

©Copyright 2016

Heather Russell

Search for long-lived particles decaying in the muon spectrometer of
the ATLAS detector at the LHC

Heather Russell

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2016

Reading Committee:

Henry Lubatti, Chair

Ann Nelson

Gordon Watts

Program Authorized to Offer Degree:
Physics

University of Washington

Abstract

Search for long-lived particles decaying in the muon spectrometer of the ATLAS detector at the LHC

Heather Russell

Chair of the Supervisory Committee:

Henry Lubatti

Department of Physics

This thesis presents the results of two searches for long-lived, weakly-interacting neutral particles. These searches are performed with 22.1 fb^{-1} of proton–proton collision data collected with the ATLAS detector at the Large Hadron Collider. The first search is for events with a single displaced muon spectrometer (MS) vertex and two, high-energy jets. The second search is for events with two displaced MS vertices, requiring no prompt activity. The individual power of these searches depends on the model they are being applied to, and two models are studied: Stealth Supersymmetry, and Standard Model Higgs boson decays to long-lived, neutral scalars. The former model is studied in the context of both searches, while the latter is only studied in the two displaced MS vertex search. No excess of events over the expected background is found in either search, so 95% confidence level limits are set. For Stealth Supersymmetry, limits are also presented for the combination of the two studied topologies.

TABLE OF CONTENTS

	Page
List of Figures	v
List of Tables	xiv
Chapter 1: Introduction	1
1.1 Some motivation	1
Chapter 2: Theory	7
2.1 The Standard Model of Particle Physics	7
2.1.1 Quarks and leptons	9
2.1.2 Gauge bosons	9
2.1.3 QCD and Hadronization	10
2.2 Limitations of the Standard Model	11
2.2.1 The hierarchy problem	11
2.3 Beyond the Standard Model physics	12
2.3.1 Supersymmetry	13
2.3.2 Stealth Supersymmetry	15
2.4 Proton physics	16
Chapter 3: The Large Hadron Collider	19
3.1 Overview	19
3.2 Bunch trains and filling schemes	20
3.3 Luminosity	22
Chapter 4: A brief description of the ATLAS detector	23
4.1 Inner detector	25
4.1.1 Pixel detector	25

4.1.2	Silicon microstrip tracker	26
4.1.3	Transition radiation tracker	27
4.2	Calorimeters	27
4.2.1	Electromagnetic calorimeter	28
4.2.2	Hadronic calorimeter	29
4.3	Muon spectrometer	30
4.3.1	Monitored drift tubes	38
4.3.2	Cathode strip chambers	39
4.3.3	Resistive plate chambers	39
4.3.4	Thin gap chambers	40
4.4	Detectors for luminosity determination	40
Chapter 5:	Triggers for long-lived particles	42
5.1	Introduction	42
5.2	Level-1 trigger system	43
5.2.1	Level-1 calorimeter trigger	44
5.2.2	Level-1 muon trigger	46
5.2.3	Level-1 topological trigger	47
5.2.4	Level-1 zero bias trigger	47
5.3	High level trigger	48
5.4	Specialized triggers	49
5.4.1	CalRatio triggers	49
5.4.2	Muon RoI Cluster trigger	49
Chapter 6:	Data and Monte Carlo simulation samples	56
6.1	Monte Carlo simulation	56
6.1.1	Signal samples	57
6.1.2	Background samples	59
6.2	Data	59
Chapter 7:	Reconstructing events	62
7.1	Standard physics objects	62
7.1.1	Inner detector track reconstruction	62
7.1.2	Jet reconstruction	63

7.1.3	Composite variables	67
7.2	Muon Spectrometer vertices	69
7.2.1	Multi-layer segment reconstruction	69
7.2.2	Tracklet reconstruction	73
7.2.3	MS vertex reconstruction	73
Chapter 8:	Identifying displaced decays	80
8.1	Identifying muon spectrometer vertices	80
8.1.1	MS vertices in the barrel-endcap overlap region	81
8.1.2	Good vertex criteria	82
8.2	Scale factors due to mismodelling in signal MC	87
8.3	MS vertex reconstruction efficiencies	93
8.4	Systematic uncertainties associated to displaced MS objects	94
8.4.1	Systematic uncertainty due to mismodelling in signal MC	94
8.4.2	Systematic uncertainties due to pileup uncertainties	95
8.4.3	Systematic uncertainties due to PDF uncertainties	96
8.4.4	Summary of systematic uncertainties	97
Chapter 9:	Search for events with two displaced vertices	100
9.1	Expected number of background events	101
Chapter 10:	Search for events with one displaced vertex and two jets	106
10.1	Description of the ABCD method for background determination	106
10.2	Variables distinguishing signal vertices from background vertices	108
10.3	Developing a signal region for Stealth SUSY events	111
10.3.1	Modelling the efficiency for Stealth SUSY events to have two jets	112
10.4	Applying the ABCD method to data	114
10.4.1	Systematic uncertainty on the ABCD method background estimation	117
Chapter 11:	Results	119
11.1	Expected number of signal events	119
11.1.1	Inclusion of systematic uncertainties	128
11.2	Expected number of signal events	129
11.3	Expected and observed limits	130

11.4 Application of two vertex search results to Higgs boson decays	140
Chapter 12: Summary	144
Bibliography	147
Appendix A: Timing of the level-1 barrel muon trigger	151
Appendix B: The quirks of MS vertex reconstruction	154
B.1 Three-hit segment seed spacing	154
B.2 Vertical plane formation	156
B.3 Fitting of parallel lines to a vertex location	158

LIST OF FIGURES

Figure Number	Page
1.1 Standard model particles and their interactions with the ATLAS detector. Clockwise from the left, the particles are muons, a jet, an electron, a photon, a neutrino, and two low-momentum charged particles. A transverse cross-section of the detector is shown. Each colored layer is a separate type of detector: starting from the inside, the yellow layer is the inner detector, green the electromagnetic calorimeter, blue the hadronic calorimeter, and gray the muon spectrometer.	4
1.2 Long-lived, neutral particles and their interactions with the ATLAS detector. Clockwise from the left, long-lived particle decays in the muon spectrometer, hadronic calorimeter, and inner detector. A transverse cross-section of the detector is shown. Each colored layer is a separate type of detector: starting from the inside, the yellow layer is the inner detector, green the electromagnetic calorimeter, blue the hadronic calorimeter, and gray the muon spectrometer.	6
2.1 The constituents of the Standard Model of Particle Physics. Each particle is labelled with its mass, charge, and spin. Adapted from [1].	8
2.2 Example of a one-loop correction (left) to the Higgs boson mass from the top quark coupling, and a higher-order correction (right).	12
2.3 Feynman diagram of the Standard Model Higgs boson decaying to a pair of long-lived scalars.	13
2.4 Feynman diagram of the Stealth SUSY process studied. The singlino, \tilde{S} , is the long-lived particle. The gravitino, \tilde{G} , has arbitrarily low mass and does not carry away a substantial amount of energy.	15
2.5 The CT14 proton parton distribution function at scale $Q = 2 \text{ GeV}$, as a function of fractional energy of the proton that a parton claims. Note that the y-axis is the distribution function multiplied by x , and the gluon PDF is divided by 5. These subtleties are simply to improve the readability of the plot. Image source: [2].	17

2.6	(a) An individual hard scatter process, with two gluons colliding and creating a $u\bar{u}$ pair. (b) The hard scatter within the context of two colliding proton beams, showing initial state and final state radiation (ISR and FSR). The beam remnant contains the protons and partons that did not participate in the hard scatter. (c) The hard scatter is shown from a detector perspective: here, the beam remnant manifests as underlying event, which leaves energy deposits in the detector in addition to the hard scatter.	18
3.1	The CERN accelerator complex. From [3].	20
3.2	An example bunch group set.	21
4.1	A digital rendering of the ATLAS detector. Each of the sub-detectors is labeled. From [4].	23
4.2	The coordinate system used to describe the ATLAS detector. The polar angle θ spans $[0, \pi]$, and the azimuthal angle ϕ spans $(-\pi, \pi]$. The pseudorapidity η is defined as $-\ln[\tan(\theta/2)]$, and spans the range $(-\infty, \infty)$. Because the detector is \sim symmetric in ϕ , $r = \sqrt{x^2 + y^2}$ is generally used to describe a displacement from the central z -axis.	24
4.3	A digital rendering of the ATLAS inner detector. Each of the tracking sub-detectors is labelled, and their radial locations are marked. From [5].	26
4.4	Number of radiation lengths in the combined inner detector and calorimeter system, as a function of $ \eta $. The uppermost blue, unlabelled region spanning $0 < \eta < 3$ is depicting the material between the outer edge of the calorimeter and the first layer of the muon spectrometer. From [4].	28
4.5	Geometrical structure of ECal modules in $\eta - \phi$. From [4].	29
4.6	Cross-section of the Muon Spectrometer. From [4].	32
4.7	Magnetic field integral for a particle traveling straight through the Muon Spectrometer (e.g. infinite momentum muon). From [4].	33
4.8	Technical drawings of the (a) barrel and (b),(c) endcap muon spectrometer. Physical differences between small and large chambers are clearly visible in both the barrel and endcaps. Note that these drawings were from the original technical design report. Additional chambers have since been added, and the chamber naming scheme in the endcaps has changed. The locations of the outer endcap wheels were altered before detector construction.	37
4.9	Schematic of an MDT chamber.	38

5.1	The minimum distance between charged particle deposits and a cluster of muon RoIs in SM multi-jet events compared to simulated signal events. “Track” here refers to well-reconstructed, offline inner detector tracks with $p_T > 5$ GeV, while “jet” refers to offline jets with $p_T > 30$ GeV and $\log_{10}(E_{\text{HAD}}/E_{\text{EM}}) < 0.5$. The nearest ΔR between a muon RoI cluster and jets and tracks in the barrel are shown in (a) and (b), while for the endcap they are shown in (c) and (d), respectively.	52
5.2	The minimum ΔR between the long-lived particle and a cluster of muon RoIs in the (a) barrel and (b) endcap.	53
5.3	The efficiency for a single long-lived particle to fire the isolated Muon RoI Cluster trigger, HLT_j30_mvtx, for two extreme Stealth SUSY mass points. Dashed lines mark the relevant detector layers. Each RPC layer has alternating large (L) and small (S) sectors in ϕ , meaning the exact L_{xy} of each trigger plane depends on the ϕ location. Barrel details are shown in (a) and endcap details in (b).	54
6.1	Distribution of the average number of interactions per bunch crossing in Monte Carlo simulation samples. This profile was generated to best represent the average interactions per bunch crossing expected in 2016, plus the existing 2015 data.	57
6.2	Integrated luminosity delivered and recorded by day in 2015 (a) and 2016 (c). The peak average interactions per bunch crossing, per day, are shown for 2015 (b) and 2016 (d). The 2015 dataset is both smaller and has a lower nominal pileup than 2016.	60
7.1	Formation of two jets from a $q\bar{q}$ pair. Straight lines indicate quarks, while the curly lines indicate gluons.	64
7.2	Transverse depiction of three event topologies with the same H_T but differing m_{eff} and H_T^{miss}	68
7.3	The numbering scheme for MDT tubes within a given multi-layer.	70
7.4	(a) The four possible fit lines for a three-hit segment seed, and (b) the parameters of the fit line used to perform a χ^2 -minimization fit.	71

