter 7

Object Reconstruction and Datasets

Before the raw data from the detector can be analyzed it must be converted into a collection of physics objects in a process called offline reconstruction (online reconstruction refers to the processing done by the trigger, none of which is retained past the final trigger decision). The same reconstruction algorithms are applied to both real data and simulated data (Monte Carlo) though the calibrations are often different in order to correct for known object-level discrepancies between the two. Event-level differences are often corrected via reweighting of Monte Carlo.

7.1 Data

The dataset used in this document is the full 2011 ATLAS dataset introduced in Section 5.3. Although the total data recorded by ATLAS during this period is 5.25 fb^{-1} , not all of the data is free of defects. ATLAS data is divided into runs, which in turn are divided into two-minute sections called luminosity blocks. Each luminosity block is capable of having its own set of data-quality flags. Unstable beams (where much of the detector remains off to protect it from damage), noisy channels, and other detector issues are logged and the affected luminosity blocks flagged semi-automatically.

Data Period	Integrated Luminosity (pb^{-1})
B	12
D	167
E	49
F	132
G	508
H	259
I	337
J	226
K	590
L	1,405
M	1,027
Total	4,713

Table 7.1: Integrated luminosity passing the Good Runs List as a function of data-taking period in 2011.

More subtle problems can be flagged offline when the monitoring histograms are examined by data-quality shifters. Depending on the severity of the issue [113, 114], the affected luminosity blocks are removed from the list of good luminosity blocks, known as the Good Run List (GRL). As the analyses of Chapters 10 and 11 use essentially the entire detector, the corresponding GRL is relatively strict and reduces the total dataset to 4.71 fb^{-1} . A summary of the usable integrated luminosity as a function of data-taking period (a group of runs collected under similar beam conditions) is shown in Table 7.1.

7.2 Monte Carlo

Monte Carlo simulation is used extensively in both the optimization and background estimation. All detector simulation was performed using the ATLAS detector simulation [115] based on GEANT4 [116], though the generators differ by physics process.

The simplified model signal samples were generated using HERWIG++ [117]. As HERWIG++ is a $2 \rightarrow 2$ generator, MADGRAPH [118] samples with zero, one, and two extra partons in the matrix element were also generated for selected signal points to verify HERWIG++ results dependent on hard radiation (Section 8.1.3).

For the background samples, the generator depends on the process and analysis. ALPGEN [119] is used to produce large ($\mathcal{O}(10^8)$) samples of $t\bar{t}$, W/Z and diboson (WW, WZ, ZZ) events. The $t\bar{t}$ and $Z \rightarrow ll$ samples are produced with up to five

additional partons at the matrix element level, while the $W \rightarrow l\nu$ and $Z \rightarrow \nu\nu$ are generated with up to six and the diboson samples with up to three. Dedicated W , Z , and $t\bar{t}$ samples with associated production of heavy flavor jets are also generated with up to three additional partons in order to boost background statistics in the signal regions. MC@NLO [120, 121, 122] is used to produce single top and additional $t\bar{t}$ samples, while SHERPA [123] is used for additional diboson samples. Both ALPGEN and MC@NLO are interfaced to HERWIG [124] for fragmentation and hadronization and to JIMMY [125] for the underlying event. MADGRAPH interfaced to PYTHIA [126] is also used to generate the $t\bar{t} + W$, $t\bar{t} + Z$, and $t\bar{t} + WW$ processes. QCD di-jet samples are generated using PYTHIA and ALPGEN but used minimally due to low statistics in high- \cancel{E}_T regions.

The Monte Carlo samples are produced using parameters tuned as described in References [127, 128]. For MC@NLO, the next-to-leading-order (NLO) PDF set CT10 [129] is used, while the leading-order (LO) generators use CTEQ6L1 [130]. All Standard Model background processes are normalized to the results of higher-order calculations when available [3], while signal sample cross sections [78, 79, 80] are computed using PROSPINO [81] at next-to-leading order + next-to-leading log (NLO+NLL) (Figure 4.8).

The MC@NLO dataset is used as the primary $t\bar{t}$ dataset for the direct sbottom analysis. The gluino-mediated analysis uses ALPGEN instead, as it provides a better description of final states with many jets. The ALPGEN diboson samples, used in the gluino-mediated analysis (though negligible), are superseded by the SHERPA samples for the direct sbottom analysis as SHERPA provides a better description for diboson final states involving jets. The Monte Carlo samples used for the background estimation, like the data (Section 7.1), were reconstructed with Release 17 of the ATLAS ATHENA framework. The search optimization, however, was conducted with samples reconstructed with ATHENA Release 16. While dictated by sample availability at the time rather than a choice, this has the advantage of making the optimization and background estimate statistically independent.

7.3 Tracks and Vertices

Inner detector track reconstruction ($p_T > 0.4$ GeV, $|\eta| < 2.5$) begins with a pre-processing stage, where silicon hits are converted to space points and TRT data is converted into calibrated drift circles. The default tracking uses an “inside-out” algorithm where the tracks are seeded using space points from the pixel detector and the first layer in the SCT. The seeds are extended through the SCT to form track candidates and outlier space points with bad fit quality are removed. The surviving track candidates are extrapolated into the TRT and matched with draft-circle information, forming extended tracks. These extended tracks are then refit (using a combinatorial Kalman filter [131]) with information from all inner detectors. Bad space points are again labeled as outliers.

The tracks are then used as inputs to electron (Section 7.4) and muon (Section 7.5) reconstruction, flavor tagging of jets (Section 7.8), and reconstruction of primary vertices. Reconstruction of primary vertices [132] proceeds via a seeding step followed by a fitting step. In the seeding step, tracks with a transverse impact parameter within 4 mm of the beam line are selected. Additionally, these tracks must have at least 4 SCT hits, at least 6 hits in the combined pixel and SCT detectors, and have well-measured transverse and longitudinal impact parameters ($\sigma(d_0) < 5$ mm and $\sigma(z_0) < 10$ mm, respectively) [133].

An iterative vertex finding algorithm [134] is then applied to the selected tracks. The first vertex is seeded by finding a global maximum in the track z coordinates computed at closest approach to the centroid of the beam spot [135]. The fitter uses the seed track and its surrounding tracks in a succession of χ^2 -based fits, with outlying tracks becoming progressively down-weighted and eventually discarded. Tracks incompatible with the vertex by more than 7 standard deviations are used to seed a new vertex, with the process continuing until no unassociated tracks remain or no additional vertices can be found. Each vertex contains a minimum of two tracks. The primary vertex with largest scalar sum of constituent track p_T^2 is designated the hard-scatter vertex of the event.

7.4 Electrons

Leptons (electrons and muons only) are used for two purposes in hadronic analyses: control regions and lepton vetoes. Electron reconstruction begins when electron clusters are created in the middle layer of the liquid-argon electromagnetic calorimeter (Section 6.3.1) using a 3 by 5 cell (in η and ϕ , respectively) sliding window. Each cell is about 0.025 by 0.025 in $\Delta\eta$ and $\Delta\phi$. The wider dimension of the window in ϕ is to account for bremsstrahlung (deceleration radiation) in the bending (transverse) plane of the inner detector. Bremsstrahlung is significantly more important for electrons due to their small mass than it is to heavier charged particles such as muons or pions. The energy threshold for the cluster is 2.5 GeV.

The candidate cluster is then matched to an inner detector track using the track nearest to the energy-weighted cluster centroid in ΔR (Section 6.1). A subtle point is that in 2011 these tracks are reconstructed according to the low-bremsstrahlung (pion or muon) hypothesis, resulting in some small “irreducible” electron inefficiency. This is a problem that becomes larger with increased luminosity. For analyses based on 2012 data it is addressed via a track refit designed specifically for electrons [136].

After track matching, the cluster is recalculated into a 3 by 7 or 5 by 5 cell cluster in the barrel and end caps, respectively. The electron definitions are applied to these recalculated calorimeter clusters and their associated tracks. The *baseline* electron definition, used for the lepton veto, is based on the 2011 “medium++” electron selections, while the *signal* electron definition is based on the “tight++” definition and is used for leptonic control regions [137, 138]. The medium++ selections require one hit in the pixel b -layer for $|\eta| < 2.01$ and two hits in the entire pixel detector for $|\eta| > 2.01$, seven hits or more in the pixel and SCT combined, $|\Delta\eta| < 0.005$ between the track and the cluster, a transverse impact parameter (d_0) less than 5 mm, and apply cuts on shower shape and hadronic leakage. Outliers are included in the hit counts for the purposes of the cuts. These selections are about 85% efficient for electrons originating in Z boson decays.

The tight++ selections are a subset of the medium++ selections with equal or tighter cuts on shower shapes, one b -layer hit for all η , $d_0 < 1$ mm, $|\Delta\phi| < 0.02$, a

difference between the number of expected and observed TRT hits less than 16, and additional requirements on the energy/momentum ratio of the electron track and the fraction of high-threshold TRT hits. The tight++ selections are about 78% efficient for electrons from Z boson decays.

The electron candidates are calibrated to match the observed ATLAS electromagnetic energy scale base on measurements of Z boson and J/Ψ meson decays to electron pairs [138]. Monte Carlo electron energies are further smeared to match the electron energy resolution observed in data. The transverse energy of the cluster is required to be greater than 10 GeV in all regions and greater than 15 GeV in the calorimeter crack region of $1.37 < |\eta| < 1.52$. Cluster η as opposed to track η is used for this and other selections, though the difference is negligible thanks to the tight $|\Delta\eta|$ cut applied as part of the electron definition. All electrons are required to have $|\eta| < 2.47$, and electrons in regions of the calorimeter with dead optical transmitters are removed.

These final selections coupled with the medium++ definition and a p_T cut of 20 GeV complete the selection of the baseline electrons used in the lepton vetoes. Signal electrons are also required to satisfy the tight++ definition and a minimum p_T of 25 GeV. Furthermore, signal electrons must be isolated. This is quantified by the requirement that the scalar sum of all track p_T within a cone of size 0.2 in ΔR divided by the electron transverse energy must be less than 10%.

7.5 Muons

Two types of muons are considered: so-called combined muons which consist of a muon spectrometer track matched to an inner detector track (Section 7.3) and segment-tagged muons which are constructed from an inner detector track matched to a muon spectrometer track segment or hit [89]. Muon spectrometer track segments are constructed in each of the three muon stations (Section 6.4) and then linked together to form muon spectrometer tracks. These tracks and segments are then extrapolated through the calorimeter to the inner detector, with both multiple scattering and energy loss in the calorimeter taken into account. The algorithm which

performs these tasks is called Muonboy [139].

The STACO (STAtistical COmbination) algorithm [89] is used to combine muon spectrometer tracks and segments with inner detector tracks. The χ^2 of the track match is defined as the difference between the inner detector and muon spectrometer track four-vectors weighted by their combined covariance matrix. After a successful match, the combined four-vector is obtained using a covariance-weighted average. The inner detector track typically dominates up to about 80 GeV in p_T in the barrel. Muons with p_T over 100 GeV are dominated by the muon spectrometer track.

Several additional selections are applied once the muon candidates are identified. The transverse momenta of muons in Monte Carlo simulation are smeared to match the observed p_T resolution in data [140]. The minimum p_T requirement (post-smearing, in the case of Monte Carlo) is 10 GeV. Along with $|\eta| < 2.4$, the muon must have at least one pixel hit, and six SCT hits. Outliers are not considered hits for the silicon trackers. For the TRT, the sum of TRT hits and outliers must be greater than 5 for the region $|\eta| < 1.9$, and the number of outliers must be less than 90% of this sum. For the region $|\eta| > 1.9$, this latter requirement is only applied when the sum is greater than 5—the sum itself has no requirement. Finally, segment-tagged muons (those without a multi-station muon spectrometer track) are required to contain one hit in the end cap ($|\eta| > 1.05$) of the muon spectrometer but no TGC hits, at least two segments, *or* at least three TGC end-cap hits in the associated segments.

This completes the selection of *baseline* muons, which are used in the lepton vetoes. *Signal* muons, which are used to select events for leptonic control regions, must have p_T greater than 20 GeV. Signal muon isolation is enforced by requiring that the scalar sum of the track energy in a cone of 0.2 in ΔR around the muon be less than 1.8 GeV.

7.6 Topological Clusters

Topological clusters (topoclusters) are the primary input for jet reconstruction (Section 7.7) and also an important component of the missing transverse energy (Section

7.9). Clusters are constructed using a “4-2-0” clustering scheme based on the root-mean-squared (RMS) noise width (σ_{noise}) measured in each cell when triggered on random events [141, 142]. Particularly “hot” or noisy cells are excluded entirely, resulting a loss of $\approx 0.1\%$ of all calorimeter cells.

The clustering begins with seed selection. Seeds are defined as those cells with $|E_{cell}|$ greater than or equal to 4 times their noise width. All surrounding cells with $|E_{cell}| \geq 2\sigma$ are then iteratively added to the seed to grow the cluster. Finally, all nearest-neighbor cells are added to the cluster, as this was shown to increase single pion energy resolution. Cluster with negative total energy (arising as a result of the liquid-argon pulse shape) are discarded. Approximately 2500 cells per collision are built into clusters. The use of thresholds based on each cell’s particular noise characteristics serves to suppress both electronic noise and noise from out-of-time and multiple interactions (pileup).

7.7 Jets

The jet objects used in this document are reconstructed from topocluster inputs (Section 7.6) using an anti- k_t clustering algorithm [143] implemented by FastJet [144] with characteristic radius R of 0.4. The anti- k_t algorithm is part of a family of clustering algorithms which define the two distance parameters $d_{i,j}$ and $d_{i,beam}$:

$$d_{i,j} = \min(k_{t,i}^{2p}, k_{t,j}^{2p}) \frac{\Delta R_{i,j}^2}{R^2} \quad (7.1)$$

$$d_{i,beam} = k_{t,i}^{2p} \quad (7.2)$$

ΔR is defined in Section 6.1 and $k_{t,i}$ is the transverse momentum (p_T) of object i . The clustering procedure iterates over all inputs i starting with the object with highest p_T . For all objects $j \neq i$, $d_{i,j}$ and $d_{i,beam}$ are calculated, and the smallest $d_{i,j}$ is selected out of all j . If the smallest $d_{i,j} < d_{i,beam}$, the corresponding object is merged with i . Otherwise, object i is removed from clustering list and added to the list of final state objects (jets). The clustering procedure is begun anew until no objects

remain in the clustering list.

Clustering algorithms of this type are *infrared* and *collinear safe*. Infrared safety refers to insensitivity of jet-finding to the emission of soft radiation. Collinear safety, meanwhile, is insensitivity to a parton splitting in two. Both of these situations and the case where neither occurs lead to the same final set of jets when using an infrared- and collinear-safe algorithm.

While all clustering algorithms in this family satisfy these safety requirements, anti- k_t ($p = -1$) has a distinct experimental advantage in that it clusters hard objects first and thus produces very regular, cone-shaped jets, as shown in Figure 7.1. On the other hand, the shape of k_t jets ($p = 1$), for example, can be very irregular and dependent on soft radiation. Regular cone-like jets are considerably easier to calibrate and correct for pileup.

The topological clusters used as inputs are scaled to the electromagnetic energy scale used for electrons and photons that was obtained from test beam data [145, 101]. For hadronic objects, the energy of these clusters is systematically underestimated due to the non-compensation of the calorimeters, dead material, and out-of-cone energy loss. This is corrected with a numerical inversion calibration based on the reconstructed jet response in Monte Carlo simulation [146]. This calibration was tested in data using jet-photon balance [147], a di- and multi-jet balance [148], and jet-track association [149] (a related topic is discussed in Appendix A). A further energy offset is applied on top of the calibration to correct for the average pileup contribution as a function of the average number of interactions for the data period $\langle\mu\rangle$, the number of primary vertices in the event, and the jet η [150]. The final kinematic jet selections for signal jets used in the analyses in Chapters 10 and 11 are $p_T > 20$ GeV and $|\eta| < 2.8$.

7.8 Heavy Flavor Tagging

The most important experimental characteristics of b -hadrons (hadrons containing a bottom quark) are their relatively long (≈ 1.5 ps) lifetimes and their large masses compared with the masses of their decay products. Coupled with substantial boosts




Figure 7.1: Example of parton-level jet reconstruction with the anti- k_t algorithm [143].

from the hard scatter or the decay of heavy particles such as the top quark, the long lifetimes mean that b -hadron decays are often displaced macroscopic distances (few mm) from the primary vertex. The large mass differences result in energetic decay products with significant angular separation even at reasonably large boosts. To a lesser extent, both of these characteristics apply to tau leptons and to mesons containing charm quarks as well. These secondary vertices and their displacements lead to the concept of b -tagging, where one attempts to either reconstruct the secondary vertex explicitly or use related variables to “tag” a jet as originating from a b -hadron. As the analyses presented in this document are searches for new particles which decay exclusively to final states involving multiple b quarks, b -tagging is a very important part of the overall strategy.

The primary b -tagging algorithm used in ATLAS is called Multi-Variate 1 (MV1). As its name implies, it is a combination (via neural network) of three separate taggers known as IP3D, SV1, and JetFitter [151, 152]. Each employ a different tagging strategy. The performance of MV1 compared to JetFitter and a combination of IP3D and SV1 is shown in Figure 7.2. Operating points are defined as cuts on the MV1 discriminant which correspond to specific efficiencies in Monte Carlo $t\bar{t}$ samples. Multiple operating points are used depending on the analysis. For instance, only the 60% efficient MV1 operating point is used in the sbottom pair production analysis (Chapter 10), while the gluino-mediated sbottom/stop analysis (Chapter 11) uses the 60%, 70% and 75% operating points depending on the signal region. The best operating point is chosen by the optimization procedure (Chapter 8).

The b -tagging algorithms use custom track selections designed to reject fake tracks, tracks from much longer-lived particles (K_s mesons, Λ baryons, etc.), and tracks from material interactions (photon conversions and hadronic interactions). These selections vary depending on the tagger. For the impact parameter taggers, at least seven silicon hits (pixel or SCT micro-strip) are required. Two of these hits must be in the pixel detector with one in the b -layer. The vertex taggers require only one pixel hit and do not require it to be in the b -layer. However, they do require that all tracks have at most one hit that is shared between multiple tracks. The minimum track p_T is 1 GeV for the impact parameter taggers, and 0.4 GeV for the vertex taggers. Track pairs




Figure 7.2: MV1 combined tagger performance compared to that of JetFitter and a combination of IP3D and SV1 [153].

whose two-track vertex (if it exists) has a mass consistent with the K_s , Λ , or photon, or which appear within material (beam pipe or inner pixel layers) are rejected. For impact parameter taggers, the minimum transverse impact parameter d_0 relative to the hard-scatter vertex must be less than 1 mm, and $|z_0| \sin \theta$ is required to be less than 1.5 mm (where z_0 is the longitudinal impact parameter of the track). Vertex taggers only require $d_0 < 3.5$ mm. The vertex tracking requirements are looser to allow for better reconstruction (and thus better rejection) of longer-lived contaminants such as K_s and Λ .

Tracks are associated to calorimeter jets with a p_T -dependent ΔR (Section 6.1) cone around the calorimeter jet axis. The cone size decreases with increasing p_T to reflect the higher collimation of high- p_T jets. For b -tagging track jets (not done for the analyses described in this document, though track jets reconstruction is described in Appendix A) the jet constituents are used directly.

IP3D Impact parameter taggers look for tracks associated with the jet which are inconsistent with the hard-scatter vertex. The impact parameter tagger which serves as an input to MV1 is called IP3D, so named because it uses both the transverse and longitudinal signed impact parameter significances. To construct the significance, the impact parameter relative to the hard-scatter vertex is divided by its error estimate. This has the effect of weighting well-measured tracks higher than more poorly-measured tracks with the same impact parameter. The “sign” of the impact parameter significance is positive if the track crosses the jet axis in front of the hard-scatter vertex, and negative if it crosses behind. Tracks from secondary decays will have predominantly positive sign, while resolution effects will populate the positive and negative sides of the distribution equally. Distributions of the signed transverse and longitudinal impact parameter significances in both data and Monte Carlo simulation are shown in Figure 7.3.

The IP3D tagger works by constructing a likelihood ratio between the b and light jet hypotheses using two-dimensional distributions of the two impact parameter significances taken from Monte Carlo simulation. Two-dimensional distributions are used to take advantage of correlations between the two variables.

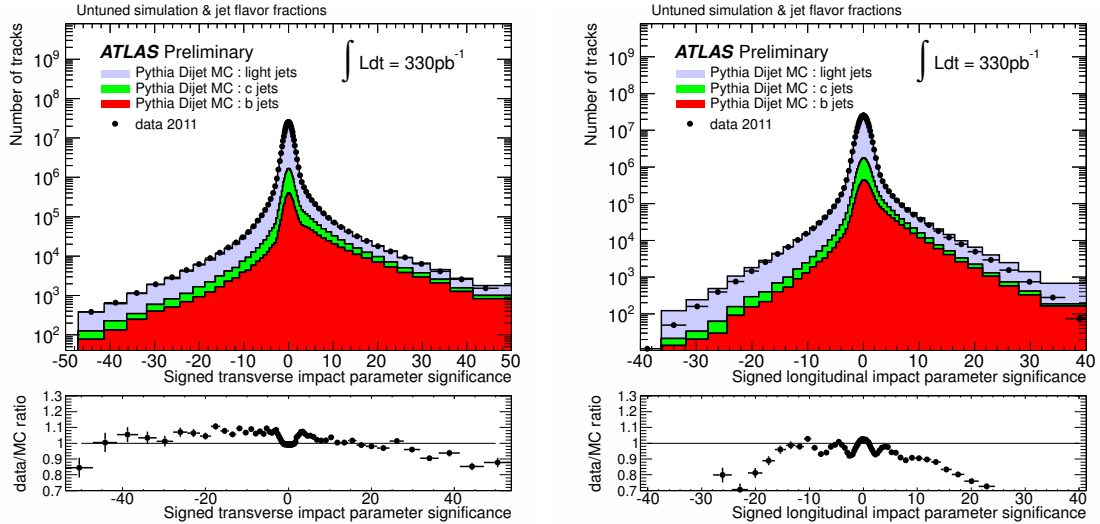


Figure 7.3: Distributions of the signed transverse (left) and longitudinal (right) impact parameter significances in both data and Monte Carlo simulation [151].

SV1 The SV1 tagger constructs all two-track vertices as well as a single inclusive vertex using all available tracks. A iterative procedure is then used to remove the most inconsistent track until the inclusive vertex χ^2 drops below a threshold. The inclusive vertex merges the b -hadron decay vertex with that of the subsequent charm decay. Discrimination between the light and b hypotheses is achieved via a likelihood ratio based on three variables: the invariant mass of the vertex, the ratio between the sum track p_T in the vertex to the total sum track p_T associated with the jet, and the number of good two-track vertices. The first two variables have significant correlations and thus a two-dimensional distribution is used. Distributions of these quantities are shown in Figure 7.4.

JetFitter The most advanced stand-alone tagger is the JetFitter tagger, which uses a Kalman filter to find a common line on which the hard-scatter, b -hadron, and c -hadron vertices lie [152]. The line serves as a constraint for vertices, and provides an approximate flight path for the b -hadron. The b - and c -hadron vertices are not necessarily merged with this procedure even though either or both may have only one associated track. Discrimination is provided by a likelihood ratio constructed from

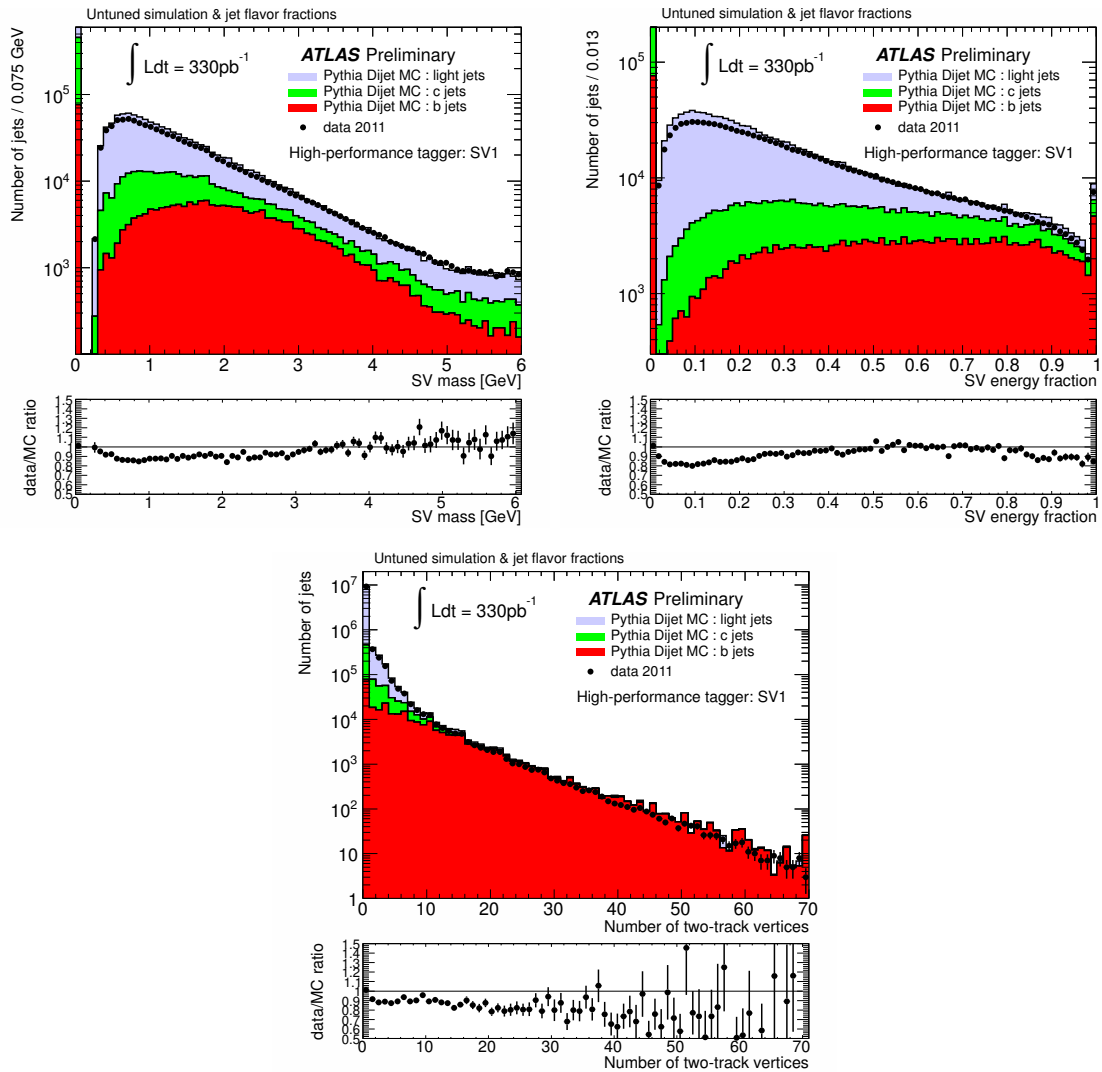


Figure 7.4: Distributions of the secondary vertex invariant mass (top left), secondary vertex p_T fraction (top right), and number of two track vertices (bottom) for the SV1 tagger in both data and Monte Carlo simulation [151].

variables similar to those used for SV1 (decay chain mass, decay chain energy fraction, number of one and multi-track vertices) as well as the flight length significances of the reconstructed secondary vertices. Example distributions of some of these quantities are shown in Figure 7.5.

7.9 Missing Transverse Energy

Missing transverse energy (\cancel{E}_T) is a complicated, compound object. It is, in fact, misnamed—it should be called missing transverse momentum. It is constructed via vector addition of all visible transverse momentum in the event followed by a sign flip. In principle, all visible objects going into the calculation should have their own calibration applied. In practice, a separate photon term is unnecessary for the analyses under consideration, and the calibration for some objects such as τ leptons is not fully validated. As a result, the ATLAS supersymmetry analyses with 2011 data use a simplified \cancel{E}_T definition:

$$\vec{\cancel{E}}_T = - \sum_{electrons} \vec{E}_T^e - \sum_{muons} \vec{p}_T^\mu - \sum_{jets} \vec{E}_T^{jet} - \sum_{clusters} \vec{E}_T^{clus} \quad (7.3)$$

The object definitions used to compute these terms are slightly different than those used for stand-alone physics objects. The electrons which enter into the calculation are similar to the baseline electrons in that they use the medium++ definition, but have an expanded η range and a lower p_T cut of 10 GeV rather than 20 GeV. The jets retain their 20 GeV p_T cut but are expanded to include all jets within $|\eta| < 4.5$. Baseline muons are used directly, and the cluster term includes all topological clusters not used in other objects. The clusters are calibrated at the electromagnetic scale.

7.10 Overlap Removal

Being calorimeter objects, all sufficiently-energetic electrons are also reconstructed as jets. Many jets, particularly those resulting from heavy flavor decays, can include energetic non-prompt electrons and muons. In both cases, overlap removal is required.

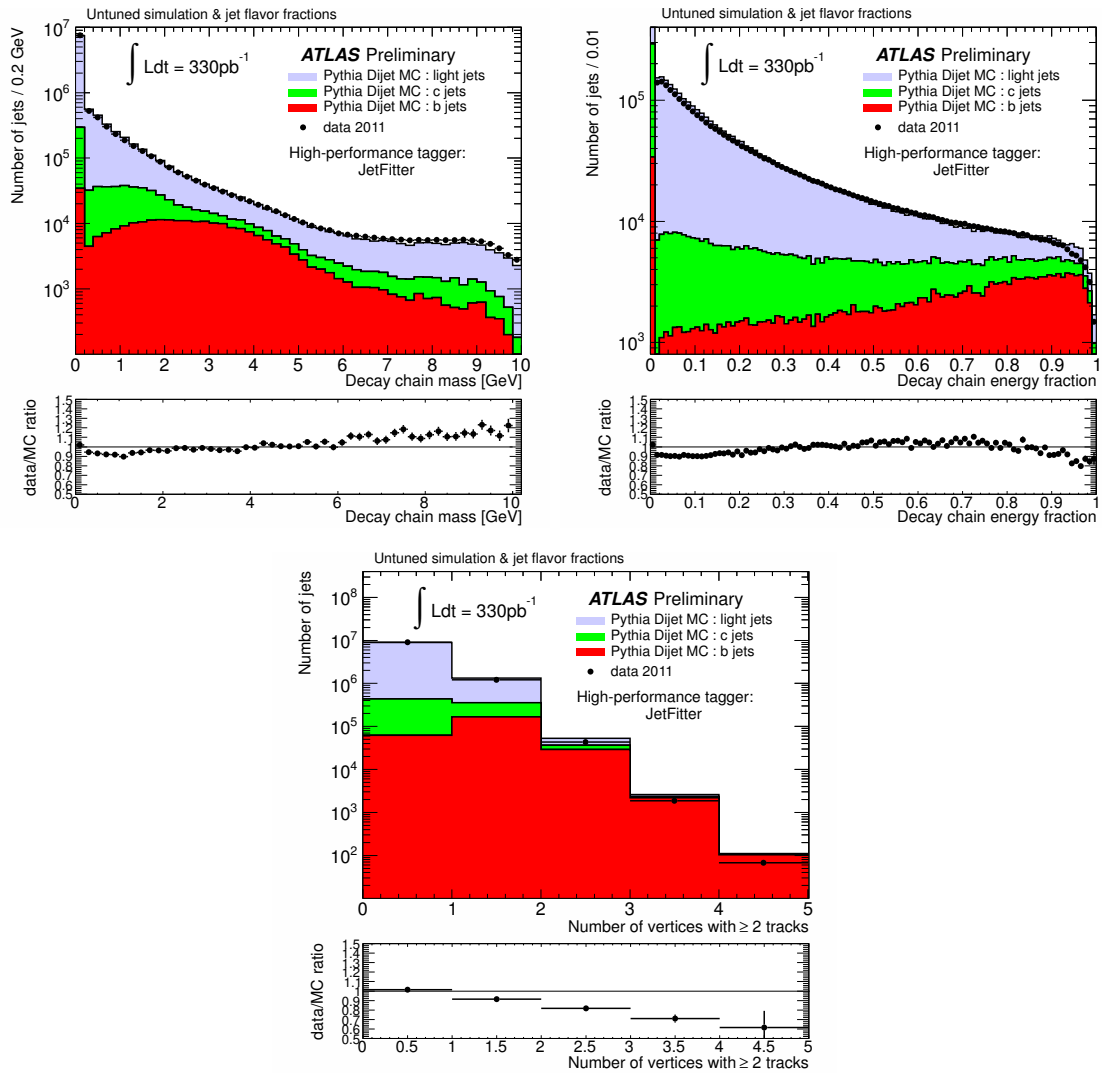


Figure 7.5: Distributions of the decay chain invariant mass (top left), decay chain p_T fraction (top right), and number of two track vertices (bottom) for the JetFitter tagger in both data and Monte Carlo simulation [151].

The first step is to check electron-jet overlap. If electron and jet four-vectors have $\Delta R < 0.2$ the jet is removed. On the other hand, if $0.2 < \Delta R < 0.4$ the electron is removed, as this implies the electron is non-isolated and thus is likely a secondary decay product. The muon is removed if $\Delta R < 0.4$ for muon-jet overlap, as muons very rarely involve significant calorimeter energy deposits and thus any jet overlap indicates a secondary muon. The overlap removal is largely redundant for signal leptons as their definitions mandate isolation.

A second type of overlap removal is Monte Carlo-specific and is performed on the ALPGEN inclusive W , Z , and $t\bar{t}$ + jets samples and the corresponding samples with dedicated heavy flavor jets. As the inclusive samples include a heavy flavor jet component, simply using both samples would result in double counting. This overlap is removed by vetoing all events in the inclusive samples which have $b\bar{b}$ or $c\bar{c}$ angles in the detector frame greater than 0.4, and doing the same in the dedicated samples with the cut reversed.

7.11 Triggers

The trigger menu evolved throughout 2011 as it was adapted to higher luminosity and pileup. This evolution usually involved higher momentum and energy thresholds. Isolation requirements were also added to the trigger in the leptonic case. For the sbottom pair production search, all signal regions used a stand-alone \cancel{E}_T trigger, while the gluino-mediated stop/sbottom search used jet + \cancel{E}_T triggers. Leptonic control regions used single lepton triggers in the case of electrons, and muon + jet triggers in the case of muons. A summary of the triggers as a function of data period is shown in Table 7.2.

7.12 Event Cleaning

Not all defects in data are removed with the Good Runs List, nor are all defects so severe that entire luminosity blocks must be vetoed. The following event-level vetoes are applied either to data only or both data and Monte Carlo. The data-only vetoes

Period	Triggers			
	\cancel{E}_T	Jet + \cancel{E}_T	Electron	Muon
B	EF_xe60_verytight_noMu	EF_j75_a4_EFFS_xe45_loose_noMu	EF_e20_medium	EF_mu18
D-I		EF_j75_a4tc_EFFS_xe45_loose_noMu		
J			EF_e22_medium	
K			EF_e22vh_medium1 or EF_e45_medium1	EF_mu18_L1J10
L-M		EF_j75_a4tc_EFFS_xe55_noMu		

Table 7.2: Triggers used in the analyses as a function of 2011 data-taking period.

have a negligible effect on the overall efficiency, and thus the integrated luminosity is not recalculated to take them into account.