8.1	(a) Fraction of vertices coming from long-lived particles that match <i>two</i> otherwise good MS vertices. The simulated MC samples shown are heavy scalar decays (H) to lighter scalar pairs (s). This behavior also visible in Stealth SUSY decays. (b) The efficiency to reconstruct a MS vertex as function of η , for the Stealth SUSY $m_{\tilde{g}} = 800$ GeV sample. Dashed lines mark the excluded region. Note that the efficiency between each pair of lines is low relative to the rest of the detector, and thus removing the overlap region does not significantly harm the overall acceptance.	81
8.2	Acceptance for selecting barrel and endcap vertices passing hit-based good vertex criteria. Simulated QCD multi-jet events are compared against signal vertices to determine selection criteria that will remove most or all vertices from background QCD events and keep most signal vertices. In each case, dashed lines show the chosen cut value. Figures (a), (b), (e), and (d) select vertices with the number of hits <i>above</i> the cut line, while Figures (c) and (d) describe and upper cut, and vertices with the number of hits <i>below</i> the cut line are selected. The cut in Figures (c) and (d) is designed to reject electronic detector noise, and not QCD multi-jet events.	84
8.3	Acceptance for selecting barrel vertices passing good vertex criteria. Simulated QCD multi-jet events are compared against signal vertices to determine selection criteria that will remove most or all vertices from background QCD events and keep most signal vertices. In each case, dashed lines show the chosen cut value.	85
8.4	Acceptance for selecting endcap vertices passing good vertex criteria. Simulated QCD multi-jet events are compared against signal vertices to determine selection criteria that will remove most or all vertices from background QCD events and keep most signal vertices. In each case, dashed lines show the chosen cut value.	86
8.5	The MS vertex reconstruction efficiency in the barrel (a) and endcaps (b), with and without the scale factor applied, for the lowest and highest gluino mass Stealth SUSY samples.	88
8.6	(a) The number of MS tracklets in the leading jet cone in MC multi-jet events with leading jet $ \eta < 0.8$ and $600 < p_T < 2000$ GeV or leading jet $p_T > 2000$ GeV. The number of MS tracklets in a cone around the LLP decay of the same size is shown for comparison. (b) The number of MS tracklets in the leading jet cone in MC multi-jet events with leading jet $1.3 < \eta < 2.5$ and $600 < p_T < 1000$ GeV or leading jet $p_T > 1000$ GeV. The lower p_T threshold is chosen for endcap jets because higher- p_T jets are preferentially at lower $ \eta $ and thus statistics are poorer for TeV-scale jets.	90

8.7	The average number of MS tracklets in the leading jet cone in dijet events in both data (filled circles) and MC (open circles). The ratio between data and MC is shown below the plot and provides a scale factor and associated uncertainty on the number of tracklets reconstructed in MC versus data, in the (a) barrel and (b) endcaps.	90
8.8	The average number of level-1 muon RoIs in the leading jet cone in dijet events in both data (filled circles) and MC (open circles). The ratio between data and MC is shown below the plot and provides a scale factor and associated uncertainty on the number of level-1 muon RoIs reconstructed in MC versus data, in the (a) barrel and (b) endcaps.	91
8.9	The MS vertex reconstruction efficiency in the barrel (a) and endcaps (b), with and without the scale factor applied, for the lowest and highest gluino mass Stealth SUSY samples.	92
8.10	The MS vertex reconstruction efficiency in the barrel (a) and endcaps (b), with all good vertex criteria and scale factors applied, for the lowest and highest gluino mass Stealth SUSY samples. Vertical dashed lines show where relevant detector structures are located.	94
8.11	The MS vertex reconstruction efficiency at nominal scale factor, $+1\sigma$, and -1σ in the (a) barrel and (b) endcaps for $m_{\tilde{g}} = 250$ GeV Stealth SUSY vertices. The overall systematic is found by rebinning the entire x -axis into a single bin, and finding the net change in efficiency.	95
8.12	The effect of pileup systematic uncertainties on the MS vertex reconstruction efficiency in the barrel, for the Stealth SUSY $m_{\tilde{g}} = 250$ GeV sample.	96
8.13	The effect of parton distribution function uncertainties on the MS vertex reconstruction efficiency in the barrel, for the Stealth SUSY $m_{\tilde{g}} = 800$ GeV sample.	97
9.1	The angular separation, ΔR , between singlinos in Stealth SUSY events and scalars in Higgs boson events.	102
10.1	Sketch of the ABCD method for background estimation.	107
10.2	Isolation of barrel MS vertices from tracks and jets in events with two jets with $p_T > 50$ GeV.	109

10.3	Two vertex-based criteria and their ability to distinguish signal vertices from punch-through jet vertices. (a) The degree of isolation of the vertex, where clearly punch-through jet vertices generally have a jet or track within $\Delta R = 0.3$ of the vertex axis. (b) The number of hits associated to a vertex. On average, most punch-through jet vertices have fewer hits than signal-like vertices, and the distribution for punch-through jet vertices falls off much faster than that for signal vertices.	110
10.4	Sub-leading jet p_T for Stealth SUSY signal MC samples, in events with a barrel MS vertex.	111
10.5	Efficiency for a long-lived particle (singlino) to leave a well-reconstructed jet in the barrel calorimeter ($ \eta < 1.5$) as a function of (a) long-lived particle transverse decay position, and (b) long-lived particle transverse momentum.	114
10.6	ABCD plane in the validation and signal regions for events with a single barrel MS vertex.	116
10.7	Event displays of a single, signal MC event from the $m_{\tilde{g}} = 800$ GeV sample. This event passes the full selection criteria for the single vertex plus two jets search. The displaced vertex is not shown, but a shower of hits shows the trajectory of a long-lived particle. (a) shows an xy -view, while (b) shows an Rz -view.	117
11.1	(a) Barrel RoI cluster trigger efficiency for barrel MS - barrel MS events. Since the RPC timing dependence is accounted for before this histogram is used, the timing dependence from MC is removed in the events in this histogram. The decays are ordered in terms of their radial decay position before entering the histogram - so the x -axis is the decay with the larger R , while the y -axis is the decay with the smaller R . (b) Barrel MS vertex reconstruction efficiency for triggered barrel MS - barrel MS events (<i>i.e</i> events in the numerator of (a)).	121
11.2	Barrel-Endcap RoI cluster trigger efficiency for barrel MS - endcap MS events. The x -axis is the decay position of the endcap decay, while the y -axis is the decay position of the barrel decay.	122
11.3	(a) Barrel MS vertex reconstruction efficiency for triggered barrel MS - endcap MS events. (b) Endcap MS vertex reconstruction efficiency for triggered barrel MS - endcap MS events.	123
11.4	(a) Endcap RoI cluster trigger efficiency for endcap MS - endcap MS events. The decays are ordered in terms of their longitudinal decay position before entering the histogram - so the x -axis is the decay with the larger $ z $, while the y -axis is the decay with the smaller $ z $. (b) Endcap MS vertex reconstruction efficiency for triggered endcap MS - endcap MS events.	124

11.5	(a) Barrel Muon RoI Cluster trigger efficiency for events with one barrel MS decay. (b) Vertex reconstruction efficiency for triggered events with one barrel MS decay, for vertices matched to the trigger cluster.	125
11.6	(a) Barrel MS vertex reconstruction efficiency for triggered barrel MS - barrel MS events. (b) Barrel MS vertex reconstruction efficiency for triggered barrel MS - endcap MS events. (c) Endcap MS vertex reconstruction efficiency for triggered barrel MS - endcap MS events.	126
11.7	(a) Barrel Muon RoI Cluster trigger efficiency for events with one barrel MS decay. (b) Vertex reconstruction efficiency for triggered events with one barrel MS decay, for vertices matched to the trigger cluster.	127
11.8	An example fit used in the signal extrapolation. The 2 MS vertex topology in the 250 GeV gluino sample is shown, with the smooth Novosibirsk fit on top of the raw output from the toy MC.	129
11.9	Expected number of signal events for each of the Stealth SUSY benchmark samples. At low masses, where a smaller fraction of events have two jets with $E_T > 150$ GeV, both the 2 MS Vertex topology and the 1 MS vertex + 2 jets topology have equal contributions to the global efficiency. At higher masses, the 2 MS vertex topology becomes nearly negligible.	131
11.10	Expected number of signal events for each of the Stealth SUSY benchmark samples, for the two vertex channel, the one vertex plus two jets channel, and the sum of the two channels.	132
11.11	Expected and observed 95% CL exclusion limits on the production cross section for Stealth SUSY events with $m_{\tilde{g}} = 250$ GeV, as a function of the singlino proper lifetime. Limits for the two vertex search are shown in (a), (b) shows the one vertex plus two jets search, and (c) shows the combination of both topologies.	133
11.12	Expected and observed 95% CL exclusion limits on the production cross section for Stealth SUSY events with $m_{\tilde{g}} = 500$ GeV, as a function of the singlino proper lifetime. Limits for the two vertex search are shown in (a), (b) shows the one vertex plus two jets search, and (c) shows the combination of both topologies.	134
11.13	Expected and observed 95% CL exclusion limits on the production cross section for Stealth SUSY events with $m_{\tilde{g}} = 800$ GeV, as a function of the singlino proper lifetime. Limits for the two vertex search are shown in (a), (b) shows the one vertex plus two jets search, and (c) shows the combination of both topologies.	135

11.14	Expected and observed 95% CL exclusion limits on the production cross section for Stealth SUSY events with $m_{\tilde{g}} = 1200$ GeV, as a function of the singlino proper lifetime. Limits for the two vertex search are shown in (a), (b) shows the one vertex plus two jets search, and (c) shows the combination of both topologies.	136
11.15	Expected and observed 95% CL exclusion limits on the production cross section for Stealth SUSY events with $m_{\tilde{g}} = 1500$ GeV, as a function of the singlino proper lifetime. Limits for the two vertex search are shown in (a), (b) shows the one vertex plus two jets search, and (c) shows the combination of both topologies.	137
11.16	Expected and observed 95% CL exclusion limits on the production cross section for Stealth SUSY events with $m_{\tilde{g}} = 2000$ GeV, as a function of the singlino proper lifetime. Limits for the two vertex search are shown in (a), (b) shows the one vertex plus two jets search, and (c) shows the combination of both topologies.	138
11.17	(a) – (c) Observed 95% CL exclusion limits on the production cross section for Stealth SUSY events, as a function of the singlino proper lifetime, for six gluino masses. Limits only considering the two vertex topology are shown in (a), the one vertex plus two jets topology in (b), and the combination of the two channels in (c). The same limits as (c) are shown in (d), but each limit is divided by the gluino production cross section at 13TeV for the corresponding gluino mass, and thus the y -axis shows the observed 95% CL exclusion limit on the branching ratio for $\tilde{g} \rightarrow \tilde{S}g, \tilde{S} \rightarrow \tilde{C}S, S \rightarrow gg$	139
11.18	Expected number of $H \rightarrow ss$ signal events produced in 22.1 fb^{-1} , for six scalar masses.	140
11.19	Observed 95% CL exclusion limits on the branching ratio for $H \rightarrow ss$, as a function of the scalar proper lifetime, for six scalar masses.	142
11.20	Expected and observed 95% CL exclusion limits on the branching ratio for $H \rightarrow ss$, as a function of the scalar proper lifetime, for six scalar masses.	143
A.1	Arrival time of hits accepted by the L1 RPC trigger in (a) data and (b) MC. 1 tick corresponds to 3.125 ns. The distributions are fitted with a gaussian distribution. Mean and standard deviation values, and their respective uncertainties, are employed for scaling the trigger efficiency on simulation and for assigning a systematic uncertainty on the scaling procedure.	152
A.2	The scale factor applied as a weight to events passing the RoI cluster trigger, as a function of the time delay as compared to a particle travelling at c	153
B.1	Examples of three-hit seeds that the algorithm should select.	155

B.2	Examples of three-hit seeds that the algorithm selects but should reject as bad.	155
B.3	Examples of three-hit seeds that the algorithm selects, but whose reflections are not allowed combinations.	156
B.4	Horizontal planes, both as intended and calculated, for a line-of-fire corresponding to $\eta = 0.5$	157

LIST OF TABLES

Table Number	Page	
4.1	Electromagnetic calorimeter coverage and resolution. Layers with a range of provided granularities have a granularity that varies with eta; a full breakdown is available in Ref. [4]. Note that there are only two layers in the endcap for $ \eta > 2.5$	30
4.2	Hadronic calorimeter coverage and resolution.	31
4.3	Muon spectrometer coverage. For MDT layers, the coverage represents the chamber extent. For trigger planes, the coverage represents the inner edge of the trigger plane. Small and large sectors alternate in ϕ , as shown in Figure 4.8. All numerical values are extracted from the technical drawings in Figure 4.8. **MDT Outer chamber values are less accurate due to the chamber locations changing after the technical designs were created.	35
5.1	Isolation criteria for the Muon RoI Cluster trigger.	51
6.1	Summary of the mass and lifetime parameters used in the Stealth SUSY signal sample generation. Every sample started with 400,000 events, but occasionally a job or two fail in one of the many simulation steps, causing the loss of a small fraction of events. Production cross-sections are 13 TeV gluino production cross-sections, with squarks decoupled [6].	58
6.2	Summary of the mass and lifetime parameters used in the $H \rightarrow ss$ signal sample generation.	58
6.3	Summary of the QCD multi-jet slice names, their leading jet (LJ) p_T range, and the number of events per slice.	59
8.1	Summary of criteria for good MS vertices in the barrel and endcap regions. .	88
8.2	Scale factors and systematic uncertainties on data-MC agreement for MS objects.	91
8.3	Summary of systematic uncertainties on the muon RoI cluster trigger in the barrel for the Stealth SUSY benchmark samples.	98
8.4	Summary of systematic uncertainties on the muon RoI cluster trigger in the endcaps for the Stealth SUSY benchmark samples.	98