Jet Cleaning Calorimeter effects can lead to spurious large energy deposits which can fake jets. Events are rejected if they contain a reconstructed jet with $p_T > 20$ GeV at any η which is flagged as bad. The bad jet criteria are formulated to look for energy spikes in the hadronic end cap, large numbers of negative-energy cells, non-collision background and cosmic rays, and coherent electromagnetic noise [154]. This veto is only applied to data.

A further requirement that the charged energy fraction (ratio of the sum of associated track p_T to calorimeter energy) of the leading signal jets ($|\eta| < 2.8$) be greater than 5% is also applied for the hadronic channels. The number of leading jets the requirement applies to depends on the signal region (three for the sbottom pair signal regions, four or six for the gluino regions). This cut is applied to both data and Monte Carlo simulation.

LAr Hole Veto A failure of multiple front-end boards of the electromagnetic barrel calorimeter occurred in May of 2011. While the problem was repaired in June 2011, 0.86 fb^{-1} of the dataset is affected by a “hole” in the liquid-argon calorimeter barrel between -0.1 and 1.5 in η and -0.9 and -0.5 in ϕ . The failure affected multiple layers of the calorimeter. The central danger of this kind of failure is fake missing transverse energy. Rather than an outright veto of any events with a jet in the hole, information from the remaining functional calorimeter layers was used to estimate the amount of energy in dead cells, and this estimate was then compared in both magnitude and direction with the \cancel{E}_T in the event. The event is only rejected if a significant fraction

of the \cancel{E}_T is consistent with arising from the mis-measured jet. This veto is only applied to data.

LAr Error Veto Some liquid-argon data-quality errors, such as excessive noise, appear and disappear on a time scale smaller than a luminosity block and therefore are better treated on a per-event basis. This veto is only applied to data.

Primary Vertex Track Multiplicity Occasionally an event is triggered for which no primary vertices can be reconstructed, or the hard-scatter vertex chosen by the maximum track $\sum p_T^2$ criterion has very few tracks. As proper choice of the primary vertex is important for both heavy flavor tagging and proper object calibration, events whose hard-scatter vertex has fewer than 5 tracks (or does not exist) are rejected. This veto is applied to both data and Monte Carlo simulation.

Muon Cleaning Events are rejected which contain a baseline muon with $\frac{\sigma_{q/p}}{|q/p|} \geq 0.2$, where q/p is the charge-momentum ratio. These bad muons are examined prior to overlap removal, and this veto is applied to both data and Monte Carlo simulation.

Cosmic Muon Veto Events which contain any baseline muons with transverse impact parameter $|d_0| > 0.2$ mm or axial impact parameter $|z_0| > 1$ mm are vetoed. Both of these impact parameters are computed relative to the reconstructed hard-scatter vertex. The veto is applied to both data and Monte Carlo simulation.

7.13 Monte Carlo Reweighting

Not all differences between data and Monte Carlo are best addressed by corrections to individual physics objects. An alternative approach is to reweight simulation events to better reproduce the efficiencies or distributions observed in data.

Lepton Reconstruction and Trigger Efficiencies η - and ϕ -dependent multiplicative scale factors are used to correct the signal electron [155] and signal muon

[156, 157] identification and reconstruction efficiencies in Monte Carlo. A similar reweighting based on η and E_T in the electron case and in the muon case η , ϕ , p_T , isolation, and type (combined or segment-tagged) is applied to account for differences in single-lepton trigger efficiency. The electron trigger scale factors also depend on which particular electron trigger is being used (Table 7.2).

Flavor Tagging Efficiency Multiplicative scale factors (one per jet considered for b -tagging) are used to correct the tagging efficiency in Monte Carlo [158, 159, 160, 153, 161]. The inputs are jet η , p_T , truth flavor (b , c , τ , or light flavor/gluon), b -tagging operating point, and the jet b -tagging weight. Jets with p_T above 200 GeV use the 200 GeV scale factors, and assume the correspondingly large uncertainties.

Pileup ATLAS Monte Carlo is generated before the data is taken using a “best guess” for relevant beam conditions. A particularly assumption is the distribution of the number of interactions per bunch crossing (μ). Monte Carlo datasets are reweighted to match the μ distribution observed in the cumulative data sample.

MC@NLO MC@NLO Monte Carlo samples contain weights of -1 and +1 on an event-by-event basis to account for interference between production processes. These are applied to the relevant Monte Carlo samples.

Chapter 8

Search Optimization

Signal region optimization has two goals: maximizing the significance of each possible new physics model considered, and ensuring sensitivity to as many models as possible. Both objectives increase the likelihood of discovery. Given practical considerations (each of the hundreds of models considered cannot have its own signal region), these goals are inevitably come into some conflict with one another and a compromise must be found. This chapter describes the optimization procedure used for both the direct sbottom and gluino-mediated sbottom/stop analyses, and how broad sensitivity to the model space was achieved while also increasing overall sensitivity dramatically.

8.1 Model Kinematics

Understanding model kinematic features and their dependence on the input parameters is key even in the case of semi-automated optimization procedures. It allows both for better-motivated choices of input variables, and equips us with the physical understanding needed to ask the all-important question when it comes to results obtained from a “black box”: does this answer make sense? Such understanding can help categorize models into groups which may be able to share a single signal region without significant loss of sensitivity for any particular model in the group.

8.1.1 A Representative Example

Consider the direct sbottom simplified model:

$$pp \rightarrow \tilde{b}\tilde{b}, \quad \tilde{b} \rightarrow b + \chi_1^0 \quad (8.1)$$

At the parton level, the final state observables are:

$$\text{Missing transverse momentum } \cancel{E}_T = (\mathbf{p}_{\chi_1^0}^A + \mathbf{p}_{\chi_1^0}^B)_T \quad (8.2)$$

$$\text{Quark transverse momenta} = (\mathbf{p}_b^A)_T, (\mathbf{p}_b^B)_T \quad (8.3)$$

where the T indicates projection into the plane transverse to the beam. Both the production mechanism and the decay mode influence the final state kinematics. Assuming production-decay correlations are small, the influence of these two processes can be considered separately. In the case of two heavy sbottoms produced at threshold, the intermediate state before sbottom decay is well approximated by two sbottoms at rest with respect to the transverse plane. The final state kinematics are then completely determined by the kinematics of the sbottom decay. In the rest frame of the sbottom, the kinematics are completely determined by energy and momentum conservation:

$$\mathbf{p}_b = -\mathbf{p}_{\chi_1^0} \quad (8.4)$$

$$m_{\tilde{b}} = E_b + E_{\chi_1^0} \quad (8.5)$$

Some simple algebra and rearrangement yields p^2 , the squared magnitude of the momenta of the daughter particles, in terms of the masses:

$$p^2 = \frac{1}{4} \left[m_b^2 - 2(m_b^2 + m_{\chi_1^0}^2) + \frac{(m_b^2 - m_{\chi_1^0}^2)^2}{m_b^2} \right] \quad (8.6)$$

As $m_b \ll m_{\chi_1^0}$, this relation can be further simplified to:

$$p = \frac{m_{\tilde{b}}^2 - m_{\chi_1^0}^2}{2m_{\tilde{b}}} \quad (8.7)$$

At this point, it is convenient to change variables from $(m_{\tilde{b}}, m_{\chi_1^0})$ to $(m_{\tilde{b}}, \Delta m)$, with Δm defined as:

$$\Delta m = m_{\tilde{b}} - m_{\chi_1^0} \quad (8.8)$$

Then Equation 8.7 becomes:

$$p = \Delta m - \frac{(\Delta m)^2}{2m_{\tilde{b}}} \quad (8.9)$$

Therefore, to first order, the final state kinematics only depend on the difference between the sbottom and neutralino masses which is a measure of the phase space available to the decay. This conclusion is robust as long as Δm is small relative to $m_{\tilde{b}}$ (or alternatively, as long as the neutralino is non-relativistic), however, even in the large Δm limit the magnitude of the quadratic term is never larger than 50% of the leading term. Models with similar mass differences should therefore have similar kinematics with some weak dependence on the sbottom mass, and thus can be optimized as a group with a single signal region. This intuitive characteristic, though more difficult to show analytically, is shared by the gluino-mediated models, where Δm is defined as $m_{\tilde{g}} - m_{\chi_1^0}$. Δm will be used as a stand-in for decay phase space and a classifier for model kinematics throughout the optimization and analyses.

A comparison between p as given by the full expression (Equation 8.6) and $p = \Delta m$ is shown in Figure 8.1. Note that even when Δm is large and thus $p = \Delta m$ is a poor approximation, the contours of constant p resemble those of constant Δm (as long as the neutralino is non-relativistic) because the derivative of p with respect to $m_{\tilde{b}}$ is small. This structure should reappear in the final optimization results provided that the contribution to the final state kinematics from the hard scatter are negligible.




Figure 8.1: Δm (left) and the characteristic momentum p (right) for the direct sbottom model in the $m_{\tilde{b}}$ and $m_{\chi_1^0}$ plane. As long as the neutralino is non-relativistic, the contours have similar shape.

8.1.2 Influence of the Hard-Scatter Process

How negligible are kinematic contributions from the hard scatter? A naive expectation might be that, to first order, the center-of-mass of the sbottom-sbottom system would still be at rest in the transverse plane, but each sbottom would acquire some equal-magnitude opposite-direction transverse momentum. This would result in a broadening of the kinematic distributions along with a shift to higher momentum and energy. A simple way to test this prediction is to compare the parton-level kinematic predictions of the toy model described above with signal Monte Carlo, which also simulates the hard scatter. An example comparison, with $m_{\tilde{b}} = 450$ GeV and $m_{\chi_1^0} = 150$ GeV ($p = 200$ GeV) is shown in Figure 8.2.

Both the b -quark p_T and \cancel{E}_T distributions are broadened in the signal Monte Carlo as compared to the toy model. The effect is tempered somewhat in the \cancel{E}_T distribution due to it being a composite quantity formed by the addition of two mostly-uncorrelated four-vectors. For the b -quark p_T , the effect is much more dramatic at the high end. However, the maxima of each distribution (at $\approx p$) are unaffected, as are, to good approximation, the soft ends of both distributions. For instance, only for b -quark p_T selections above 150 GeV would there be significant differences in acceptance. The same conclusion holds for the \cancel{E}_T distributions out to 400 GeV. The expectation that this distortion of the decay-only spectra arises primarily from the

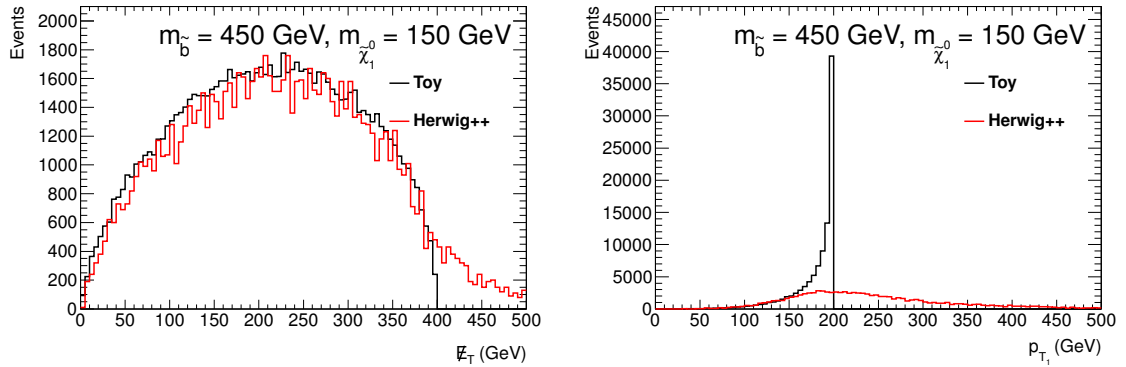


Figure 8.2: \cancel{E}_T (left) and leading b -jet p_T (right) for the direct sbottom model with $m_{\tilde{b}} = 450$ GeV and $m_{\tilde{\chi}_1^0} = 150$ GeV ($p = 200$ GeV) in the $m_{\tilde{b}}$ and $m_{\tilde{\chi}_1^0}$ plane. Black is the toy model described in Section 8.1.1 with no corrections for the hard-scatter process, while red is the Monte Carlo prediction from HERWIG++. The Monte Carlo distributions are broadened compared to the toy model, mostly at the hard end of the spectrum.

back-to-back momenta of the two sbottoms with the sbottom system center-of-mass remaining at rest in the transverse plane can be confirmed by applying the Monte Carlo sbottom transverse momentum distribution to the toy model. This is shown in Figure 8.3.

8.1.3 Compressed Scenarios

In the small- Δm limit, the phase space available for the decay is also small. The neutralino and bottom quark are produced approximately at rest in the frame of the sbottom. This is a very large experimental problem, as even with a large back-to-back boost of the sbottoms the two neutralinos in the final state will essentially be anti-aligned with one another, producing very little observable \cancel{E}_T . This type of final state is almost invisible experimentally and is known as a *compressed scenario* [162, 163].

Radiation of additional hard partons during the hard-scatter process changes this by boosting the rest of the event in the direction opposite to the hard parton. This both increases the momenta of the neutralinos and bottom quarks in the detector frame and tends to bring the neutralinos away from anti-alignment. Emission of




Figure 8.3: \cancel{E}_T (left) and leading b -jet p_T (right) for the direct sbottom model with $m_{\tilde{b}} = 450$ GeV and $m_{\chi_1^0} = 150$ GeV ($p = 200$ GeV) in the $m_{\tilde{b}}$ and $m_{\chi_1^0}$ plane. Black is the toy model described in Section 8.1.1 with a symmetric application of the HERWIG++ sbottom transverse momentum distribution, while red is the Monte Carlo prediction from HERWIG++. For models with large p , such as this one, the sbottom transverse momentum distribution accounts for all of the difference seen in Figure 8.2.

relatively hard radiation is more likely for production processes involving large momentum transfer (higher sbottom or gluino mass). The kinematic difference is more dramatic for the \cancel{E}_T , both due to the angular dependence as well as the fact that the momentum of a particle with velocity β is proportional to its mass, which is much larger in the neutralino case. This effect can be seen in the \cancel{E}_T and b -quark p_T distributions for $m_{\tilde{b}} = 750$ GeV and $m_{\chi_1^0} = 700$ GeV ($p \approx 50$ GeV) shown in Figure 8.4.

In summary, radiation of additional partons, often considered a nuisance in hadron collider physics, in this case makes searching for otherwise undetectable models viable. It does, however, imply a dedicated signal region (in the sbottom case this region would involve an additional hard jet from radiation, soft b -jets, and large \cancel{E}_T). A steep price may be paid in effective cross section by requiring an additional hard object in the event. The \cancel{E}_T enhancement from heavy neutralinos should also be visible in the optimization results in cases of small Δm in both the gluino-mediated sbottom/stop and direct sbottom models.

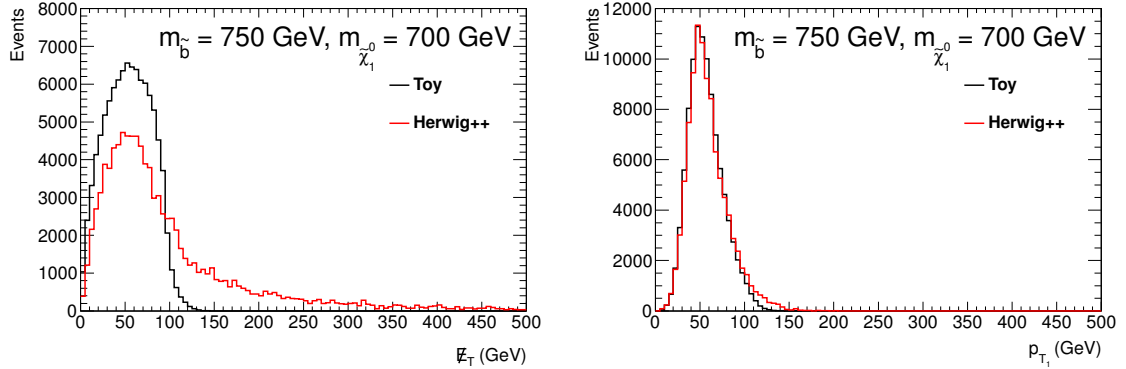


Figure 8.4: \cancel{E}_T (left) and leading b -jet p_T (right) for the direct sbottom model with $m_{\tilde{b}} = 750$ GeV and $m_{\chi_1^0} = 700$ GeV ($p \approx 50$ GeV) in the $m_{\tilde{b}}$ and $m_{\chi_1^0}$ plane. Black is the toy model described in Section 8.1.1 with a symmetric application of the HERWIG++ sbottom transverse momentum distribution, while red is the Monte Carlo prediction from HERWIG++. The \cancel{E}_T shows large differences due to the combined effect of small Δm , large $m_{\chi_1^0}$, and the toy model’s lack of additional hard-scatter partons. The leading b -jet p_T distribution is much less affected.

8.2 Optimization Procedure

The optimization process is a simple approach to developing cut-based signal regions suitable for counting experiments such as the analyses described in Chapters 10 and 11. The figure of merit used to rank candidate signal regions is the significance Z as given by:

$$Z = \frac{s}{\sqrt{b + \sigma_b^2}} \quad (8.10)$$

where s is the expected signal yield, b the expected background yield, and σ_b the expected absolute background systematic error on the background. This formula has a straightforward qualitative interpretation as a comparison between the signal yield and the total background error, both statistical and systematic. More rigorously, it can be justified as the small $\frac{s}{b}$, small $\frac{\sigma_b^2}{b}$ limit of the significance from the profile likelihood test on the Asimov dataset [164]:

$$Z_A = \frac{s}{\sqrt{b + \sigma_b^2}} \left(1 + \mathcal{O}\left(\frac{s}{b}\right) + \mathcal{O}\left(\frac{\sigma_b^2}{b}\right) \right) \quad (8.11)$$

Equation 8.10 is linear in s . Therefore, the signal region with the highest significance will remain the most sensitive signal region independent of the effective signal cross section, provided the yield is small. This is not the case for the dependence on integrated luminosity—the optimization must be redone when the dataset size changes.

The optimization process begins with the application of a set of preselection cuts related to data quality, lepton vetoes, trigger momentum and energy thresholds and QCD background rejection. The data-quality and baseline-lepton criteria are described in Chapter 7. The lepton veto applies to any event with a baseline lepton. For the jet + \cancel{E}_T trigger (Section 7.11) used for the gluino-mediated analysis, the leading jet p_T is required to be greater than 130 GeV, and the $\cancel{E}_T > 160$ GeV. For the \cancel{E}_T -only trigger used for the direct sbottom analysis, the preselection cut is $\cancel{E}_T > 150$ GeV. The QCD preselections require that the difference in azimuthal angle (ϕ) between each of the leading 3 jets in the event and the \cancel{E}_T be larger than 0.4. This selection is very effective at rejecting events with fake \cancel{E}_T arising from a single mis-measured jet.

All additional selections are varied within the optimization. The set of variables considered is tailored to each analysis. While the lists of variables are extensive, they are restricted to quantities that can be justified by the kinematics of the simplified model considered. Moreover, invariant-mass and angular variables are avoided in favor of less-sensitive quantities such as the scalar sum of the p_T of the jets and \cancel{E}_T , known as the effective mass (m_{eff}). The goal of this restriction is to limit dependence on the gluino phase-space matrix element assumption in the gluino-mediated case.

For n variables, an n -dimensional cut grid is then constructed, with appropriate granularity chosen for each variable by examination of the one-dimensional distributions. The significance (including b -tagging and jet energy scale systematic uncertainties calculated using the prescription for Monte Carlo described in Chapter 9) is calculated for each point in the grid with signal and background Monte Carlo, as well as numerical derivatives of the significance with respect to all variables. The b -tagging and jet energy scale uncertainties are the dominant detector-level systematic errors, and the only ones considered in the optimization. As every cut on a given variable is paired with every other cut on every other variable in the input set, correlations

between variables are automatically taken into account.

For each analysis, the cut grid contains $\approx 10^6$ points. Therefore, for $\mathcal{O}(100)$ signal samples, there tends to be a unique best point for each model. The definition of “best” is the point with highest significance, subject to the restriction that the expected background yield b_{exp} be more than five events. This minimum threshold was chosen to ensure that not only the total yields in the signal regions but also the shape of the distributions can be compared reliably between data and Monte Carlo, and that Equation 8.10 is a good approximation for the significance (Figure 8.5). Such a minimum background yield also helps ensure sufficient Monte Carlo statistics are used in the significance calculations and avoids bias in the calculation of the systematic uncertainties.

100 signal regions are of course impractical, and as expected, models with similar kinematics have similar best selections. The differences are often due to statistical fluctuations in the Monte Carlo samples. The derivatives are used to identify the variables upon which the significance has only weak dependence. Cuts on such variables are then consolidated to an average value or eliminated. The number of signal regions is iteratively reduced until a reasonable number (≈ 5) is reached. This reduction is not an automated process. The goal is final selections which maintain at least 50% of “best” sensitivity over the entire model space. In other words, sensitivity which is both broad and deep. Statistical independence between signal regions is considered advantageous but not mandatory.

8.3 Variables

The variables considered in the optimization for both the direct sbottom and gluino-mediated analyses are:

- **Number of jets** - Minimum or maximum number of jets. Subject to a p_T threshold.
- **Jet p_T threshold** - Minimum p_T requirement for the jets counted in the number-of-jets calculation.




Figure 8.5: Comparison between the significance as calculated using the $\frac{s}{\sqrt{b+\sigma_b^2}}$ approximation, the significance using the Asimov dataset (Z_A), and the significance calculated using toy Monte Carlo [164]. All methods agree well in the small $\frac{s}{b}$ and small $\frac{\sigma_b^2}{b}$ limit.

- **Number of b -jets** - Minimum or maximum number of b -jets for a chosen operating point. Subject to a p_T threshold.
- **b -jet p_T threshold** - Minimum p_T requirement for b -jets.
- **b -jet operating point** - Minimum MV1 (Section 7.8) selection used to identify b -jets. The available operating points correspond to b -tagging efficiencies of 60%, 70%, and 75% in $t\bar{t}$ Monte Carlo simulation.
- **Leading jet p_T threshold** - Minimum p_T requirement for the leading jet in the event.
- \cancel{E}_T - Minimum \cancel{E}_T threshold.
- H_T - Minimum scalar sum of the p_T of the n leading jets.
- m_{eff} - Minimum scale sum of H_T and \cancel{E}_T . Designed to reject $t\bar{t}$ and vector boson + jets background. High m_{eff} signatures are often characteristic of SUSY [165].
- $\frac{\cancel{E}_T}{m_{eff}}$ - Minimum cut applied to reject QCD multi-jet background.

Additional variables considered only for the direct sbottom optimization are:

- **Leading jet anti-tag** - The leading jet is required to *not* be b -tagged. Designed for rejection of $t\bar{t}$ background for signal regions targeting direct sbottom signatures with hard radiation, as in this case the leading jet should be a gluon/light flavor jet.
- **Leading b -jet p_T** - Maximum p_T of leading b -tagged jet. Designed to remove $t\bar{t}$ background in cases where soft b -jets are expected from the signal (small Δm).
- $\Delta\phi(b_0, \cancel{E}_T)$ - Maximum difference in ϕ between the leading b -tagged jet and the \cancel{E}_T . Designed to reject $t\bar{t}$ background in cases where the sbottom system recoils against a hard radiation jet.
- $\Delta\phi(j_0, \cancel{E}_T)$ - Minimum difference in ϕ between the leading jet and the \cancel{E}_T . Designed to reject $t\bar{t}$ background in cases where the sbottom system recoils against a hard radiation jet (the leading jet).

- $\Delta\phi(j_1, \cancel{E}_T)$ - Maximum difference in ϕ between the second-leading jet and the \cancel{E}_T . Designed to reject $t\bar{t}$ background in cases where the sbottom system recoils against a hard radiation jet (the leading jet).
- $H_{T,n}$ - Maximum scalar sum of all jets *excluding* the n leading jets. Rejects $t\bar{t}$ when fewer jets are expected for signal.
- m_{CT} - Minimum contransverse mass [166, 167]. $t\bar{t}$ has a kinematic endpoint at ≈ 140 GeV. Intended for signal models with large Δm , as direct sbottom signal has an endpoint at $\approx 2\Delta m$.
- m_T^t - Minimum top transverse mass. This mass is reconstructed assuming the \cancel{E}_T represents the transverse W momentum, with combinatorics resolved according to the combination which yields the maximum m_T^t for the event.

The majority of this second list are designed for signal regions for direct sbottom events with additional hard radiation. Several were suggested by Reference [168].

8.4 Direct Sbottom Signal Regions

The final direct sbottom signal regions after optimization are given in Table 8.1. No benefit is found for a separate b -jet p_T threshold, so it remains 30 GeV. The tightest b -tagging operating point (60% efficiency) was optimal for all regions.

All signal regions require two b -tagged jets to reject QCD and vector boson + jets and a common $\cancel{E}_T/m_{\text{eff}}$ cut to reject QCD multi-jet events. As expected from the discussion in Section 8.1, the signal region most sensitive for a given signal point depends largely on Δm as seen in Figure 8.6.

When Δm is large (> 250 GeV), the most sensitive region is the one denoted SR1. SR1 is similar to the signal region of the previous iteration of the analysis [2], in that it relies on hard cuts on the leading jet p_T and m_{CT} to reject $t\bar{t}$, the primary background after requiring two b -tagged jets. It is divided into three subregions with progressively harder m_{CT} cuts, though only the subregion with the hardest cut (> 250 GeV) is usually relevant. There is also a veto on the third jet for $p_T > 50$ GeV.

Description	Direct Sbottom Signal Regions			
	SR1	SR2	SR3a	SR3b
$\cancel{E}_T/m_{\text{eff}}$	> 0.25 , m_{eff} computed with leading 3 jets			
\cancel{E}_T	> 150 GeV	> 200 GeV	> 150 GeV	> 250 GeV
Leading jet p_T (j_0)	> 130 GeV	> 60 GeV	> 130 GeV	> 150 GeV
Jet p_T	> 50 GeV	> 60 GeV	> 30 GeV	
n jets	exactly 2	at least 2	at least 3	
n b -jets	leading 2 jets		2^{nd} - and 3^{rd} -leading jets	
m_{CT}	$> 150, 200, 250$ GeV	> 100 GeV	-	
$H_{T,n}$	-	< 50 GeV, $n = 2$	< 50 GeV, $n = 3$	
$\Delta\phi(\cancel{E}_T, j_0)$	-		> 2.5	
Leading jet anti-tag	-		yes	
Leading b -jet p_T	-		< 110 GeV	

Table 8.1: Summary of the optimization-derived event selection for the direct sbottom signal regions.

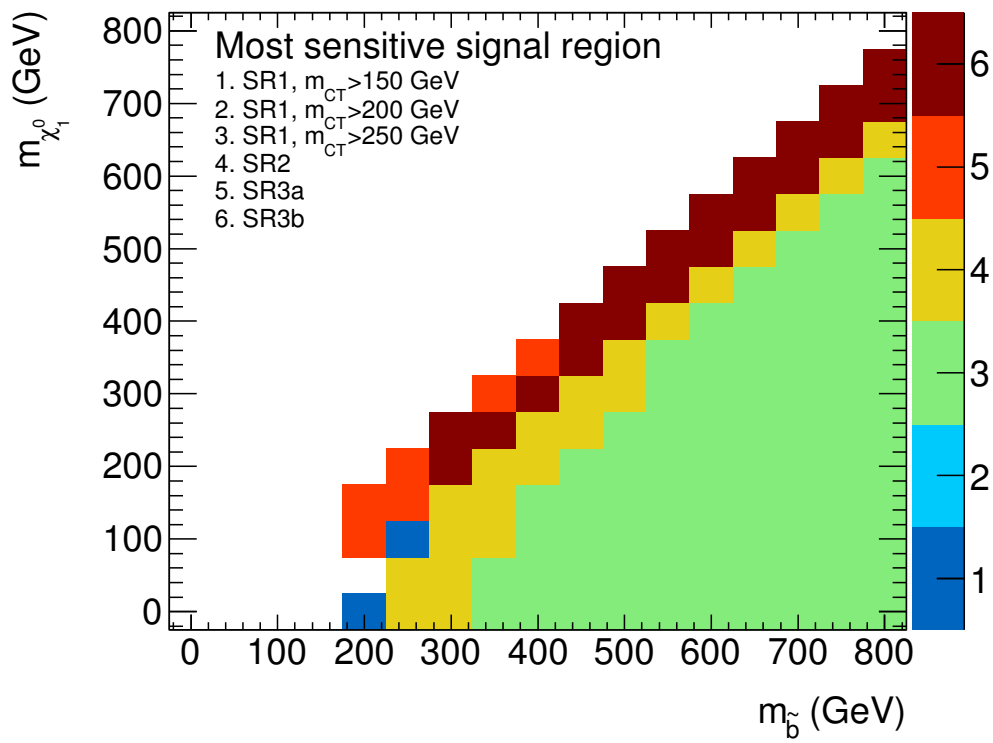


Figure 8.6: Most sensitive direct sbottom pair production signal region as a function of the sbottom and neutralino masses.

For Δm between 100 GeV and 250 GeV, SR2 becomes the most sensitive signal region. This region is similar to SR1 with significantly relaxed leading jet p_T and m_{CT} cuts. This is necessary to boost signal acceptance as m_{CT} has an endpoint for the direct sbottom topology as well [166]:

$$m_{CT}^{max} = \frac{m_b^2 - m_{\chi_1^0}^2}{m_{\tilde{b}}} \approx 2\Delta m \quad (8.12)$$

The final approximation in Equation 8.12 is valid in the small- Δm limit. The SR2 cut relaxation drastically reduces the $t\bar{t}$ rejection, so the \cancel{E}_T cut is raised slightly to compensate. This does not affect signal efficiency as much as the jet p_T thresholds (Section 8.1.3). The quantity $H_{T,2}$ is relied upon for primary rejection of $t\bar{t}$.

Near the kinematically forbidden line ($m_{\tilde{b}} = m_b + m_{\chi_1^0}$), the signal kinematics are soft, and the \cancel{E}_T in particular is small for final states with no hard radiation. Here, the two variants of SR3 dominate. They require a hard leading jet, anti-tagged, and that the second- and third-leading jets be b -tagged. The p_T thresholds on these b -tagged jets are low as the jets are expected to be soft. Various other selections shown in Table 8.1 are applied to SR3 to enforce the desired topology of one hard parton recoiling against the sbottom-pair system. The SR3b region has higher leading jet p_T and \cancel{E}_T cuts and is best when $m_{\chi_1^0}$ is large (expected from the discussion in Section 8.1.3). The requirement of an additional hard parton also causes an enormous loss in signal acceptance compared to SR2 and SR1 (Figure 8.7) as predicted in Section 8.1.3.

Quantifying how much is gained by the optimization compared to the signal regions from the previous version of the analysis [2] is an important benchmark. While this will be somewhat apparent in the final limits, the limits are also very dependent on integrated luminosity and the dataset used here (4.71 fb^{-1}) is ≈ 1.5 times larger. A luminosity-independent comparison can be done by comparing the predicted significance for the new best signal region (highest significance) for each point with that of the old best signal region. The results of such a comparison are shown in Figure 8.8.

An improvement of over 100% (more than a factor of two) is typical in the bulk




Figure 8.7: Signal acceptance, in percent, as a function of mass plane for the four direct sbottom signal regions. Note that the scale is different for each plot.




Figure 8.8: Relative sensitivity increase in percent of the optimized 4.71 fb^{-1} direct sbottom signal regions over the signal regions from the previous version of the analysis [2]. The most sensitive signal region is used at each point (Figure 8.6). Note that the scale is logarithmic.

of the mass plane (SR1), while in the transition region between the two types of final state signatures (SR2) the sensitivity increase is a more modest 20%. The improvement along the forbidden line (SR3) is often several thousand percent. This is chiefly because the previous version of the sbottom search was almost entirely insensitive to this region of the mass plane.

8.5 Gluino-mediated Sbottom/Stop Signal Regions

The optimization for the gluino-mediated analysis is very different from the sbottom case, though the methods used are the same. The central feature of the gluino-mediated topologies is that there are four b -quarks in the final state. While the sbottom analysis has to deal with irreducible $t\bar{t}$ and $Z \rightarrow \nu\nu + b\bar{b}$ backgrounds, the gluino-mediated analysis can, in principle, simply require more b -tagged jets (the Standard Model backgrounds which have four real b -quarks and \cancel{E}_T in the final state, such as $t\bar{t}b\bar{b}$, have very low cross sections). As the final states are often dramatic (large amounts of jet activity and \cancel{E}_T), the effective mass (m_{eff}) and similar variables are also very powerful tools. The final signal region selections are shown in Table 8.2.

All gluino-mediated signal regions require three b -tagged jets instead of the two used in the previous iteration [3], making this the first new physics search to do so. The difficulties with three b -tags have to do with the background estimation (Chapter 11). From an optimization perspective, the fact that $t\bar{t}$ is the dominant background and contains only two real b -jets make the advantage of three or more tags seemingly obvious. This simplistic view, however, ignores the subtle balance between signal efficiency, background rejection, and systematic errors. Four b -tagged jets, for instance, is worse than three for sensitivity due to the higher systematic error and the loss in signal efficiency incurred by requiring an additional tag. The three-tag signal regions are not as affected by efficiency loss due to the signal topologies containing a “spare” real b -jet.

The four-jet signal regions (SR4-L, SR4-M, and SR4-T, with the trailing letters representing “loose”, “medium”, and “tight”, respectively) were optimized for the

Description	Gluino-mediated Sbottom and Stop Signal Regions				
	SR4-L	SR4-M	SR4-T	SR6-L	SR6-T
Leading jet p_T (j_0)	> 130 GeV				
$\cancel{E}_T/m_{\text{eff}}$	> 0.2, m_{eff} computed with leading 4 or 6 (n) jets				
\cancel{E}_T	> 160 GeV			> 200 GeV	
Jet p_T	> 50 GeV				
n jets	at least 4			at least 6	
b -jet p_T	> 30 GeV, not restricted to the leading n jets				
b -tagging operating point	60%		70%		75%
n b -jets	at least 3				
m_{eff}	> 500 GeV	> 700 GeV	> 900 GeV	> 700 GeV	> 900 GeV

Table 8.2: Summary of the optimization-derived event selection for the gluino-mediated sbottom and stop signal regions.

gluino-mediated sbottom simplified models, while the six-jet signal regions were optimized for gluino-mediated stop. The hybrid gluino-mediated decay with 2 top quarks and two bottom quarks in the final state also does well with the six-jet signal regions. The most-sensitive signal regions as a function of the gluino and neutralino masses are shown in Figure 8.9.

The contours roughly resemble those of Δm , as this a measure of the phase space available for the gluino decay. Deviations from the Δm expectation appear for signal points near the forbidden line, particular those with high gluino mass. This is consistent with the discussion of compressed regions and hard-scatter influence in Sections 8.1.2 and 8.1.3.

As discussed earlier in this section, requiring three b -tagged jets is expected to significantly reduce the dominant $t\bar{t}$ background, and therefore one might expect the improvement in sensitivity relative to the two-tag analysis [3] to be dramatic. This is expectation is confirmed in Figure 8.10. The relative sensitivity improvement is typically between 200% and 400% for the sbottom case (three-fold and five-fold sensitivity




Figure 8.9: Most sensitive gluino-mediated sbottom (left) and stop (right) signal region as a function of the gluino and neutralino masses.

increase, respectively), and between 600% and 1200% (seven-fold and thirteen-fold sensitivity increase, respectively) for the stop case.