8.5	Summary of MS vertex systematic uncertainty in the barrel for the Stealth SUSY benchmark samples.	98
8.6	Summary of MS vertex systematic uncertainty in the endcaps for the Stealth SUSY benchmark samples.	99
9.1	The numbers of events in the 13 TeV dataset passing the Muon RoI Cluster trigger without isolation criteria, and events passing said trigger also with a matched vertex. These values are necessary to calculate the background prediction. The notation is explained in the text. Uncertainties are statistical only.	103
9.2	The probabilities needed to compute the background from events containing MS vertices from non-signal processes. $P[BV_x BClus]$ and $P[EV_x EClus]$ are the probabilities of finding one MS vertex in events passing the Muon RoI Cluster trigger that have a single cluster that is matched to a vertex in the barrel or endcaps, respectively. The probability of finding an MS vertex not matched to a cluster of muon RoIs is given by $P[V_x \text{no MS trig}]$	103
10.1	Summary of the criteria used to select events for the two jets – one vertex search. All selection criteria are also applied to signal MC events when determining the amount of expected signal events in the dataset.	112
10.2	Summary of systematic uncertainties on the event-level probability for two jets with $E_T > 150$ GeV, in events with one MS vertex and a second LLP with $ \eta < 1.5$, for the Stealth SUSY benchmark samples.	115
10.3	Summary of systematic uncertainties on the event-level probability for two jets with $E_T > 150$ GeV, in events with one MS vertex and a second LLP with $ \eta > 1.5$, for the Stealth SUSY benchmark samples.	115
10.4	Event counts in each of the four regions of the ABCD plane, and the expected number in region A, for both the validation and signal regions.	116
10.5	Event counts using varied x - and y -axis cuts in each of the four regions of the ABCD plane, and the expected number in region A, for both the validation and signal regions. The largest discrepancy between A and B^*C/D is 35%, and is from the second variation, with cuts values $t \min(\Delta R(\text{vertex, jet/trk})) > 0.35$ and $n\text{Hits} > 2000$	118
11.1	Ranges of long-lived scalar proper decay lengths excluded at 95% CL assuming a 10%, or 1% BR, for $m_H = 125$ GeV.	141

ACKNOWLEDGMENTS

Completing this thesis was only possible with the help of many people. Not all of you might be mentioned here, but you are all certainly appreciated.

First, to all the professors who went out of their way to help me throughout the years, at the University of Victoria, the Perimeter Institute, and the University of Washington. The entire displaced hadronic jets analysis team, former and present members: I've learned so much from working with all of you. Special thanks is necessary to Daniel Blackburn, who was so patient teaching me about ATLAS software when I first joined, and to Rachel Rosten for always being a good, non-judgemental sounding board for all my ideas, both good and horrible. And of course, to my supervisor, Henry Lubatti: you've given me an immense amount of support over the past five years, and for that I will be eternally grateful.

To Alison Eliot, who embarked on a Swiss adventure with me, and made me appreciate that if it's cheaper than a beer, it doesn't cost *that* much. We had so many excellent adventures, and walking across a country with you was definitely one of the highlights of the past few years. Geneva isn't quite the same without you. Thanks is also necessary to my friends from Seattle, Jackie and Stella, who came all the way to Geneva to say hi.

Two people who always went out of their way to help me deserve my profound thanks. First, Mark Stockton, for always answering my questions about ATLAS, no matter how random they seemed, and teaching me to cycle up mountains. And second, Sam Harper, who has taught me so much about programming and particle

physics, and how to ski up mountains when there's too much snow to cycle. My code would be a disaster without you.

My family certainly deserves thanks here for supporting me from the start. A special thanks to my brother, Andrew Russell, who taught me how to code. Hopefully I've taught you some physics along the way.

Lastly, to Thomas Blake: I'm not sure I would have made it this far without you. You've given me an incredible amount of support, reminded me when I needed it that physics really is pretty cool, and made sure I could always come hide under your couch. I don't have enough words to express my gratitude. Even if you are a snowboarder.

DEDICATION

To my grandfather, Nikita Kusnezov, who showed me that anything is possible.

Chapter 1

INTRODUCTION

For millennia, humanity could satisfy their curiosity with tangible inventions. Discoveries. Experiments.

The intangible was left to the imagination - or, in some cases, was simply so intangible that not a single person had asked the question.

It might be nice to think that, one day, a brilliant mind awoke and described a beautiful image of the universe that none had imagined. While this paints a romantic picture of scientific discovery, it does a disservice to those without whom such progress would not have been possible.

Each experiment, each paper, each presentation - they all contribute to a guided but mostly random walk, until eventually one is close enough to stumble across a solution. But oftentimes a solution only opens a new box of questions, and then the process starts again.

Sometimes the goal is known, but predominantly, search parties are simply sent out in as many directions as can be fathomed, in hopes that one might eventually stumble upon a yet-unknown goal.

This is one of those journeys.

1.1 Some motivation

If one truly wants to explain the current state of being, it makes sense to start at the beginning. Particle physics, however, has two beginnings. One was when Dmitri Mendeleev¹ discovered that organizing elements by their properties and mass resulted in a period arrange-

1

and others!

ment, leading to the development of the first periodic table. The other beginning started much, much earlier, when a singularity gave birth to the universe in what is commonly described as the Big Bang.

The matter we interact with on a day-to-day basis is representative of only a small fraction of the matter in the universe. Complex molecules, minerals – these can all be broken into smaller pieces: elements. Elements are what we commonly think of as pure substances. But to explain their properties and interactions one must look deeper. A pure sample of any element - let's look at gold - can be cut into smaller and smaller pieces and still retain the properties of that element. However, this is only true on a macroscopic scale. That is to say, the scale on which we see and interact with objects. Eventually, the pieces of gold will be small enough that one more cut will reach the point where the resulting pieces are no longer gold. The smallest possible division that can be made of an element is called an atom.

The actual composition of an atom was the subject of much investigation for quite some time, and went through many iterations and models until scientists converged on something that accurately describes the interactions that occur on an atomic scale.

The model most people think of when they hear “atom” is Bohr’s model. It consists of a solid, positively charged nucleus made of protons (positive charge) and neutrons (neutral) that is orbited by negatively charged electrons.

The exact number of protons in the nucleus defines the element, while the number of neutrons defines the isotope of that element. The number of electrons is usually equal to the number of protons, unless the atom is ionized - *i.e.* either positively or negatively charged.

This almost could have been the end of the story - except nuclear decay experiments showed atoms could change into others via the decay of a neutron. The following decades brought a rich array of new particles and interactions, and eventually led to the development of the Standard Model of Particle Physics. With the addition of the Higgs boson, the Standard Model describes most of the phenomena we observe, but it does leave many unanswered questions.

The current state of the art technology for studying high-energy particles and their in-

teractions is at the European Organization for Nuclear Research (CERN), where the Large Hadron Collider (LHC) accelerates and collides protons below 100 m of rock at in the foothills of the Jura Mountains. Four detectors are dotted along the collider, and are designed to detect the particles that emerge from high-energy proton collisions.

When two protons collide at the energies generated at the LHC, it's the constituent quarks and gluons that interact. The energies of the colliding particles create new particles: other quarks, gluons, leptons, or even a Higgs boson. To understand what happens in these collisions, the outgoing particles need to be identified and studied: this is the main purpose of the ATLAS detector.

Many things can be searched for in proton–proton collisions. Known Standard Model processes can be precisely measured, to ensure their parameters and rates agree with theoretical predictions and previous experiments. Theoretical particles whose existence could explain unanswered questions, like the nature of dark matter, can be searched for in the plethora of data.

Most of the time, these studies all involve a careful analysis of the Standard Model particles that can be measured with the ATLAS detector, such as electrons, photons, hadrons (the “stable” state of quarks and gluons), and muons. The detector therefore has four distinct parts, carefully designed such that these different particles can be easily distinguished and their energies well-measured:

- The inner detector is a tracking detector, designed to provide spatial information about where charged particles have travelled. It is inside a magnetic field, so the trajectories of charged particles are bent as they travel, providing information about their momentum.
- The electromagnetic calorimeter is designed to absorb and measure the energy of electromagnetically interacting particles: namely, electrons and photons. Hadrons deposit some of their energy here as well.
- The hadronic calorimeter is designed to absorb and measure the energy of hadrons.

Generally, the energy from all particles except muons should be absorbed by the end of the hadronic calorimeter.

- The muon spectrometer is a tracking detector that provides spatial information about anything that is not absorbed by the first three detector layers. From the point of view of the Standard Model, the only detectable particles reaching the muon spectrometer should be muons. There are also magnetic fields that bend the muons to provide momentum information.




Figure 1.1: Standard model particles and their interactions with the ATLAS detector. Clockwise from the left, the particles are muons, a jet, an electron, a photon, a neutrino, and two low-momentum charged particles. A transverse cross-section of the detector is shown. Each colored layer is a separate type of detector: starting from the inside, the yellow layer is the inner detector, green the electromagnetic calorimeter, blue the hadronic calorimeter, and gray the muon spectrometer.

Figure 1.1 shows how each particle type interacts with each part of the detector. Starting on the left are muons: they leave tracks in the inner detector and some energy depositions in

the calorimeters. However, since they are not absorbed in the calorimeters, they also leave tracks in the muon spectrometer.

Next is the result of a quark or gluon: when produced in a high-energy collision, these particles tend to shower into many hadrons, producing a collimated group of particles that leave tracks in the inner detector and deposit their energy in the calorimeters.

Continuing clockwise after the jet is an electron: it is a charged particle, so it is bent in the inner detector, leaves a track, and then deposits its energy in the electromagnetic calorimeter. The following particle is a photon, which is not charged and thus leaves no track in the inner detector, but it still deposits its energy in the electromagnetic calorimeter. The dashed line is a neutrino, which does not interact with any of the detector, and its energy escapes detection.

Last is the signature left by any low-energy charged particles, which dissipate before reaching the calorimeters. When two protons interact, most of the time the result is low-energy particles. These cause a *background* over which high-energy objects can be searched for.

Most searches for physics beyond the standard model involve looking for some combination of the above particles at specific energies. However, there is also the possibility that new physics might look somewhat *different*. Instead of a new particle being created in the proton-proton collision and immediately decaying into standard model particles, what if it travelled first? And what if these new particles *didn't interact with the detector*?

This would leave a unique and striking signature in the detector: a burst of activity far away from the collision, with nothing in between. Figure 1.2 shows some examples of signatures these types of particles could leave, if they decayed in various sub-detectors: The sudden appearance of many tracks in the muon spectrometer, a jet with no tracks or energy in the electromagnetic calorimeter, or many tracks in the inner detector all coming from the same location - with no link to the collision point. This thesis outlines a search for these so-called long-lived, neutral particles.




Figure 1.2: Long-lived, neutral particles and their interactions with the ATLAS detector. Clockwise from the left, long-lived particle decays in the muon spectrometer, hadronic calorimeter, and inner detector. A transverse cross-section of the detector is shown. Each colored layer is a separate type of detector: starting from the inside, the yellow layer is the inner detector, green the electromagnetic calorimeter, blue the hadronic calorimeter, and gray the muon spectrometer.

Chapter 2

THEORY

The first truly fundamental particle was discovered in 1896, when Sir Joseph Jon Thomson showed that cathode rays were actually negatively charged particles with an unprecedentedly high charge-to-mass ratio. In the following century, our understanding of the fundamental forces and particles of nature progressed to what is now referred to as the Standard Model of Particle Physics.

The Standard Model, or SM, accurately describes many of the physical phenomena that have been observed. Its true value, however, lies in its predictive nature. Many particles have been discovered because they were predicted by the SM, which provides validation of the model as a theory of fundamental physics. The most recent of these discoveries was the Higgs boson, whose discovery was simultaneously announced by the ATLAS and CMS experiments and the Large Hadron Collider on 4 July 2012[7, 8].

2.1 The Standard Model of Particle Physics

The Standard Model describes the fundamental forces (strong, weak, and electromagnetic)¹ and particles of nature. The SM is a relativistic quantum field theory described by a Lagrangian \mathcal{L} . The model has 19 free parameters, which have been experimentally determined. Together with the Lagrangian, this set of parameters describes the dynamics the SM must obey. The SM must also obey two sets of symmetries: one associated with the Lagrangian, and another with the gauge group that describes the interactions of the SM. Symmetries of

1

A discerning reader will have noticed that gravity is omitted in this list. Unifying the gravitational force with the Standard Model, and thus the principles of quantum mechanics, is a difficult task that has been the sole focus of many research careers, but no solution has yet been discovered

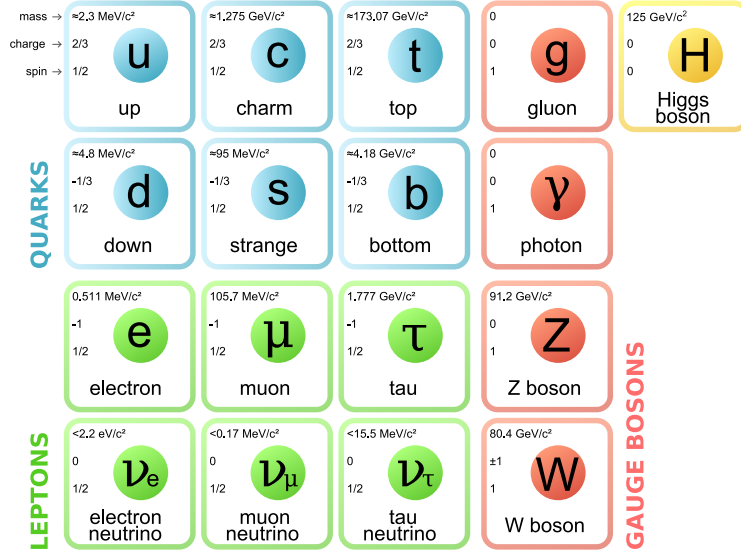


Figure 2.1: The constituents of the Standard Model of Particle Physics. Each particle is labelled with its mass, charge, and spin. Adapted from [1].

the Lagrangian are space–time symmetries, while those associated with the gauge transformations are local symmetries. These symmetries restrict the possible physical manifestations of the SM and thus dictate the particles and interactions that have been observed in nature.

The underlying symmetry of the SM can be written as:

$$SU(3)_C \times SU(2)_L \times U(1)_Y. \quad (2.1)$$

The first term, $SU(3)_C$, is the gauge group of quantum chromodynamics (QCD). QCD is the theory of the strong interaction, and all particles that have color charge (as indicated by the subscript C) are bound by the symmetries of SU(3). The second and third terms describe the symmetries of electroweak interactions.

The SM is made up three categories of fundamental particles: fermions, gauge bosons, and the Higgs boson. Fermions are spin-1/2 particles, and are the principal component of matter. Gauge bosons are force mediators, and the Higgs boson is the excitation of the Higgs field, which is what gives mass to the other fundamental particles.

2.1.1 Quarks and leptons

The fermions are traditionally split into two categories, quarks and leptons, each of which consist of three generations. The properties of each generation are fundamentally the same, aside from the particle masses. Particles belonging to the lightest generation of fermions (the up quark, down quark, and electron) are what comprise the atoms that make up matter. Protons are made of two up quarks and a down quark (uud), neutrons are made of an up quark and two down quarks (udd), and atoms are comprised of a nucleus made of protons and neutrons surrounded by an electron cloud.

The remaining nine fermions are less common, but have been observed both directly and indirectly. Muons (μ) are a common constituent of cosmic rays at the surface of the Earth, because they have a very small probability of interacting with matter and thus can pass through much of the atmosphere undisturbed. Electron neutrinos ($\bar{\nu}_e$) can be inferred from the energy balance in nuclear beta decay or electron capture. The other fermions are not observed outside of dedicated experiments, though that does not make them any less important to the Standard Model.

2.1.2 Gauge bosons

The gauge bosons are spin-1 particles that mediate the weak, electromagnetic, and strong forces. At high energies, well above the mass scale of the weak interaction ($\mathcal{O}(100 \text{ GeV})$, the observed mass of the W and Z bosons), the electromagnetic and weak forces are unified into the electroweak interaction. Here, the two forces are indistinguishable. At lower energies, the two forces separate because of a mechanism called electroweak symmetry breaking [9]. This mechanism is why the W^\pm and Z bosons have a non-zero mass and are split from the mediator of the electromagnetic force, the massless photon γ .

2.1.3 QCD and Hadronization

The mediator of the strong force is the massless gluon, which binds together quarks into *hadrons*. No single quark or gluon exists as a free object, unbound to anything else. This is referred to as *confinement*, and is a consequence of color charge: only combinations of particles with no net color charge can exist as an isolated state. In general, quarks will form either quark–anti-quark states known as mesons, or three (anti-)quark states known as (anti-)baryons. Quarks can also form higher multiplicity bound states: In recent years, both tetraquark and pentaquark states have been confirmed [10, 11, 12, 13, 14].

The structure of the $SU(3)_C$ group implies three color charges: red (r), green (g), and blue (b). Quarks exist in the fundamental representation of $SU(3)$ and thus carry a single color charge, (r , g , or b), while anti-quarks carry an anti-color charge (\bar{r} , \bar{g} , or \bar{b}). Gluons, however, exist in the adjoint representation of $SU(3)$, and thus have eight possible color charges. These can have many different representations, but they are always a combination of a color and an anti-color, such as $r\bar{b}$ or $(g\bar{b} + \bar{g}b)/\sqrt{2}$.

When a pair of quarks are produced in a high-energy collision or decay, they can simply travel in opposite directions. However, when a pair of quarks ($q\bar{q}$) or gluons (gg) are produced additional interactions must occur to ensure they travel as color singlets. If a pair of quark's momenta are sufficiently collimated, they can bind together to produce a meson. If instead two quarks are produced with momenta in opposite directions, it will likely become energetically favorable for an additionally $q\bar{q}$ pair to be produced from the vacuum to reduce the energy necessary to maintain the distance between the quarks. These new quarks will combine with the original pair to form two, colorless mesons that will proceed to travel apart. This is the simplest example of the *hadronization* process.

In general, especially at higher energies, there are many $q\bar{q}$ pairs produced and thus many hadrons are formed. These will generally be collimated around the original quark or gluon produced in the high-energy collision, and the sum of their momenta should be that of the original parton from the hard scatter. In order to study the parameters of the hard scatter,

it becomes interesting to know the properties of the original parton. The set of collimated hadrons resulting from hadronization is known as a *jet*.

2.2 *Limitations of the Standard Model*

The standard model has evolved with our understanding of subatomic phenomena. Some particles were discovered first, then the SM was adapted to accommodate them. Others, such as the massive vector bosons and the Higgs boson, were first predicted, and their subsequent discoveries validated the theories that predicted them.

The only *discovery* inconsistent with the predictions of the SM is that neutrinos oscillate. However, there are still many phenomena that are outside the scope of the current manifestation of the SM. These include dark matter, dark energy, unification with gravity, and the matter–anti-matter asymmetry in today’s universe. Many models have been proposed to reconcile these observations with the SM, and much work has been done to prove or disprove these models’ existence. This thesis presents the results of a search for two specific models. The first, Stealth Supersymmetry, is outlined in Section 2.3.2. The second is a model in which the SM Higgs boson decays to pairs of long-lived, neutral scalars, which are discussed briefly in Section 2.3.

In addition to the physical phenomena that have not yet been integrated into the standard model, there is one problem with the SM that motivates the search for Stealth Supersymmetry: the hierarchy problem.

2.2.1 *The hierarchy problem*

The hierarchy problem manifests most clearly in the calculation of the Higgs boson mass (m_H). The higher order corrections to m_H from the top quark (Figure 2.2) result in a predicted mass equal to the ultraviolet cutoff - but the Higgs boson has been observed with a mass of 125 GeV!

The simplest solution to the hierarchy problem is fine tuning, in which the parameters in the equations describing the Standard Model have to match to an extreme number of signif-




Figure 2.2: Example of a one-loop correction (left) to the Higgs boson mass from the top quark coupling, and a higher-order correction (right).

ificant digits in order to accurately describe the observed Higgs boson mass. One argument in favor of fine-tuning is the anthropogenic principle, which says that the parameters in the standard model are the exact values they are simply because that is what they needed to be in order for the universe to be the way it is today. The anthropogenic principle appeals to many people, but one can also argue that it is a lazy solution to a real problem.

A much more elegant solution to this problem arises with supersymmetry: each Standard Model particle has a superpartner, and each loop contribution from a quark is *exactly* cancelled out the contribution from its superpartner. Supersymmetry will be discussed in Section 2.3.1.

2.3 Beyond the Standard Model physics

There are two ways of making a discovery. The first is trying to discover physics that has been predicted, generally in a well-motivated way. The other is by searching for new phenomena that are not necessarily predicted by any theoretical models, but are also not ruled out by current constraints.

The accuracy to which the Higgs boson couplings have been measured does not rule out the Higgs boson coupling to new physics at appreciable rates - up to 30%! Thus, this is a clear place where undiscovered phenomena could be lurking. It is fairly simple to introduce a new scalar field that couples to the Higgs boson. This allows the Higgs boson to decay into

pairs of these scalars, which will subsequently decay back to SM particles through the Higgs boson coupling. Providing the decay width of the new scalar is sufficiently small, the lifetime can be macroscopic. This provides the framework for a simplified model where the Higgs boson decays to two long-lived, neutral scalars or pseudoscalars, which proceed to decay to SM particles[15], as shown in Figure 2.3.




Figure 2.3: Feynman diagram of the Standard Model Higgs boson decaying to a pair of long-lived scalars.

More complex models are possible, though the signatures are often very similar and thus any limits set with the simplified model are easily reinterpreted as limits on other models. Examples of other scenarios that predict long-lived particles in Higgs boson decays are Hidden Valley models [16, 17, 18], neutral naturalness [19, 20], and models aiming to explain baryogenesis and the baryon–anti-baryon asymmetry [21].

2.3.1 Supersymmetry

One might ask the question, “Why do some particles have integer spin, but others have half-integer spin?” Supersymmetry, or SUSY, introduces another symmetry to the SM: between particles of different spin. Every fermion gains a bosonic superpartner of the same flavor, charge, and mass, denoted by adding an “s” to the start of the particle name. For example, the electron’s superpartner is a selectron. Every boson gains a fermionic superpartner of the same charge and mass, denoted by adding “ino” to the end of the particle name. For example, the gluon’s superpartner is a gluino. The superpartner symbols are the same as

the SM particle, but with an added tilde over top.

Like most symmetries we encounter, supersymmetry is also a broken symmetry. Below some SUSY-breaking scale F , the SM particles and their superpartners are distinguishable by mass: the mass symmetry between SM particles and their SUSY partners is broken. There are many experimental constraints from squark and gluino searches on the allowable mass scale of SUSY particles [22]. Recent LHC results have pushed the mass limits very high, to the point where either some fine tuning is necessary to explain the mass discrepancies, or something additional is necessary to explain why the superpartners have not been found.

In SUSY, each particle has an additional property called R -parity, and this must be conserved for all interactions and decays. There is also a special class of SUSY models that aim to circumvent this restriction by allowing for R -parity violating interactions [23]. The R -parity of a particle is determined by its baryon number, lepton number, and spin: $(-1)^{3B+L+2s}$, and must be conserved in any interaction, such that the product of the R -parities of the outgoing particles must equal the R -parity of the incoming particle. Standard Model particles carry an R -parity of $+1$, while superpartners have an R -parity of -1 . Consequently, any decaying superpartner must decay to a standard model partner *and a superpartner*. This means that the lightest supersymmetric particle (LSP) cannot decay: there are no lighter SUSY particles for it to decay to, while conserving R -parity.

In many SUSY models the LSP is a neutralino, meaning it is heavy, stable, and non-interacting. This has two benefits. First, the LSP is a dark matter candidate. Second, it leaves a very distinct signature when it is produced in a high energy proton–proton collisions: a large amount of energy escaping the detector. Most SUSY searches setting strong limits on gluino or squark masses rely on the presence of heavy, stable neutralinos, and thus it becomes interesting to circumvent these limits by constructing a model that does not have the same LSP structure.




Figure 2.4: Feynman diagram of the Stealth SUSY process studied. The singlino, \tilde{S} , is the long-lived particle. The gravitino, \tilde{G} , has arbitrarily low mass and does not carry away a substantial amount of energy.