One might ask why no specialized signal regions were constructed for the compressed region (near the $m_{\tilde{g}} = 2m_q + m_{\chi_1^0}$ line) as they were in the direct sbottom search optimization. While a case can be made, the fact that the b -tagged jets in the gluino-mediated signal regions are not required to be the leading jets makes these signal regions much more sensitive to final states with hard radiation than are the “normal-topology” direct sbottom signal regions SR1 and SR2.

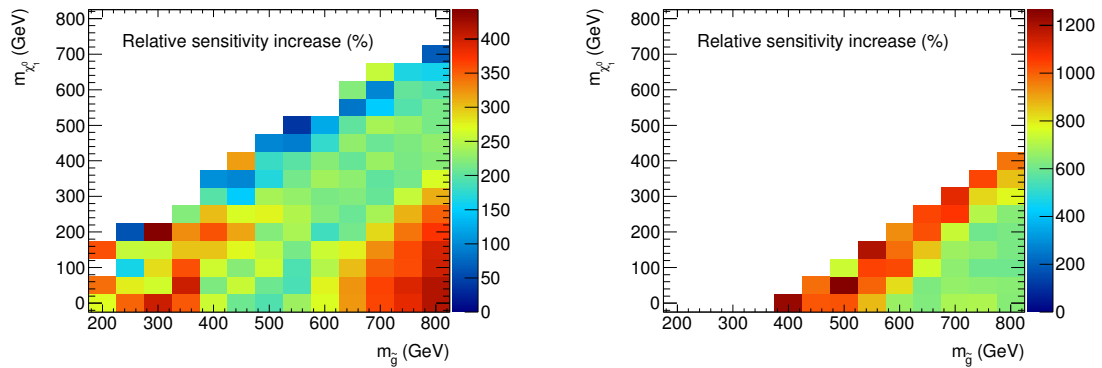


Figure 8.10: Relative sensitivity increase in percent of the optimized 4.71 fb^{-1} gluino-mediated sbottom (left) and stop (right) signal regions over the signal regions from the previous version of the analysis [3]. The most sensitive signal region is used at each point (Figure 8.9).

Chapter 9

Background Estimation Techniques

Background contributions to the signal regions and their uncertainties are estimated in different ways depending on the accuracy of the Monte Carlo description of the process and the process' relative importance to the total background estimation.

9.1 Monte Carlo-based Predictions

Some backgrounds contribute little to the final signal region yield and/or are not feasible to measure via control region. The alternative is to obtain a background estimate directly from Monte Carlo simulation. For the direct sbottom analysis, the primary backgrounds in this category are W +jets and dibosons. Both are very effectively suppressed by the lepton vetoes, the two b -tagged jets, and the high \cancel{E}_T requirements (Table 8.1). Other backgrounds such as $t\bar{t}+W/Z/WW/b\bar{b}$ have very low cross sections and are further suppressed by the cuts designed to reject $t\bar{t}$ generally. For the gluino-mediated analysis, the category of small backgrounds expands to also include Z +jets.

9.2 Monte Carlo Systematic Uncertainties

Even for non-negligible backgrounds, accurate Monte Carlo predictions and quantitative estimates of the systematic uncertainties on those predictions are essential, as

these numbers are important inputs to the semi-data-driven background estimation methods. The general estimation method is similar for all event-dependent systematic errors. A comparison is made between Monte Carlo and data (or, for the generator uncertainties, between one Monte Carlo and another) to quantify any discrepancy, e.g. differences in b -jet efficiency. In some cases, like b -tagging, a correction is derived (Section 7.13) and the systematic uncertainty is expressed as uncertainties on the scale factors. In other cases, the discrepancy itself is taken to be the systematic error.

Signal yield predictions can only be made using Monte Carlo, and the predicted signal systematic errors are estimated in the same way as done for background Monte Carlo estimates. These uncertainties are used in the limit computations in Chapter 12.

9.2.1 Detector Uncertainties

Once detector systematic errors are appropriately parameterized by physics object (for instance, the efficiency to tag a jet as a function of b -tagging operating point, jet p_T , jet η , and jet flavor), the yields for a given region are recalculated by rerunning the analysis with all objects shifted up by one standard deviation, and also with all objects shifted down. In the case of event weights (i.e. b -tagging and pileup reweighting), the Monte Carlo events in the region do not change with these shifts, but their weights do, and thus so do the yields. In other cases, such as the jet energy scale, the event kinematics are modified (the jets are shifted up and down in p_T) and thus the region of interest is populated by a somewhat different set of events after the “shifted” runs. The differences in event yield between the shifted runs and the nominal run quantify the systematic uncertainty for the region. While the “up” and “down” systematic errors can be asymmetric, in these analyses all uncertainties are symmetrized conservatively by taking whichever deviation is larger.

Heavy Flavor Tagging As discussed in Section 7.13, the discrepancy between data and Monte Carlo in b -tagging efficiency is corrected via a scale factor dependent on b -tagging operating point, jet p_T , jet η , jet flavor, and whether the jet is tagged. The event scale factor is the product of the individual jet scale factors, whose uncertainties

range between 5% and 20% per b -jet (larger for higher p_T) [158, 159, 160, 153, 161]. The shifted event scale factor is the product of the shifted jet scale factors. This is often the dominant systematic uncertainty due to the large number of b -tagged jets in the signal regions. The uncertainties for b -jets, c/τ -jets, and light jets are computed independently and are uncorrelated.

Jet energy scale Jet p_T are shifted by the standard deviation of the jet energy scale [146, 147, 148, 149] and parameterized as a function of jet p_T and η . This variation is also propagated to the \cancel{E}_T [142] and is often the second-largest source of systematic error (between 10% and 20%). Jet energy scale uncertainty is most important in regions with very hard jets (i.e. high m_{eff}).

Jet energy resolution Jet p_T are smeared (broadened) by the standard deviation of the jet energy resolution, parameterized as a function of jet p_T and η [169]. This variation is also propagated to the \cancel{E}_T , and is usually very small. Note that this “shift” can only be applied in one direction.

\cancel{E}_T cluster scale and resolution Similar to the jet versions, these uncertainties are due to the scale and resolution of the unclustered energy in the calorimeter. They are almost always negligible relative to other uncertainties.

Pileup reweighting A one-sided uncertainty is generated by shifting the pileup weights (Section 7.13) higher by 10%. This uncertainty is typically relatively small.

Luminosity A flat 3.9% luminosity uncertainty [170] is applied for the 4.71 fb⁻¹ 2011 dataset.

Lepton trigger and reconstruction There are small uncertainties associated with the lepton reconstruction efficiencies and the single-lepton trigger scale factors (Section 7.13). They were found to be always negligible and thus are not included in the final results.

9.2.2 Theoretical/Generator Uncertainties

The theoretical/generator uncertainties vary between 10% and 40% and tend to be dominated by the uncertainties associated with $t\bar{t}$, as it is the dominant or co-dominant background for all signal regions. This uncertainty mostly cancels in the semi-data-driven background estimation methods.

Cross section Variations in PDF set and renormalization/factorization scale are used to assign a 4% baseline cross section uncertainty to vector boson + jets and diboson production and an 8% uncertainty to $t\bar{t}$ and single top production. A similar procedure is used to estimate a signal cross section uncertainty, though this is only used in the limit calculations.

Initial and final state radiation For $t\bar{t}$, the uncertainty due to initial and final state radiation (ISR/FSR) is assessed using **AcerMC** [171] samples with variations of the **PYTHIA** parameters related to the ISR branching phase space and the FSR low- p_T limit (the point below which partons are assumed not to radiate). These variations are chosen to produce jet activity consistent with the data. This uncertainty tends to be the dominant generator uncertainty for $t\bar{t}$ background, and can be as large as 40% for large- m_{eff} regions. For the direct sbottom signal samples with low Δm , the default **HERWIG++** samples are compared with **MADGRAPH** samples generated with zero, one, and two additional partons in the matrix element. This quantifies the differences in ISR description between the two generators and is used as a systematic in the direct sbottom limits calculated in Chapter 12.

Generator For the $t\bar{t}$ background, **ALPGEN** and **MC@NLO** generators have been compared with each other and with **POWHEG** [172] interfaced to **HERWIG** and **JIMMY** to compare the generator dependence of the simulated $t\bar{t}$ distributions. The uncertainty is typically between 5% and 15% and is often one of the dominant theoretical/generator systematic uncertainties.

Parton shower The uncertainty due to the parton shower model in $t\bar{t}$ Monte Carlo is evaluated via comparison of POWHEG interfaced to HERWIG/JIMMY with POWHEG interfaced to PYTHIA. This uncertainty is always much less than the generator uncertainty.

MLM matching (ALPGEN) An additional multiplicative uncertainty of 24% per additional parton (added in quadrature) is assigned to the ALPGEN samples due to uncertainties in the MLM matching procedure [173].

Finite n -partons (ALPGEN) A comparison between ALPGEN yields using an inclusive three-part sample (additional partons come from the parton shower) and the default five- or six-parton inclusive sample is made to evaluate the dependence on the number of partons generated in the matrix element. It was found to be negligible.

Heavy flavor Additional conservative uncertainties ranging from 30% to 100% are applied to the $W/Z/t\bar{t} + b\bar{b}$ samples. For W/Z +jets, these are based on dedicated studies. For $t\bar{t} + b\bar{b}$, an ad-hoc 100% uncertainty was assigned. As the $t\bar{t} + b\bar{b}$ yield never exceeds 10% in any signal region, the effect of this uncertainty on the final results is minimal.

$t\bar{t} + W/Z$ A uncertainty of 70% based on PDF uncertainties and variation of renormalization/factorization scales is applied [174].

9.3 Transfer Factors

For the dominant backgrounds ($t\bar{t}$ and $Z \rightarrow \nu\nu$ +jets for the direct sbottom analysis, $t\bar{t}$ only for the gluino-mediated analysis), use of the Monte Carlo predictions directly would result in large overall systematic uncertainties on the background estimate. These uncertainties can be reduced significantly using control regions enriched in the dominant backgrounds as constraints. The semi-data-driven “transfer factor” approach was used in the 2.05 fb^{-1} analyses [2, 3] as the primary method of estimation for the dominant backgrounds. In the 4.71 fb^{-1} analyses, it is used as a cross-check.

For a given background i , the transfer factor is a Monte Carlo-based number defined as the ratio between the yield in the signal region (SR) and the yield in the control region (CR):

$$T_i = \frac{b_{i,SR}^{MC}}{b_{i,CR}^{MC}} \quad (9.1)$$

The semi-data-driven predicted contribution from background i in the signal region can then be written as:

$$b_{i,SR}^T = T_i(n_{CR} - \sum b_{\neq i,CR}) \quad (9.2)$$

where n_{CR} is the measured total yield in the control region in data and $\sum b_{\neq i,CR}$ is the predicted yield from backgrounds other than background i . The control region is usually chosen such that $\sum b_{\neq i,CR}$ is small, but also so that it is as kinematically similar its corresponding signal region as possible without suffering from significant signal contamination or low statistics. This latter consideration is to ensure that the extrapolation encapsulated by the transfer factor T_i is kept to a minimum.

The advantage of this method over the direct use of Monte Carlo is both a more accurate prediction and correspondingly lower systematic uncertainties. As the systematic uncertainties are correlated, they tend to cancel wholly or in part in the transfer factor T_i . As long as $\sum b_{\neq i,CR}$ is much less than n_{CR} , even large relative uncertainties associated with $\sum b_{\neq i,CR}$ are small when propagated to $n_{CR} - \sum b_{\neq i,CR}$. If the number of events in the control region is large, the Poisson error on n_{CR} is small, and as a result the total uncertainty of $b_{i,SR}^T$ is small. The quantity $\sum b_{\neq i,CR}$ can be taken from Monte Carlo or from another data-driven or semi-data-driven technique.

9.4 Profile Likelihood

In cases where a background is constrained by multiple control regions or a single control region constrains multiple dominant backgrounds, a simultaneous fit based on the profile likelihood method can be advantageous. This is the primary method used for the dominant backgrounds in the 4.71 fb^{-1} analyses.

Neglecting systematic uncertainties, the likelihood function is constructed as a product of Poisson probabilities for all control regions. For a single dominant background i and N control regions j , the likelihood function is given by Equation 9.3:

$$L(\mu_i) = \prod_j^N \frac{(\mu_i b_{i,j}^{MC} + \sum b_{\neq i,j})^{n_j}}{n_j!} e^{-(\mu_i b_{i,j}^{MC} + \sum b_{\neq i,j})} \quad (9.3)$$

where μ_i is scale factor for background i , $b_{i,j}$ is the predicted contribution of background i to control region j from Monte Carlo, and $\sum b_{\neq i,j}$ is the sum of all other contributions to control region j . For a single control region, maximizing the likelihood functions is trivial and the result equivalent to the transfer factor method:

$$\hat{\mu}_i = \frac{n_{CR} - \sum b_{\neq i,CR}}{b_{i,CR}^{MC}} \quad (9.4)$$

$$n_{i,SR}^{PL} = \hat{\mu}_i b_{i,SR}^{MC} \quad (9.5)$$

If there are multiple control regions, however, solving for $\hat{\mu}_i$ (the value of μ_i which maximizes the likelihood) is slightly less trivial, and $n_{i,SR}^{PL}$ is influenced by all control regions with appropriate statistical weight. Equation 9.3 can be expanded to consider two scale factors for two co-dominant backgrounds i and k :

$$L(\mu_i, \mu_k) = \prod_j^N \frac{(\mu_i b_{i,j}^{MC} + \mu_k b_{k,j}^{MC} + \sum b_{\neq i/k,j})^{n_j}}{n_j!} e^{-(\mu_i b_{i,j}^{MC} + \mu_k b_{k,j}^{MC} + \sum b_{\neq i/k,j})} \quad (9.6)$$

$L(\mu_i, \mu_k)$ can then be maximized to find $\hat{\mu}_i$ and $\hat{\mu}_k$ simultaneously, assuming at least two control regions. Systematic uncertainties can be incorporated into the likelihood function with additional Gaussian factors. Generalizing Equation 9.6 to arbitrary numbers of scale factors μ_i by converting i to an index, and adding the systematic factors gives:

$$L(\mu_i) = \prod_j^N \frac{(\sum \mu_i b_{i,j})^{n_j}}{n_j!} e^{-(\sum \mu_i b_{i,j})} \prod_k^M e^{-s_k^2/2} \quad (9.7)$$

where s_k is a parameter measuring the deviation due to systematic k from nominal in units of standard deviation. The $b_{i,j}$ are now defined as:

$$b_{i,j} \equiv b_{i,j}^{MC} + \sum_k^M s_k t_{i,j,k} \sigma_{i,j,k} \quad (9.8)$$

$\sigma_{i,j,k}$ is the standard deviation of systematic k for background i in control region j , and $t_{i,j,k}$ is a parameter with absolute value one which represents the sign of the relationship between s_k and the deviation in control region j of background i . The likelihood function can then be maximized with respect to both the μ_i and the s_k to obtain better estimates of μ_i . In practice, the fitted values of s_k are often examined as additional goodness-of-fit measures. Values of s_k significantly different from zero can indicate something wrong with either the fit or inconsistency in the input parameters.

9.5 Matrix Method

The so-called *matrix method* is a (potentially) fully data-driven and somewhat generic procedure for background estimation based on the two definitions of a physics object, *loose* (L) and *tight* (T). The downside of the method is that it introduces correlations between the background estimate and the signal region yield. This is why it is used in these analyses only for small backgrounds not well-modeled by Monte Carlo (the QCD multi-jet contribution to the one- and two-lepton control regions from jets faking leptons) or as a cross-check (verification of the total background estimate for the gluino-mediated analysis). In the one- and two-lepton QCD multi-jet estimation procedures loose and tight leptons are used (Sections 7.4 and 7.5), while for the global cross-check in the gluino-mediated analyses the physics object of interest is the jet and tight jets are those which are b -tagged.

The method is simplest when the number of physics object in question can be only one or zero, as is the case for the one-lepton control region QCD multi-jet estimate. Two versions of the control region are constructed: the default one, requiring exactly one tight (signal) lepton, and a modified region which instead requires one loose lepton. The tight control region is a subset of the loose control region, and the

number of events in each region can be written:

$$\epsilon N_R + f N_F = N_T \quad (9.9)$$

$$N_R + N_F = N_L \quad (9.10)$$

where N_R and N_F are the number of real and fake leptons in the loose control region, respectively, and ϵ and f are their corresponding efficiencies to pass the tight selection criteria. These efficiencies are ideally taken from data, as is the case for the fake lepton estimates, or can be estimated from Monte Carlo, as is done with the gluino-mediated cross-check.

These two equations form a linear system which can be written as:

$$\begin{pmatrix} \epsilon & f \\ 1 & 1 \end{pmatrix} \begin{pmatrix} N_R \\ N_F \end{pmatrix} = \begin{pmatrix} N_T \\ N_L \end{pmatrix} \quad (9.11)$$

The matrix can be then inverted to solve for the unknown N_R and N_F in terms of the known N_T , N_L , ϵ , and f :

$$\frac{1}{\epsilon - f} \begin{pmatrix} 1 & -f \\ -1 & \epsilon \end{pmatrix} \begin{pmatrix} N_T \\ N_L \end{pmatrix} = \begin{pmatrix} N_R \\ N_F \end{pmatrix} \quad (9.12)$$

The generalization to multiple objects is straightforward, though the size of the matrix increases rapidly. For instance, for 3 jets, any of which could be b -tagged, there are eight equations, each with eight terms.

$$N_{TTT} = \epsilon_1 \epsilon_2 \epsilon_3 N_{RRR} + \epsilon_1 \epsilon_2 f_3 N_{RRF} + \dots \quad (9.13)$$

9.6 Jet Smearing

The QCD multi-jet background to the signal regions in both analyses is not well-reproduced by PYTHIA leading-order Monte Carlo due to the dominant contribution from states with many jets, the high \cancel{E}_T cuts, and the high QCD rejection of the

$\Delta\phi(\text{jet}, \cancel{E}_T)$ and $\frac{\cancel{E}_T}{m_{eff}}$ cuts. For any other background, the rejection is such that its contribution to the signal region could be ignored. However, the QCD multi-jet cross section is enormous (≈ 0.1 bn [175], as opposed to the $t\bar{t}$ cross section of 168 pb).

The method used in these and other ATLAS SUSY searches is called *jet smearing* [176], and is based on the fact that QCD events have no true \cancel{E}_T and that the fake \cancel{E}_T arises from jet energy mis-measurement. Multi-jet events from data are selected with low \cancel{E}_T significance S (< 0.6 GeV^{1/2}), as defined by:

$$S \equiv \frac{\cancel{E}_T}{\sqrt{\sum E_T}} \quad (9.14)$$

The sum is over all transverse energy in the event, both clustered and unclustered. These events are collected using a hierarchy of single jet triggers with prescales varying inversely with p_T threshold. Small S implies that the jets in these events are very close to their true p_T . The p_T of these jets can then be smeared by the jet response function:

$$R \equiv \frac{p_T^{reco}}{p_T^{true}} \quad (9.15)$$

Single seed events can be re-smearred many times to generate large “simulated” QCD samples with significant \cancel{E}_T . The $\Delta\phi(\text{jet}, \cancel{E}_T)$ cuts tend to reject the events with one mis-measured jet, but do not reject events with multiple mis-measured jets. These are the type of QCD multi-jet events that populate the signal regions. A region similar to the signal region but with the $\Delta\phi(\text{jet}, \cancel{E}_T)$ cut reversed will be highly enriched in QCD with one mis-measured jet. Such a region can be used as a control region to normalize the contribution to the signal region. The smeared data sample acts as the “Monte Carlo” sample in the transfer factor estimation method (Section 9.3).

Measuring the jet response function R , however, is not trivial. The strategy is to begin with a Monte Carlo-based correction, which reflects some non-Gaussian features of the response function. Corrections are done separately for b -tagged and light-flavor jets to reflect differences in response caused primarily by b -jets decaying semi-leptonically. Data-driven corrections are then applied.

The first data-driven correction is to the Gaussian width using di-jet p_T balance. For the gluino-mediated analyses, events are selected using single jet triggers with a leading jet $p_T > 130$ GeV, second-leading jet $p_T > 40$ GeV, and no additional jet with $p_T > 30$ GeV. The p_T cuts are relaxed slightly for the direct sbottom analysis. To reduce contamination from $Z \rightarrow \nu\nu + jj$ events, the \cancel{E}_T is required to align with one of the two leading jets ($\Delta\phi(\text{jet}, \cancel{E}_T) < 0.3$). The asymmetry A of such an event is defined as:

$$A \equiv \frac{p_T^1 - p_T^2}{p_T^1 + p_T^2} \quad (9.16)$$

The Gaussian width of A in a bin of average $p_T^1 + p_T^2$ denoted by $\langle p_T^1 + p_T^2 \rangle$ is given by:

$$\sigma_A = \frac{\sqrt{\sigma_1^2 + \sigma_2^2}}{\langle p_T^1 + p_T^2 \rangle} \quad (9.17)$$

If both jets are assumed to have $p_T^{\text{true}} \simeq \frac{\langle p_T^1 + p_T^2 \rangle}{2}$ and the same Gaussian width σ_{p_T} , then Equation 9.17 can be written as:

$$\sigma_A \simeq \frac{\sigma_{p_T}}{\sqrt{2}p_T^{\text{true}}} = \frac{\sigma_R}{\sqrt{2}} \quad (9.18)$$

The Gaussian width of the response in Monte Carlo tends to be narrower than in data. It is broadened until it agrees with σ_R .

While the asymmetry distribution is sensitive to the Gaussian width of the response, up and down fluctuations in p_T cannot be distinguished. For corrections to the low and high tails of the response distribution, events with three or more jets are used. Called ‘‘Mercedes’’ events due to their resemblance to the three-prong Mercedes logo, they are selected for the gluino-mediated analyses with the p_T and angular selections in Table 9.1. The p_T cuts are again relaxed for the sbottom analysis, though the angular selections remain the same.

The response of the jet parallel or anti-parallel to the \cancel{E}_T is given by:

Mercedes Angular and p_T Selections for the Gluino-Mediated Analysis	
Parallel	Anti-parallel
≥ 3 jets with $p_T > 40$ GeV, leading jet $p_T > 130$ GeV	
$\cancel{E}_T > 30$ GeV	
$\Delta\phi(j_1, \cancel{E}_T) < \pi - \Delta\phi(j_N, \cancel{E}_T)$	$\pi - \Delta\phi(j_N, \cancel{E}_T) < \Delta\phi(j_1, \cancel{E}_T)$
$\Delta\phi(j_1, \cancel{E}_T) < 0.1$	$\pi - \Delta\phi(j_N, \cancel{E}_T) < 0.1$
$\Delta\phi(j_2, \cancel{E}_T) > 0.5$	$\pi - \Delta\phi(j_{N-1}, \cancel{E}_T) > 0.5$

Table 9.1: Selections for “Mercedes” events after the single jet trigger. Jets are ordered by azimuthal distance from the $\vec{\cancel{E}}_T$ — j_1 is closest, and j_N is farthest away.

$$R_2 \equiv \frac{\vec{p}_T \cdot (\vec{p}_T + \vec{\cancel{E}}_T)}{|\vec{p}_T + \vec{\cancel{E}}_T|^2} \quad (9.19)$$

as $p_T^{true} \approx |\vec{p}_T + \vec{\cancel{E}}_T|$. After this correction for the tails, the resulting response distribution is used to smear the low- S seed events. Appropriate number of b -jets are selected depending on signal region. Systematic uncertainties are obtained considering statistics and variations in b -tagging, jet energy scale, Gaussian width, and distribution tails. The uncertainty is dominated by the b -tagging and tail variations, and results in $\approx 100\%$ systematic uncertainty on the QCD yield in the signal regions. As the relative contribution of the QCD multi-jet estimate is small, this large uncertainty has negligible effect on the final results.

Chapter 10

Direct Sbottom Search

We now use the optimized signal regions from Section 8.4 and the background estimation methods from Chapter 9 to perform the direct sbottom production search. The general topology for this search is two b -jets + \cancel{E}_T . The dominant background for this search is $t\bar{t}$, however, steps are taken to reduce this background via m_{CT} cuts as well as removing events with jet activity beyond the two required b -jets. As a result, the Z +jets background is co-dominant (primarily $Z \rightarrow \nu\nu + b\bar{b}$). These backgrounds are estimated via one- and two-lepton control regions.

10.1 Top Control Regions

The top control regions require exactly 1 signal lepton (electron or muon). A transverse mass window between 40 GeV and 100 GeV is used to select semi-leptonic W bosons.

$$m_T = \sqrt{2p_T^{lep} \cancel{E}_T (1 - \cos \Delta\phi)} \quad (10.1)$$

This also selects W bosons from direct W production, but their contribution to the control regions is small due to the b -tagging and high \cancel{E}_T requirements and is therefore estimated using Monte Carlo. The selections for the one-lepton control regions beyond the trigger and event cleaning are listed in Table 10.1. They are

CR1L_SR1	CR1L_SR2	CR1L_SR3
exactly 1 tight (signal) electron or muon		
exactly 2 jets with $p_T > 50$ GeV		≥ 3 jets with $p_T > 20$ GeV
$p_T(j_0) > 130$ GeV	$p_T(j_0, j_1) > 60$ GeV	$p_T(j_0) > 130$ GeV
$\cancel{E}_T > 90$ GeV	$\cancel{E}_T > 120$ GeV	
2 leading jets b -tagged		leading jet anti-tag, 2^{nd} - and 3^{rd} -leading jets b -tagged
$40 \text{ GeV} < m_T < 100 \text{ GeV}$		

Table 10.1: Definition of the control regions adopted for top background estimation, after single lepton trigger and common event cleaning.

similar to the signal region selections with the addition of exactly one signal lepton and somewhat lower kinematic thresholds in order to increase statistics.

Comparisons between the Standard Model expectation from Monte Carlo and the measured control region yields as a function of m_T are shown in Figure 10.1 and Table 10.2. The QCD multi-jet contribution to the leptonic control regions is estimated with the matrix method (Section 9.5) and found to be completely negligible. Table 10.2 also includes the SM background expectation after the profile likelihood fit.

10.2 Z Control Regions

For the Z control regions, two opposite-sign same-flavor signal leptons (electrons or muons) are required, and the primary figure of merit is the dilepton invariant mass:

$$m_{ll} = \sqrt{2p_1 p_2 (1 - \cos \theta)} \quad (10.2)$$

where θ is the angle between the two leptons. The control region m_{ll} window is $40 \text{ GeV} < m_{ll} < 140 \text{ GeV}$. This is further divided into a smaller window around the Z mass ($75 \text{ GeV} < m_{ll} < 105 \text{ GeV}$) and two sideband regions enriched in dileptonic $t\bar{t}$. The sidebands are used as an additional constraint on the $t\bar{t}$ normalization inside the profile likelihood fit (Section 9.4).

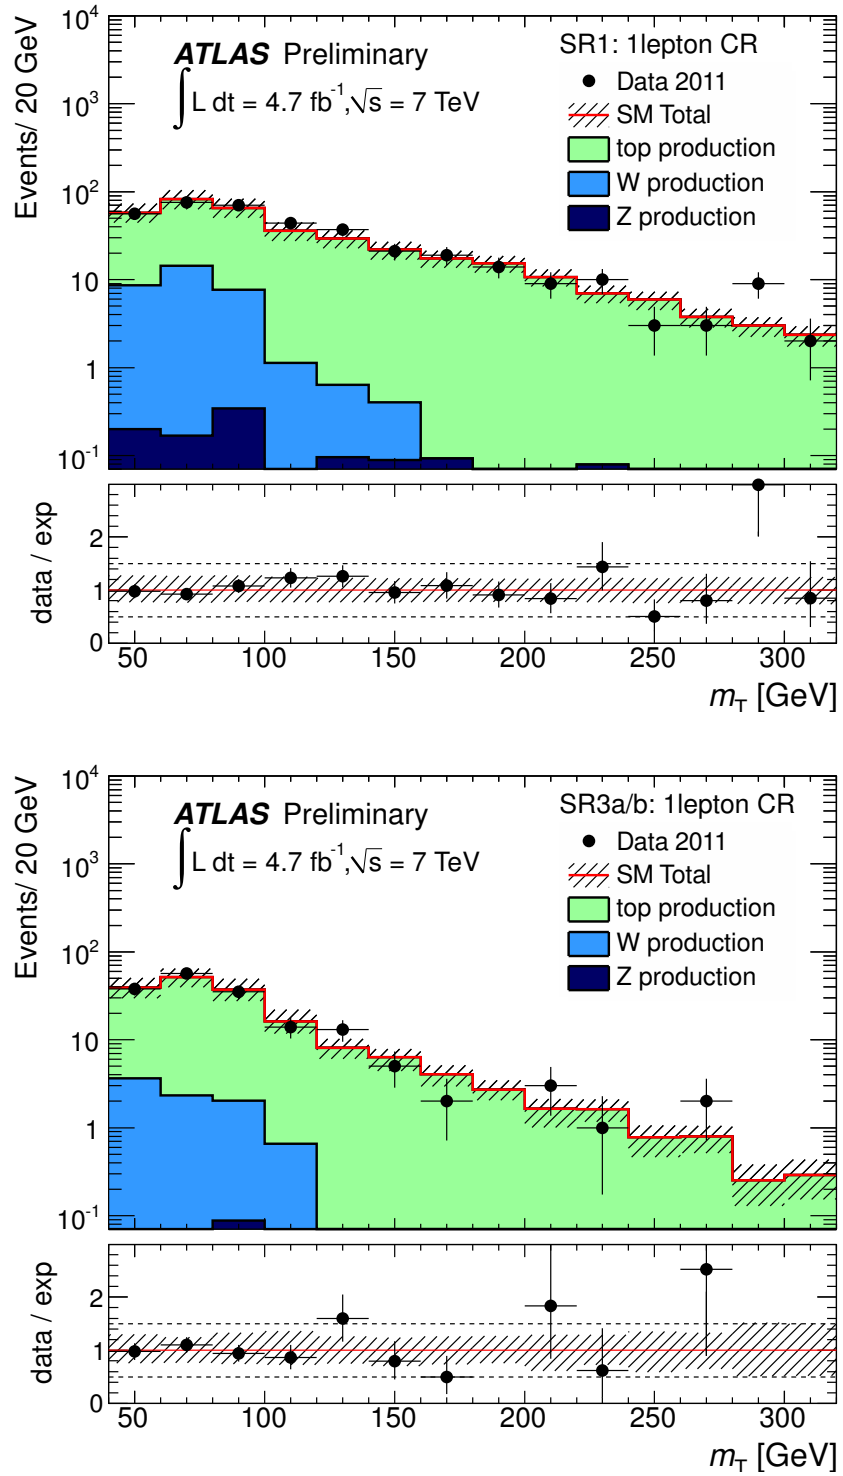


Figure 10.1: The m_T distribution before the upper selection on m_T in CR1L_SR1 (top) and in CR1L_SR3 (bottom). The slashed band includes both detector and theoretical systematic uncertainties. The SM prediction is normalized according to Monte Carlo expectations.

Region	CR1L_SR1	CR1L_SR2	CR1L_SR3
Observed	202	109	130
Fitted SM	205 ± 15	113 ± 10	129 ± 12
Fitted Top	167 ± 16	87 ± 11	114 ± 14
Fitted Z	1.2 ± 0.4	0.70 ± 0.31	0.14 ± 0.12
Fitted W	31 ± 17	20 ± 13	12 ± 8.2
Fitted Others	5.9 ± 1.4	5 ± 1.5	2.3 ± 1.4
MC SM	210	123	134
MC Top	173	92	120
MC Z	0.7	0.44	0.09
MC W	30	24.5	11.3
MC Others	6.8	6.2	2.4

Table 10.2: For each one-lepton ($t\bar{t}$) control region, the observed event yield is compared to the prediction obtained from the profile likelihood fit and to the expectation from Monte Carlo. The errors on the profile likelihood fit results include all statistical and systematic uncertainties, and thus are partially correlated.

The kinematic selections are similar to those of the signal regions and one-lepton control regions, but slightly lower to boost statistics. However, the quantity used for the \cancel{E}_T cut is not the same. As we are attempting to mimic the kinematic selections of signal regions whose Z contributions are largely from $Z \rightarrow \nu\nu$, the p_T of the two leptons are added to the \cancel{E}_T to produce the “lepton-corrected” \cancel{E}_T . The selections for the two-lepton control regions after the trigger and event cleaning are listed in Table 10.3. Note that SR1 and SR2 use the same two-lepton control region.

Comparisons between the Standard Model expectation from Monte Carlo and the measured control region yields as a function of m_{ll} are shown in Figure 10.2 and Table 10.4.

10.3 QCD Control Regions

The QCD multi-jet control regions are the same as the signal regions (Section 8.4) except that the $\Delta\phi$ cut on the leading three jets is reversed ($\Delta\phi < 0.4$) and no \cancel{E}_T/m_{eff} , m_{CT} , or $H_{T,n}$ cuts are applied. Thus, all SR1 variants share a control

CR2L_SR1/CR2L_SR2	CR2L_SR3
exactly 2 tight (signal) leptons, e^-e^+ or $\mu^-\mu^+$	
exactly 2 jets with $p_T > 50$ GeV)	≥ 3 jets with $p_T > 20$ GeV
lepton-corrected $\cancel{E}_T > 90$ GeV	
2 leading jets b -tagged	leading jet anti-tag, 2 nd - and 3 rd -leading jets b -tagged
-	leading b -tagged jet $p_T < 100$ GeV
$40 \text{ GeV} < m_{ll} < 100 \text{ GeV}$	

Table 10.3: Definition of the control regions adopted for the Z background estimation, after single lepton trigger and common event cleaning.

Region	CR2L_SR1/CR2L_SR2	CR2L_SR3
Observed	211	67
Fitted SM	208 ± 16	68 ± 8.3
Fitted Top	141 ± 15	48 ± 7
Fitted Z	63 ± 14	19 ± 8
Fitted W	0	0
Fitted Others	2.6 ± 0.5	1.2 ± 0.2
MC SM	177	66
MC Top	137	52
MC Z	37	13
MC W	0	0
MC Others	2.9	1.2

Table 10.4: For each two-lepton (Z and $t\bar{t}$) control region, the observed event yield is compared to the prediction obtained from the profile likelihood fit and to the expectation from Monte Carlo. The errors on the profile likelihood fit results include all statistical and systematic uncertainties considered, and thus are partially correlated.

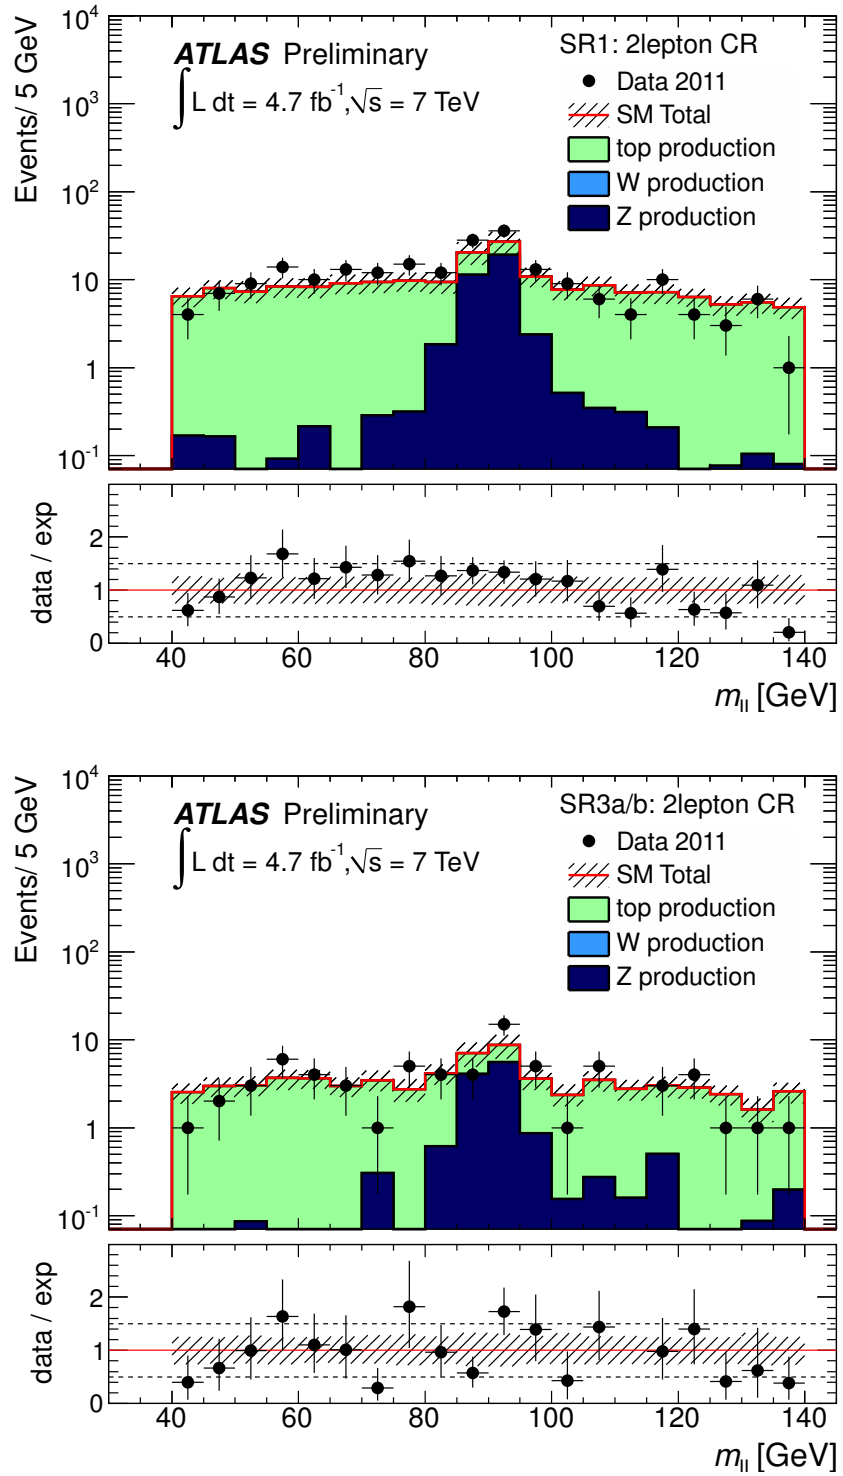


Figure 10.2: The dilepton invariant mass (m_{ll}) distribution before the upper selection on m_{ll} in CR2L-SR1/CR2L-SR2 (top) and in CR2L-SR3 (bottom). The slashed band includes both detector and theoretical systematic uncertainties. The SM prediction is normalized according to Monte Carlo expectations.

region. Also, both SR3a and SR3b share the region corresponding to SR3a. A jet-smearred data sample (Section 9.6) with two b -tagged jets is normalized to these control regions (a separate normalization is used for SR1, SR2, and SR3). The shapes of the normalized jet-smearred sample match the data well (Figure 10.3).

10.4 Results

The final background estimate is obtained from a profile likelihood fit as described in Section 9.4. No shape information is used—each region was considered a single bin, with the $t\bar{t}$ normalization, Z +jets normalization, and systematic uncertainty parameters floating. The QCD multi-jet contribution is fixed. The net result of the fit is a significant upward scaling of the Z -jets background over the Monte Carlo-based estimate. This effect was also seen in the previous version of the analysis [2]. The $t\bar{t}$ estimate is largely unaffected by the fit. Both the fit and Monte Carlo-based estimates, as well as the measured signal region yields in data, are shown in Table 10.5.

Signal region yields in data are well-within one standard deviation of the background estimates, with the exception of SR2 which has a deficit in data of approximately one standard deviation. All experimental, theoretical, and statistical uncertainties are combined into the errors shown in Table 10.5. Comparisons of representative signal region distributions between data and Monte Carlo scaled to the fitted yields are shown in Figures 10.4 and 10.5. The shapes of the distributions seem to agree to the extent that can be expected by statistics.


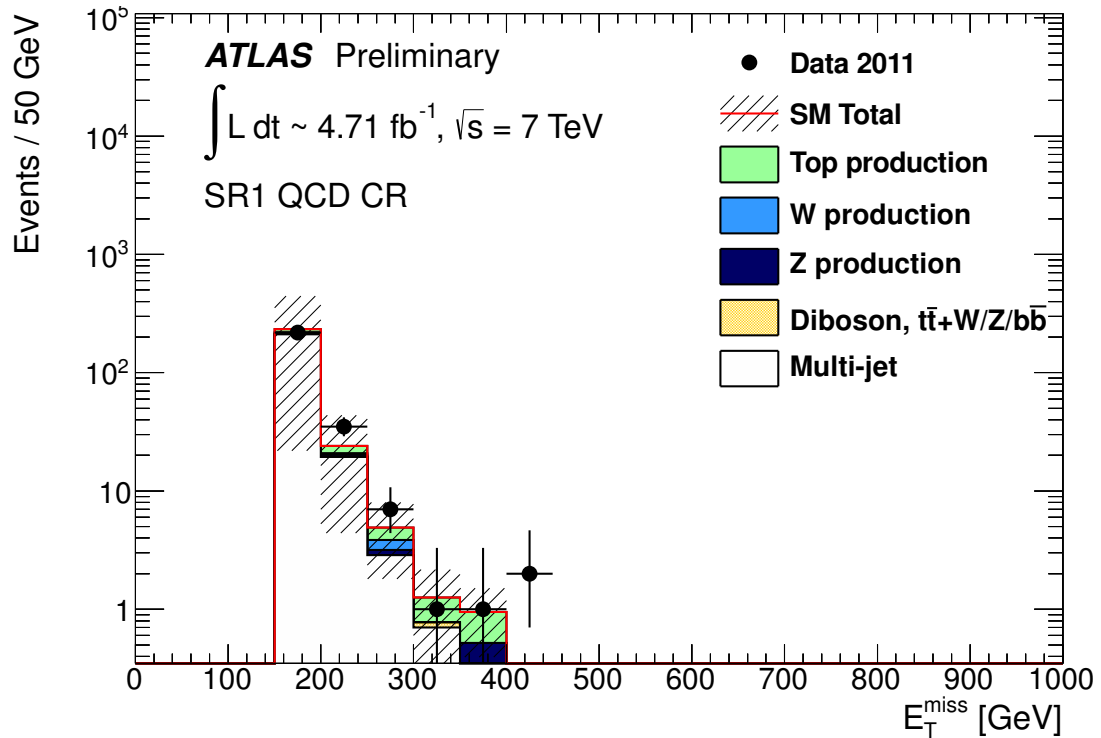


Figure 10.3: The \cancel{E}_T distribution in the SR1 QCD control region (top) and the leading jet p_T distribution in the SR3 QCD control region (bottom). The slashed band includes both detector and theoretical systematic uncertainties.

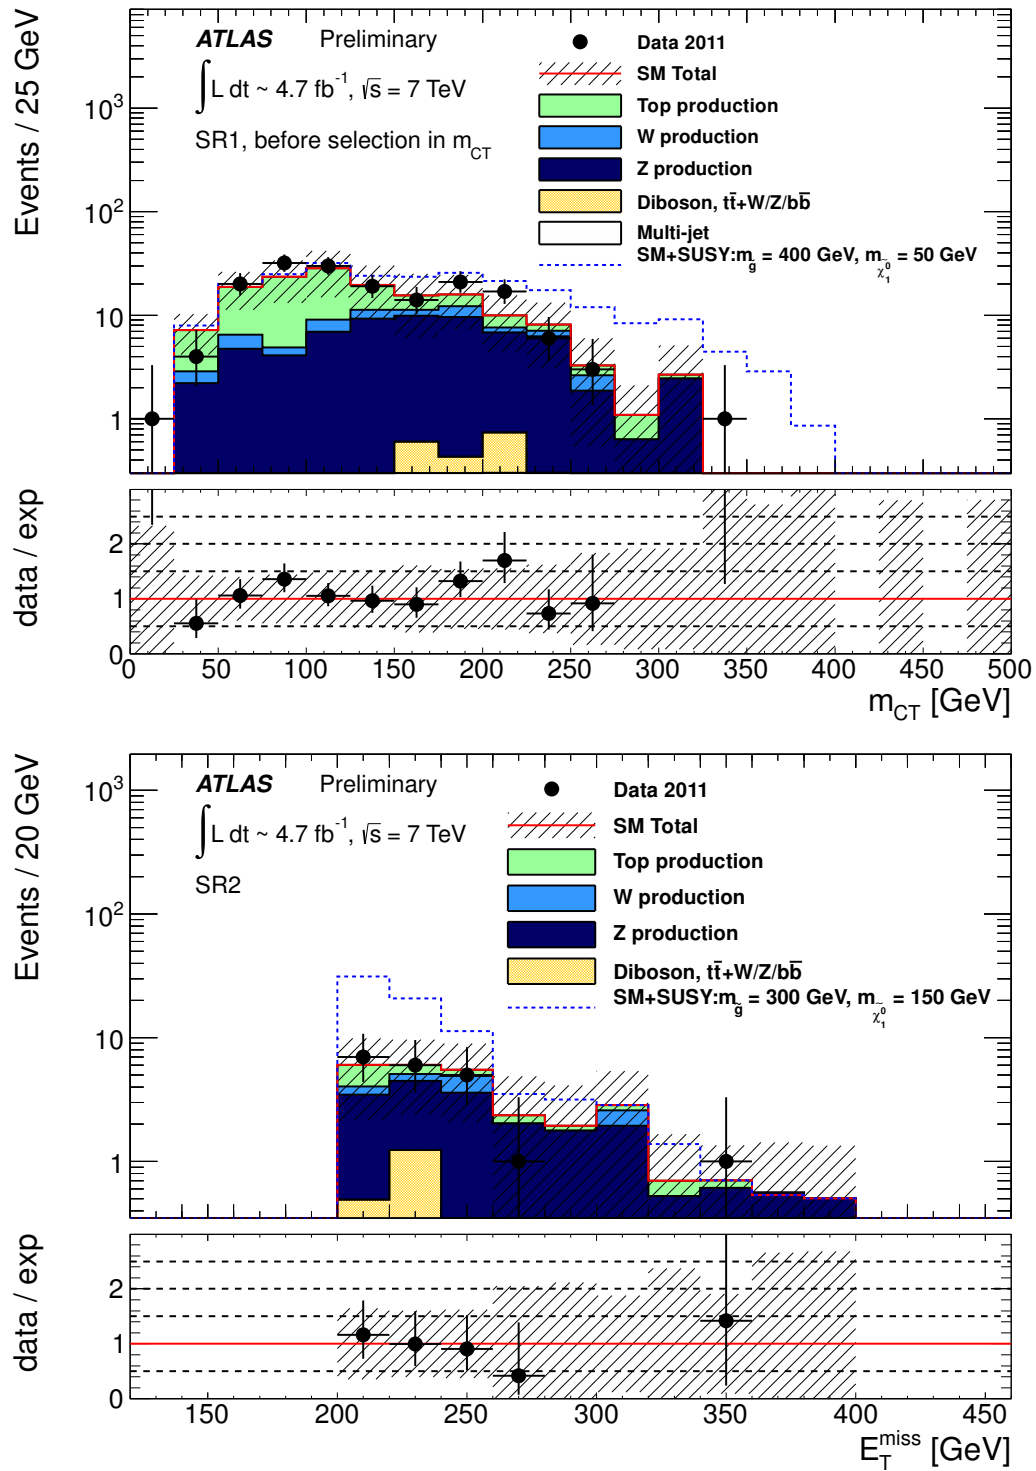


Figure 10.4: Top: m_{CT} distribution in SR1 before the selection on m_{CT} . Bottom: E_T^{miss} distribution in SR2. The slashed band includes both detector and theoretical systematic uncertainties. All Monte Carlo distributions are normalized to the results of the profile likelihood fit.

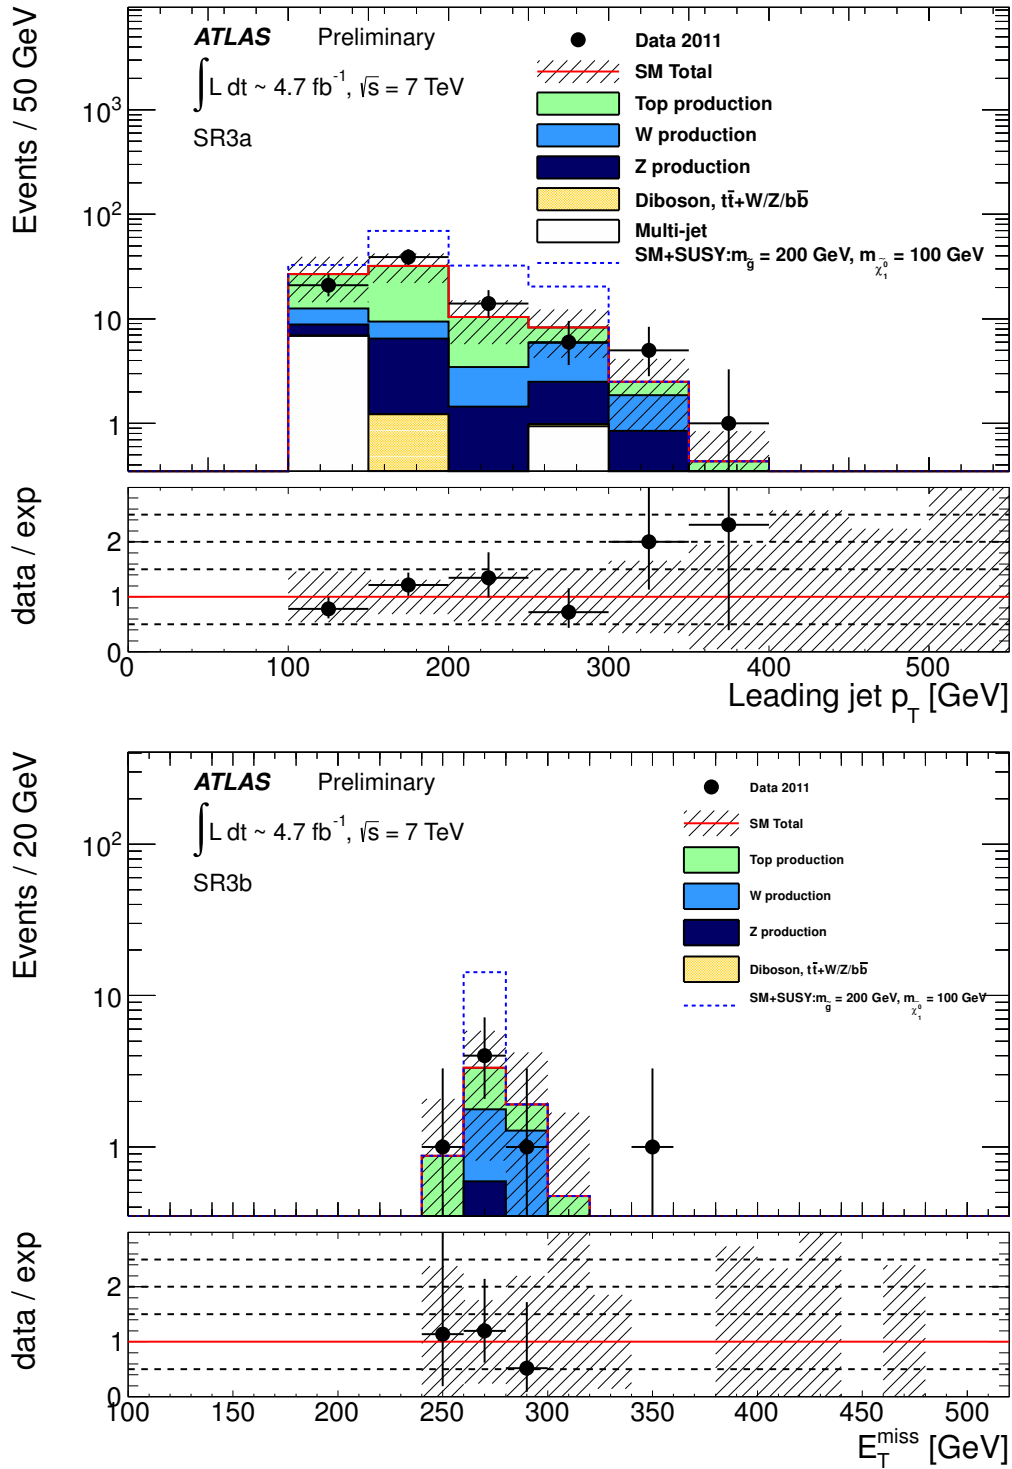


Figure 10.5: Top: Leading jet p_T distribution in SR3a. Bottom: E_T^{miss} distribution in SR3b. The slashed band includes both detector and theoretical systematic uncertainties. All Monte Carlo distributions are normalized to the results of the profile likelihood fit.

Region	SR1			SR2	SR3	
m_{CT} (SR1)	> 150 GeV	> 200 GeV	> 250 GeV		SR3a	SR3b
Observed	62	27	4	20	86	7
Fitted SM	56 ± 11	24.9 ± 5.8	6.9 ± 2.3	27 ± 7	81 ± 14	8.0 ± 2.7
Fitted Top	13 ± 3	5 ± 1	1.5 ± 0.5	4.8 ± 1.2	47.8 ± 9.5	4.1 ± 1.2
Fitted Z	35 ± 10	16 ± 5	4.1 ± 1.7	17 ± 6	11.1 ± 4.5	1.3 ± 0.9
Fitted W	6.2 ± 3.8	2.3 ± 1.1	0.8 ± 0.6	3.1 ± 1.5	13 ± 8	2.4 ± 2.0
Fitted Others	2.2 ± 0.6	1.2 ± 0.4	0.5 ± 0.2	2.5 ± 0.8	1.6 ± 0.3	0.2 ± 0.1
QCD Multi-jet	0.5 ± 0.5	0.4 ± 0.4	0.07 ± 0.07	0	7.9 ± 4.5	0
MC SM	44 ± 17	20 ± 8	6.7 ± 3.0	22 ± 9	79 ± 21	7.5 ± 2.5
MC Top	13 ± 5	6.1 ± 2.9	1.5 ± 0.8	4.7 ± 2.2	51 ± 14	4.3 ± 1.6
MC Z	22 ± 15	10 ± 7	3 ± 2	11 ± 7	7.3 ± 4.9	0.8 ± 0.8
MC W	6.1 ± 4.0	2.2 ± 1.5	1 ± 1	3.5 ± 2.1	11 ± 7	2.2 ± 2.0
MC Others	2.5 ± 1.0	1.4 ± 0.8	0.6 ± 0.4	2.9 ± 0.7	1.5 ± 0.4	0.2 ± 0.1
QCD Multi-jet	0.5 ± 0.5	0.4 ± 0.4	0.07 ± 0.07	0	7.9 ± 4.5	0

Table 10.5: For each signal region, the observed event yield is compared with the prediction obtained from the profile likelihood fit. The contribution of each SM process to each signal region from Monte Carlo Simulation is also shown. The errors on the profile likelihood fit results include all statistical and systematic uncertainties, and thus are partially correlated.

Chapter 11

Gluino-mediated Sbottom/Stop Search

This iteration of the gluino-mediated sbottom and stop analysis represents a vast improvement over the previous version [3] due primarily to the move to signal regions with three b -tags instead of two. $t\bar{t}$, the primary background, only contains two real b -jets. The gluino-mediated signals, however, contain four. In order to produce three b -tagged jets in a basic $t\bar{t}$ event, either a charm quark hadron or a tau lepton must be tagged, but these have much lower tagging efficiency than do real b -jets.

The difficulty with three b -tags is related to its advantage—there is very little Standard Model background. This significantly complicates data-driven background estimation techniques. Monte Carlo-based estimates, meanwhile, incur large b -tagging systematics and often lack statistics. These problems were thought so insurmountable that this analysis represents, to our knowledge, the first three b -tag search ever published.

Part of what makes three b -tags possible here is the 4.71 fb^{-1} integrated luminosity of the data. Relative to the 2.05 fb^{-1} analysis, this increases the expected number of background events in the signal regions and enables a robust background estimate. While this increase makes zero-lepton three-tag signal regions viable, the same cannot be said for one-lepton three-tag control regions due to low statistics. In the gluino-mediated stop model, these control regions would also have significant

signal contamination. A three-tag control region is simply not an option.

Our solution is to essentially repurpose the zero-lepton two-tag signal regions from the 2.05 fb^{-1} analysis, already known to be dominated by $t\bar{t}$, as the $t\bar{t}$ control regions for the three-tag analysis. The downside of this approach relative to using hypothetical control region with the same b -tagged jet multiplicity is that the cancellation of systematic uncertainties (Section 9.3) is incomplete. In particular, residual b -tagging uncertainty from the third tag as well as jet energy scale uncertainty from the b -jet energy scale remain, and account for the majority of the detector-related contribution to the total uncertainty of the background estimate.

As discussed previously, the background estimates for the gluino-mediated third-generation squark production signal regions (Section 8.5) are completely dominated by $t\bar{t}$, with the third tag coming from a charm quark or τ lepton from W decay, or occasionally from $t\bar{t} + c\bar{c}/b\bar{b}$ in the high- m_{eff} regions. Therefore, only top and QCD control regions are necessary, and Monte Carlo-based predictions are used for all other backgrounds ($W/Z + b\bar{b}$, single top, and $t\bar{t} + W/Z$). As the two-tag $t\bar{t}$ control regions are zero-lepton regions, the QCD multi-jet background is small but non-negligible. QCD control regions are therefore needed for both the signal regions and the top control regions.

11.1 Top Control Regions

The top control regions are the same as the signal regions (Section 8.5), with three exceptions. The \cancel{E}_T and m_{eff} cuts are uniformly 160 GeV and 500 GeV, respectively, to minimize signal contamination, and exactly two b -tagged jets are required instead of three or more. This means that SR4-L and SR4-M share a top control region. The control region selections are summarized in Table 11.1. It should be noted that $t\bar{t} + b\bar{b}$ is separated from the other $t\bar{t}$ and estimated using a dedicated Monte Carlo sample. This is done because the three-tag to two-tag ratio in Monte Carlo differs significantly between $t\bar{t} + b\bar{b}$, which has four real b -jets.

Comparisons between the Standard Model expectation from Monte Carlo and jet smearing and the measured control region yields are shown in Figures 11.1 and 11.2

lepton (e/μ) veto, leading jet $p_T > 130$ GeV, exactly 2 b -jets, $\frac{\cancel{E}_T}{m_{eff}} > 0.2$, $\Delta\phi > 0.4$ for leading 3 jets, $\cancel{E}_T > 160$ GeV, $m_{eff} > 500$ GeV			
Region	n jets	b -tagging efficiency	corresponding SR
CR4-60	≥ 4	60%	SR4-L, SR4-M
CR4-70	≥ 4	70%	SR4-T
CR6-70	≥ 6	70%	SR6-L
CR6-75	≥ 6	75%	SR6-T

Table 11.1: Definition of the four control regions used to estimate the $t\bar{t}$ background, after the jet + \cancel{E}_T trigger and common event cleaning.

and Table 11.2.

Region	$t\bar{t}$ +jets	Others	Total SM	Observed
CR4-60	329 ± 92	66 ± 26	395 ± 115	402
CR4-70	489 ± 125	102 ± 37	590 ± 160	515
CR6-70	38 ± 11	7 ± 3	45 ± 13	46
CR6-75	40 ± 12	10 ± 4	50 ± 15	52

Table 11.2: Expected Standard Model event yield and observed data events in the four $t\bar{t}$ control regions. The multi-jet estimate, included in the category “Others” and accounting for $\approx 20\%$ of the total yield, was estimated with the jet-smearing technique. All other background yields were taken from Monte Carlo. The errors include all detector-related and theoretical systematic uncertainties.

11.2 QCD Control Regions

QCD multi-jet control regions are defined for each signal region and top control region. The $\Delta\phi$ cuts are reversed ($\Delta\phi < 0.4$), and the $\frac{\cancel{E}_T}{m_{eff}}$ cuts removed. All m_{eff} and \cancel{E}_T cuts are lowered to 500 GeV and 160 GeV, respectively, meaning SR4-L and SR4-M also share a QCD control region and each three-tag control region has a two-tag counterpart. Jet-smearred data samples (Section 9.6) with two and three b -tagged jets are normalized to these control regions separately. The shapes of the normalized jet-smearred samples match the data well (Figure 11.3).

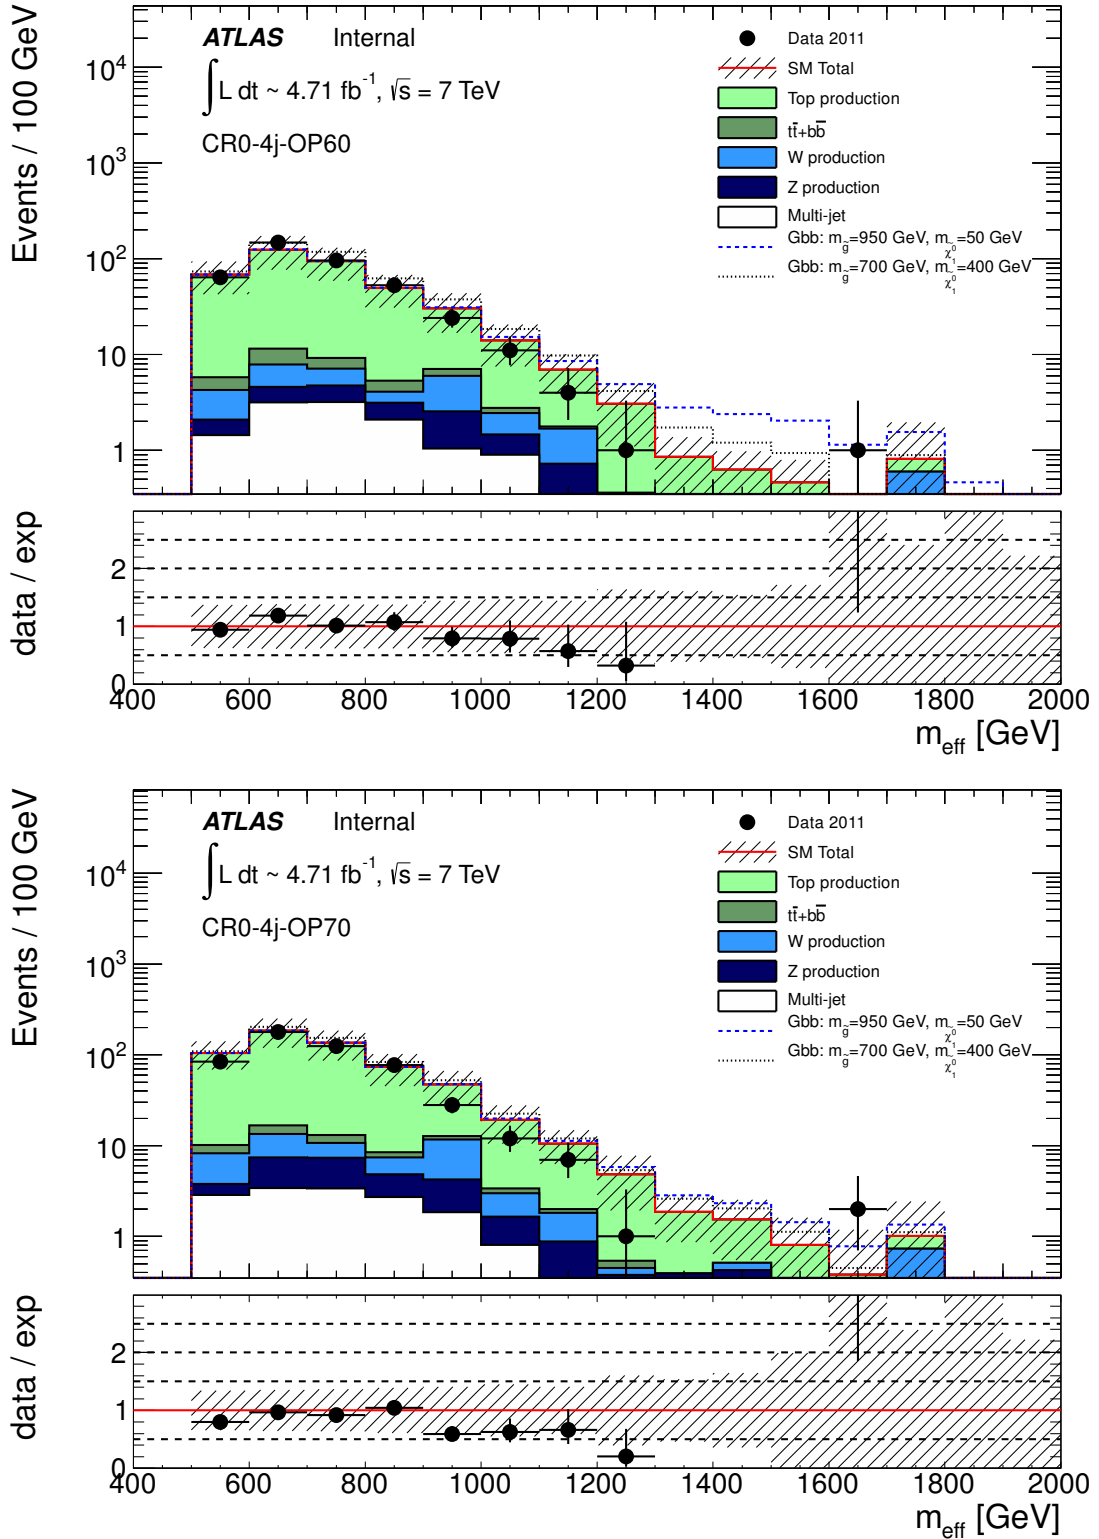


Figure 11.1: Comparison of the distributions of the effective mass (m_{eff}) between data and Monte Carlo expectation in the two four-jet $t\bar{t}$ control regions CR4-60 (top) and CR4-70 (bottom). The slashed band includes all detector-related and theoretical systematic uncertainties.

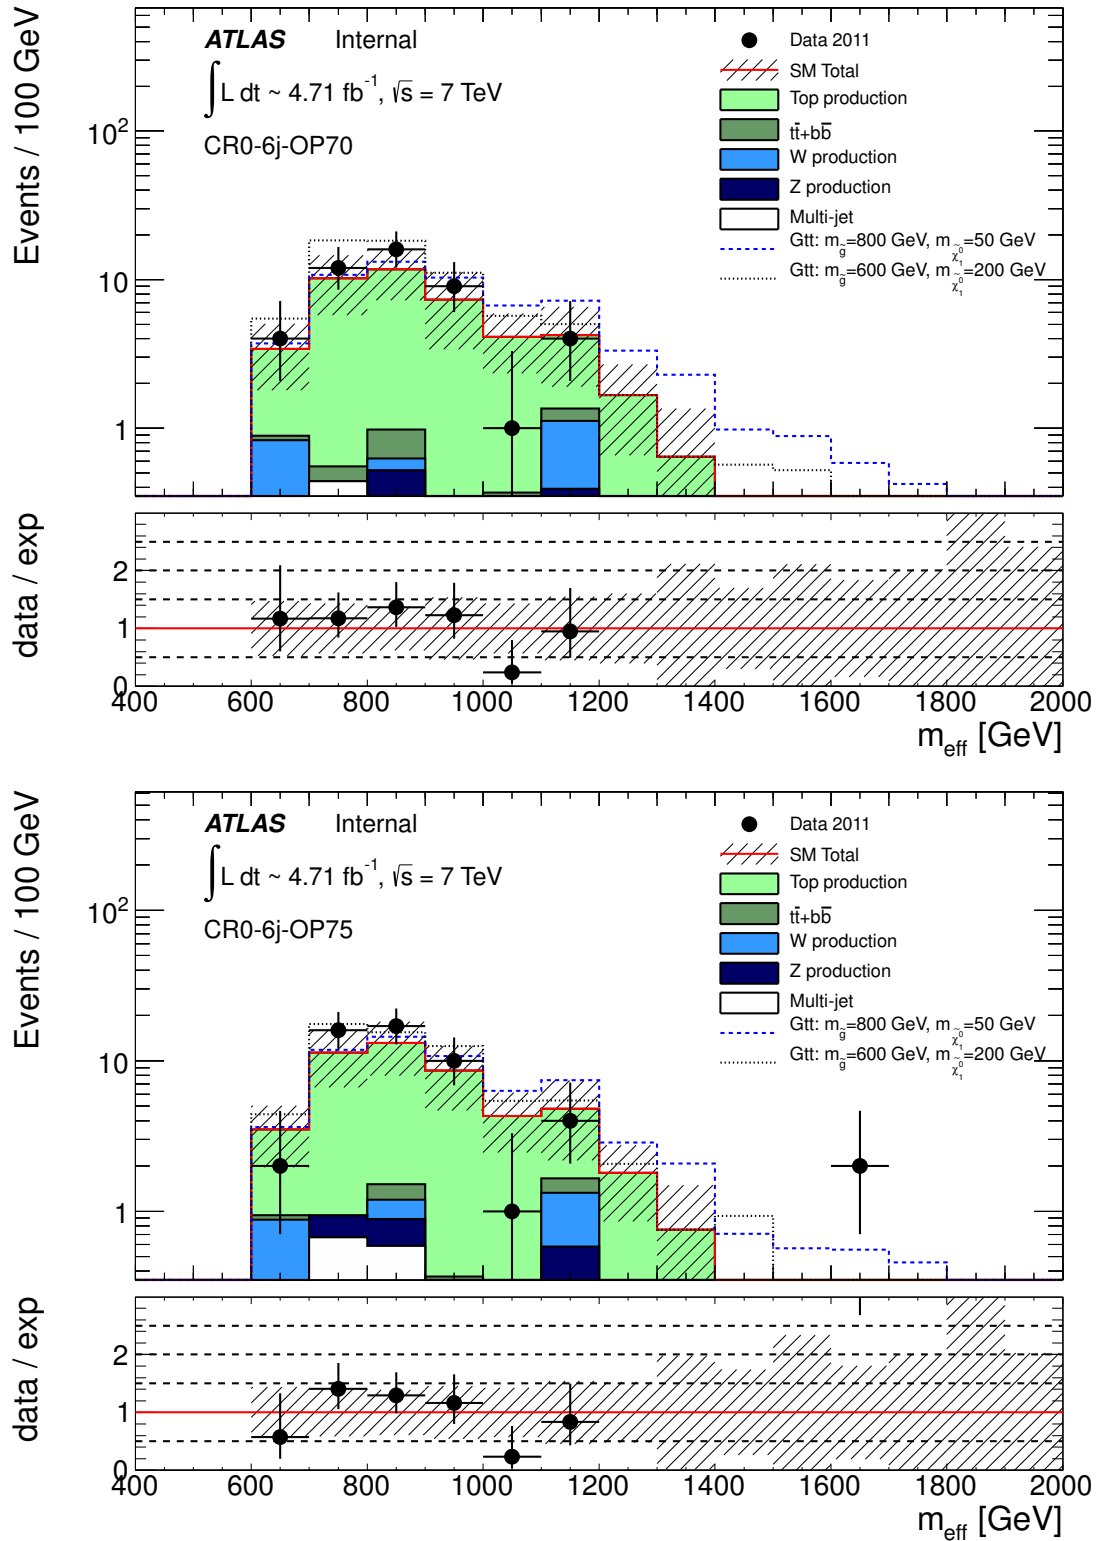


Figure 11.2: Comparison of the distributions of the effective mass (m_{eff}) between data and Monte Carlo expectation in the two six-jet $t\bar{t}$ control regions CR6-70 (top) and CR6-75 (bottom). The slashed band includes all detector-related and theoretical systematic uncertainties.