2.3.2 Stealth Supersymmetry

Stealth SUSY models [24, 25] are a subclass of R -parity conserving SUSY models that do not have large energy escaping the detector (missing transverse momentum, or E_T^{miss}). This can be accomplished in many different ways, but this search explores a model that involves adding a hidden (stealth) sector singlet superfield S at the electroweak scale, which has a superpartner singlino \tilde{S} . By weakly coupling the hidden sector to the MSSM, the mass splitting between S and \tilde{S} (δM) is small, given low-scale SUSY breaking². The SUSY decay chain ends with the singlino decaying to a singlet plus a gravitino \tilde{G} , where the gravitino carries off very little E_T^{miss} and the singlet promptly decays to two gluons. The effective decay processes are $\tilde{g} \rightarrow \tilde{S}g$ (prompt), $\tilde{S} \rightarrow S\tilde{G}$ (not prompt), and $S \rightarrow gg$ (prompt). This results in one prompt gluon jet and two displaced gluon jets per gluino decay. A Feynman diagram of this process is shown in Figure 2.4.

The decay width (and, consequently, the lifetime) of the singlino is determined by both the mass splitting δM and the SUSY breaking scale F : $\Gamma_{\tilde{S} \rightarrow S\tilde{G}} \sim m_{\tilde{S}}(\delta M)^4/\pi F^2$ [24]. The SUSY breaking scale F is not a fixed parameter, and thus the singlino has the possibility

2

It should be noted that high-scale SUSY breaking can also be consistent with small mass splitting and Stealth SUSY, though this requires a more complex model and is not considered in this search.

of traveling an appreciable distance through the detector, leading to a significantly displaced vertex. Since R -parity is conserved, each event necessarily produces two gluinos, resulting in two displaced vertices.

2.4 Proton physics

The Large Hadron Collider (LHC) is a proton–proton collider, and thus it is important to understand the structure of the proton. Recall that a proton is made up of three quarks: uud . However, an astute reader will probably have noticed that the sum of these quark masses, as shown in Figure 2.1, is less than 10 MeV: barely a tenth the mass of a proton! In fact, most of the proton mass actually comes from gluons binding the constituent quarks together. These gluons can also spontaneously create more gluons, or sea quarks (quark–anti-quark pairs of any flavor). Each of these quarks, anti-quarks, or gluons is referred to as a *parton*.

When two protons collide, it is not just the three quarks in each that can interact, but also the gluons and the sea quarks. The probability that a proton contains any given quark or gluon at a given energy and time is represented by a set of Parton Distribution Functions (PDFs). Many different PDFs exist, but they all contain input from many experimental results and precise theoretical calculations. An example of one PDF set is shown in Figure 2.5. Notice that the most likely constituents at any given energy are u -quarks, d -quarks, and gluons.

The LHC collides *bunches* of protons every 25 ns, at a centre of mass energy of 13 TeV. Much of the time, the protons in each bunch pass right by or through each other without interacting. When two partons do have a high energy inelastic collision, this is referred to as a *hard scatter*. Consider the case where two incoming gluons interact, and produce a $u\bar{u}$ pair (Figure 2.6a).

Before the gluons interact, they can radiate soft photons and gluons: this is called *initial state radiation*, or ISR. Radiation can also occur in the same manner, but from the outgoing partons (in this case, the u or \bar{u}), in which case it is *final state radiation*, or FSR. Both ISR and FSR are soft compared to the energy of the hard scatter, such that the gluons or photons




Figure 2.5: The CT14 proton parton distribution function at scale $Q = 2 \text{ GeV}$, as a function of fractional energy of the proton that a parton claims. Note that the y-axis is the distribution function multiplied by x , and the gluon PDF is divided by 5. These subtleties are simply to improve the readability of the plot. Image source: [2].

have little impact on the event. Particles can also radiate with a much higher fraction of the parton energy, in which case they will impact the topology of the final event.

In addition to the hard scatter, the *beam remnant*, containing both the partons and protons that did not interact, does not simply disappear: their interactions are collectively referred to as the *underlying event* (Figures 2.6b and 2.6c).

The two protons involved in the hard scatter *fragment*: once two partons interact, the protons are no longer stable hadrons. The partons rearrange to form stable objects with different momenta, generally producing activity mostly aligned with the direction of the incoming particles. Occasionally, another set of partons will interact in an elastic or inelastic collision, causing more transverse activity.

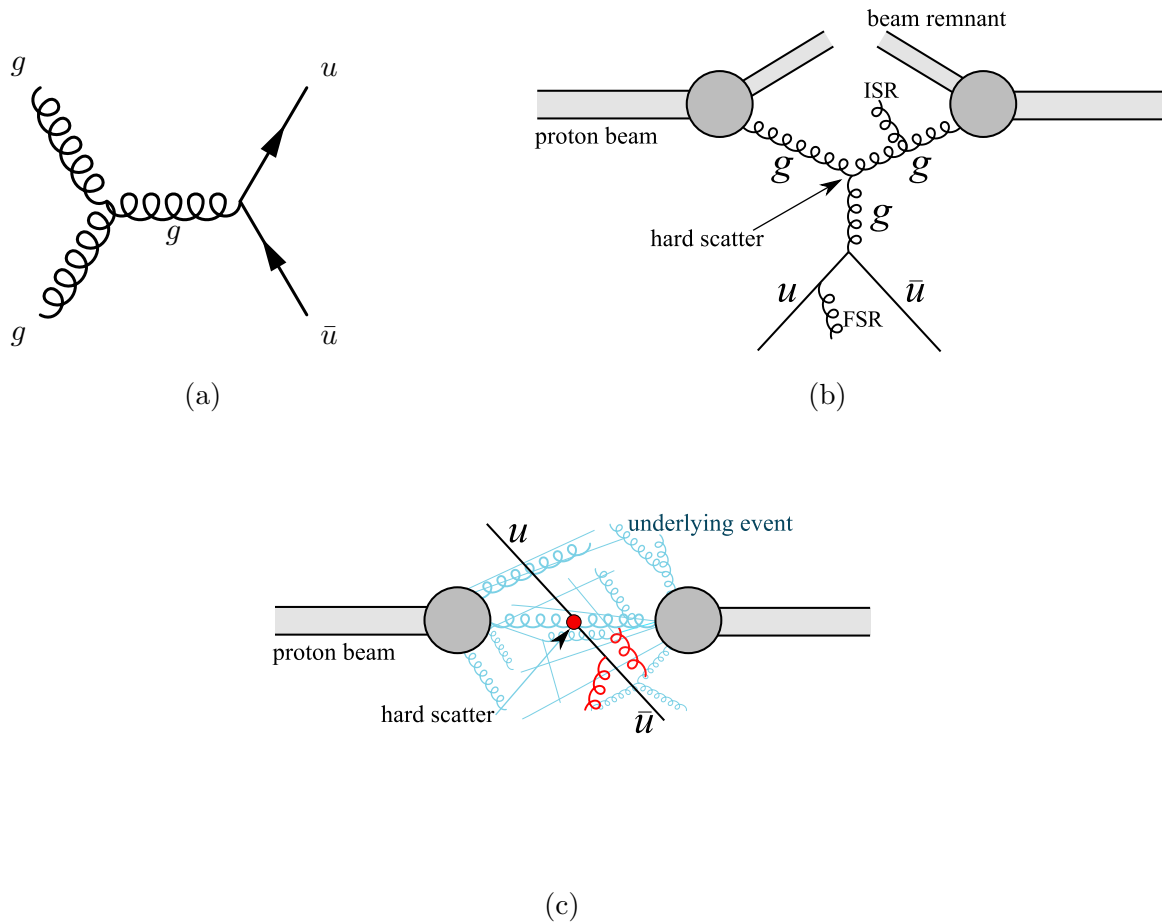


Figure 2.6: (a) An individual hard scatter process, with two gluons colliding and creating a $u\bar{u}$ pair. (b) The hard scatter within the context of two colliding proton beams, showing initial state and final state radiation (ISR and FSR). The beam remnant contains the protons and partons that did not participate in the hard scatter. (c) The hard scatter is shown from a detector perspective: here, the beam remnant manifests as underlying event, which leaves energy deposits in the detector in addition to the hard scatter.

The remainder of the protons that did not produce the hard scatter will often continue to travel through the beam pipe, but occasionally there will also be lower energy parton collisions. The contributions to underlying event from these collisions are referred to as *pileup*. The amount of pileup in an event is directly proportional to the number of protons per colliding bunch.

Chapter 3

THE LARGE HADRON COLLIDER

3.1 Overview

The evolution of particle accelerators arguably starts with the study of cathode rays by Johann Hittorf with what later became known as Crooks tubes, the predecessor to cathode ray tubes. Since then, accelerators have developed along with the complexity and energy of the phenomena they are designed to study.

The ATLAS detector, described in Chapter 4, is used to measure and record the results of high-energy proton collisions. The final acceleration and collision of the proton beams is performed by the Large Hadron Collider (LHC). However, before the protons reach the LHC, they are sent through a series of smaller accelerators to incrementally increase the beam energy. The full chain is shown in Figure 3.1, and makes use of many parts of the old accelerators dating back to the 1950s.

The life of a proton designated for collision starts in a small bottle of hydrogen gas. Molecular hydrogen is passed through an electric field and the hydrogen nuclei - protons - are extracted. They are then sent into LINAC 2, a linear accelerator that accelerates protons to 50 MeV. The protons then pass in six bunches to the Proton Synchrotron Booster, which accelerates them to 1.4 GeV before sending them into the Proton Synchrotron. This accelerator boosts the energy of the protons to 25 GeV before they are directed into the Super Proton Synchrotron (SPS). The SPS accelerates the protons to 450 GeV, before they are split into two beams and delivered into the Large Hadron Collider (LHC).

The LHC is a 27 km ring made up of a series of RF cavities for acceleration and magnets for focusing and bending the beams. In 2015 and 2016, the LHC accelerated the proton

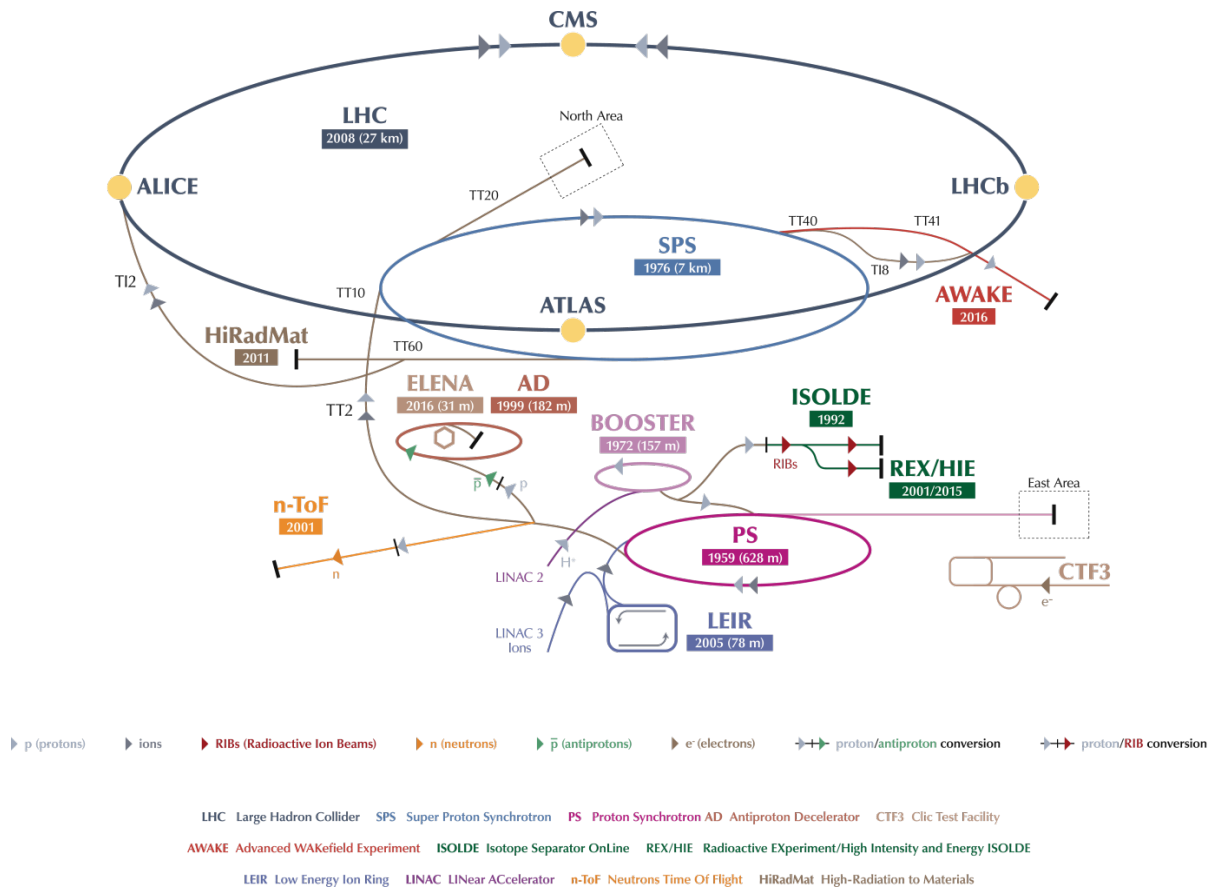


Figure 3.1: The CERN accelerator complex. From [3].

beams to an energy of 6.5 TeV each, only 0.5 TeV below the design energy.