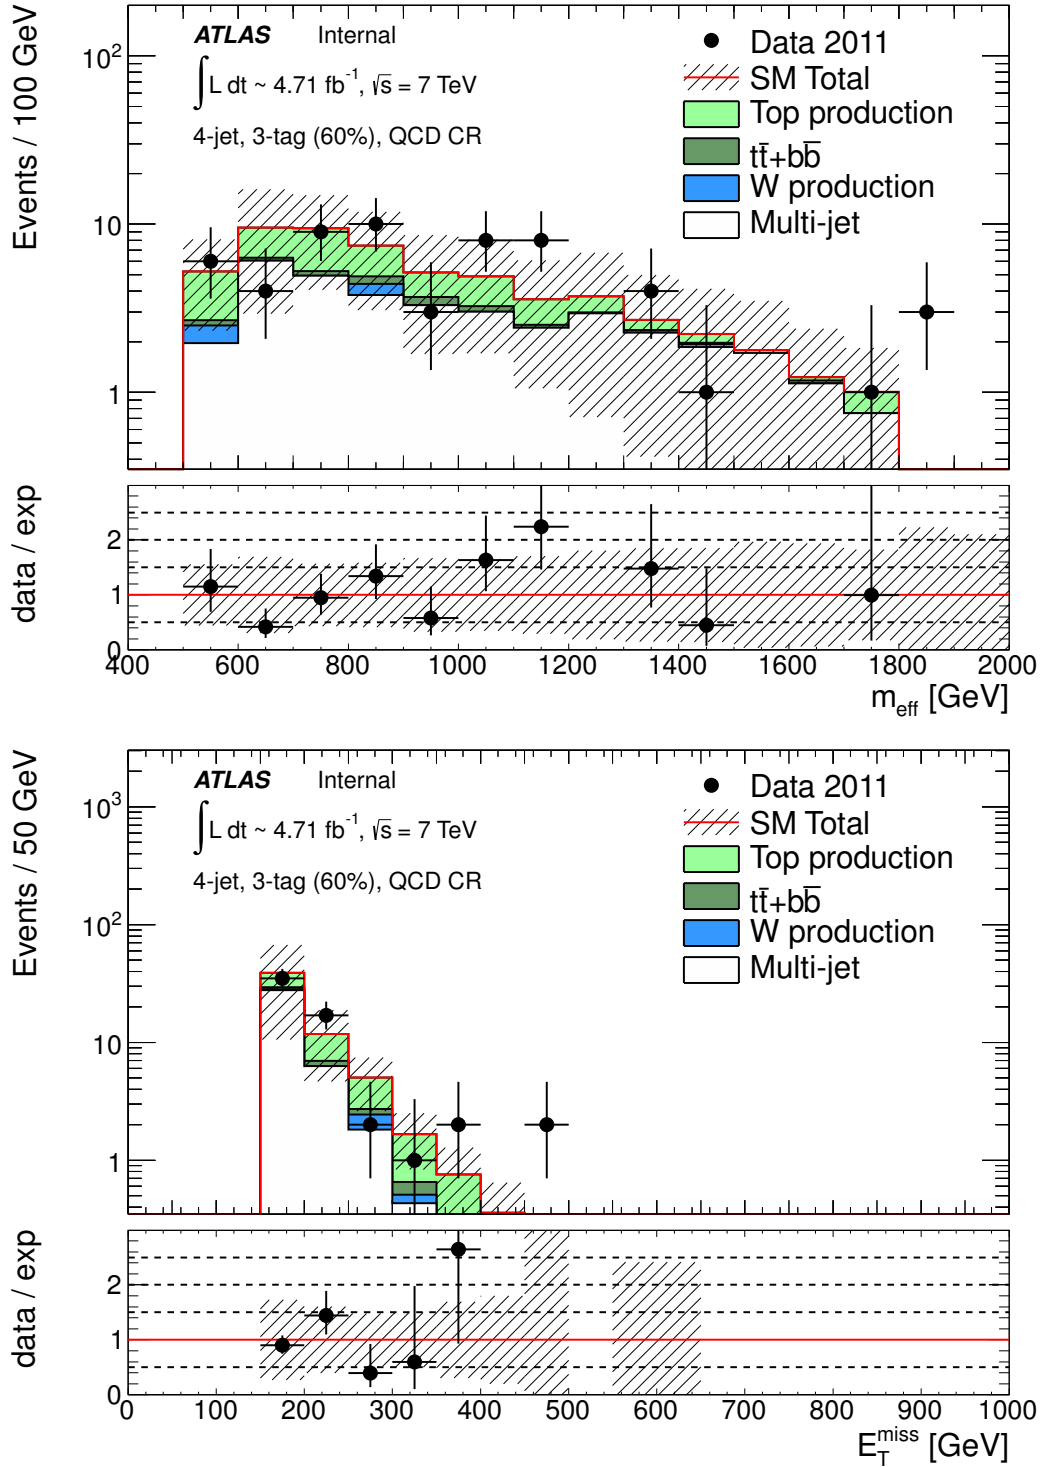


Figure 11.3: Comparison of the distributions of the effective mass (m_{eff} , top) and missing transverse momentum (E_T^{miss} , bottom) between the data and Monte Carlo expectation in QCD multi-jet control region corresponding to three b -tags at 60% b -tagging efficiency. The slashed band includes all detector-related and theoretical systematic uncertainties.

11.3 Results

The final background estimation is obtained via a profile likelihood fit as described in Section 9.4. No shape information was used—each region was considered a single bin. The $t\bar{t}$ normalization, Z +jets normalization, and systematic uncertainty parameters were free to float. The fit has almost zero net effect and is completely consistent with the Monte Carlo-based estimate, albeit with lower errors. Both the fit-based and Monte Carlo-based estimates, as well as the measured signal region yields in data are shown in Table 11.3.

Region	$t\bar{t}$ +jets (MC)	Others	Total SM	Observed
SR4-L	33.3 ± 7.9 (32.6 ± 15.4)	11.1 ± 4.9	44.4 ± 10.0	45
SR4-M	16.4 ± 4.1 (16.1 ± 8.4)	6.6 ± 2.9	23.0 ± 5.4	14
SR4-T	9.7 ± 2.1 (11.4 ± 5.4)	3.8 ± 1.6	13.3 ± 2.6	10
SR6-L	10.3 ± 3.3 (10.0 ± 6.2)	2.4 ± 1.4	12.7 ± 3.6	12
SR6-T	8.3 ± 2.4 (7.9 ± 5.3)	1.6 ± 1.1	9.9 ± 2.6	8

Table 11.3: Comparison between the results of the profile likelihood fit and the numbers of observed events in the five signal regions. The $t\bar{t}$ event yield predicted by the MC simulation is quoted in parentheses, and all other non-QCD backgrounds are taken from the Monte Carlo in both cases. QCD multi-jet events contribute less than 5% of the total background in all signal regions. The errors include all experimental, theoretical, and statistical uncertainties.

The SR4-L and both SR6 observed signal region yields agree well with the background estimates. The SR4-M and SR4-T yields in data are slightly low compared to the background estimates, but in agreement at the level of one standard deviation or better when the Poisson error on the data is taken into account. All experimental, theoretical, and statistical uncertainties are combined into the errors shown in Table 11.3. Comparisons of representative signal region distributions between data and Monte Carlo scaled to the fitted yields are shown in Figures 11.4 and 11.5. The

shapes of the distributions also appear to agree, with the caveat that in the several of the signal regions the number of events is too low for reliable comparison.

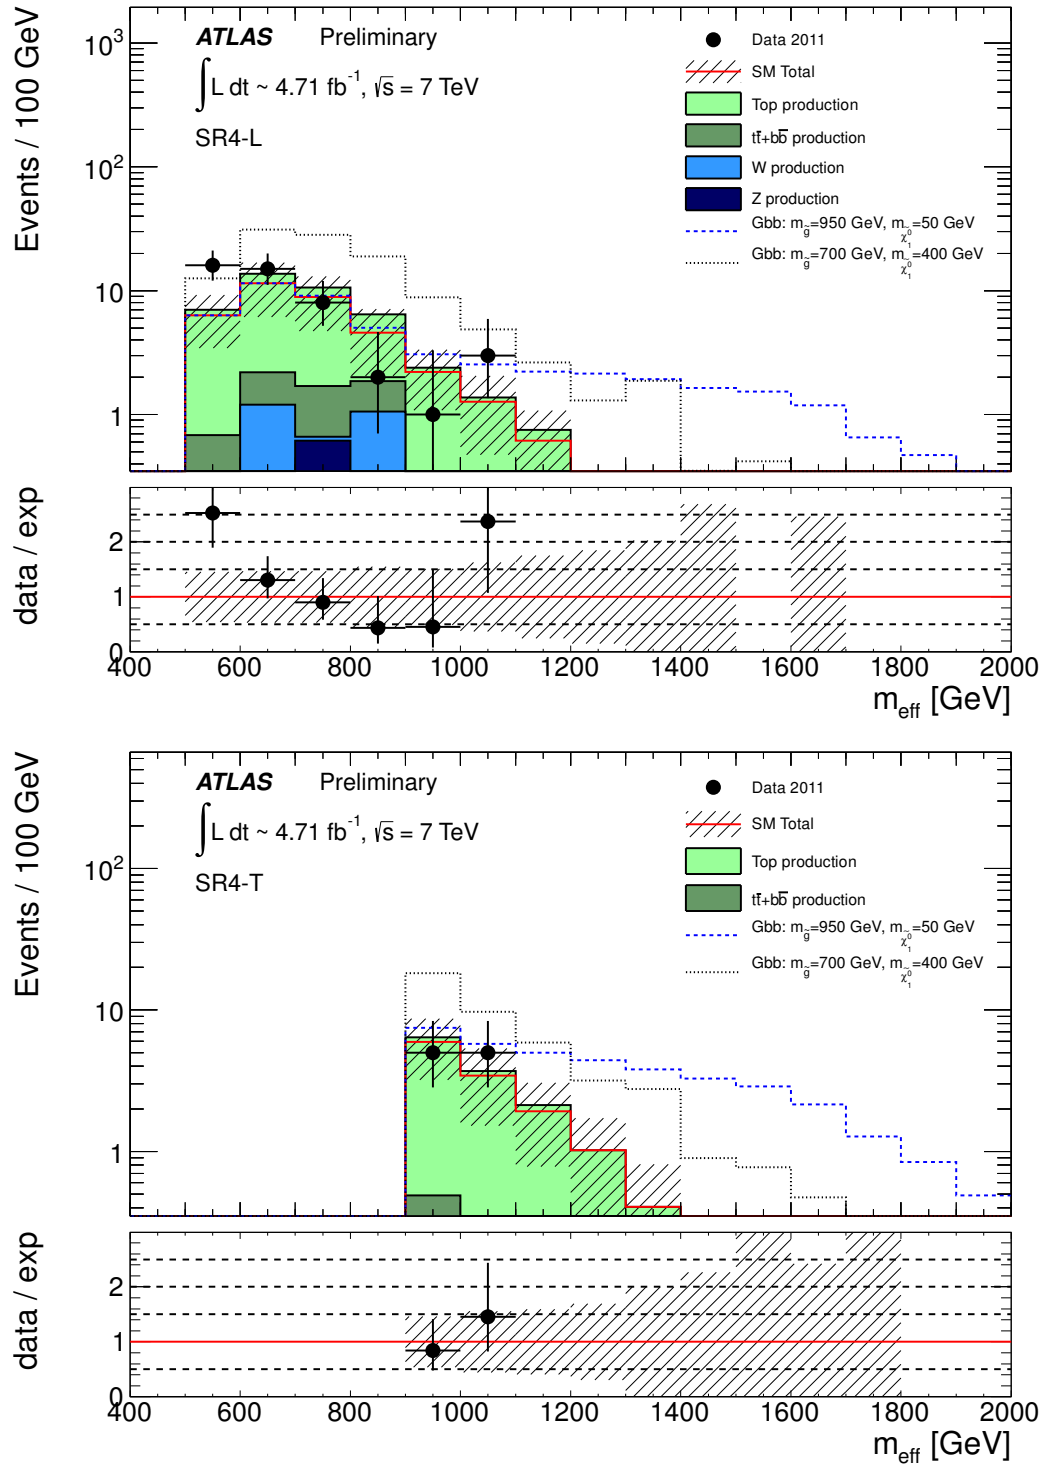


Figure 11.4: Comparison of the distributions of the effective mass (m_{eff}) between data and Standard Model expectation in the four-jet signal regions SR4-L and SR4-M (top) and SR4-T (bottom). The Monte Carlo distributions have been scaled to the yields from the profile likelihood fit. The slashed band includes all detector-related and theoretical systematic uncertainties.

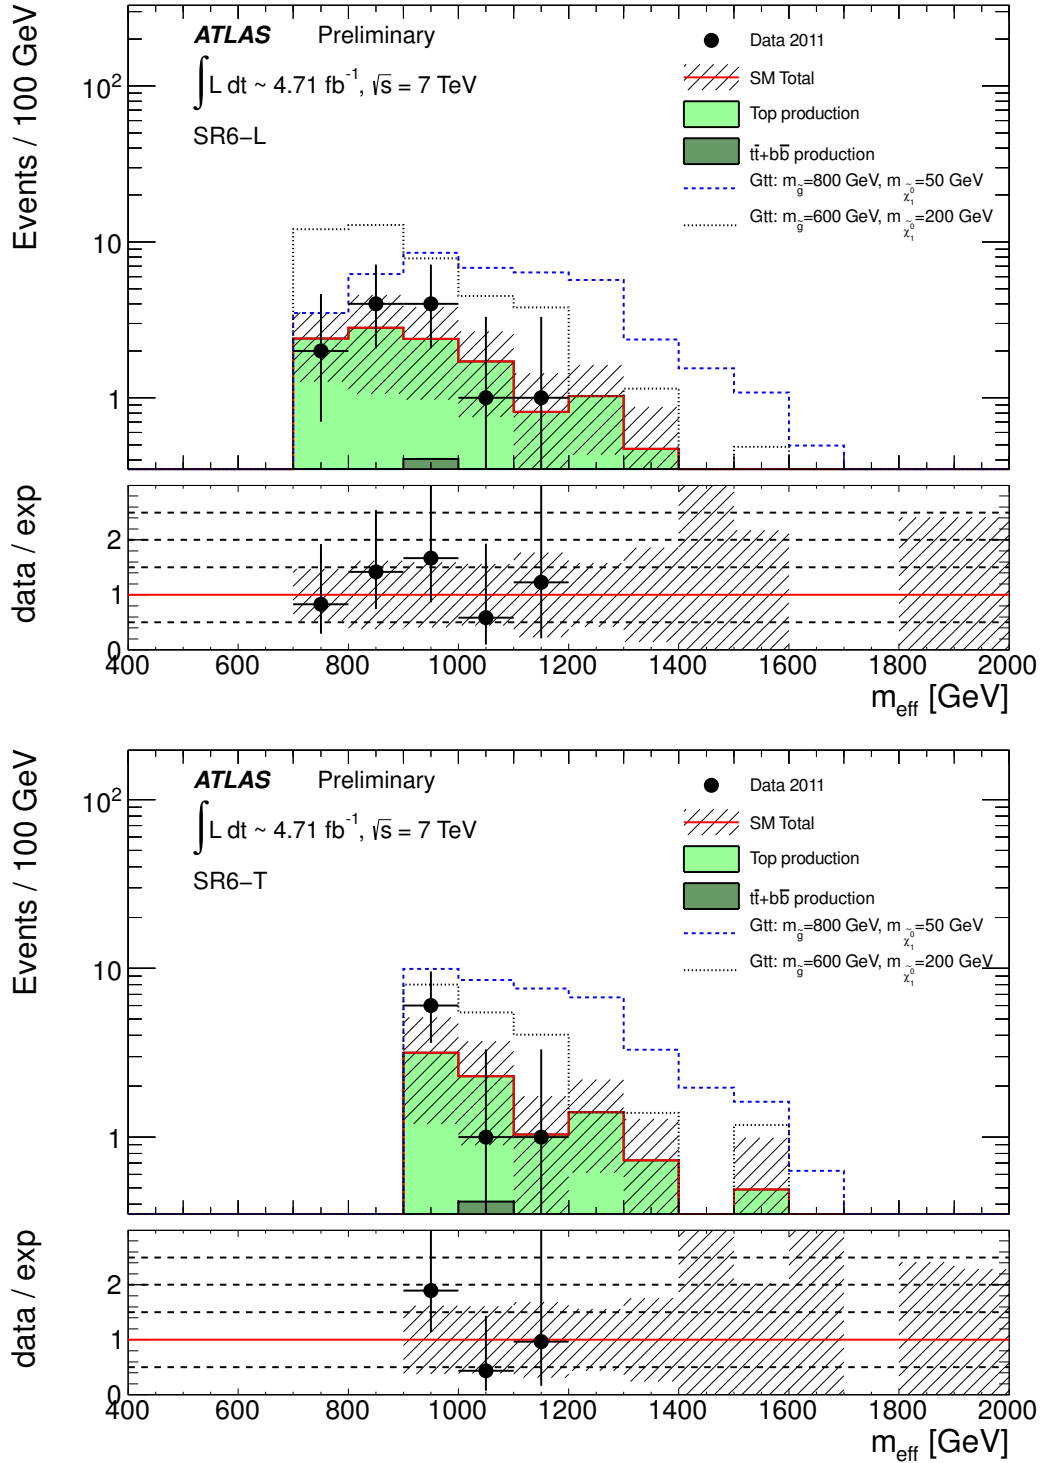


Figure 11.5: Comparison of the distributions of the effective mass (m_{eff}) between data and Standard Model expectation in the two six-jet signal regions SR6-L (top) and SR6-T (bottom). The Monte Carlo distributions have been scaled to the yields from the profile likelihood fit. The slashed band includes all detector-related and theoretical systematic uncertainties.

Chapter 12

Interpretation of Results

No significant excess is seen in the signal regions of either the direct sbottom or gluino-mediated sbottom/stop analyses (Section 10.4 and 11.3). We quantify this agreement via limits on the contribution of new physics to the signal regions.

12.1 Profile Likelihood Ratio Test and Limits

Limits are set using an extension of the profile likelihood method described in Section 9.4. The full likelihood function for the background estimate (Equation 9.7) can be written:

$$L(b, \theta | n, \theta^0) = P_{CR}(n | b, \theta) \cdot P_{Sys}(\theta^0 | \theta) \quad (12.1)$$

where b and θ represent the background estimates and systematic uncertainties (nuisance parameters), respectively. n are the measured yields in data, and θ^0 is the value around which θ varies (zero). P_{CR} is a product of control region Poisson distributions, and P_{Sys} a product of normal (Gaussian) distributions, one for each systematic uncertainty parameter. This function should be read as the likelihood that b and θ represent the true background and systematic parameters, given n and θ^0 fixed.

For the limits, the likelihood function is expanded to:

$$L(\mu, b, \theta | n, \theta^0) = P_{SR}(n | \mu, b, \theta) \cdot P_{CR}(n | \mu, b, \theta) \cdot P_{Sys}(\theta^0 | \theta) \quad (12.2)$$

where μ is now the *signal strength*, with $\mu = 0$ the background-only hypothesis and $\mu = 1$ the nominal signal yield. P_{SR} has the form of a Poisson probability distribution for the signal region, and the μ -dependence of P_{CR} is due to signal contamination of the control regions and is small by design. We will drop the dependence on n and θ^0 going forward as these are fixed. The test static used is the profile log likelihood ratio:

$$\Lambda(\mu) \equiv -2 \ln \left(\frac{L(\mu, \hat{\hat{b}}, \hat{\hat{\theta}})}{L(\hat{\mu}, \hat{b}, \hat{\theta})} \right) \quad (12.3)$$

The parameters $\hat{\mu}$, \hat{b} , and $\hat{\theta}$ are the unconditional maximum likelihood estimators: they are the choices for μ , b , and θ which maximize L . The parameters $\hat{\hat{b}}$ and $\hat{\hat{\theta}}$ are the conditional maximum likelihood estimators, which are the values of b and θ which maximize L for a given value of μ . For $\hat{\mu} < 0$ $\hat{\mu} = 0$ is substituted, as the addition of signal can only increase the overall yield. Also, for $\hat{\mu} > \mu$ $\Lambda(\mu) = 0$, as we do not want to regard data with $\hat{\mu} > \mu$ as being incompatible with μ . By the Neyman-Pearson lemma, the likelihood ratio test is the most powerful test to distinguish between two hypotheses.

The p -value is the probability of the data to be equally or more incompatible with signal strength μ ($\mu = 1$ for the model-dependent limits). It is found by integrating the probability distribution of the test statistic from $\Lambda(\mu)$ to ∞ . This probability distribution is constructed assuming that the data have signal strength $\mu' = \mu$. While there exist asymptotic formulas for this probability distribution with large ranges of validity [164], we use toy Monte Carlo. The toys are generated by “reversing” the likelihood function and writing it as a probability density using μ' and the maximum likelihood estimators of b and θ (assuming μ') as the parameters. The probability distribution of $\Lambda(\mu)$ can then be generated from this Monte Carlo dataset.

If the p -value or something similar to it (Section 12.2) is less than a given threshold, the data is said to exclude the μ model at a confidence level corresponding to the

threshold. In this case, the chosen p -value threshold is 0.05 for 95% confidence level.

So-called model-independent limits are computed similarly, but in this case μ is varied until the p -value matches the threshold. The signal contributions to the control regions are assumed to be zero. This special value of μ can be translated into a maximum contribution to the signal region (maximum number of signal events) at 95% confidence level, or equivalently, a maximum visible cross section:

$$\sigma_{vis} = \sigma \cdot BR \cdot \alpha \cdot \epsilon \quad (12.4)$$

σ is the production cross section, BR is the branching ratio to the final state of interest, α is the signal acceptance (fiducial volume, cuts), and ϵ is the signal efficiency (b -tagging, lepton identification, resolution effects, etc.).

12.2 CL_s

In these analyses, it is the CL_s value [177] which is compared to the 95% confidence level exclusion threshold, not the p -value. The CL_s value is defined as:

$$CL_s \equiv \frac{p_{s+b}}{1 - p_b} \quad (12.5)$$

p_{s+b} is the p -value discussed in Section 12.1. p_b , on the other hand, is defined as the integral from 0 to $\Lambda(\mu)$ of the $\Lambda(\mu)$ distribution produced assuming $\mu' = 0$. The quantity $1 - p_b$ is therefore a measure of the incompatibility of the data with the background-only hypothesis due to under-fluctuations. The motivation behind this normalization of the signal + background hypothesis p -value (p_{s+b}) is that if the background fluctuates down, with the unmodified p -value one can exclude signal hypotheses for which there is *no sensitivity*. The CL_s method prevents this by inflating the CL_s value when $1 - p_b$ is small. It also results in more conservative limits generally.

12.3 Model-dependent Limits

Model-dependent limits are set using the production cross sections shown in Figure 4.8 and assuming 100% branching ratio to the final states considered in Chapter 4. Signal regions without overlap such as SR1/SR2 with SR3 were combined for limit-setting, while for signal regions with overlap (such as SR1 and SR2) the most sensitive region for a given signal point according to Monte Carlo simulation was chosen.

Limits at 95% confidence level on direct sbottom production with sbottom decay to a bottom quark and a neutralino are shown in Figure 12.1 as a function of $m_{\tilde{b}}$ and $m_{\chi_1^0}$. Relative to the previous version of the analysis (the dark grey contour), the exclusion limits are extended ≈ 100 GeV along the $m_{\tilde{b}}$ axis, and ≈ 50 GeV in the $m_{\chi_1^0}$ direction. In many ways this plot understates the improvement near the $\tilde{b} \rightarrow b + \chi_1^0$ forbidden line. While the 95% confidence level limit does not extend to it for cross section and acceptance reasons (as illustrated in Chapter 8) this analysis is at least somewhat sensitive to this region, whereas the previous analysis [2] was completely insensitive.

For the gluino-mediated analyses with off-shell third-generation squarks, the model-dependent exclusion limits are shown as contours in the $m_{\tilde{g}}/m_{\chi_1^0}$ plane. Limits for the $\tilde{g} \rightarrow b\bar{b} + \chi_1^0$ decay mode are shown in Figure 12.2. The SR4 signal regions were considered, and the most sensitive chosen for each point. Along the $m_{\tilde{g}}$ axis, the exclusion is extended to 1.05 TeV, an increase of ≈ 150 GeV compared with the previous version analysis [3] (the orange contour). This is the first > 1 TeV exclusion limit for this model. In the $m_{\chi_1^0}$ direction, the exclusion is increased by an amount between 50 GeV and 200 GeV depending on $m_{\tilde{g}}$. The region with $\Delta m < 50$ GeV is not excluded due to the need for a hard jet from radiation and corresponding lower signal acceptance.

The limits for the $\tilde{g} \rightarrow t\bar{t} + \chi_1^0$ decay mode are shown in Figure 12.3. The SR6 signal regions were considered, and the most sensitive chosen for each point. Relative to the previous version of the analysis [3] (orange contour), the extension of the exclusion limit along the $m_{\tilde{g}}$ axis is enormous, from 750 GeV to almost 1 TeV. However, in this case the previous analysis the leptonic rather than the hadronic version, and was




Figure 12.1: Exclusion contours at 95% confidence level with the CL_s method in the $m_{\tilde{b}}/m_{\tilde{\chi}_1^0}$ plane for direct sbottom pair production, with the sbottom decaying to a bottom quark and neutralino. The theoretical uncertainty on the sbottom cross section is shown as a yellow band. All other systematic uncertainties are included in the nominal limit.




Figure 12.2: Exclusion contours at 95% confidence level with the CL_s method in the $m_{\tilde{g}}/m_{\tilde{\chi}_1^0}$ plane for off-shell gluino-mediated sbottom production ($pp \rightarrow \tilde{g}\tilde{g}, \tilde{g} \rightarrow b\bar{b} + \tilde{\chi}_1^0$). The theoretical uncertainty on the gluino cross section is shown as a yellow band. All other systematic uncertainties are included in the nominal limit.




Figure 12.3: Exclusion contours at 95% confidence level with the CL_s method in the $m_{\tilde{g}}/m_{\chi_1^0}$ plane for off-shell gluino-mediated stop production ($pp \rightarrow \tilde{g}\tilde{g}, \tilde{g} \rightarrow t\bar{t} + \chi_1^0$). The theoretical uncertainty on the gluino cross section is shown as a yellow band. All other systematic uncertainties are included in the nominal limit.

severely limited due to statistics and the resulting inability to b -tag multiple jets. A more interesting comparison is with the $\sqrt{s} = 7$ TeV inclusive multi-jet SUSY search [178] (brown contour) which analyzed the same data but did not incorporate b -tagging. The analysis with b -tagged jets described in Chapter 11 still excludes ≈ 100 GeV further along both axes, making it by far the most sensitive $\sqrt{s} = 7$ TeV search for $\tilde{g} \rightarrow t\bar{t} + \chi_1^0$. As each additional 50 GeV in gluino mass is approximately a factor of two less in production cross section, near the $m_{\tilde{g}}$ axis this corresponds to about a factor of four in increased sensitivity for three-tag search over the multi-jet search.

Limits for the $\tilde{g} \rightarrow t\bar{b} + \chi_1^-$ decay mode are also computed and are shown in Figure 12.4. The SR6 signal regions are most sensitive to this final state. The excluded region is comparable to that of the $\tilde{g} \rightarrow b\bar{b} + \chi_1^0$ decay mode. The improvement relative to




Figure 12.4: Exclusion contours at 95% confidence level with the CL_s method in the $m_{\tilde{g}}/m_{\chi_1^0}$ plane for off-shell gluino-mediated sbottom/stop production with an alternate decay mode ($pp \rightarrow \tilde{g}\tilde{g}$, $\tilde{g} \rightarrow t\bar{b} + \chi_1^{-1}$, $\chi_1^{-1} \rightarrow l^- \chi_1^0$ where the lepton from the chargino decay is lost). The theoretical uncertainty on the gluino cross section is shown as a yellow band. All other systematic uncertainties are included in the nominal limit.

the previous version of the analysis [3] (orange contour) is the most dramatic of all gluino-mediated topologies.

12.4 Model-independent Limits

Model-independent limits were computed as described in Section 12.1 for each signal region in both analyses, and expressed in Table 12.1 in terms of the maximum number of signal events from new physics at 95% confidence level. The limit on the visible cross section σ_{vis} (Equation 12.4) can be obtained via division by 4.71 fb^{-1} .

Signal Region	95% CL Upper Limit	
	Observed	Expected
SR1 ($m_{CT} > 150$ GeV)	28.9	25.6
SR1 ($m_{CT} > 200$ GeV)	16.9	15.6
SR1 ($m_{CT} > 250$ GeV)	5.22	6.94
SR2	10.8	14.4
SR3a	36.9	34.3
SR3b	7.45	8.04
SR4-L	23.8	23.4
SR4-M	8.6	12.8
SR4-T	7.1	9.2
SR6-L	9.6	10.1
SR6-T	7.1	8.3

Table 12.1: Observed and expected 95% confidence level upper limits on the non-Standard Mode contributions (number of events) to all signal regions. Systematic uncertainties on the SM background estimates are included.

12.5 Discussion and Caveats

As discussed in Chapter 4, an advantage of two-parameter simplified models such as the ones considered here is that the entire parameter space is visible in the limit plots, and it is often very obvious where the analyses are less sensitive (i.e., near the kinematically forbidden lines). Similar limit plots for models with more than two parameters can easily hide weaknesses. Despite this being obvious, this caveat is often overlooked. Illustrating this point well are the exclusion limits of on-shell gluino-mediated decay modes in the $m_{\tilde{g}}/m_{\tilde{q}}$ plane. A $\tilde{g} \rightarrow \tilde{b}\tilde{b}, \tilde{b} \rightarrow b + \chi_1^0$ model is shown in Figure 12.5, and a $\tilde{g} \rightarrow \tilde{t}\tilde{t}, \tilde{t} \rightarrow t + \chi_1^0$ model is shown in Figure 12.6. Both have $m_{\chi_1^0}$ fixed to 60 GeV. This is a very favorable choice for discovery, as almost every point in the mass plane has decay signatures with significant visible energy and \cancel{E}_T . As a result, the exclusion contours are almost vertical and determined essentially by the cross section alone.

There is no dishonesty in these limits, but for other heavier choices of $m_{\chi_1^0}$, the mass plane coverage is significantly worse. There is no way to discern this fact from these plots, and it is easy to assume that an 800 GeV gluino is simply excluded when




Figure 12.5: Exclusion contours at 95% confidence level with the CL_s method in the $m_{\tilde{g}}/m_{\tilde{b}}$ plane for on-shell gluino-mediated sbottom production ($pp \rightarrow \tilde{g}\tilde{g}, \tilde{g} \rightarrow \tilde{b}\tilde{b}, \tilde{b} \rightarrow b + \chi_1^0$) assuming $m_{\chi_1^0} = 60$ GeV. The theoretical uncertainty on the gluino cross section is shown as a yellow band. All other systematic uncertainties are included in the nominal limit.

this is *not* generally true. These figures do confirm, at least for the large phase space case, the qualitative assertion made in Chapter 4 that signal regions optimized for the off-shell case would also be sensitive to the on-shell case. The on-shell limit positions in terms of $m_{\tilde{g}}$ are similar to the exclusion contours near the $m_{\tilde{g}}$ axis (large Δm) in the corresponding off-shell cases (Figures 12.2 and 12.3).

A more subtle issue affecting all simplified model limit plots is the assumption of 100% branching ratio. Only a fraction of supersymmetric models have a single dominant decay mode. Even if the real branching ratio for the desired decay mode is as high as 80%, the effective cross section is only 64% of the number used in the model-dependent limit calculations. This translates into roughly ≈ 50 GeV in terms




Figure 12.6: Exclusion contours at 95% confidence level with the CL_s method in the $m_{\tilde{g}}/m_{\tilde{t}}$ plane for on-shell gluino-mediated sbottom production ($pp \rightarrow \tilde{g}\tilde{g}, \tilde{g} \rightarrow \tilde{t}\tilde{t}, \tilde{t} \rightarrow t + \chi_1^0$) assuming $m_{\chi_1^0} = 60$ GeV. The theoretical uncertainty on the gluino cross section is shown as a yellow band. All other systematic uncertainties are included in the nominal limit.




Figure 12.7: Mass spectrum of a pMSSM natural supersymmetric model which evades all limits as of June 2012 [179]. The decay modes of the sbottom in this model are shown in Figure 12.8.

of $m_{\tilde{g}}$ or $m_{\tilde{b}}$. A 50% branching ratio for the desired decay mode would suppress the effective cross section by a factor of four.

A study applying all relevant collider and cosmological constraints to the nineteen-parameter pMSSM (“phenomenological” MSSM) illustrates this point nicely [179]. Included in the constraints are this version of the gluino-mediated analysis and the previous iteration of the direct sbottom search [2]. The mass spectrum of one model from the paper which is relatively natural (light stop) and evades all current constraints is shown in Figure 12.7.

In this model, one stop and one sbottom are relatively light with masses of ≈ 400 GeV, as desired from naturalness. However, the gluino is approaching 3.6 TeV. This means that the gluino production modes with their high LHC cross sections are not available for discovery at $\sqrt{s} = 7$ TeV. Moreover, three of the four neutralinos and

Light Sbottom Decays




Figure 12.8: Sbottom decay modes of a pMSSM natural supersymmetric model which evades all limits as of June 2012 [179]. The mass spectrum of the model is shown in Figure 12.7.

both charginos are lighter than the lightest squark, which indicates that there are likely several possible squark decay modes.

The sbottom decay modes in this model are shown in Figure 12.8. Not only is the majority of the decay width assigned to decays other than $\tilde{b} \rightarrow b + \chi_1^0$, most of the other decay modes are cascades involving vector bosons, which can potentially contaminate the one- and two-lepton control regions. The branching ratio to the mode targeted by the direct sbottom analysis is 19%, corresponding to a 3.6% effective cross section—a factor of twenty-five in effective cross section suppression.

One of the strengths of the simplified model approach and the models used in these analyses in particular is that we can, in a rough sense, apply the limits shown

in Figure 12.1 to the specific model shown in Figures 12.7 and 12.8 and see the effect of this suppression. If the branching ratio were 100%, the relevant point would be $m_{\tilde{b}} = 400$ GeV and $m_{\chi_1^0} = 100$ GeV, and the model would be excluded at 95% confidence level. Using Figure 4.8 or the factor-of-two-per-50 GeV rule of thumb, a factor of 25 reduction in cross section corresponds to an increase of about 250 GeV in sbottom mass, from 400 GeV to 650 GeV. We expect the kinematics to be roughly the same for the same value of Δm , which in this case is 300 GeV. Therefore, the point in Figure 12.1 which best represents this particular full SUSY model is $m_{\tilde{b}} = 650$ GeV and $m_{\chi_1^0} = 350$ GeV. This point is very much not excluded by the direct sbottom analysis. In fact, it is not even within the range of the exclusion plot axes.

It is possible that the various direct stop analyses [27, 28, 29, 30, 31, 32, 33, 34] may significantly constrain this model. The dominant decay mode is $\tilde{b} \rightarrow t + \chi_1^-$, $\chi_1^- \rightarrow l^- + \chi_1^0$, but the mass difference between the chargino and neutralino is very small, so the lepton from the chargino decay is likely to be lost. Therefore, the final state looks very much like the direct stop simplified model final state ($\tilde{t} \rightarrow t + \chi_1^0$). However, there is still a factor of 4 in effective cross section to contend with, and direct stop searches are often difficult due to the limited ways one can reduce the $t\bar{t}$ background.

This iteration of the direct sbottom and gluino-mediated sbottom/stop analyses has increased the search sensitivity enormously relative to previous versions, as Figures 12.1, 12.2, 12.3 and 12.4 make clear. However, even these limits need to be viewed in proper context, as illustrated by the preceding discussion. Natural supersymmetry may be wounded but it is certainly not dead, and there remain plenty of potential hiding places.

Chapter 13

Conclusion and Prospects

The search optimization strategies (Chapter 8) based on simplified models (Chapter 4) and the analyses (Chapters 10 and 11) presented in this document have broadly and significantly increased the experimental sensitivity to two supersymmetric processes relevant to natural supersymmetry (Chapter 3): direct sbottom squark pair production followed by sbottom decay to a bottom quark and neutralino, and gluino pair production with gluino decay to two bottoms or tops and a neutralino via on- or off-shell sbottoms and stops. Combined with the direct stop pair production searches [27, 28, 29, 30, 31, 32, 33, 34], these represent some of the best-motivated of new physics searches. The gluino-mediated analysis is also the first new physics search to successfully use signal regions requiring three b -tagged jets.

No significant excesses are observed in the full 2011 $\sqrt{s} = 7$ TeV 4.71 fb^{-1} ATLAS dataset, leading to model-dependent and model-independent limits at 95% confidence (Chapter 12). It should be emphasized that the models used for the model-dependent limits are simplified models with 100% branching ratios to a single decay mode. These limits are therefore maxima, and the relevant limit for any full supersymmetric model is less restrictive—in some cases significantly so.

This model-based uncertainty is avoided almost entirely in the optimization by use of the simplified models approach, which parameterizes a single decay chain in terms of the relevant masses and leaves the product of the cross section and branching ratio a free parameter. The optimization is aided by careful characterization of the

production and decay processes of interest in order to identify the primary drivers of the final state kinematics. The signal regions are chosen to cast a wide net over all kinematic regimes and thereby minimize the chance that new physics is missed, as well as maximize sensitivity to the degree possible and practical. The 4.71 fb⁻¹ sbottom search ATLAS conference note and gluino-mediated analysis paper can be found at References [4] and [5], respectively.

The search for natural supersymmetry in hadronic channels continues with these methods at $\sqrt{s} = 8$ TeV. Both the direct sbottom pair production search and the three b -tag gluino-mediated sbottom and stop searches have recently been updated with the first 12.8 fb⁻¹ of the 2012 dataset [6, 7]. While the optimization for these analyses has been redone for the higher \sqrt{s} and integrated luminosity (with correspondingly large jump in sensitivity), the general search strategy has remained the same.

Significant effort is also being put into new channels and strategies. A one-lepton and three b -tag analysis would be equally sensitive to gluino-mediated stop as the current hadronic search, and is planned for the full 2012 dataset. At $\sqrt{s} = 7$ TeV and 4.71 fb⁻¹, such an analysis would have suffered from extremely low event yields. For heavy gluinos, the possibility of boosted top quarks arises and leads to large radius “top jets” which may be best analyzed using jet substructure techniques. An analysis following this approach is also being pursued.

Partner quark and the lightest neutralino is not the only allowed squark decay mode in supersymmetry. For truly complete model coverage one needs to ensure sensitivity to all other major decay modes including those with multiple steps. This process has begun and will continue throughout the imminent LHC shutdown period. The next LHC data will be taken at the design energy of 14 TeV and will open up an entirely new range of gluino and squark masses to explore.

The LHC era has just begun, and while significant efforts have been made to find supersymmetry, much work is yet to be done. Natural supersymmetry remains one of the best-motivated beyond-the-Standard Model theories. I look forward to the day, not too long from now, when we will learn definitively whether it exists.

Appendix A

Track Jets

Before committing full time to heavy flavor supersymmetry analysis in 2010, my contributions to ATLAS were dominated by studies aimed at understanding and improving detector performance, with particular emphasis on jets, tracks, and applications utilizing both together. The project that yielded the most dividends to the collaboration as a whole was three-dimensional (3D) track jet reconstruction. This appendix also serves to illustrate some of the difficulties posed by pileup to event reconstruction, a topic largely neglected in this iteration of the heavy flavor supersymmetry analyses (Chapters 10 and 11) due to the small pileup-dependence of the control and signal region selections and the fact that both analyses are counting experiments. With ever higher luminosity and moves to shape-based analysis, such avoidance will no longer be possible.

Pre-2008, track jets in ATLAS were constructed by simply passing track four-vectors through the standard jet reconstruction. While for many purposes this was satisfactory, the method ignores the additional information tracks provide compared to calorimeter clusters. In the presence of multiple interactions, track impact parameters relative to the reconstructed vertices in the event can be used to ensure that track jets only contain tracks from the same interaction, and associate the resulting jets to their respective vertices. The inherent robustness of track jets constructed in this way to distortions from pileup makes them useful both for cross-checks of calorimeter-based objects and directly as primary physics objects.




Figure A.1: 3D (left) vs. 2D (right) track jets in a simulated di-jet event with pileup. The minimum track jet p_T cut in both cases is 5 GeV. A line with length 1 cm in the above displays corresponds to a track with p_T of 0.5 GeV. Note that the vertical and horizontal scales are different. 5 jets pass the 5 GeV minimum p_T cut in the 3D case, and each is clearly associated with a single origin interaction. In the 2D case 13 jets pass, and most of these are composite objects with significant contributions from multiple interactions and therefore dubious physics value.

An event display illustrating the difference and advantage of 3D track jets relative to two-dimensional (2D) track jets in the presence of pileup is shown in Figure A.1. It is clear that ignoring the tracks' longitudinal impact parameters when constructing track jets, as occurs in the 2D case, produces jets made from tracks originating from different interactions. This distorts the jets from the hard-scatter and makes unambiguous jet-to-vertex assignment difficult. It also allows minimum-bias tracks to push several low energy track jets from the hard-scatter interaction above threshold, resulting in extra fake low- p_T jets. A more statistical measure of the distortion is shown in Figure A.2, where the large excess of 2D jets as compared to 3D jets is due to jets built from tracks coming from different interactions.

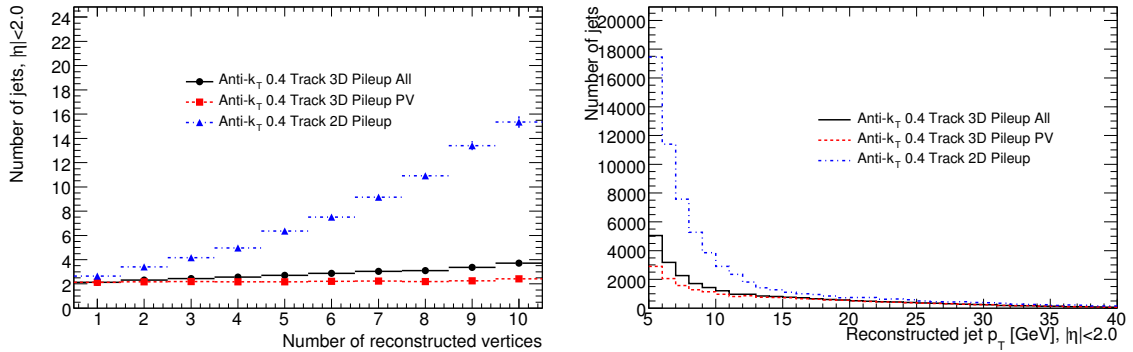


Figure A.2: Average number of 2D and 3D anti- k_t track jets vs. number of reconstructed vertices (left) and reconstructed p_T distribution (right) for simulated di-jet events with pileup ($2 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$ luminosity). PV refers to 3D track jets matched to the primary vertex of the event (only possible with 3D, where all track jet tracks share a common origin vertex). In the 2D case, the removal of the requirement that all jet constituents (tracks) come from the same interaction allows a large number of “mixed origin” jets to pass the 5 GeV p_T cut in events with many reconstructed vertices.

A.1 Monte Carlo and Calorimeter Jets

The Monte Carlo datasets used in this study (the work predates first LHC data) are PYTHIA di-jet samples consisting of 30,000 events without pileup. 10,000 events with pileup were also used, corresponding to a luminosity of $2 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$ with 25 ns bunch-spacing (this is to be compared with the 50 ns bunch-spacing in the 2011 dataset). Each event has an average of 4.6 additional pileup interactions. Reference calorimeter jets are reconstructed using topological clusters (Section 7.6) as inputs with the anti- k_t algorithm with $R = 0.4$ (Section 7.7).

A.2 Reconstruction

Jet reconstruction in ATLAS is mediated by an instance of `JetAlgorithm`, a process which is passed an ordered sequence of `JetAlgTools` which select inputs, do the actual jet-finding, perform cuts on the resulting jets, calibrate them, and set the jets’ energy signal state (a method by which both calibrated and uncalibrated four momenta can

Parameter	Default Value	Description
UseVtxSeeding	True	Mode selection
DeltaZRange	5.0	Δz cut (mm)
TrackJetMinPt	0.0	Minimum jet p_T cut (MeV)
TrackJetMinMulti	2	Minimum track multiplicity cut
TrackParticleContainerName		Track container
VxContainerName		Vertex container
TrackSelector		TrackSelectorTool ToolHandle
JetFinder		JetAlgToolBase ToolHandle

Table A.1: Configurable algorithm parameters for JetTrackZClusterTool

be stored). When the JetAlgTools have finished, the resulting JetCollection is written to StoreGate (the service which handles both transient and persistent data storage) by the JetAlgorithm. JetTrackZClusterTool, the algorithm at the heart of track jet reconstruction, takes the place of the jet finder and input selector in the JetAlgTool sequence. A jet finder must still be instantiated and configured, but it is passed to JetTrackZClusterTool via parameters (Table A.1) rather than included directly in the JetAlgorithm sequence.

JetTrackZClusterTool has two operational modes, defined by the *UseVtxSeeding* parameter. The default and recommended setting is *True*. In this mode, the tool retrieves the first vertex in the vertex container specified with *VxContainerName* and iterates through all tracks in the input track container. Each track and vertex are passed to a configured TrackSelectorTool (specified with JetTrackZClusterTool's *TrackSelector* ToolHandle), which applies track quality and impact parameter cuts with respect to the vertex under consideration. If the track passes the TrackSelectorTool and has a z impact parameter (z_0) within *DeltaZRange* of the vertex z-position, it is then added to the input collection for that particular vertex and removed from the list of available tracks. The z_0 cut is applied in case the user does not wish to use the TrackSelectorTool to make cuts relative to the vertex. If the user does elect to use the TrackSelectorTool for this (recommended) then *DeltaZRange* should be set to something large to ensure it does not interfere (for instance, 5 mm).

This process is repeated for each vertex in the container. Note that any track-vertex association ambiguities are resolved in favor of whichever vertex is nearer to the beginning of the container. This is the desired outcome as the reconstructed vertex container is by default ordered with vertices with higher $\sum p_T^2$ and track multiplicity first. The first vertex in the container is defined as the hard-scatter vertex of the event.

The result of the vertex iteration is the creation of separate track input collections for each reconstructed vertex in the event. `JetTrackZClusterTool` then calls the configured jet finder `JetAlgTool` (specified with the `JetFinder ToolHandle`) on each input collection independently, which produces a collection of track jets for each vertex. These track jets are associated to their respective vertices via a persistent link set to the appropriate vertex in the `VxContainer` specified with `VxContainerName`. Track jets which contain at least `TrackJetMinMulti` track constituents and have $p_T > \text{TrackJetMinPt}$ are then added to the final output `JetCollection`.

The `UseVtxSeeding = False` mode is provided to construct track jets in a way unbiased by vertex reconstruction. Instead of vertices being used to seed track z-clusters, the tool selects as the first seed the highest p_T track in the event which passes the `TrackSelectorTool` cuts. The tool then iterates through the input track container and adds all selected tracks whose z impact parameters lie within `DeltaZRange` of the seed track. The second cluster is then seeded with the highest p_T track remaining, and the process is repeated until all tracks have either failed the selection cuts or been assigned to a z-cluster. An attempt is made to match each z-cluster with a reconstructed vertex found within `DeltaZRange` of the seed track's z_0 (again vertices towards the beginning of the container have priority), however if no match is found the tool simply moves on. The rest of the tool's operation in this mode is identical to the `UseVtxSeeding = True` case, except that vertex associations are only made if the parent cluster had been matched to a vertex. Diagrams of the two operational modes are shown in Figure [A.3](#).

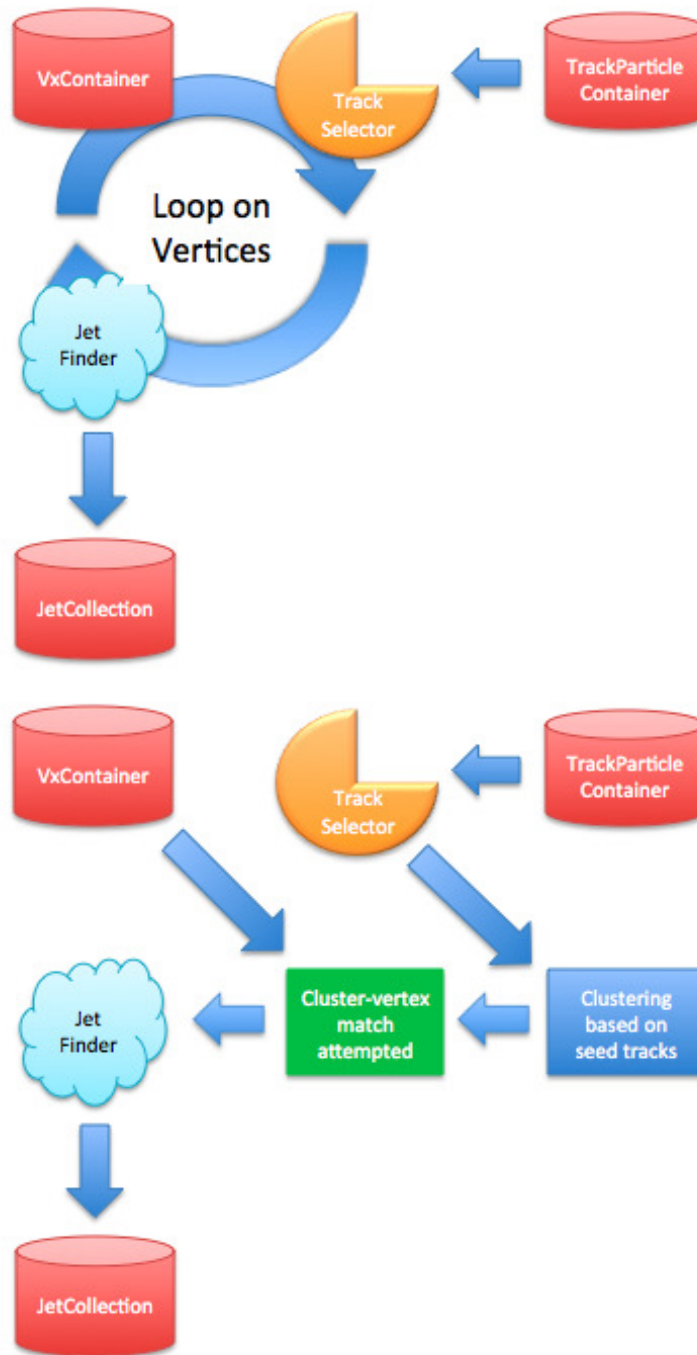


Figure A.3: Diagrams representing the two operational modes of JetTrackZClusterTool. The setting of *UseVtxSeeding* to *True* (top) uses the reconstructed vertices of the event to cluster tracks and allows cuts to be applied relative to them. The *False* setting (bottom) instructs the tool to perform its own clustering based on the high p_T tracks in the event. The default mode is *True*.

A.3 Optimization

Di-jet events without pileup were used to optimize the track jet algorithm. As one of the potential uses for track jets is b -tagging, we adopt the standard impact parameter tagger track selection cuts (Section 7.8) with a few modifications. Only seven silicon hits are required, removing the additional b -tagging requirement of at least two pixel-layer hits, one being in the innermost layer (b -layer). This change was made because requiring a b -layer hit negatively impacted the track jet efficiency (fraction of truth jets matched to reconstructed track jets) without improving the purity (fraction of reconstructed track jets matched to truth). The impact parameter cuts are $|d_0| < 1$ mm and $|z_0 - z_{pv}| \sin \theta < 1.5$ mm. The minimum track p_T cut is 0.5 GeV, compared to the b -tagging cut of 1 GeV.

We use the `UseVtxSeeding = True` operation mode and use the `InDetDetailedTrackSelector TrackSelectorTool` (a standard `TrackSelectorTool` used by both the inner detector alignment and b -tagging) to perform the track selection. The anti- k_t algorithm (Section 7.7) with $R = 0.4, 0.5,$ and 0.6 is used to find the track jets. All plots made with reference to MC truth jets use anti- k_t truth jets constructed with $R = 0.4$. We justify considering larger R for track jet reconstruction than those used in truth by noting that as we are only clustering the charged energy in the track jets (as opposed to all truth particles in the truth), using $R = 0.4$ is not any more consistent a priori than using 0.5 or 0.6 , and larger R could yield better performance by compensating for the lack of neutral particles. The minimum track multiplicity per jet is two (another compromise between efficiency and purity), and minimum track jet p_T cuts of 5 GeV, 5.5 GeV, and 6 GeV have been applied to $R = 0.4, 0.5$ and 0.6 respectively to ensure a minimum purity level of 75% in bins of reconstructed p_T , as shown in Figure A.4. The matching criterion used between truth and reconstructed jet is $\Delta R < 0.4$.




Figure A.4: Anti- k_t track jet purity vs. reconstructed p_T (left) and reconstructed η (right) for different R simulated di-jet events without pileup ($|\eta| < 2.0$). In the right plot, p_T cuts of 5 GeV, 5.5 GeV, and 6 GeV have been applied to $R = 0.4$, 0.5 and 0.6 respectively to ensure a minimum purity level of 75% in bins of reconstructed p_T .

A.4 Performance

Efficiencies for track jets as a function of the truth jet p_T with respect to all truth jets are shown in the top-left plot of Figure A.5. Notice that none of the three R settings seems to have a clear advantage here over the others, and that the efficiencies for all three drop precipitously below 20 GeV. In order to explain the origin of this inefficiency, we define two truth jet/truth event selection criteria:

- *Findable* truth jets are defined as those which have two or more selected tracks within $R = 0.4$ of the truth jet axis and whose combined p_T is greater than 5 GeV.
- *Isolated* events are those in which all truth jets have $\Delta R > 1.0$ with respect to the nearest truth jet.

These selections attempt to separate inefficiencies related directly to jet finding (merging/splitting, p_T cut) from underlying problems such as low selected track efficiency or truth jets with low track multiplicity/energy. The largest efficiency increase is obtained by requiring truth jets to be findable (compare the top-left and bottom-left plots of Figure A.5). As a significant relaxation of the track selection cuts only

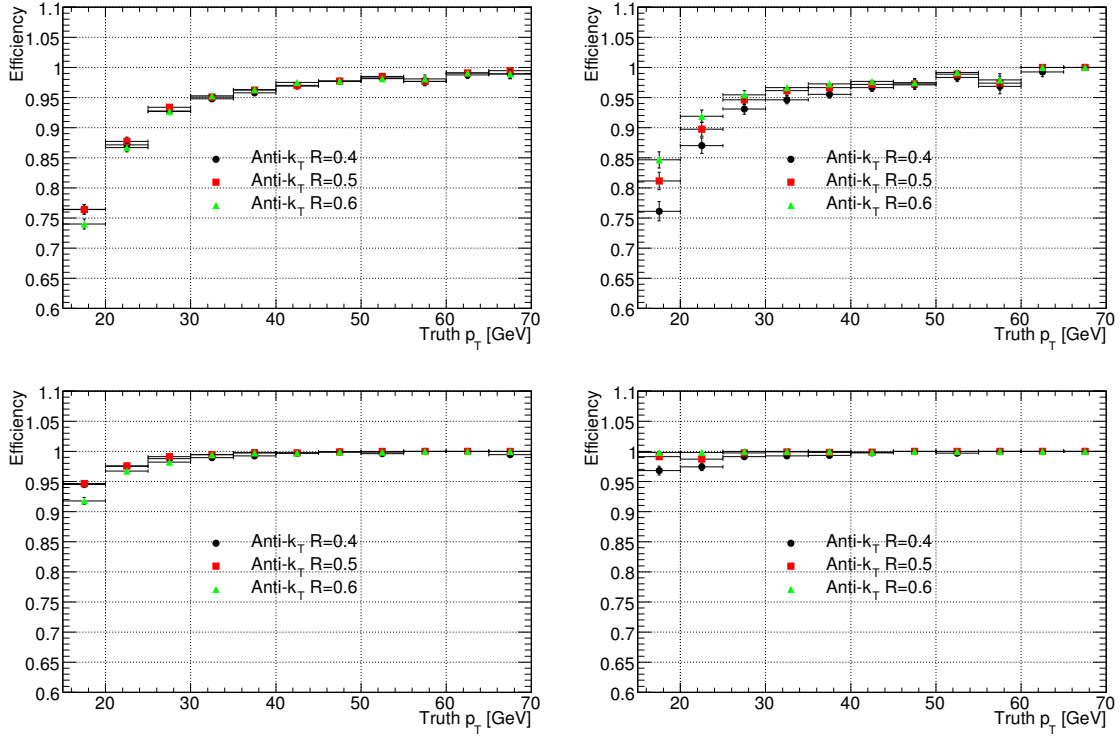


Figure A.5: Anti- k_t track jet efficiency vs. truth jet p_T for different R simulated di-jet events without pileup ($|\eta| < 2.0$). Clockwise from top-left are efficiencies for all jets, isolated jets, findable and isolated jets, and all findable jets. Reconstructed p_T cuts of 5 GeV, 5.5 GeV, and 6 GeV have been applied to $R = 0.4, 0.5$ and 0.6 respectively to ensure a minimum purity level of 75% in bins of reconstructed p_T . All efficiencies are flat with respect to η .




Figure A.6: Anti- k_t $R = 0.4$ truth jets selected by the findable jet and isolated event criteria vs. truth p_T in simulated di-jets without pileup ($|\eta| < 2.0$).

marginally reduces this difference, we conclude that a large fraction of truth jets below 20 GeV simply do not have enough associated reconstructed track energy/multiplicity to pass the track jet p_T and multiplicity cuts.

The isolated event cut is applied to determine the effect of merging on the track jet efficiency. By comparing the bottom-left and bottom-right plots in Figure A.5 it is clear that a large fraction of the residual inefficiency for 0.5 and 0.6 findable jets is due to merging. An example of this is shown in the top event display of Figure A.7, where two adjacent $R = 0.4$ truth jets are merged into a single $R = 0.6$ track jet, resulting in one of the truth jets remaining unmatched. Near 100% efficiency with respect to findable, isolated jets over the entire p_T range can be obtained for $R = 0.5$ and 0.6 by lowering the reconstructed p_T cuts from 5.5 and 6 GeV, respectively, to 5 GeV, the findable criterion. The residual inefficiency of $R = 0.4$ with respect to findable, isolated jets is largely due to pathological track distributions in the jet cone resulting in split jets (bottom-left and bottom-right event displays in Figure A.7). Figure A.5 can therefore be understood as follows. While increasing the R helps collect more tracks and compensate for small-scale pathological track distributions, it also produces more fakes and merged jets. The increase in fakes is addressed by raising the minimum p_T cut to ensure a minimum of 75% purity. The end result is similar efficiencies with respect to all truth jets, as seen in the top-left plot of Figure A.5. A comparison to the efficiency of anti- k_t 0.4 calorimeter jets is given in Figure A.8.

Jet multiplicity per event, track multiplicity per track jet, and jet η and p_T distributions for simulated di-jet events are shown in Figure A.9. Where applicable, the corresponding anti- k_t $R = 0.4$ calorimeter jet distributions are shown as well.

The ϕ resolution ($\phi_{reco} - \phi_{truth}$) of calorimeter and track jets in representative p_T bins are shown in Figure A.10. Track jets appear to have similar ϕ resolution to calorimeter jets at high p_T and better resolution at low p_T . The resolution decreases with increasing R . In contrast with calorimeter jets, track jet ϕ resolution is clearly non-Gaussian when binned only in truth p_T due to its additional dependence on track multiplicity and the p_T of the constituent tracks. Plots of η resolution show similar performance. Combining this with the efficiency measurements, there appears to be

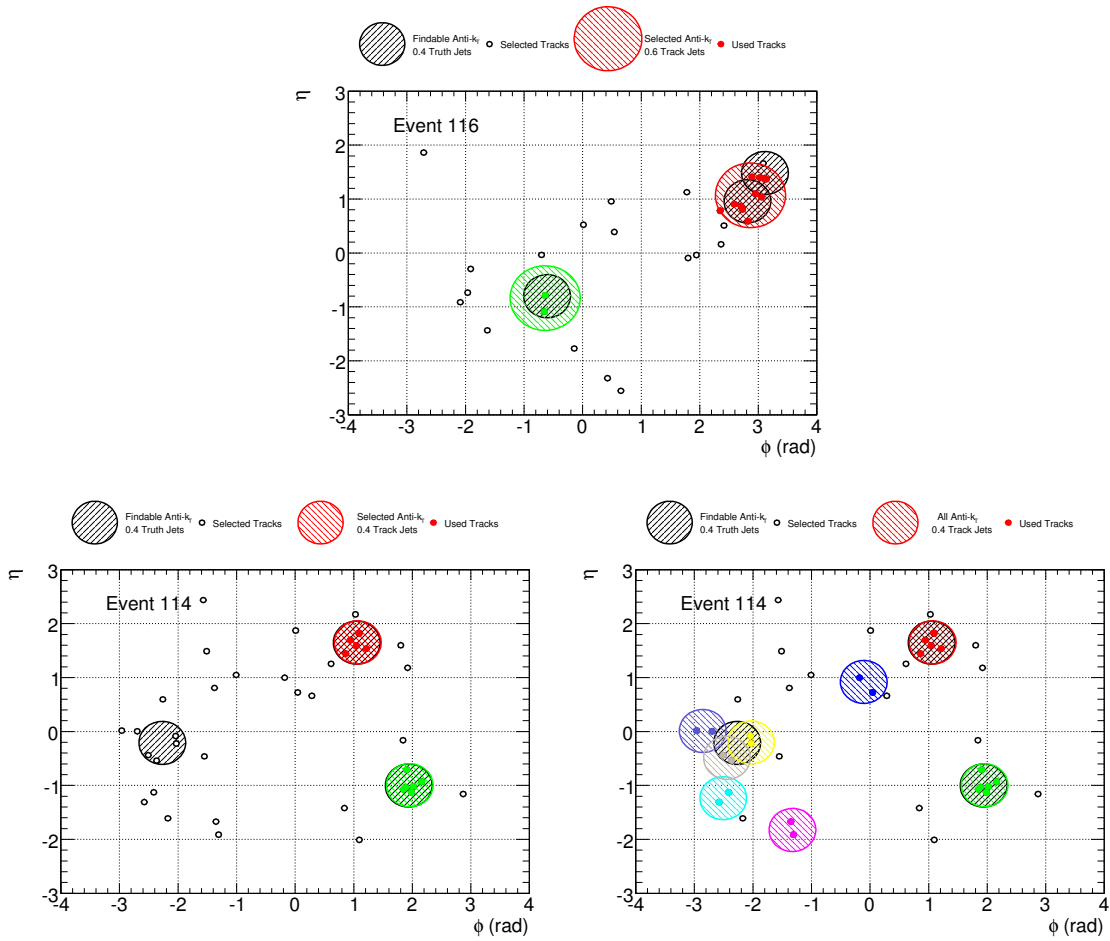


Figure A.7: Examples of events with track jet efficiency problems with respect to findable jets. Track jets and tracks are shown in the η/ϕ plane. In the top figure, two adjacent $R = 0.4$ truth jets are merged into a single $R = 0.6$ track jet. In the bottom-left event display, tracks corresponding to a single truth jet are highly separated with respect to the jet axis, resulting in two low-energy $R = 0.4$ track jets, neither of which passes the $5 \text{ GeV } p_T$ cut (bottom-right).




Figure A.8: Anti- k_t track jet and calorimeter jet efficiencies in simulated di-jet events without pileup (reconstructed jets with $|\eta| < 2.0$).




Figure A.9: Anti- k_t jet multiplicity per event (top-left), track multiplicity per track jet (top-right), and jet p_T (bottom-left) and η (bottom-right) distributions for simulated di-jet events without pileup. Both calorimeter and track jets are included where appropriate. The additional structure seen for calorimeter jets as a function of η is due to calorimeter structure (gaps) and dead material beyond the tracker.

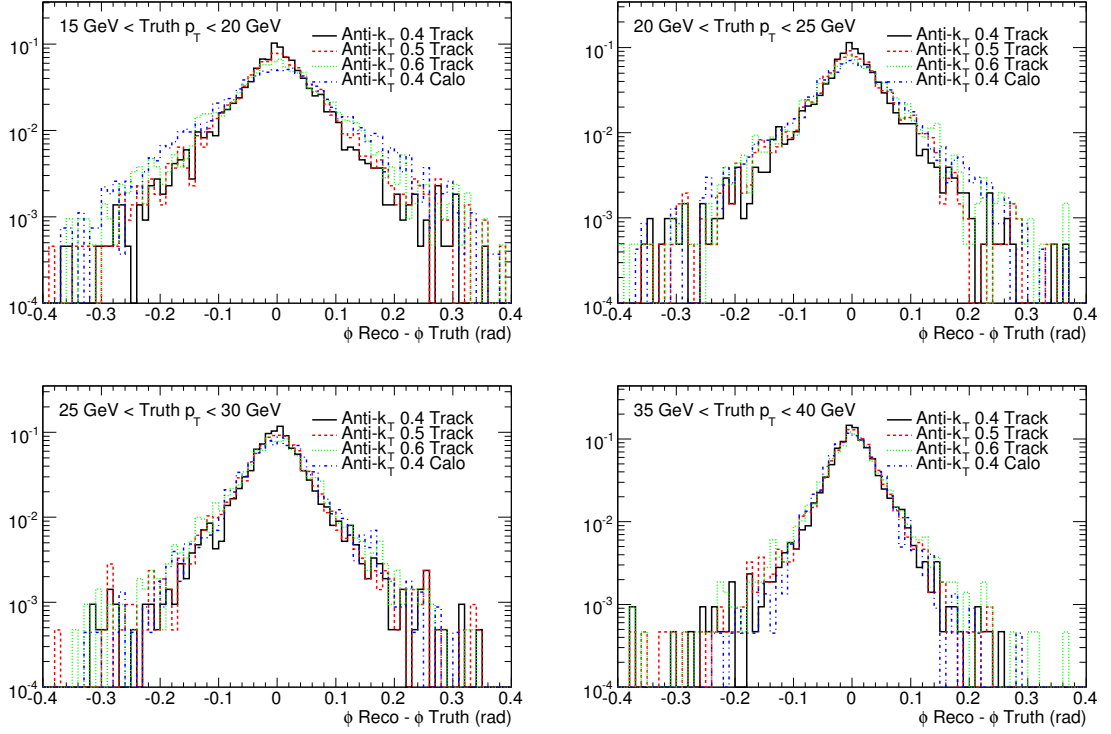


Figure A.10: Anti- k_t track jet and calorimeter jet ϕ resolution ($\phi_{reco} - \phi_{truth}$) in simulated di-jet events without pileup (reconstructed jets with $|\eta| < 2.0$) for representative truth p_T bins $15 \text{ GeV} < p_T < 20 \text{ GeV}$ (top-left), $20 \text{ GeV} < p_T < 25 \text{ GeV}$ (top-right), $25 \text{ GeV} < p_T < 30 \text{ GeV}$ (bottom-left), and $35 \text{ GeV} < p_T < 40 \text{ GeV}$ (bottom-right). The area under the curves has been normalized to 1 for better direct comparison.

very little if any benefit obtained from using track jets with R larger than 0.4 with $R = 0.4$ truth jets. Thus the remainder of the discussion shall be limited to $R = 0.4$ track jets only.

A.5 Pileup

Pileup (event contamination due to additional minimum-bias interactions superimposed on top of the hard-scatter event) is a serious challenge for analysis at the LHC. With calorimeter or 2D track jets, pileup can distort an event both by creating new jets and by distorting jets from the hard-scatter in both direction and energy. In

contrast, 3D track jets are largely immune from distortion caused by the addition of pileup. As shown in Figure A.11, the kinematic distributions of the track jets associated with the hard-scatter vertex with pileup are essentially identical to those from the non-pileup sample. In addition, we see that the number of primary vertex jets is constant with increasing numbers of reconstructed vertices, which indicates minimal minimum-bias contamination. This is in sharp contrast to calorimeter jets, where there is no straightforward calorimeter-based method of separating out minimum-bias from hard-scatter signal contributions. As shown in Figure A.12, a significant increase in jet multiplicity is observed at low jet p_T (caused by pure minimum-bias and contaminated low-energy jets from the hard-scatter vertex).

The ϕ resolution of 3D track jets is also unaffected by pileup, as shown in Figure A.13. 3D track jet η resolution is similarly unaffected. While there are other track-based methods to correct the energy of calorimeter jets for pileup and reject minimum-bias jets, 3D track jets offer a straightforward and calorimeter-independent way of correcting the *direction* of these jets.

A.6 Jet Angular Resolution

Track jets are constructed independently from the calorimeter using track information only. Calorimeter jets are constructed independently of the tracker. Thus, the uncertainties in jet-axis direction for the two types of jets are uncorrelated, and the measurements can be combined to improve the overall ϕ and η resolution for matched jets. Good Gaussian fits for track jet resolutions require binning in reconstructed p_T , truth p_T , track multiplicity, and to a small extent η . In addition to requiring large numbers of jets, the binning in truth p_T requires using the matched calorimeter p_T as a proxy. For simplicity we instead use the RMS (root mean square) of the resolution distribution as a measure of its width and bin in reconstructed track jet p_T only. Calorimeter jets are binned in reconstructed p_T and η .

The combination scheme works as follows. The selected anti- k_t calorimeter jets are iterated through and matched to track jets using the criterion $\Delta R < 0.4$. If multiple jets satisfy this condition, the jet closest to the calorimeter jet is chosen.