3.2 Bunch trains and filling schemes

During the acceleration process, the original six proton bunches are each split into more and more bunches using a multiple splitting process (see [26] for details), until they reach the LHC. The exact splitting depends on the filling scheme being used and the technical limitations of the various systems along the acceleration path. The final structure injected into the LHC is called a *bunch train*. Multiple bunch trains are injected into each beam, and their structure is referred to as a *filling scheme*. There are total of 3564 possible *buckets* that can contain a proton bunch. Each beam has the same filling scheme describing which

of the buckets are filled with bunches and which are not. Buckets that do not end up filled with a bunch are generally referred to as *empty bunches*.

The 2015+2016 dataset considered in this thesis uses data collected with a bunch spacing of 25 ns. In principle, this means that the ATLAS detector could see a collision (or the possibility for one) every 25 ns. However, not every bunch ends up being filled with protons. For example, after each bunch train a set of empty bunches is left to allow for enough time ($3\mu\text{s}$) to turn on the magnets that direct the beams to the beam dump. This set of empty bunches is referred to as the *abort gap*.

When the beams collide, they do so in a predictable way based of the filling scheme. Each time two bunches pass by in the detector, this is called a *bunch crossing*, and the pattern of bunch crossings is referred to as a *bunch group set*.




Figure 3.2: An example bunch group set.

When two bunches filled with protons cross, the bunch crossing is considered *paired*. For example, the filling scheme used for Figure 3.2 contains 2208 paired bunches. When two unfilled bunches cross, the bunch crossing is *empty*. If the empty bunch is within five bunch crossings of a paired bunch, it falls into the category empty before/after paired. In practice, the distance from a paired crossing is important because it means any residual activity from a collision should have dissipated.

If one bunch is filled but the other is not, the crossing is *unpaired*. If the unpaired bunch

crossing is at least seven bunch crossings away from any paired crossing, it is considered *isolated*.

3.3 Luminosity

In order to know how many physics events to expect during a given running period, it is necessary to know the rate at which protons have the opportunity to collide. This rate is called instantaneous luminosity, and is determined by the exact parameters of the beam in the following way:

$$L = \frac{N_b^2 n_b f_{\text{rev}} \gamma_r}{4\pi \epsilon_n \beta^*} F. \quad (3.1)$$

The first two values, N_b and n_b are the number of protons per bunch and the number of filled bunches per beam, respectively. The next two values describe the energy of the beam: f_{rev} is the frequency with which the beams travel the ring and γ_r is the relativistic factor accounting for the near speed-of-light travel of the beams.

The denominator contains the normalized transverse beam emittance, ϵ_n and the beta function at time of collision, β^* . The transverse beam emittance is a measure of how wide the beam is in the transverse (xy) plane, while the beta function is a measure of the beam width in the longitudinal direction. The entire function is then multiplied by a geometric factor F that accounts for the crossing angle at time of collision (the beams do not collide fully head-on, but at a non-negligible angle).

The actual production rate for a given process is then $N_{\text{events}} = L \times \sigma_{\text{prod}}$, where σ_{prod} is the production cross-section for the desired physics process. The full, integrated luminosity actually used in any analysis is measured carefully by the detectors themselves, and not simply calculated from the delivered beam parameters [47]. The luminosity detectors are described in Section 4.4.

Chapter 4

A BRIEF DESCRIPTION OF THE ATLAS DETECTOR

The ATLAS (**A Toroidal Large Hadron Collider ApparatuS**) experiment is a multi-purpose detector designed for the detection of a few significant physics processes. The first of these was the discovery of the Higgs boson, which was in fact discovered in 2012 [7]. However, the Standard Model is still not fully descriptive of the physical processes we have, thus spawning many other theories (see Section 2.2). Thus, the discovery of these processes, such as heavy W' or Z' bosons or supersymmetric partners to SM particles, were also an important consideration in the ATLAS detector design.




Figure 4.1: A digital rendering of the ATLAS detector. Each of the sub-detectors is labeled. From [4].




Figure 4.2: The coordinate system used to describe the ATLAS detector. The polar angle θ spans $[0, \pi]$, and the azimuthal angle ϕ spans $(-\pi, \pi]$. The pseudorapidity η is defined as $-\ln[\tan(\theta/2)]$, and spans the range $(-\infty, \infty)$. Because the detector is \sim symmetric in ϕ , $r = \sqrt{x^2 + y^2}$ is generally used to describe a displacement from the central z -axis.

The wide range of signatures from the aforementioned processes, including clean diphoton events from $H \rightarrow \gamma\gamma$ and multi-jet, high E_T^{miss} SUSY events, ensured the ATLAS detector has multiple subsystems, a rendering of which is shown in Figure 4.1. In addition to the original design, the detector underwent upgrades between the 2012 and 2015 data-taking campaigns. The full, upgraded detector was used to collect the data used in this thesis, and is thus the detector “version” described in this chapter.

The ATLAS detector is oriented along the LHC ring such that the z -axis follows the beam-line, the x -axis points inwards to the centre of the LHC ring, and the y -coordinate points vertically upwards. The detector is approximately symmetric both around the z -axis (symmetric in ϕ) and upon reflection in the $x - y$ plane. In the $x - y$ plane, the azimuthal angle is denoted by ϕ , spans $(-\pi, \pi]$ and starts from the x -axis, while the polar angle θ spans the range $[0, \pi]$. These variables are all depicted in Figure 4.2.

The detector is split into three sections: a middle *barrel* sandwiched between two *endcaps*.

The detector layers in the barrel are cylindrical, and designed to detect particles traveling perpendicularly away from the beam pipe (particles with a large transverse component to their momentum). The endcaps, while still symmetric in ϕ , are designed to detect particles with a larger longitudinal momentum traveling somewhat along the z -axis. Consequently, layers of the endcaps are wheel-shaped, sandwiching the barrel.

4.1 Inner detector

The innermost sub-detector, the inner detector (ID), is a tracking system designed to accurately reconstruct charged tracks from particles down to 100 MeV. The ID, shown in Figure 4.3, has three subsystems, each located concentrically around the beam line, inside a 2T axial magnetic field created by a solenoid. Closest to the beam line is the pixel detector, followed by the silicon microstrip tracker (SCT) and then the transition radiation tracker (TRT). An additional pixel layer, the insertable B-layer (IBL), was recently added for Run 2. The pixel and SCT detectors both provide discrete, precision measurements for tracking, while the TRT provides a continuous extension of tracks used for improved resolution.

4.1.1 Pixel detector

The pixel detector provides barrel coverage in the region $0 < |z| < 400.5$ mm, $50.5 < R < 122.5$ mm and $495 < |z| < 650$ mm, $88.8 < R < 149.6$ mm in the endcaps. There are over 80 million readout channels, from 47232 pixels on each of 1744 pixel sensors. Each sensor is $50 \times 400 \mu\text{m}^2$, and the fine granularity provides a hit accuracy of 10μ in $(R - \phi)$ and $115 \mu\text{m}$ in z (10μ in $(R - \phi)$ and $115 \mu\text{m}$ in R) in the barrel (endcaps).

The IBL, which was partially constructed at the University of Washington, has recently been added between the innermost pixel layer and the beam pipe. This serves two important purposes: first, to increase tracking efficiency, especially for b -tagging, and second to protect the inner pixel detector layer from the radiation effects from increasing luminosity. The IBL provides coverage in the range $|z| < 332$ mm and $31.0 < R < 40.0$ mm, with resolution similar to the pixel layers.




Figure 4.3: A digital rendering of the ATLAS inner detector. Each of the tracking sub-detectors is labelled, and their radial locations are marked. From [5].

4.1.2 Silicon microstrip tracker

The SCT covers the range $299 < R < 514$ mm, $0 < |z| < 749$ mm in the barrel and $275 < R < 560$ mm, $839 < |z| < 2735$ mm in the endcaps. The barrel is comprised of four cylindrical layers, each with two strips. In each layer, one strip is aligned along z (the beam line), and the other is offset by 40 mrad, with a pitch of $80 \mu\text{m}$. There are nine SCT layers in the endcap, with similar strip pairs.

Three-dimensional space points (SPs) are created by fitting the pulse registered in each strip. Since this procedure results in measurements with a large uncertainty, SPs are additionally constrained by requiring the line between the two strips to be pointing to the beam spot¹. The resolution in $(R \times \phi)$ is $17 \mu\text{m}$ and in both the barrel and the endcaps, while in the barrel the resolution in z is $580 \mu\text{m}$ and in the endcaps the resolution in R is $580 \mu\text{m}$.

1

In cosmic reconstruction, this constraint is removed in favor of a perpendicular line, which simply fits to the closest line between the strips

4.1.3 Transition radiation tracker

The TRT is made up of drift tubes (straws) and transition radiation material. When a particle traverses the TRT, it averages 36 reconstructed hits on track. The barrel consists of 73 layers of 4 mm diameter straws, in the region $563 < R < 1066$ mm, $0 < |z| < 712$ mm. The 144 cm long straws are oriented parallel with the beam pipe. Because of their length, no z information is available, but they provide a resolution of $130 \mu\text{m}$ in $(R \times \phi)$.

In the endcaps, there are 160 layers of 37 cm long straws, covering the region $644 < R < 1004$ mm, $848 < |z| < 2710$ mm (corresponding to $|\eta| < 2.0$). The per-straw resolution is the same as the barrel, though the straws are oriented vertically and thus provide $130 \mu\text{m}$ resolution in z .

4.2 Calorimeters

The calorimeter system is split into two parts: an electromagnetic calorimeter (ECal) designed to accurately detect electrons and photons, and a hadronic calorimeter (HCal) designed to absorb energy from hadronic showers. The ECal generally absorbs all electromagnetic energy, since the final decay products in an electromagnetic shower will always be electrons, which the calorimeter proceeds to measure. However, hadronic energy is not as simple, and is not fully contained and measured for two main reasons: Through weak interactions, hadrons can decay into undetectable neutrinos or minimally interacting muons, and hadrons can lose some of their energy to inducing nuclear interactions in the detector material. The energy of hadronic jets is corrected with a scale factor to account for these unmeasurable components.

In combination, the two calorimeters provide coverage in the region $|\eta| < 4.9$ (Figure 4.4) with between 10 and 20 interaction lengths of material. Both calorimetry systems are constructed of alternating layers of absorbers and active sampling material. The absorbers serve to cause particles traveling through the detector to shower, after which the energy of the particle shower can be measured by the sampling material. This allows the calorimeter to

measure the entire shower within a smaller volume; however, the exact energy of the shower must be estimated, since some of the original particle's energy ends up deposited in the absorbers.




Figure 4.4: Number of radiation lengths in the combined inner detector and calorimeter system, as a function of $|\eta|$. The uppermost blue, unlabelled region spanning $0 < |\eta| < 3$ is depicting the material between the outer edge of the calorimeter and the first layer of the muon spectrometer. From [4].

4.2.1 Electromagnetic calorimeter

The electromagnetic calorimeter (ECal) is divided into one barrel and two endcap calorimeters, each with three layers and a presampler. The ECal is made of liquid argon (LAr) as the sampling material and lead plates as the absorber. Copper electrodes are located inside the folds to detect electrons. The lead plates are folded into an accordion shape, which allows for both full coverage and symmetry in ϕ . Figure 4.5 shows the accordion structure and the cell geometry in $\eta - \phi$ in the barrel.