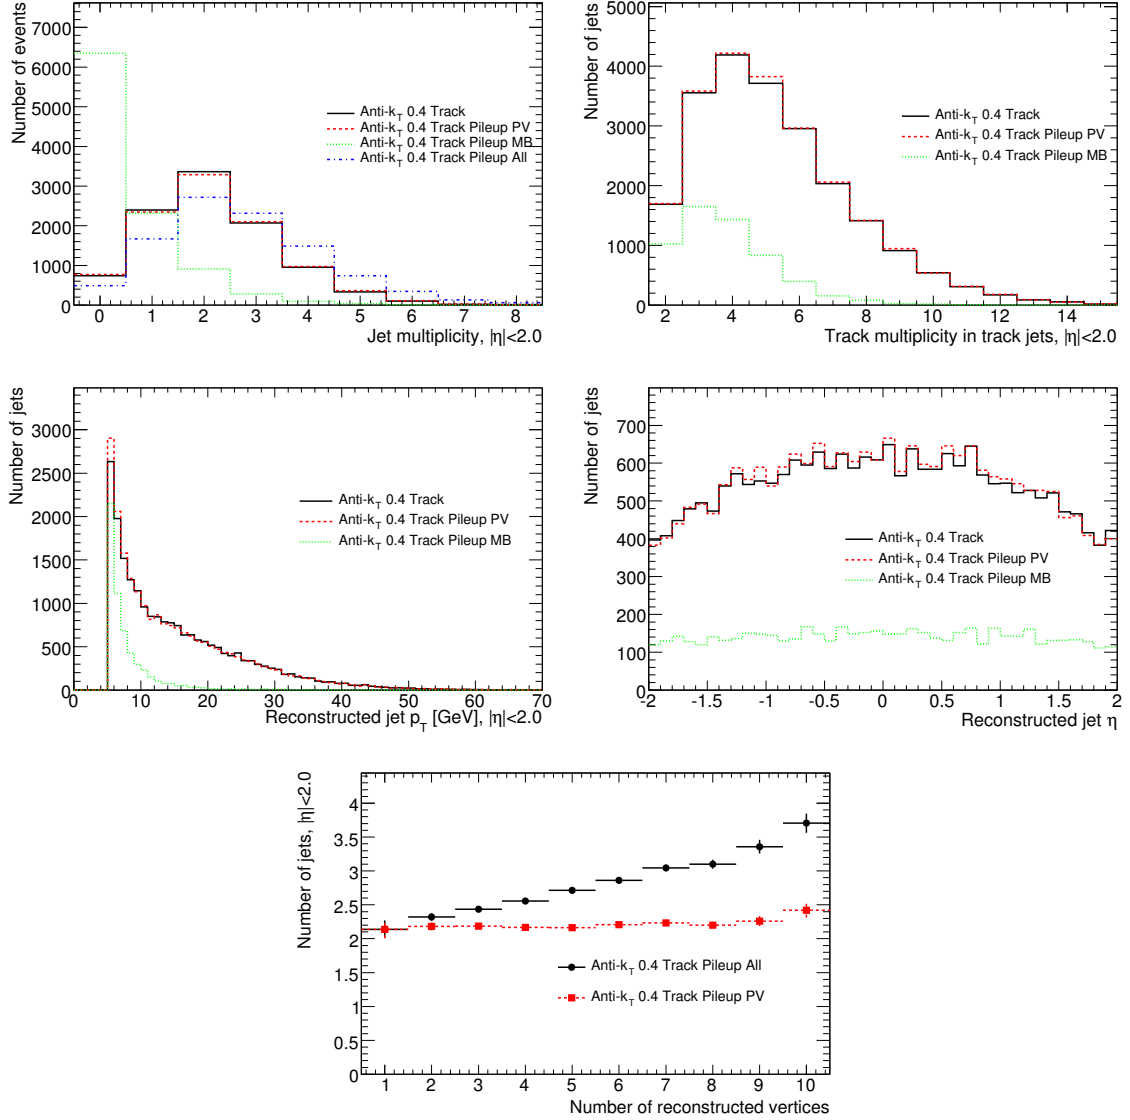


Figure A.11: Anti- k_t track jet multiplicity per event (top-left), track multiplicity per track jet (top-right), track jet p_T (middle-left) and η (middle-right) distributions, and number of track jets vs. number of reconstructed vertices (bottom) for simulated di-jet events with pileup ($2 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$ luminosity). PV refers to track jets matched to the primary vertex of the event, while MB refers all other jets (minimum-bias/pileup jets).

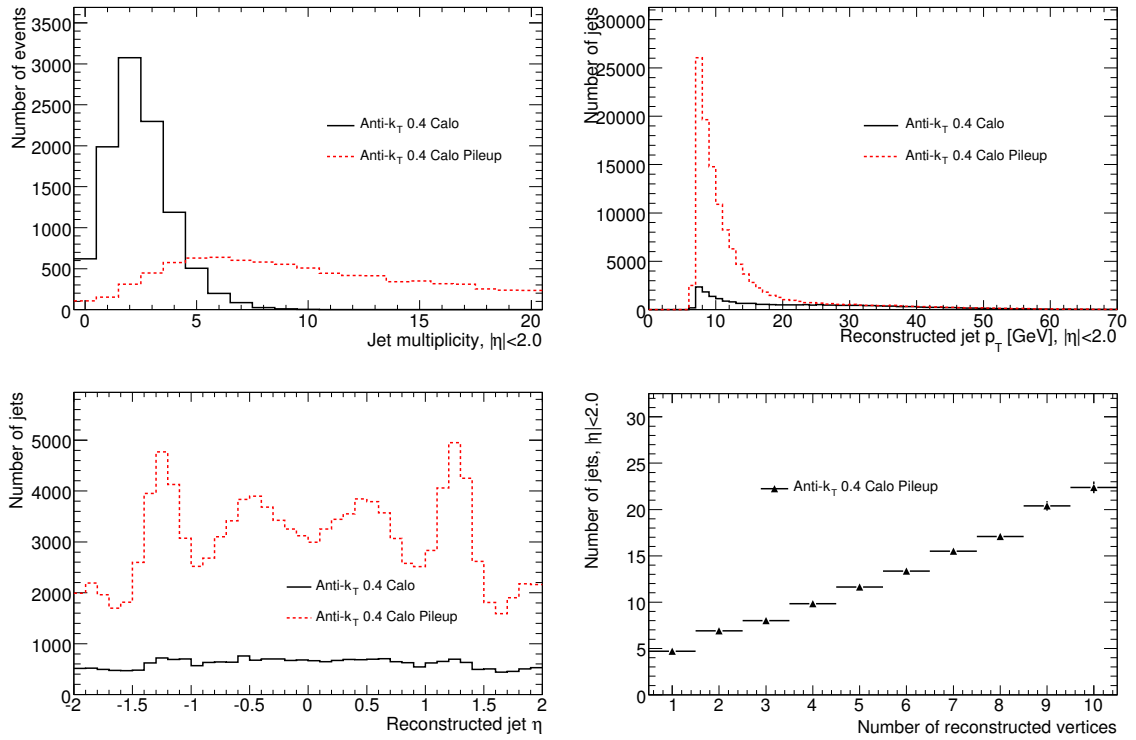


Figure A.12: Anti- k_t calorimeter jet multiplicity per event (top-left), jet p_T (top-right) and η (bottom-left) distributions, and number of jets vs. number of reconstructed vertices (bottom-right) for simulated di-jet events with pileup ($2 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$ luminosity).




Figure A.13: Anti- k_t track jet and calorimeter jet ϕ resolution ($\phi_{reco} - \phi_{truth}$) in simulated di-jets (reconstructed $|\eta| < 2.0$) with pileup ($2 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$ luminosity) for truth p_T bins $15 \text{ GeV} < p_T < 20 \text{ GeV}$ (top-left), $20 \text{ GeV} < p_T < 25 \text{ GeV}$ (top-right), $25 \text{ GeV} < p_T < 30 \text{ GeV}$ (bottom-left), and $35 \text{ GeV} < p_T < 40 \text{ GeV}$ (bottom-right). The area under the curves has been normalized to 1 for better direct comparison.

The RMS widths for the matched calorimeter jets and track jets are used to compute weights ($\frac{1}{\text{RMS}^2}$). These weights are normalized and multiplied by the unit three vectors corresponding to each jet axis. These weighted vectors are then added, and the original calorimeter jet four vector magnitude is restored.

The ϕ resolution of these “combination” jets for representative truth p_T bins, as well as the corresponding calorimeter jet and track jet resolutions, is shown in Figure [A.14](#) for the central eta bin ($|\eta| < 0.3$). The ϕ resolution (as measured by RMS) as a function of truth jet p_T is shown in Figure [A.15](#). It is clear from these figures that the combination scheme both suppresses the tails of the resolution distributions and substantially reduces their width (up to 50%), even at reasonably high truth p_T . Studies with the goal of large-scale implementation of a track-based angular correction to ATLAS calorimeter jets continue.

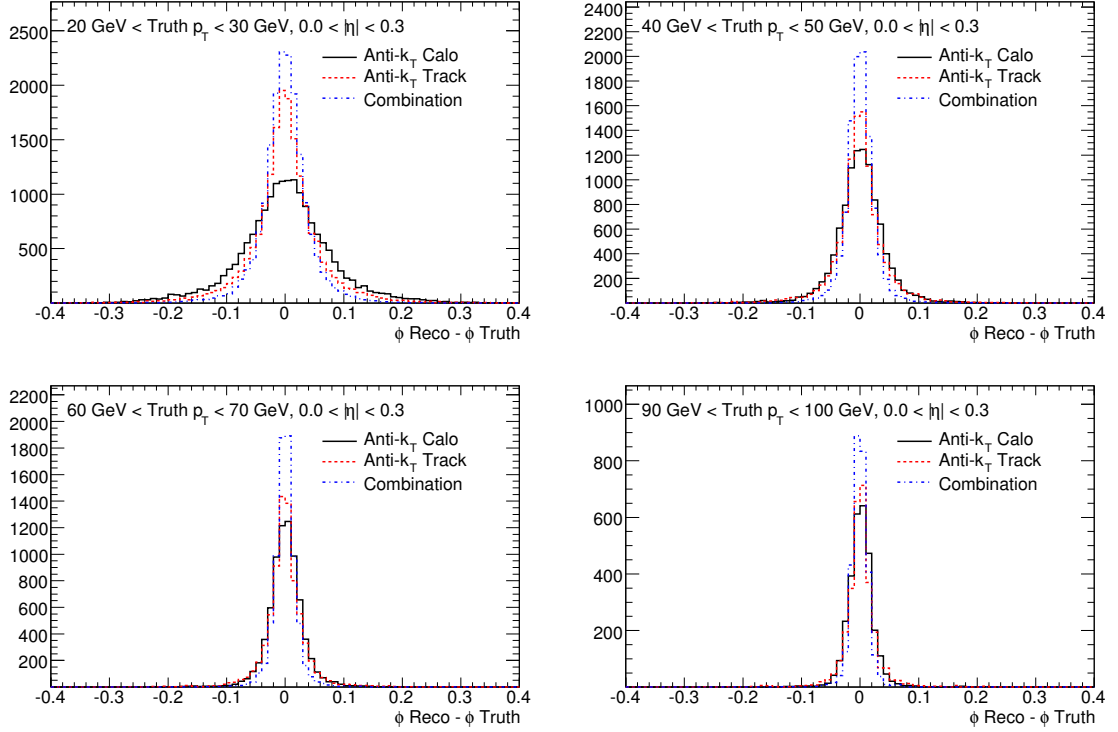


Figure A.14: Anti- k_t $R = 0.4$ track jet, calorimeter jet, and combination jet resolution ($\phi_{reco} - \phi_{truth}$) in simulated di-jets (reconstructed $|\eta| < 0.3$) for truth p_T bins $20 \text{ GeV} < p_T < 30 \text{ GeV}$ (top-left), $40 \text{ GeV} < p_T < 50 \text{ GeV}$ (top-right), $60 \text{ GeV} < p_T < 70 \text{ GeV}$ (bottom-left), and $90 \text{ GeV} < p_T < 100 \text{ GeV}$ (bottom-right).




Figure A.15: Anti- k_t $R = 0.4$ track jet, calorimeter jet, and combination jet resolution ($\phi_{reco} - \phi_{truth}$) RMS in simulated di-jets (reconstructed $|\eta| < 0.3$) as a function of truth p_T .

Appendix B

PROOF, the Parallel ROOT Facility

ROOT [180] is an object-oriented C++ class framework intended for large-scale analysis of event-like data. Written specifically for high-energy physics (HEP), it has been a key component of the computing infrastructure of every major HEP experiment since the mid-1990s. The standard analysis procedure in ROOT is sequential processing of events held in within classes called trees. In ATLAS, tree-based analysis with ROOT is by far the most popular method of doing physics.

Parallel processing of large datasets is traditionally accomplished via batch systems, where stand-alone jobs with specified event ranges are run independently on many different batch nodes. Each node typically corresponds to a single CPU core on a batch computer, which retrieves its assigned data from a server over a network connection. Monitoring, job management, and merging of output is typically handled by the user.

PROOF [181], the Parallel ROOT Facility, is ROOT’s solution for parallel processing tree-based datasets. Instead of running multiple independent jobs on single cores, PROOF jobs can be considered single jobs running on multiple cores simultaneously. There are two related variants: PROOF-Lite, which enables use of multiple cores on a single machine in a single interactive job, and “full” PROOF, which does the same using a cluster of machines and has the option of a queue to handle batch-style job

submission.

PROOF’s “single job” approach enables active job monitoring, management, and load-balancing. This last feature can result in a significant decrease in total job time over the equivalent batch submission, and all of these features help simplify the process of utilizing all cores on a single machine for analysis. However, for large datasets (tens or hundreds of millions of events), the primary advantage of a full PROOF cluster is its dual role as both dataset processor and storage. In an ideal configuration, each machine in a cluster has significant disk capacity, and these machines are grouped into an XRootD storage node [182]. A dataset is registered on the cluster by resolving the absolute storage locations of each file and cataloguing its contents. When a PROOF job is run on this dataset, each machine processes its own local files first, with transfer of data from other nodes only occurring once the local data has been exhausted.

Local input/output (I/O), in particular when RAID arrays are used, is significantly faster than network I/O. As analysis jobs tend to be I/O-limited, cluster jobs can execute in a fraction of the time an equivalent batch submission takes to complete. With jobs dominated by CPU time, the PROOF I/O configuration has no downside compared to batch, and the load-balancing is still a significant advantage.

B.1 SLAC Cluster

The first iteration of the SLAC PROOF cluster was constructed during the end of 2010 and the beginning of 2011. It had eight total machines, one of which was and remains the designated master node for the cluster and seven worker machines whose 18 terabyte (usable) twelve-disk RAID 6 arrays provide storage, I/O speed, and redundancy. Each machine has eight 2.4 gigahertz cores and 24 gigabytes of memory. Reserving one core on each worker node for the operating system, the original cluster provided 49 analysis cores and approximately 119 terabytes of dataset storage.

In early 2012, the number of machines in the cluster was doubled to sixteen, with all the new machines becoming worker nodes. This raised the total number of processing cores to 105 and the total usable dataset storage space to 255 terabytes. This version of the cluster provided much of the computing horsepower behind the

optimization and analyses presented Chapters 8, 10, and 11. The current iteration of the PROOF cluster was brought online in October of 2012 with the addition of twenty new worker machines. Each new node has twelve 2.6 gigahertz cores and 25 terabytes of usable storage, bringing the total cluster to 325 processing cores and 755 terabytes of dataset storage.

B.2 Development

The SLAC cluster is (and has been since its sixteen-machine iteration) the largest and most powerful PROOF cluster in existence. It is also has one of the heaviest workloads. This makes SLAC ground zero for tests of PROOF's ability to scale, both with regard to stability and performance. Intermittent stability issues are particularly problematic, as PROOF's single job model is not well-suited to partial retries of failed sub-jobs and rerunning the entire job can be time-consuming. The failure of the official developers to reproduce and address these issues has motivated significant PROOF development activities at SLAC. The results of this work have been continuously merged back into the main PROOF code base.

Sub-mergers by host The first PROOF development project addressed merging performance. The traditional PROOF method of merging output objects was a direct merge performed on the master node with a single core. This method scales linearly with the number of worker cores, and fails to exploit any of the parallelization of the cluster. A new method called sub-merging was recently introduced by the PROOF development team as a way to address this issue. With sub-merging, some worker nodes (for n workers, roughly \sqrt{n}), become sub-mergers which pre-merge output from a certain number of workers before passing it on to the master. As these pre-merges happen more-or-less in parallel, the total CPU time for the merging is proportional to $2 \times \sqrt{n}$ rather than n .

Use of sub-mergers sped up the merging process considerably. However, the method of choosing which worker nodes to promote was solely based on the order

in which worker nodes completed processing. For clusters made up of multi-core machines (the norm), this would often result in most of the sub-mergers being placed on a single physical machine. In this case, the sub-mergers would still be faster than the direct merge in terms of CPU time because multiple cores would be used. However, they would lose ground in real time, as this one physical machine would have to collect output from every other machine in the cluster. This would quickly saturate the merging machine's network connection. Even when the sub-mergers were more or less evenly spread in terms of the physical cluster machines, the workers were assigned to sub-mergers on a first-come first-serve basis, with no regard to their relative physical location.

The SLAC solution was to add an option to promote exactly as many sub-mergers as exist physical cluster machines (made generic through use of hostnames), and have each sub-merger handle the output from their machine's workers only. This completely eliminates any network traffic between nodes during merging, which is not only more stable but much faster due to the vast speed and latency differences between local and network I/O. In tests on the SLAC cluster, this feature reduced merging time by about a factor of 10 on average.

Single-file merge A special case for merging which typically arises only when dealing with signal grids (running many samples as one job with a different output file for each signal sample) is the single-file merge. This occurs when a sub-merger or master merger only receives a single output of a given type. In the case of the sub-merger, it only needs to pass the file on to the master node, whereas in the case of the master, the output file is complete as-is. This situation, however, was not foreseen by the PROOF development team, and therefore the file was still being opened and all objects contained within uncompressed and recompressed as would be done during merging. These unnecessary operations wasted a large amount of CPU time. The solution, of course, is to detect this situation and simply copy the compressed file to its new location, which is what is now done thanks to SLAC's efforts.

Crashes due to network disconnections The sub-merger and single-file optimizations were motivated by stability as well as performance concerns, particularly in the case of the sub-mergers. It had been noted in the early days of the SLAC cluster that the more network activity during a job, the more likely a cluster crash. These crashes were catastrophic, occurring in the PROOF server daemon itself rather than individual worker threads and bringing down all concurrent jobs, even those submitted by other users. Through log file examination, it became clear that the crashes were almost always preceded by a network disconnection. Strategies to avoid and deal with such crashes such as sub-mergers by host, cron jobs to restart crashed PROOF servers, and simply not running concurrent jobs were developed. While these strategies were effective in that they made the cluster usable, the underlying problem remained difficult to reproduce and thus hard to troubleshoot. Brookhaven National Laboratory's 96-core cluster exhibited similar intermittent symptoms, indicating that the problem was not site-specific.

The problem was finally identified in June 2012 (one and a half years after the cluster had been originally constructed) via the accidental discovery of a set of analysis jobs that always crashed when run concurrently and a month-long period of dedicated debugging. At issue was not the network disconnections themselves, which seem to be unavoidable in periods of high load, but the PROOF daemon's response to a disconnection. In short, a disconnection results in a method being called which "recycles" the connection object. This recycling involves deletion of several sub-objects used to communicate over the connection. Unfortunately, other components of the PROOF daemon store pointers to some of these sub-objects, and are not notified when they are deleted. These "dangling" pointers, when accessed post-disconnection, cause segmentation faults and bring down the entire PROOF server.

The solution was, of course, to find and remove the dangling pointers on disconnection. While the implementation of the fix was fairly complicated, the desired result was achieved. Disconnection-related PROOF server crashes seem to have been completely eliminated.

Recovery from network disconnections The solution to the network disconnection crashes ensured that concurrent jobs would not be affected by disconnections in other jobs. It did not, however, save the job which had experienced the disconnection. That particular job became unresponsive and had to be manually killed. PROOF server code to ensure proper reconnection theoretically was in place, but due to the disconnection crashes it had never been tested and did not work as intended.

There were three main problems to be fixed. The first was how to distinguish erroneous disconnections from the normal disconnections involved in the end of a job. Several methods were explored. The best was based on intercepting a “destroy” signal sent to the worker processes during job termination. Without a destroy signal, it is assumed the disconnection is random and that a reconnection attempt will be made. The second issue was that the worker processes must be protected during the the reconnection period (300s) from the automatic cleanup methods whose job it is to remove orphan sessions. Finally, connections have unique identifier numbers and these numbers are imprinted on the worker threads during their construction. After a successful reconnection, the new connection will only be able to talk to the worker sessions if it has the same identifier number as the old connection did. As these identifiers are recycled and assigned based on the lowest number available, this almost never happened in the original code.

With solutions in place for all three problems, disconnections are now fully recoverable, and, in fact, go totally unnoticed during cluster operation. While poorly-written user code can still generate cluster-wide disruptions, the PROOF daemons themselves are now stable on the timescale of several months. This has greatly increased the reliability and usefulness of the cluster as a whole.

B.3 Cluster Configuration Optimization

The PROOF software is only one component of cluster performance. Understanding and optimizing the cluster storage configuration is just as important.

Output disk on master node In the first iteration of the SLAC cluster, output files were placed on the same XRootD cluster as the input (dataset files). However, it is the master node which performs the final merging. A much faster (and less network-intensive) system was to add a reasonably-sized disk to the master node itself and store the output there. This completely eliminates the second network copy of the output back to the worker machines.

Dataset rebalancing PROOF's central advantage is that most of the data processed by each node is on that node's local disk. If datasets are imbalanced, some nodes run out of data long before the others. The nodes with data remaining must then both continue processing locally and ship data to idle workers over the network. This is a performance disaster for I/O-dominated jobs.

The methods used to download to the cluster typically do a reasonable job of evenly spreading the data over the available storage nodes. However, with the large cluster expansions which have now occurred twice, all datasets downloaded before the expansion become extremely imbalanced. This imbalance can then be propagated to new downloads, as XRootD attempts to fill up empty disks before adding to partially-full ones.

The cluster expansions prompted development of a rather complicated script to rebalance cluster datasets. As a first step, a list of files in terms of their XRootD URLs is resolved in terms of each file's absolute location in the cluster. After the file name resolution, the script has several operation modes. The simplest distributes the files "evenly" according to given machine weights (the weights reflect differences in disk capacity between the eight and twelve-core machines) in the minimum number of file moves. This mode essentially assumes that all files in the list are approximately the same size, which is not a good assumption for heterogenous datasets (i.e. an inclusive W +jets dataset including files with different number of additional partons). A slightly more nuanced option is the "spread" mode, which uses the file names, regular expressions, and multiple passes to first spread out the individual datasets (defined as those files which share a dataset or run number, i.e. $W + 3$ partons), and then ensure that the total inclusive dataset is evenly spread as well. This is the

most common mode for background Monte Carlo and collision data, as most files in an individual dataset or run are nearly the same size.

A unique operation mode intended for Monte Carlo signal grids is called “isolation” mode. Signal grids consist of many individual datasets, usually with only a couple files each. Most of these files are of similar size. In addition, each individual dataset has its own separate output file (there is no single merged output). Like other large datasets, a signal grid’s files need to be spread out over all cluster storage nodes. However, from a merging perspective, it is advantageous to ensure all files for each individual dataset are stored on a single machine. With this configuration, instead of the master node merging hundreds of output files in sequence with a single core, all merging operations occur on the sub-mergers in parallel. The master node simply copies the pre-merged output files to their final destination (Section B.2). Before dataset isolation and single-file merge, a 250-point signal grid which took 9 minutes to process could take upwards of 9 hours to merge. With these optimizations the merge time has been reduced to around 5 minutes.

All modes attempt to minimize the total number of file moves. Despite this, rebalancing of a single large dataset can involve moving around many tens of terabytes. Unfortunately, XRootD has no third-party copy implementation. If one uses an interactive machine to copy files from one cluster machine to another, all the data is funneled through the interactive machine via network connection. This is painfully inefficient and takes far too long to be practical. As a work-around, the rebalancing script has an option to generate failure-safe copy scripts designed to be executed on the cluster storage nodes themselves. Each copy script deals only with those files originating on the particular host where it is run. This means all operations are direct point-to-point copies. When executed in parallel using cluster management tools, rebalancing can proceed at a rate of approximately 1 terabyte per minute. At this rate a maximally imbalanced SLAC cluster could be rebalanced in about 6 hours.

Appendix C

ProofAna, an analysis framework

The construction of the SLAC PROOF cluster (Section B.1) motivated a reevaluation of the supersymmetry analysis codebase, as at least a partial rewrite would be necessary to use PROOF. The very early SLAC work on SUSY analyses had been done using a framework known as FlatAna (Flat ntuple Analysis). While possessing some useful features, it was not PROOF-compatible and had some major performance-related drawbacks. What began as a side project to make FlatAna compatible with PROOF soon morphed into a completely new project called ProofAna (for up-to-date information, see Reference [183]). Two years later, ProofAna has been the backbone of the heavy flavor SUSY effort at SLAC and is now used for several different analyses within the SLAC group and at other institutions.

C.1 Design Goals

ProofAna was written from the ground up with several design goals in mind.

Simplicity Most physicists learn programming on the job, and are far from being experts. One ProofAna goal was creation of a simple, intuitive pseudocode-like object-oriented framework usable by experts and non-experts alike.

Speed As mentioned previously, most physicists are not expert programmers. In many cases, there is no exposure to best practices and people use what works, even if the method is very inefficient in terms of their own time and computing time. As computing resources are shared, such inefficiencies affect everyone. The ProofAna framework has been heavily profiled to remove inefficiencies, and provides numerous helper functions to avoid common performance pitfalls such as those involving data retrieval.

Flexibility Not all analyzers are beginners, and a framework which can only perform certain specified tasks will not be very useful in the long term. ProofAna supports running multiple analyses in parallel on the same data, analysis “chains” which execute in sequence on a single event (and whose execution can be stopped by any analysis in the chain), and seamless submission to the SLAC PROOF cluster, batch system, local computer in single core and multi-core modes, and the GRID distributed analysis system [184]. The ProofAna object-based event data model (Section C.3) is enormously flexible and can be customized at run-time to the user’s needs.

Modularity ProofAna uses a “plug-in” style of analysis management where users can add or remove analyses at will with no changes to the rest of the package required. The ProofAna build system will automatically detect, compile, link the user code during the build process. Many auxiliary classes are needed for ATLAS ROOT analysis, and ProofAna utilizes the RootCore package build system to manage these packages in a relatively simple manner (though this is one area where the current system could be significantly improved).

Ease of Maintenance There was and remains no appetite for devotion of significant time to ProofAna code maintenance, though some maintenance is of course unavoidable. The ProofAna framework and its object-oriented event data model (Section C.3) in particular were designed to minimize maintenance by concentrating on development of a few generically useful classes rather than a menagerie of related but specific classes. The result has been a remarkably stable core class structure, with

few updates and no major overhauls.

C.2 Class Structure

Most of the core ProofAna code is contained within a few important classes. All other commonly-used classes derive from these.

ProofAna The confusingly-named `ProofAna` class inherits from the ROOT `TSelector` class, and thus provides the interface to the PROOF system. Run-time signals such as notification of the opening of a new file or the processing of a new event are sent to `ProofAna` via `TSelector` methods. `ProofAna` is charged with reading the job configuration, preparing “EventBuilder” and analysis objects and output files, running the analyses on each event, and managing output files and output trees (if applicable).

EventBuilderBase `EventBuilderBase` is the base class for EventBuilders, whose job it is to construct object-based events from flat ntuple inputs. Typically, these utilize numerous auxiliary classes to calibrate and otherwise select physics objects, and are specific to the type of input ntuple. This conversion step is necessary because while flat ntuples are convenient for disseminating data (each “branch” is an object attribute, such as number of muon hits or jet p_T , and one does not have to load branches not needed for the given analysis), an object-oriented approach (i.e. a muon is represented by a four vector with attributes) is much better for doing actual analysis.

The main purpose of the `EventBuilderBase` class is to provide a set of methods to load ntuple branches which are easy to use and maximize performance. This is done via a set of templated functions which take the branch name as an argument and the branch type as the template parameter. Only branches which are requested in the EventBuilder are loaded, automatically ensuring maximum I/O efficiency. Subsequent analyses which rebuild the event and use the same branches can use the already-loaded data. Loading unneeded ntuple branches is by far the single most costly error (in terms of processing time) committed by ROOT analyzers.

`EventBuilderBase` solves this problem elegantly.

AnalysisBase Analyses are classes designed to operate on object-based events. Jobs can be organized as single analysis classes placed after an `EventBuilder`, or as a chain of analysis classes executed one after another (the head of the chain must always be an `EventBuilder`). The `AnalysisBase` base class implements job control methods used by the `ProofAna` class to steer execution, as well as methods to book histograms, trees, and other objects in output files.

MomentObj The `MomentObj` class is truly the heart and soul of `ProofAna`. All physics object classes, the `Event` class itself, and the `Config` class used for job scheduling inherit from `MomentObj`, and the majority of their functionality is contained within it. `MomentObj` is essentially a heterogeneous container class—it is an associative array (map) with string-compatible keys which index attributes of the object. These attributes can be any plain old data (POD) type, such as integers, floating point numbers, or strings, as well as other objects as long as they inherit from the `ROOT` class `TObject`. In the object case, an indexing key can refer to a vector of objects. As `MomentObj` itself inherits from `TObject`, objects inheriting from `MomentObj` can be nested. This feature is key for the event data model discussed in Section [C.3](#).

Two features of `MomentObj` deserve mention. The first is automatic type conversion between compatible types. While type-safety purists will likely object, this feature is often very convenient for analyzers in practice. The second is memory management for sub-objects via reference-counting. When a pointer to an object is added to `MomentObj`, `MomentObj` assumes ownership over it by default. The same pointer can be added elsewhere in the `MomentObj` structure as many times as desired. These “object links” allow track objects, for instance, to reside in their own named container and also be linked to matched jet objects within the same `Event` without duplicating the track objects themselves. “Circular” links, a uncommon but serious problem with reference-counting schemes, are dealt with by manually denoting a “weak” link.

Keys While convenient for analyzers, string-based keys are absolutely terrible for performance. Comparisons between integers are a single CPU operation, while comparisons between strings must proceed character-by-character until a difference is encountered, and must compare the entirety of two strings to confirm an exact match. This problem is alleviated (to an extent somewhat dictated by the adherence of the user to ProofAna coding guidelines) by use of the `Key` class. A `Key` corresponding to a unique string is a *singleton*—all keys in all ProofAna modules associated with this same string refer to the same location in memory. Therefore, if one desires to compare the strings associated with two `Key` objects, one can instead compare the string memory addresses (an integer operation). Maps utilizing `Key` classes, such as the one in `MomentObj` and the one in `EventBuilderInterface` which manages ntuple branch names, are ordered by string memory address rather than string value, which enables fast binary searches. String concatenation is facilitated by a fast integer-based map lookup.

It should be noted that the `Key` singleton structure is purely a run-time attribute. If a `Key` object is written to disk, it is written in string form. Also, `Key` construction involves a string lookup by necessity. The user-dependent performance aspect enters here, as there is an enormous performance difference between initializing a `Key` once and initializing a new `Key` for every event. Still, performance gains are realized in both cases in the low-level ProofAna functions, which are all designed to use the `Key` class efficiently. Multiple `Key` classes (for instance, for ntuple branches and for `MomentObj`) leverage type safety to prevent coding mistakes. These classes, being identical, are differentiated by a dummy template parameter to avoid code duplication.

C.3 Event Data Model

The object-oriented ProofAna event data model is based on the `MomentObj` class and designed for maximum flexibility and minimal maintenance. There is a performance cost to this. Many analysis frameworks include different classes for different physics objects, each with attributes specific to the object they are intended to represent. For a single analysis, this can be efficient. However, for a general-purpose framework

the attribute list must be exhaustive. Not only does this bloat the event size, but it also means someone must continually maintain and update the attribute list for all physics objects.

This problem is avoided using `MomentObj` and the ability to dynamically add and remove object attributes. This comes at the cost of map lookups which tend to be fast thanks to the `Key` class (Section C.2). The complete ProofAna event data model is formed by three classes inheriting from `MomentObj`. The `Event` class has large numbers of helper methods designed to aid in writing analyses. Methods to automatically compute event weights based on dataset information or flags set during `EventBuilder` execution are also included. As its name implies, it is the top-level container class for all physics objects. Named object vectors in the event class hold either `Particle` or `Point` objects. These also inherit from `MomentObj`, but include a four-vector or a three-vector member object, respectively. All physics objects (electrons, muons, jets, vertices) can be represented fully in arbitrary detail by these two classes, and they can serve as base classes for more complicated objects as well (such as jets with multiple calibrations). As `Particle` and `Point` have all `MomentObj` features, they can also include object vectors with links to associated objects such as jet constituents or tracks.

Maintenance of the event data model consists essentially of maintenance of `MomentObj` and the `Event` helper methods, neither of which are analysis-specific. Whole or partial events can be fully persistified in object form to ROOT files, and there is a special `EventBuilder` which can read trees of `Event` objects as an input dataset for further analysis. Such `Event` trees can be created very simply with methods found in `AnalysisBase`.

C.4 Heavy Flavor Supersymmetry Analysis

The heavy flavor supersymmetry analyses in Chapters 10 and 11 utilize many of ProofAna's features. There is an `EventBuilder` specific to the standard ATLAS SUSY group ntuples, which applies quality selections and builds the object-based event structure. Variations in event-building related to systematic errors (Section 9.2) can

be turned on and off in the EventBuilder via analysis configuration flags. The analysis class contains cut flows, histograms, and also writes out an `Event` tree for the signal and control regions for quick re-histogramming. The use of a single EventBuilder and analysis (cut flow) class for all systematic variations is noteworthy only because this is so often not done in physics analysis, despite the maintenance nightmare and propensity for error that duplicate code introduces.

The job is configured in a run script which instantiates various `Config` objects, one for each analysis variation and one for ProofAna as a whole (the variations only apply to Monte Carlo simulation). The ProofAna `Config` object contains the dataset information (such as cross section and luminosity, which are necessary for proper event weighting) as well as flags controlling the merging of output objects. Each analysis `Config` object includes the name of the analysis class to use (all the same in this case), a unique name for the directory in which output objects are stored, and flags which set the variation one wishes to run. For heavy flavor supersymmetry, jet energy scale (up and down), jet energy resolution (up only), and \cancel{E}_T cluster scale (up and down) and resolution (up and down) are all included as separate event construction and analysis jobs, as these variations modify the event itself. The nominal event, heavy flavor tagging variations (b -tagging efficiency up and down, c -tagging efficiency up and down, and light tagging efficiency up and down), and pileup reweighting analyses are all run as a single `AnalysisChain` where the event is reconstructed once and passed from analysis to analysis where different weights are applied.

All of this configuration is contained within the run script, and no recompilation is required for changes. When the run script is executed, the `Config` objects are constructed and saved to a configuration `ROOT` file along with any configured auxiliary classes needed by the EventBuilder. These objects are shipped to every worker node, whether there is only one local core or an entire PROOF cluster. On each worker node, a `ProofAna` object reads in the analysis `Config` objects and creates and configures heavy flavor SUSY analysis class objects, as well as a SUSY EventBuilder. It also creates the desired output file(s) and directory objects in each output file corresponding to the analyses' names.

When event execution begins, the EventBuilder reconstructs the event according

to the first analysis' configuration flags. As this is the first analysis, all the ntuple data needed is loaded as it is requested. The `Event` object is passed to the first analysis, which runs its cut flows and books and fills its histograms in the corresponding output directory. If this analysis is part of a chain, the event is then passed to the next analysis in the chain, which does the same. If not, the `Event` is deleted, and the `EventBuilder` is run again for the next analysis in the list. This time, however, most if not all the branches used are already loaded, so data retrieval is extremely fast. This cycle continues until all analyses have been run, and then starts anew for the next event. When all events have been processed, the output files are merged.

The final output files contain one histogram-filled directory for each analysis run. A macro is used to convert this collection of varied histograms to a single histogram with per-bin systematic error estimates. A complete analysis rerun (approximately 500 million events) on the sixteen-machine SLAC PROOF cluster (Section B.1) took approximately 4 hours. On the current cluster it can be done in about one hour. The process could be made faster still by removing rejected events from the dataset (skimming), but there is hardly any need given the speed of the cluster hardware and the efficiency of the ProofAna framework.

Bibliography

- [1] ATLAS Collaboration. Search for supersymmetry in pp collisions at $\sqrt{s} = 7$ TeV in final states with missing transverse momentum, b -jets and no leptons with the ATLAS detector. ATLAS-CONF-2011-098, <http://cdsweb.cern.ch/record/1369212>, July 2011.
- [2] G. Aad et al. Search for Scalar Bottom Quark Pair Production with the ATLAS Detector in pp Collisions at $\sqrt{s} = 7$ TeV. *Physical Review Letters*, **108(18):181802**, May 2012.
- [3] G. Aad et al. Search for supersymmetry in pp collisions at $\sqrt{s} = 7$ TeV in final states with missing transverse momentum and b -jets with the ATLAS detector. *Physical Review D*, **85(11):112006**, June 2012.
- [4] ATLAS Collaboration. Search for scalar bottom pair production in final states with missing transverse momentum and two b -jets in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS Detector. ATLAS-CONF-2012-106, <http://cdsweb.cern.ch/record/1472685>, August 2012.
- [5] G. Aad et al. Search for top and bottom squarks from gluino pair production in final states with missing transverse energy and at least three b -jets with the ATLAS detector. *The European Physical Journal C*, **72(10):2174**, October 2012.
- [6] ATLAS Collaboration. Search for direct sbottom production in event with two b -jets using 12.8 fb^{-1} of pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS Detector. ATLAS-CONF-2012-165, <http://cdsweb.cern.ch/record/1497668>, December 2012.

- [7] ATLAS Collaboration. Search for gluino pair production in final states with missing transverse momentum and at least three b -jets using 12.8 fb^{-1} of pp collisions at $\sqrt{s} = 8 \text{ TeV}$ with the ATLAS Detector. ATLAS-CONF-2012-145, <http://cdsweb.cern.ch/record/1493484>, November 2012.
- [8] G. Aad et al. Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Physics Letters B*, **716(1):1–29**, September 2012.
- [9] S. Chatrchyan et al. Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Physics Letters B*, **716(1):30–61**, September 2012.
- [10] M. Regis. Particle Dark Matter: the state of the art. [arXiv:1008.0506](https://arxiv.org/abs/1008.0506), August 2010.
- [11] L. Canetti, M. Drewes, and M. Shaposhnikov. Matter and antimatter in the universe. *New Journal of Physics*, **14(9):095012**, September 2012.
- [12] H. Miyazawa. Baryon Number Changing Currents. *Progress of Theoretical Physics*, **36(6):1266–1276**, December 1966.
- [13] P. Ramond. Dual Theory for Free Fermions. *Physical Review D*, **3(10):2415–2418**, May 1971.
- [14] Y. Golfand and E. Likhtman. Extension of the Algebra of Poincare Group Generators and Violation of P invariance. *Journal of Experimental and Theoretical Physics Letters*, **13(8):323**, April 1971.
- [15] A. Neveu and J. Schwarz. Quark Model of Dual Pions. *Physical Review D*, **4(4):1109–1111**, August 1971.
- [16] A. Neveu and J. Schwarz. Factorizable dual model of pions. *Nuclear Physics B*, **31(1):86–112**, August 1971.

- [17] J. Gervais and B. Sakita. Field theory interpretation of supergauge in dual models. *Nuclear Physics B*, 34(2):632–639, November 1971.
- [18] D. Volkov and V. Akulov. Is the neutrino a goldstone particle? *Physics Letters B*, 46(1):109–110, September 1973.
- [19] J. Wess and B. Zumino. A lagrangian model invariant under supergauge transformations. *Physics Letters B*, 49(1):52–54, March 1974.
- [20] J. Wess and B. Zumino. Supergauge transformations in four dimensions. *Nuclear Physics B*, 70(1):39–50, February 1974.
- [21] S. Troitsky. Unsolved problems in particle physics. *Physics-Uspekhi*, 55(1):72–95, January 2012.
- [22] J. Alwall, P. Schuster, and N. Toro. Simplified models for a first characterization of new physics at the LHC. *Physical Review D*, 79(7):075020, April 2009.
- [23] D. Alves, E. Izaguirre, and J. Wacker. Where the Sidewalk Ends: Jets and Missing Energy Search Strategies for the 7 TeV LHC. [arXiv:1102.5338](https://arxiv.org/abs/1102.5338), February 2011.
- [24] D. Alves et al. Simplified Models for LHC New Physics Searches. [arXiv:1105.2838](https://arxiv.org/abs/1105.2838), May 2011.
- [25] G. Aad et al. The ATLAS Experiment at the CERN Large Hadron Collider. *Journal of Instrumentation*, 3(8):S08003, August 2008.
- [26] L. Evans and P. Bryant. LHC Machine. *Journal of Instrumentation*, 3(08):S08001, August 2008.
- [27] ATLAS Collaboration. Search for direct top squark pair production in final states with one isolated lepton, jets, and missing transverse momentum in $\sqrt{s} = 8$ TeV pp collisions using 13.0 fb⁻¹ of ATLAS data. ATLAS-CONF-2012-166, <http://cdsweb.cern.ch/record/1497732>, December 2012.

- [28] ATLAS Collaboration. Search for a supersymmetric top-quark partner in final states with two leptons in $\sqrt{s} = 8$ TeV pp collisions using 13 fb^{-1} of ATLAS data. ATLAS-CONF-2012-167, <http://cdsweb.cern.ch/record/1497733>, December 2012.
- [29] G. Aad et al. Search for Direct Top Squark Pair Production in Final States with One Isolated Lepton, Jets, and Missing Transverse Momentum in $\sqrt{s} = 7$ TeV pp Collisions Using 4.7 fb^{-1} of ATLAS Data. *Physical Review Letters*, [109\(21\):211803](#), November 2012.
- [30] G. Aad et al. Search for a Supersymmetric Partner to the Top Quark in Final States with Jets and Missing Transverse Momentum at $\sqrt{s} = 7$ TeV with the ATLAS Detector. *Physical Review Letters*, [109\(21\):211802](#), November 2012.
- [31] G. Aad et al. Search for a heavy top-quark partner in final states with two leptons with the ATLAS detector at the LHC. *Journal of High Energy Physics*, [2012\(11\):94](#), November 2012.
- [32] G. Aad et al. Search for light scalar top-quark pair production in final states with two leptons with the ATLAS detector in $\sqrt{s} = 7$ TeV proton-proton collisions. *The European Physical Journal C*, [72\(11\):2237](#), November 2012.
- [33] G. Aad et al. Search for scalar top quark pair production in natural gauge-mediated supersymmetry models with the ATLAS detector in pp collisions at $\sqrt{s} = 7$ TeV. *Physics Letters B*, [715\(1-3\):44–60](#), August 2012.
- [34] ATLAS Collaboration. Search for light top squark pair production in final states with leptons and b -jets with the ATLAS detector in $\sqrt{s} = 7$ TeV proton-proton collisions. [arXiv:1209.2102](https://arxiv.org/abs/1209.2102), September 2012.
- [35] A. Pich. QUANTUM CHROMODYNAMICS. [arXiv:9505231 \[hep-ph\]](https://arxiv.org/abs/9505231), May 1995.
- [36] M. Peskin and D. Schroeder. *An Introduction To Quantum Field Theory*. Westview Press, 1995.

- [37] K. Wilson. Confinement of quarks. *Physical Review D*, 10(8):2445–2459, October 1974.
- [38] R. Gupta. Introduction to Lattice QCD. [arXiv:9807028 \[hep-lat\]](#), July 1998.
- [39] Yang-Mills Millenium Problem, http://www.claymath.org/millennium/Yang-Mills_Theory.
- [40] S. Glashow. Partial-symmetries of weak interactions. *Nuclear Physics*, 22(4):579–588, February 1961.
- [41] S. Weinberg. A Model of Leptons. *Physical Review Letters*, 19(21):1264–1266, November 1967.
- [42] S. Novaes. Standard Model: An Introduction. [arXiv:0001283 \[hep-ph\]](#), January 2000.
- [43] F. Englert and R. Brout. Broken Symmetry and the Mass of Gauge Vector Mesons. *Physical Review Letters*, 13(9):321–323, August 1964.
- [44] P. Higgs. Broken Symmetries and the Masses of Gauge Bosons. *Physical Review Letters*, 13(16):508–509, October 1964.
- [45] G. Guralnik, C. Hagen, and T. Kibble. Global Conservation Laws and Massless Particles. *Physical Review Letters*, 13(20):585–587, November 1964.
- [46] T. Nakano and K. Nishijima. Charge Independence for V-particles. *Progress of Theoretical Physics*, 10(5):581–582, November 1953.
- [47] M. Gell-Mann. The interpretation of the new particles as displaced charge multiplets. *Il Nuovo Cimento*, 4(S2):848–866, April 1956.
- [48] S. Dawson. Introduction to Electroweak Symmetry Breaking. [arXiv:9901280 \[hep-ph\]](#), January 1999.

- [49] E. Majorana. Teoria simmetrica dellelettrone e del positrone. *Il Nuovo Cimento*, 14(4):171–184, April 1937.
- [50] T. Yanagida. Horizontal Symmetry and Masses of Neutrinos. *Progress of Theoretical Physics*, 64(3):1103–1105, June 1980.
- [51] R. Mohapatra and G. Senjanović. Neutrino Mass and Spontaneous Parity Non-conservation. *Physical Review Letters*, 44(14):912–915, April 1980.
- [52] J. Schechter and J. Valle. Neutrino masses in $SU(2) \otimes U(1)$ theories. *Physical Review D*, 22(9):2227–2235, November 1980.
- [53] M. Auger et al. Search for Neutrinoless Double-Beta Decay in ^{136}Xe with EXO-200. *Physical Review Letters*, 109(3):032505, July 2012.
- [54] F. P. An et al. Improved Measurement of Electron Antineutrino Disappearance at Daya Bay. [arXiv:1210.6327](https://arxiv.org/abs/1210.6327), October 2012.
- [55] Y. Abe et al. Indication of Reactor ν_e Disappearance in the Double Chooz Experiment. *Physical Review Letters*, 108(13):131801, March 2012.
- [56] N. Cabibbo. Unitary Symmetry and Leptonic Decays. *Physical Review Letters*, 10(12):531–533, June 1963.
- [57] M. Kobayashi and T. Maskawa. CP-Violation in the Renormalizable Theory of Weak Interaction. *Progress of Theoretical Physics*, 49(2):652–657, February 1973.
- [58] B. Pontecorvo. Neutrino Experiments and the Problem of Conservation of Leptonic Charge. *Soviet Physics JETP*, 26, May 1968.
- [59] Z. Maki, M. Nakagawa, and S. Sakata. Remarks on the Unified Model of Elementary Particles. *Progress of Theoretical Physics*, 28(5):870–880, November 1962.
- [60] S. Martin. A Supersymmetry Primer. [arXiv:9709356](https://arxiv.org/abs/9709356) [hep-ph], September 1997.

- [61] M. Peskin. Supersymmetry in Elementary Particle Physics. [arXiv:0801.1928](#), January 2008.
- [62] M. Perelstein. Little Higgs models and their phenomenology. *Progress in Particle and Nuclear Physics*, 58(1):247–291, January 2007.
- [63] M. Shifman. Large Extra Dimensions: Becoming Acquainted with an Alternative Paradigm. *International Journal of Modern Physics A*, 25(02n03):199–225, January 2010.
- [64] C. Macesanu. The Phenomenology of Universal Extra Dimensions at Hadron Colliders. *International Journal of Modern Physics A*, 21(11):2259–2296, April 2006.
- [65] S. Coleman and J. Mandula. All Possible Symmetries of the S Matrix. *Physical Review*, 159(5):1251–1256, July 1967.
- [66] A. Masiero and O. Vives. Flavour and CP violation in supersymmetry. *New Journal of Physics*, 4(1):4, February 2002.
- [67] M. Drees, R. Godbole, and P. Roy. *Theory and Phenomenology of Sparticles: An Account of Four-Dimensional $N=1$ Supersymmetry in High Energy Physics*. World Scientific Publishing Company, 2005.
- [68] H. Baer, V. Barger, P. Huang, A. Mustafayev, and X. Tata. Radiative natural SUSY with a 125 GeV Higgs boson. [arXiv:1207.3343](#), July 2012.
- [69] H. Haber, R. Hempfling, and A. Hoang. Approximating the radiatively corrected Higgs mass in the Minimal Supersymmetric Model. [arXiv:9609331 \[hep-ph\]](#), September 1996.
- [70] L. Hall, D. Pinner, and J. Ruderman. A Natural SUSY Higgs Near 125 GeV. [arXiv:1112.2703](#), December 2011.
- [71] G. Bertone. The moment of truth for WIMP dark matter. *Nature*, 468(7322):389–93, November 2010.

- [72] W.-M. Yao et al. Review of Particle Physics. *Journal of Physics G: Nuclear and Particle Physics*, **33**(1), July 2006.
- [73] A. Chamseddine, R. Arnowitt, and P. Nath. Locally Supersymmetric Grand Unification. *Physical Review Letters*, **49**(14):970–974, October 1982.
- [74] R. Barbieri, S. Ferrara, and C. Savoy. Gauge models with spontaneously broken local supersymmetry. *Physics Letters B*, **119**(4-6):343–347, December 1982.
- [75] L. Ibáñez. Locally supersymmetric SU(5) grand unification. *Physics Letters B*, **118**(1-3):73–78, December 1982.
- [76] ATLAS Collaboration. Search for squarks and gluinos with the ATLAS detector using final states with jets and missing transverse momentum and 5.8 fb⁻¹ of $\sqrt{s} = 8$ TeV proton-proton collision data. ATLAS-CONF-2012-109, <http://cdsweb.cern.ch/record/1472710>, August 2012.
- [77] B. Dutta, T. Kamon, A. Krislock, N. Kolev, and Y. Oh. Determination of nonuniversal supergravity models at the Large Hadron Collider. *Physical Review D*, **82**(11):115009, December 2010.
- [78] W. Beenakker et al. Squark and Gluino Hadroproduction. *International Journal of Modern Physics A*, **26**(16):2637–2664, June 2011.
- [79] W. Beenakker et al. Supersymmetric top and bottom squark production at hadron colliders. *Journal of High Energy Physics*, **2010**(8):98, August 2010.
- [80] W. Beenakker, M. Krämer, T. Plehn, M. Spira, and P. Zerwas. Stop production at hadron colliders. *Nuclear Physics B*, **515**(1-2):3–14, March 1998.
- [81] W. Beenakker, R. Hoepker, and M. Spira. PROSPINO: A Program for the Production of Supersymmetric Particles in Next-to-leading Order QCD. [arXiv:9611232 \[hep-ph\]](https://arxiv.org/abs/9611232), November 1996.
- [82] C. Csáki, J. Heinonen, and M. Perelstein. Testing gluino spin with three-body decays. *Journal of High Energy Physics*, **2007**(10):107, November 2007.

- [83] M. M. Nojiri and J. Shu. Two jets and missing E_T signature to determine the spins of new particles. *Journal of High Energy Physics*, 2011(6):47, June 2011.
- [84] C. Smith. How the LHC came to be. *Nature*, 448(7151):281–4, July 2007.
- [85] M. Lamont. LHC: status and commissioning plans. [arXiv:0906.0347](https://arxiv.org/abs/0906.0347), June 2009.
- [86] S. Chatrchyan et al. The CMS experiment at the CERN LHC. *Journal of Instrumentation*, 3(08):S08004, August 2008.
- [87] K. Aamodt et al. The ALICE experiment at the CERN LHC. *Journal of Instrumentation*, 3(08):S08002, August 2008.
- [88] A. Alves et al. The LHCb Detector at the LHC. *Journal of Instrumentation*, 3(08):S08005, August 2008.
- [89] G. Aad et al. Expected Performance of the ATLAS Experiment - Detector, Trigger and Physics. [arXiv:0901.0512](https://arxiv.org/abs/0901.0512), December 2008.
- [90] G. Aad et al. Charged-particle multiplicities in pp interactions measured with the ATLAS detector at the LHC. *New Journal of Physics*, 13(5):053033, May 2011.
- [91] ATLAS Collaboration. The ATLAS Pixel Insertable B-Layer (IBL). ATL-INDET-PROC-2010-034, <http://cdsweb.cern.ch/record/1303015>, October 2010.
- [92] ATLAS Collaboration. ATLAS Insertable B-Layer Technical Design Report. ATLAS-TDR-019, <http://cdsweb.cern.ch/record/1291633>, September 2010.
- [93] A. La Rosa. ATLAS IBL Pixel Upgrade. [arXiv:1104.1980](https://arxiv.org/abs/1104.1980), April 2011.
- [94] F. Campabadal et al. Design and performance of the ABCD3TA ASIC for readout of silicon strip detectors in the ATLAS semiconductor tracker. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators*,

- Spectrometers, Detectors and Associated Equipment*, 552(3):292–328, November 2005.
- [95] E. Abat et al. The ATLAS TRT Barrel Detector. *Journal of Instrumentation*, 3(02):P02014, February 2008.
- [96] E. Abat et al. The ATLAS TRT electronics. *Journal of Instrumentation*, 3(06):P06007, June 2008.
- [97] E. Abat et al. Combined performance studies for electrons at the 2004 ATLAS combined test-beam. *Journal of Instrumentation*, 5(11):P11006, November 2010.
- [98] M. Aharrouche et al. Measurement of the response of the ATLAS liquid argon barrel calorimeter to electrons at the 2004 combined test-beam. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 614(3):400–432, March 2010.
- [99] J. Pinfold et al. Performance of the ATLAS liquid argon endcap calorimeter in the pseudorapidity region $2.5 < |\eta| < 4.0$ in beam tests. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 593(3):324–342, August 2008.
- [100] M. Aharrouche. The ATLAS liquid argon calorimeter: Construction, integration, commissioning and combined test beam results. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 581(1-2):373–376, October 2007.
- [101] G. Aad et al. Readiness of the ATLAS Tile Calorimeter for LHC collisions. *The European Physical Journal C*, 70(4):1193–1236, December 2010.
- [102] P. Adragna et al. Testbeam studies of production modules of the ATLAS Tile Calorimeter. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 606(3):362–394, July 2009.

- [103] M. Aharrouche et al. Study of the response of ATLAS electromagnetic liquid argon calorimeters to muons. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 606(3):419–431, July 2009.
- [104] ATLAS Collaboration. *ATLAS liquid-argon calorimeter : Technical Design Report*. CERN, Geneva, 1996.
- [105] I. Koletsou. The ATLAS Liquid Argon Calorimeter at the LHC. ATL-LARG-PROC-2010-001, <http://cdsweb.cern.ch/record/1247962>, March 2010.
- [106] G. Aad et al. Commissioning of the ATLAS Muon Spectrometer with cosmic rays. *The European Physical Journal C*, 70(3):875–916, October 2010.
- [107] ATLAS Collaboration. Monitored Drift Tubes Chambers for Muon Spectroscopy in ATLAS. ATL-MUON-94-044, <http://cdsweb.cern.ch/record/685844>, June 1994.
- [108] A. Aloisio et al. The Muon Spectrometer Barrel Level-1 Trigger of the ATLAS Experiment at LHC. *IEEE Transactions on Nuclear Science*, 53(4):2446–2451, August 2006.
- [109] H. Ten Kate. Superconducting magnet system for the ATLAS detector at CERN. *IEEE Transactions on Applied Superconductivity*, 9(2):841–846, June 1999.
- [110] R. Ruber et al. Ultimate Performance of the ATLAS Superconducting Solenoid. *IEEE Transactions on Applied Superconductivity*, 17(2):1201–1204, June 2007.
- [111] P. Veldre et al. Manufacturing and Integration Progress of the ATLAS Barrel Toroid Magnet at CERN. *IEEE Transactions on Applied Superconductivity*, 14(2):491–494, June 2004.
- [112] H. Pernegger. First test results of a high-speed beam conditions monitor for the ATLAS experiment. *IEEE Transactions on Nuclear Science*, 52(5):1590–1594, October 2005.

- [113] P. Waller. ATLAS Data Quality Monitoring: Experience with First Collision Data. ATL-DAPR-PROC-2010-001, <http://cdsweb.cern.ch/record/1310319>, November 2010.
- [114] J. Frost. Commissioning of ATLAS Data Quality Infrastructure with First Collision Data. ATL-SOFT-PROC-2011-005, <http://cdsweb.cern.ch/record/1321627>, January 2011.
- [115] G. Aad et al. The ATLAS Simulation Infrastructure. *The European Physical Journal C*, 70(3):823–874, September 2010.
- [116] S. Agostinelli et al. Geant4 - a simulation toolkit. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3):250–303, July 2003.
- [117] M. Bähr et al. Herwig++ physics and manual. *The European Physical Journal C*, 58(4):639–707, November 2008.
- [118] J. Alwall, M. Herquet, F. Maltoni, O. Mattelaer, and T. Stelzer. MadGraph 5: going beyond. *Journal of High Energy Physics*, 2011(6):128, June 2011.
- [119] M. Mangano, F. Piccinini, A. Polosa, M. Moretti, and R. Pittau. ALPGEN, a generator for hard multiparton processes in hadronic collisions. *Journal of High Energy Physics*, 2003(07):1, July 2003.
- [120] S. Frixione and B. Webber. Matching NLO QCD computations and parton shower simulations. *Journal of High Energy Physics*, 2002(06):29, June 2002.
- [121] S. Frixione, P. Nason, and B. Webber. Matching NLO QCD and parton showers in heavy flavour production. *Journal of High Energy Physics*, 2003(08):7, August 2003.
- [122] S. Frixione, E. Laenen, P. Motylinski, and B. Webber. Single-top production in MC@NLO. *Journal of High Energy Physics*, 2006(03):92, March 2006.

- [123] T. Gleisberg et al. Event generation with SHERPA 1.1. *Journal of High Energy Physics*, 2009(02):7, February 2009.
- [124] G. Corcella et al. HERWIG 6: an event generator for hadron emission reactions with interfering gluons (including supersymmetric processes). *Journal of High Energy Physics*, 2001(01):10, January 2001.
- [125] J. M. Butterworth, J. R. Forshaw, and M. H. Seymour. Multiparton interactions in photoproduction at HERA. *Zeitschrift für Physik C: Particles and Fields*, 72(4):637–646, December 1996.
- [126] T. Sjöstrand, S. Mrenna, and P. Skands. PYTHIA 6.4 physics and manual. *Journal of High Energy Physics*, 2006(05):26, May 2006.
- [127] ATLAS Collaboration. First tuning of HERWIG/JIMMY to ATLAS data. ATL-PHYS-PUB-2010-014, <http://cdsweb.cern.ch/record/1303025>, October 2010.
- [128] ATLAS Collaboration. Charged particle multiplicities in pp interactions at $\sqrt{s} = 0.9$ and 7 TeV in a diractive limited phase-space measured with the ATLAS detector at the LHC and new PYTHIA6 tune. ATLAS-CONF-2010-031, <http://cdsweb.cern.ch/record/1277665>, July 2010.
- [129] H. Lai et al. New parton distributions for collider physics. *Physical Review D*, 82(7):074024, October 2010.
- [130] J. Pumplin et al. New Generation of Parton Distributions with Uncertainties from Global QCD Analysis. *Journal of High Energy Physics*, 2002(07):12, July 2002.
- [131] ATLAS Collaboration. Performance of the ATLAS Inner Detector Track and Vertex Reconstruction in the High Pile-Up LHC Environment. ATLAS-CONF-2012-042, <http://cdsweb.cern.ch/record/1435196>, March 2012.

- [132] G. Piacquadio, K. Prokofiev, and A. Wildauer. Primary vertex reconstruction in the ATLAS experiment at LHC. *Journal of Physics: Conference Series*, **119(3):032033**, July 2008.
- [133] ATLAS Collaboration. Performance of primary vertex reconstruction in proton-proton collisions at $\sqrt{s} = 7$ TeV in the ATLAS experiment. ATLAS-CONF-2010-069, <http://cdsweb.cern.ch/record/1281344>, July 2010.
- [134] R. Frühwirth, W. Waltenberger, and P. Vanlaer. Adaptive Vertex Fitting. CMS-NOTE-2007-008, <http://cdsweb.cern.ch/record/1027031>, March 2007.
- [135] ATLAS Collaboration. Characterization of Interaction-Point Beam Parameters Using the pp Event-Vertex Distribution Reconstructed in the ATLAS Detector at the LHC. ATLAS-CONF-2010-027, <http://cdsweb.cern.ch/record/1277659>, May 2010.
- [136] ATLAS Collaboration. Improved electron reconstruction in ATLAS using the Gaussian Sum Filter-based model for bremsstrahlung. ATLAS-CONF-2012-047, <http://cdsweb.cern.ch/record/1449796>, May 2012.
- [137] ATLAS Collaboration. Expected electron performance in the ATLAS experiment. ATL-PHYS-PUB-2011-006, <http://cdsweb.cern.ch/record/1345327>, April 2011.
- [138] G. Aad et al. Electron performance measurements with the ATLAS detector using the 2010 LHC proton-proton collision data. *The European Physical Journal C*, **72(3):1909**, March 2012.
- [139] R. Nicolaidou et al. Muon identification procedure for the ATLAS detector at the LHC using Muonboy reconstruction package and tests of its performance using cosmic rays and single beam data. *Journal of Physics: Conference Series*, **219(3):032052**, April 2010.
- [140] ATLAS Collaboration. Muon Momentum Resolution in First Pass Reconstruction of pp Collision Data Recorded by ATLAS in 2010. ATLAS-CONF-2011-046, <http://cdsweb.cern.ch/record/1338575>, March 2011.

- [141] W. Lampl et al. Calorimeter Clustering Algorithms. ATL-LARG-PUB-2008-002, <http://cdsweb.cern.ch/record/1099735>, April 2008.
- [142] G. Aad et al. Performance of missing transverse momentum reconstruction in proton-proton collisions at $\sqrt{s} = 7$ TeV with ATLAS. *The European Physical Journal C*, 72(1):1844, January 2012.
- [143] M. Cacciari, G. P. Salam, and G. Soyez. The anti- k_t jet clustering algorithm. *Journal of High Energy Physics*, 2008(04):63, April 2008.
- [144] M. Cacciari and G. P. Salam. Dispelling the N^3 myth for the k_t jet-finder. *Physics Letters B*, 641(1):57–61, September 2006.
- [145] M. Aleksa et al. ATLAS Combined Testbeam : Computation and Validation of the Electronic Calibration Constants for the Electromagnetic Calorimeter. ATL-LARG-PUB-2006-003, <http://cdsweb.cern.ch/record/942528>, April 2006.
- [146] ATLAS Collaboration. Jet energy measurement with the ATLAS detector in proton-proton collisions at $\sqrt{s} = 7$ TeV. [arXiv:1112.6426](https://arxiv.org/abs/1112.6426), December 2011.
- [147] ATLAS Collaboration. Determination of the ATLAS jet energy measurement uncertainty using photon-jet events in proton-proton collisions at $\sqrt{s} = 7$ TeV. ATLAS-CONF-2011-031, <http://cdsweb.cern.ch/record/1337781>, March 2011.
- [148] G. Aad et al. Measurement of multi-jet cross sections in proton-proton collisions at a 7 TeV center-of-mass energy. *The European Physical Journal C*, 71(11):1763, November 2011.
- [149] ATLAS Collaboration. ATLAS jet energy scale uncertainties using tracks in proton-proton collisions at $\sqrt{s} = 7$ TeV. ATLAS-CONF-2011-067, <http://cdsweb.cern.ch/record/1349308>, May 2011.
- [150] ATLAS Collaboration. In-situ jet energy scale and jet shape corrections for multiple interactions in the first ATLAS data at the LHC. ATLAS-CONF-2011-030, <http://cdsweb.cern.ch/record/1337780>, March 2011.

- [151] ATLAS Collaboration. Commissioning of the ATLAS high-performance b -tagging algorithms in the 7 TeV collision data. ATLAS-CONF-2011-102, <http://cdsweb.cern.ch/record/1369219>, July 2011.
- [152] G. Piacquadio and C. Weiser. A new inclusive secondary vertex algorithm for b -jet tagging in ATLAS. *Journal of Physics: Conference Series*, 119(3):032032, July 2008.
- [153] ATLAS Collaboration. Measurement of the b -tag Efficiency in a Sample of Jets Containing Muons with 5 fb^{-1} of Data from the ATLAS Detector. ATLAS-CONF-2012-043, <http://cdsweb.cern.ch/record/1435197>, March 2012.
- [154] 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 $\sqrt{s} = 7 \text{ TeV}$. ATLAS-CONF-2010-038, <http://cdsweb.cern.ch/record/1277678>, July 2010.
- [155] ATLAS Collaboration. Observation of inclusive electrons in the ATLAS experiment at $\sqrt{s} = 7 \text{ TeV}$. ATLAS-CONF-2010-073, <http://cdsweb.cern.ch/record/1281364>, July 2010.
- [156] ATLAS Collaboration. A measurement of the muon reconstruction efficiency in 2010 ATLAS data using J/Ψ decays. ATLAS-CONF-2012-125, <http://cdsweb.cern.ch/record/1474642>, August 2012.
- [157] ATLAS Collaboration. Muon reconstruction efficiency in reprocessed 2010 LHC proton-proton collision data recorded with the ATLAS detector. ATLAS-CONF-2011-063, <http://cdsweb.cern.ch/record/1345743>, April 2011.
- [158] ATLAS Collaboration. b -Jet Tagging Efficiency Calibration using the System8 Method. ATLAS-CONF-2011-143, <http://cdsweb.cern.ch/record/1386703>, October 2011.
- [159] ATLAS Collaboration. b -jet tagging calibration on c -jets containing D^{*+} mesons. ATLAS-CONF-2012-039, <http://cdsweb.cern.ch/record/1435193>, March 2012.

- [160] ATLAS Collaboration. Measurement of the Mistag Rate with 5 fb^{-1} of Data Collected by the ATLAS Detector. ATLAS-CONF-2012-040, <http://cdsweb.cern.ch/record/1435194>, March 2012.
- [161] ATLAS Collaboration. Measuring the b -tag efficiency in a top-pair sample with 4.7 fb^{-1} of data from the ATLAS detector. ATLAS-CONF-2012-097, <http://cdsweb.cern.ch/record/1460443>, July 2012.
- [162] H. Dreiner, M. Krämer, and J. Tattersall. How low can SUSY go? Matching, monojets and compressed spectra. [arXiv:1207.1613](https://arxiv.org/abs/1207.1613), July 2012.
- [163] H. Baer, A. Box, E.-K. Park, and X. Tata. Implications of compressed supersymmetry for collider and dark matter searches. *Journal of High Energy Physics*, 2007(8):60, August 2007.
- [164] G. Cowan, K. Cranmer, E. Gross, and O. Vitells. Asymptotic formulae for likelihood-based tests of new physics. *The European Physical Journal C*, 71(2):1–19, February 2011.
- [165] I. Hinchliffe, F. E. Paige, M. D. Shapiro, J. Söderqvist, and W. Yao. Precision SUSY measurements at CERN LHC. *Physical Review D*, 55(9):5520–5540, May 1997.
- [166] D. Tovey. On measuring the masses of pair-produced semi-invisibly decaying particles at hadron colliders. *Journal of High Energy Physics*, 2008(4):34, April 2008.
- [167] G. Polesello and D. Tovey. Supersymmetric particle mass measurement with the boost-corrected contranverse mass. *Journal of High Energy Physics*, 2010(3):30, March 2010.
- [168] E. Alvarez and Y. Bai. Reach the Bottom Line of the Sbottom Search. [arXiv:1204.5182](https://arxiv.org/abs/1204.5182), April 2012.

- [169] ATLAS Collaboration. Jet energy resolution in proton-proton collisions at $\sqrt{s} = 7$ TeV recorded in 2010 with the ATLAS detector. [arXiv:1210.6210](#), October 2012.
- [170] ATLAS Collaboration. Luminosity Determination in pp Collisions at $\sqrt{s} = 7$ TeV using the ATLAS Detector in 2011. ATLAS-CONF-2011-116, <http://cdsweb.cern.ch/record/1376384>, August 2011.
- [171] B. Kersevan and E. Richter-Was. The Monte Carlo event generator AcerMC versions 2.0 to 3.8 with interfaces to PYTHIA 6.4, HERWIG 6.5 and ARIADNE 4.1. *Computer Physics Communications*, November 2012.
- [172] S. Frixione, P. Nason, and C. Oleari. Matching NLO QCD computations with parton shower simulations: the POWHEG method. *Journal of High Energy Physics*, 2007(11):70, November 2007.
- [173] S. Hoeche et al. Matching Parton Showers and Matrix Elements. [arXiv:0602031 \[hep-ph\]](#), February 2006.
- [174] G. Aad et al. Search for Gluinos in Events with Two Same-Sign Leptons, Jets, and Missing Transverse Momentum with the ATLAS Detector in pp Collisions at $\sqrt{s} = 7$ TeV. *Physical Review Letters*, 108(24):241802, June 2012.
- [175] M. Perelstein. Introduction to Collider Physics. [arXiv:1002.0274](#), February 2010.
- [176] G. Aad et al. Search for squarks and gluinos using final states with jets and missing transverse momentum with the ATLAS detector in protonproton collisions. *Physics Letters B*, 701(2):186–203, July 2011.
- [177] A. L. Read. Presentation of search results: the CL s technique. *Journal of Physics G: Nuclear and Particle Physics*, 28(10):2693–2704, October 2002.
- [178] G. Aad et al. Hunt for new phenomena using large jet multiplicities and missing transverse momentum with ATLAS in 4.7 fb^{-1} of $\sqrt{s} = 7$ TeV proton-proton collisions. *Journal of High Energy Physics*, 2012(7):167, July 2012.

- [179] M. Cahill-Rowley, J. Hewett, A. Ismail, and T. Rizzo. The Higgs Sector and Fine-Tuning in the pMSSM. [arXiv:1206.5800](#), June 2012.
- [180] R. Brun and F. Rademakers. ROOT: An object oriented data analysis framework. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 389(1-2):81–86, April 1997.
- [181] M. Ballintijn, R. Brun, F. Rademakers, and G. Roland. The PROOF Distributed Parallel Analysis Framework based on ROOT. [arXiv:0306110 \[physics\]](#), June 2003.
- [182] XRootD home page, <http://xrootd.slac.stanford.edu/index.html>.
- [183] B. Butler and M. Swiatlowski. ProofAna documentation wiki page, <http://twiki.cern.ch/twiki/bin/viewauth/AtlasSandboxProtected/ProofAna>.
- [184] T. Maeno. PanDA: distributed production and distributed analysis system for ATLAS. *Journal of Physics: Conference Series*, 119(6):062036, July 2008.