Figure 4.5: Geometrical structure of ECal modules in $\eta - \phi$. From [4].

The dimensions of each ECal layer are listed in Table 4.1. Most of the energy from an electromagnetic shower is deposited in layer 2, which comprises the most radiation lengths (see Figure 4.5). The thin presampler, which only covers the region $|\eta| < 1.8$, aims to correct for any energy lost by the incident particle as it traveled from the interaction point to the ECal. Layer 1 is finely grained in order to be able to provide an accurate η measurement, while the outermost Layer 3 measures the tail of the shower not absorbed by the first three layers.

4.2.2 Hadronic calorimeter

The hadronic calorimeter (HCal) is made up of a number of distinct regions and materials. The barrel ($|\eta| < 0.8$) and extended barrel ($0.8 < |\eta| < 1.7$) regions are made with a plastic scintillator tiles and steel absorbers. Signal (i.e., light) in the scintillator tiles is detected with wavelength-shifting fibers, which are bundled into photomultiplier tubes (PMTs).

Table 4.1: Electromagnetic calorimeter coverage and resolution. Layers with a range of provided granularities have a granularity that varies with eta; a full breakdown is available in Ref. [4]. Note that there are only two layers in the endcap for $|\eta| > 2.5$.

Layer	Region	Coverage	η range	Granularity ($\Delta\eta \times \Delta\phi$)
Presampler	Barrel	$1.42 < R < 1.44$ m	$ \eta < 1.52$	0.025×0.1
	Endcap	$3.68 < z < 3.72$ m	$1.5 < \eta < 1.8$	0.025×0.1
Layer 1	Barrel	$1.48 < R < 1.58$ m	$ \eta < 1.4$	$0.025/8 \times 0.1$
		$1.48 < R < 1.58$ m	$1.4 < \eta < 1.475$	0.025×0.025
	Endcap	$3.75 < z < 3.80$ m	$1.375 < \eta < 2.5$	$(0.025/8 - 0.025) \times (0.025 - 0.1)$
		$3.75 < z < 4.20$ m	$2.5 < \eta < 3.2$	0.1×0.1
Layer 2	Barrel	$1.58 < R < 1.84$ m	$ \eta < 1.35$	0.025×0.0245
	Endcap	$3.80 < z < 4.18$ m	$1.375 < \eta < 2.5$	$(0.025 - 0.1) \times (0.025 - 0.1)$
		$4.20 < z < 4.25$ m	$2.5 < \eta < 3.2$	0.1×0.1
Layer 3	Barrel	$1.84 < R < 1.99$ m	$ \eta < 1.35$	0.05×0.0245
	Endcap	$4.18 < z < 4.25$ m	$1.5 < \eta < 2.5$	0.05×0.0245

The endcaps (HEC) use LAr like the ECal, though the absorber is copper, and they do not have the same accordion geometry. Despite this, because of the LAr usage, the ECal and HEC are often collectively referred to as the LAr calorimeters.

The HCal is intended to absorb all energy from jets. This is of course not the case: a small fraction of the time, some of the particles *punch through* the calorimeters, and end up interacting with the muon system. Punch-through is an important systematic uncertainty for jet measurement. However, for this analysis the most important consequence of punch-through is that in rare cases, the activity in the muon system can resemble that from a displaced decay. Jet punch-through is discussed further in Chapter 8.

4.3 Muon spectrometer

From the point of view of the Standard Model, the main particles of interest that are not fully absorbed by the calorimeter system are muons. The muon spectrometer is designed

Table 4.2: Hadronic calorimeter coverage and resolution.

Layer	Region	Coverage	η range	Granularity ($\Delta\eta \times \Delta\phi$)
Tile layer 1	Barrel	$2.28 < R < 2.60$ m	$ \eta < 1.0$	0.1×0.1
	Extended barrel	$2.28 < R < 2.60$ m	$0.8 < \eta < 1.7$	0.1×0.1
Tile layer 2	Barrel	$2.60 < R < 3.45$ m	$ \eta < 1.0$	0.1×0.1
	Extended barrel	$2.60 < R < 3.14$ m	$0.8 < \eta < 1.7$	0.1×0.1
Tile layer 3	Barrel	$3.45 < R < 3.87$ m	$ \eta < 1.0$	0.2×0.1
	Extended barrel	$3.14 < R < 3.87$ m	$0.8 < \eta < 1.7$	0.1×0.1
LAr layer 1	Endcap	$4.35 < z < 4.63$ m	$1.5 < \eta < 3.2$	$0.1 \times 0.1, 1.5 < \eta < 2.5$ $0.2 \times 0.2, 2.5 < \eta < 3.2$
LAr layer 2	Endcap	$4.63 < z < 5.10$ m	$1.5 < \eta < 3.2$	$0.1 \times 0.1, 1.5 < \eta < 2.5$ $0.2 \times 0.2, 2.5 < \eta < 3.2$
LAr layer 3	Endcap	$5.13 < z < 5.59$ m	$1.5 < \eta < 3.2$	$0.1 \times 0.1, 1.5 < \eta < 2.5$ s $0.2 \times 0.2, 2.5 < \eta < 3.2$
LAr layer 4	Endcap	$5.59 < z < 6.05$ m	$1.5 < \eta < 3.2$	$0.1 \times 0.1, 1.5 < \eta < 2.5$ $0.2 \times 0.2, 2.5 < \eta < 3.2$

to provide location and momentum measurements of muons both in the trigger and offline reconstruction. A large air-core toroid magnet bends charged particles traveling in the region $|\eta| < 1.4$, while in there is a smaller toroidal magnet in each endcap, covering the region $1.6 < |\eta| < 2.7$. The field in the region between the toroids is created by the overlap of the edges of the fields from the barrel and endcap toroids. The overall toroid structure is configured such that muons traversing the covered regions will encounter a field perpendicular to their direction of motion, thus bending their trajectories proportional to their momentum.

It is important to note that these toroids are designed to bend muons traveling into the muon spectrometer, and are not necessarily placed in a way such that the decay products of long-lived particles will encounter a magnetic field. The endcap toroid is located between the inner and middle layers of the endcap muon system (see Figure 4.6), and thus tracks in the middle and outer chambers themselves will not be bent.

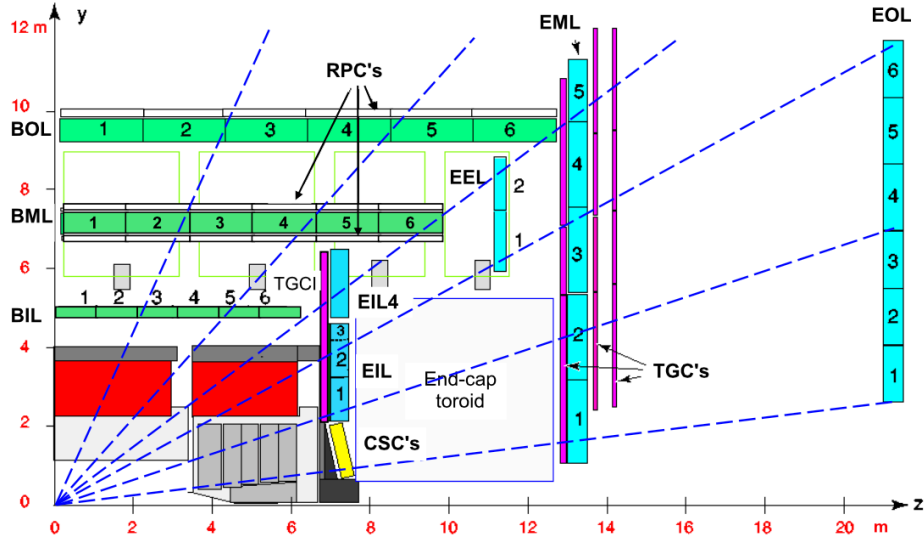


Figure 4.6: Cross-section of the Muon Spectrometer. From [4].

The muon system is split into a barrel and two endcaps, and made up of four sub-detectors. Two of the sub-detectors are for offline measurements: monitored drift tubes (MDTs) and cathode strip chambers (CSCs), and the remaining two, resistive plate chambers





Figure 4.7: Magnetic field integral for a particle traveling straight through the Muon Spectrometer (e.g. infinite momentum muon). From [4].


(RPCs) and thin gap chambers (TGCs) are used both as a trigger system and for additional hit information for offline reconstruction. Each of these is described in detail in the following subsections. Coverages of each of the sub-detectors in the barrel and endcap regions are detailed in Table 4.3. In addition to MDT Inner, Middle, and Outer layers, there are also “Extra” chambers. The locations of these are not detailed because they are not used in the analysis described in this thesis. Also, there are caveats to the locations and dimensions of some chambers when they would otherwise interfere with detector support structures.

Table 4.3: Muon spectrometer coverage. For MDT layers, the coverage represents the chamber extent. For trigger planes, the coverage represents the inner edge of the trigger plane. Small and large sectors alternate in ϕ , as shown in Figure 4.8. All numerical values are extracted from the technical drawings in Figure 4.8. **MDT Outer chamber values are less accurate due to the chamber locations changing after the technical designs were created.


Barrel			
Layer	Type	Coverage	η range
MDT Inner	Small	$4383 < R < 4667$ mm	$ \eta < 1.05$
	Large	$4718 < R < 5134$ mm	$ \eta < 1.05$
MDT Middle	Small	$7888 < R < 8252$ mm	$ \eta < 1.05$
	Large	$6861 < R < 7372$ mm	$ \eta < 1.05$
MDT Outer	Small	$10275 < R < 10786$ mm	$ \eta < 1.05$
	Large	$9222 < R < 9733$ mm	$ \eta < 1.05$
RPC1	Small	$R = 7758$ mm	$ \eta < 1.05$
	Large	$R = 6729$ mm	$ \eta < 1.05$
RPC2	Small	$R = 8282$ mm	$ \eta < 1.05$
	Large	$R = 7372$ mm	$ \eta < 1.05$
RPC3	Small	$R = 10158$ mm	$ \eta < 1.05$
	Large	$R = 9733$ mm	$ \eta < 1.05$
Endcaps			
Layer	Type	Coverage	η range
MDT Inner	Small	$7023 < z < 7390$ mm	$1.0 < \eta < 2.7$
	Large	$7409 < z < 7776$ mm	$1.0 < \eta < 2.7$
MDT Middle	Small	$13265 < z < 13632$ mm	$1.0 < \eta < 2.7$
	Large	$13660 < z < 14024$ mm	$1.0 < \eta < 2.7$
MDT Outer**	Small	$21260 < z < 21580$ mm	$1.0 < \eta < 2.7$
	Large	$21680 < z < 22000$ mm	$1.0 < \eta < 2.7$
TGC1		$ z = 13.3, z = 13.4$ m	$1.05 < \eta < 2.4$
TGC2		$ z = 14.6, z = 14.7$ m	$1.05 < \eta < 2.4$
TGC3		$ z = 15.0, z = 15.1$ m	$1.05 < \eta < 2.4$



(a)



(b)



(c)

Figure 4.8: Technical drawings of the (a) barrel and (b),(c) endcap muon spectrometer. Physical differences between small and large chambers are clearly visible in both the barrel and endcaps. Note that these drawings were from the original technical design report. Additional chambers have since been added, and the chamber naming scheme in the endcaps has changed. The locations of the outer endcap wheels were altered before detector construction.

4.3.1 Monitored drift tubes

The monitored drift tubes (MDTs) are small, aluminum tubes with diameter 29.970 mm filled with a mixture of 93% argon, 7% carbon dioxide, and a tiny amount of water. The inner pressure is held at a constant 3 bar, and in the centre of each tube there is a 50 μm diameter tungsten-rhenium wire held at a potential of 3080 V.

When a charged particle travels through an MDT, it will ionize the gas in the tube. Freed electrons will then be attracted to the anode wire, and a charge deposition is collected. The drift time is defined as the time it takes for electrons freed at the distance of closest approach to reach the wire. Drift times are converted into drift radii using RT functions. RT functions depend on gas and magnetic field conditions and thus each chamber is assigned a separate function.




Figure 4.9: Schematic of an MDT chamber.

Sets of MDTs are arranged into *chambers*. In general, each chamber has two *multilayers*, and each multilayer is made of three or four layers of MDT tubes. The multilayers are separated by a rigid spacer. A cross-section of such a chamber is shown in Figure 4.9.

In the barrel, MDT chambers are arranged into three cylindrical layers: inner, middle, and outer. Because the chambers are rectangular, they are staggered in each layer between small and large chambers, meaning the inner radius of each layer varies periodically around

ϕ . This staggering is illustrated in Figure 4.8a.

4.3.2 Cathode strip chambers

When two proton bunches collide, most of the resulting energy stays relatively collimated with the beam pipe. Thus, the sub-detectors located at large η must be able to withstand a much higher rate than those located at a more transverse position, or farther from the interaction point. MDT chambers can only withstand rates of up to 150 Hz/cm², and thus a separate system made of cathode strip chambers, which can withstand rates up to 1000 Hz/cm², is located in the innermost layer of the muon spectrometer endcap ($2.0 < \eta < 2.7$).

Cathode strip chambers (CSCs) are segmented into two wheels, each of eight chambers. The two wheels are offset, such that no coverage is lost between chambers. Each chamber contains wires oriented radially, which induce charges on cathodes oriented perpendicular to the wires. The particle position is determined by analyzing the signal induced on neighboring cathode strips.

4.3.3 Resistive plate chambers

In the barrel muon system, the middle MDT layer is sandwiched between resistive plate chambers (RPCs). A third layer of RPCs is located on the outside of the outer MDT layer of large chambers, and the inside of the outer MDT later of small chambers. Each of these is referred to as an RPC *station*, and from inside out the stations are referred to as RPC1, RPC2, and RPC3.

RPC chambers consist of two, parallel resistive plates made from phenolic-melaminic plastic laminate. The plates are held 2 mm apart, and the gap between them is filled with the gas mixture C₂H₂F₄/Iso-C₄H₁₀/SF₆ (94.7/5/0.3). The electric field between the two plates is held at 9.8 kV.

Each RPC station is made from two RPC chambers, such that a charged particle traveling through a chamber should leave two hits.

4.3.4 *Thin gap chambers*

In the muon system endcaps, the trigger system is instead made from thin gap chambers (TGCs). There are nine layers of chambers arranged into four sectors. One pair of layers is located on the inside of the inner MDT plane, a triplet of layers is on the inside of the middle MDT plane (M1), and two pairs of layers (M2 and M3) on the outside of the middle MDT plane.

The TGC chambers are constructed similar to the cathode strip chambers. Parallel wires 1.8 mm apart are sandwiched between graphite cathodes 2.8 mm apart, such that the graphite-wire distance is less than the wire-wire distance. The chambers are filled with CO₂ and n-C₅H₁₂ in a 55%-45% mixture.

4.4 *Detectors for luminosity determination*

Besides being able to record and reconstruct events, the luminosity recorded by the detector must also be known. This is done with two main detector systems: LUCID and the beam conditions monitor (BCM). Having two separate detectors allows for a lower uncertainty on the luminosity measurement, and these are also complemented by luminosity algorithms with the standard sub-detectors.

Luminosity measurement with a Cherenkov Integrating Detector (LUCID) is made of sixteen, 1.5 m long aluminum tubes. These were originally filled with C₄F₁₀, but in 2012 the gas was removed. There is one LUCID module located on either side of the ATLAS detector, each at a distance of 17 m from the interaction point. The detector was designed such that photomultiplier tubes (PMTs) detect Cherenkov photons caused by particles ionizing the gas. There is a layer of quartz separating the tubes and PMTs, and this can serve to create Cherenkov photons for detection as well.

The BCM is located much closer to the interaction point, and its primary purpose is to detect if there are beam issues requiring an abort to protect main detector. However, they also provide valuable luminosity information. The BCM consists of four diamond sensors

located on either side of the interaction point at $|z| = 184$ mm, $R = 55$ mm (corresponding to $|\eta| = 4.2$). They are oriented in a cross pattern around the beam, with two sensors paired vertically along the y -axis (BCM V) and the other two horizontally along the x -axis (BCM H). The BCM pairs provide independent luminosity measurements, while a third measurement is made by grouping all four sensors together, and is referred to as BCM T .

Chapter 5

TRIGGERS FOR LONG-LIVED PARTICLES

5.1 Introduction

The LHC is currently running with a bunch spacing of 25 ns, which means that the input rate of events to ATLAS is 40 MHz. It is impossible with current technology and financial limitations to store and process all of these events, so a *trigger* system is employed in order to decide which events are interesting enough to keep, and which can be skipped. A rate of 40 MHz is also too high to fully process events before deciding to store them or not, so the trigger system is broken up into two distinct steps.

First, a hardware-based trigger is used to make a fast, coarse selection of candidate events to record. This is referred to as the *Level-1* (L1) trigger, and it is designed to output events at a rate of 100 kHz. Anything selected by the Level-1 trigger is then passed on to the *High Level Trigger* (HLT), which reconstructs events in a much more detailed manner. Events from the HLT are selected at a maximum rate of 1.5 kHz, although this number is luminosity-dependent and thus decreases over the course of a run.

The events are selected by each level of the trigger system according to a list of trigger *items* and chains. These are defined separately at L1 and HLT, and each HLT chain is generally seeded by a specific L1 item (though a single L1 item can seed many HLT chains). Sometimes, the types of events that one wants to select occur with too large a rate to write out every occurrence. Triggers selecting such events are assigned a *prescale*, which defines the fraction of events passing the criteria that are actually kept. At Level-1 the prescales are applied *after* the trigger decision is made. If there is an additional, separate prescale at HLT, only the fraction of events equal to the prescale fraction are reconstructed at the HLT,

saving precious CPU time¹. Each L1 item and each HLT chain has an associated prescale. A given set of L1 items, HLT chains, and their associated prescales comprise a *trigger menu*.

ATLAS employs a number of triggers designed to select events with standard physics objects, such as leptons and jets. While these triggers are well-designed and efficient for selecting the topologies they are designed for, they are not specialized for selecting events with long-lived neutral particles, which often lack the signatures of traditional decays. Using standard triggers causes two problems in a displaced analysis. First, it restricts the search to topologies that have a prompt object in addition to the displaced object if the displaced decay is beyond the calorimeters. Second, it restricts the search to high energy decays or events with high energy objects, since in order to have an unprescaled trigger the threshold must be sufficiently high to have an acceptable rate. Using a prescaled trigger would lower the sensitivity of the analysis by the prescale factor, which would correspondingly decrease the sensitivity of any search by a factor equal to the prescale.

5.2 Level-1 trigger system

The L1 trigger system reduces the rate from an input of 40 MHz to a maximum of 100 kHz², which is the maximum rate that the detector readout and HLT farm can deal with. Much of this rate reduction is helped by the fact that not every single bunch crossing happens with two filled bunches - so only a fraction of the bunch crossings have possibly interesting physics activity³. A typical 2016 bunch group structure has 2200 paired bunch crossings out of a total of 3563 total bunch crossings, thus reducing the input rate from 40 MHz to 25 MHz.

¹

e.g. If 100 events pass L1 and the HLT prescale is 5, then only $100/5 = 20$ events are sent to the HLT. If the HLT passes on average 10% of events, then 2 events pass the trigger.

²

The Run-2 L1 output rate was limited by various detector readout systems, and thus the maximum output rate has ranged from 50 kHz to 100 kHz. The majority of recorded data has been taken with an average L1 output rate of 75 kHz.

³

A small fraction of events from empty or unpaired bunch crossings are also kept for non-collision background studies, though these triggers are generally comparatively low rate and often highly prescaled

The remaining rate reduction to 100 kHz is done by selecting events passing L1 items. Up to 512 L1 items can be defined, each of them selecting a specific type of event. They contain single item thresholds (such as a 20 GeV muon, L1_MU20), combinations of items (such as a 6 GeV muon and a 20 GeV jet, L1_MU6_J10), or more complicated item combinations determined using the topological trigger (see Section 5.2.3). The L1 trigger currently uses information from two main parts of the detector - the calorimeters (L1Calo trigger) and the muon system (L1Muon trigger), in addition to smaller subsystems used for luminosity and minimum-bias physics studies.

Information about the types of energy deposits from the L1Calo and L1Muon systems are passed to the *Central Trigger Processor* (CTP), which then determines if any of the L1 item thresholds have been reached. The CTP also performs the combination of items, such as the aforementioned L1_MU6_J10. Combinations can also be multiplicities of the same type of trigger, such as L1_2MU10, which is used in the specialized trigger described in Section 5.4.2.

5.2.1 Level-1 calorimeter trigger

The level-1 calorimeter trigger (L1Calo) compiles three types of L1 items: jet, electromagnetic (EM), and tau. All three items use *trigger towers* of granularity $(\Delta\eta \times \Delta\phi) = (0.1 \times 0.1)^4$ made of calorimeter cell energy deposits. The towers are formed from energy deposited in any layer of the electromagnetic or hadronic calorimeters, along lines of constant η .

Level-1 jet regions of interest (RoIs) are formed with the Jet/Energy Module (JEM), using 2×2 sums of trigger towers called *jet elements*. A sliding window algorithm [27] looks for 2×2 , 3×3 , and 4×4 detector regions that contain a local maxima in energy. To eliminate the possibility of overlapping windows, a local maxima of size 2×2 jet elements is required to be centered within the window. Identified regions of interest are denoted by L1_J, followed by the minimum energy threshold the RoI has passed. For example, a 50 GeV L1 jet RoI would be called L1_J50.

4

The granularity is slightly larger in high η regions, see Reference [27] for details.

Level-1 EM and tau RoIs are formed in much the same way as jet RoIs, though because of the narrower signatures from electrons, photons, or taus, the original 0.1×0.1 trigger towers are used. For these algorithms, the calorimeter is defined by all possible 4×4 clusters of trigger towers, including those that overlap. For each cluster possibility, the following steps are performed for EM RoIs:

- Calculate the energy sum of all possible 2×1 groups using the centre 4 towers in the cluster.
- Determine if any pair passes one of the EM energy thresholds.
- Sum the energy EM deposited in the 12 outer trigger towers of the cluster, and determine if this sum is less than the allowable energy threshold, for isolation.
- Sum the hadronic energy in the centre 2×2 region of the cluster, and determine if this sum is less than the allowable energy threshold for hadronic leakage of EM activity.
- If both isolation requirements are met and an EM threshold is passed, create an RoI to pass to the CTP.

For tau RoIs, the procedure is similar:

- Calculate the energy sum of all possible 2×1 groups using the centre 4 towers in the cluster.
- Add the hadronic energy in the centre 2×2 region of the cluster to each of the EM sums, and determine if this sum passes one of the tau energy thresholds.
- Sum the energy EM deposited in the 12 outer trigger towers of the cluster, and determine if this sum is less than the allowable energy threshold, for isolation.

- If the isolation requirement is met and a TAU threshold is passed, create an RoI to pass to the CTP.

The L1 RoIs are denoted L1_EM and L1_TAU, followed by the minimum energy threshold passed.

5.2.2 Level-1 muon trigger

There are two types of L1 muon items: low- p_T and high- p_T . Low- p_T muon RoIs have looser hit requirements than high- p_T muons. The looser requirements result in a higher signal efficiency, since it allows for triggering on muons that might not otherwise travel as far as the third trigger plane. However, it also allows for more noise, especially in the endcaps.

In the barrel, the middle layer of the MS is comprised of MDT chambers sandwiched between RPC trigger layers in the barrel, as described in section 4.3. A third trigger layer is located on the outer MDT layer. Each trigger layer has two RPC planes and thus reports two hits (a *doublet*) when traversed by a charged particle. Each L1 muon trigger is seeded by a doublet from the middle, or pivot, RPC plane. The trajectory of the hit is then projected backwards towards the IP, depending on the muon p_T . If a matching hit is found in at least one of the inner RPC planes, the muon p_T is computed and if it is greater than 4 GeV, the trigger result is passed to the CTP.

For a muon to be considered high- p_T , there must be hits in all three trigger planes: if at least one hit in the outer RPC doublet matches a low- p_T trigger, the p_T is computed. The item is passed to the CTP if it passes a high- p_T threshold of 10 GeV. To discriminate high- p_T , 10 GeV muons from low- p_T , the high- p_T muons are referred to as MU11.

There are six preset L1 muon thresholds, and each muon RoI is described by the highest threshold it passes. In 2016 the low- p_T muon thresholds were 4, 6, and 10 GeV. The three high- p_T muon thresholds for 2016 were 10, 15, and 20 GeV. However, even if a low- p_T muon has a p_T of 25 GeV, it is a MU10 item, because it does not meet the requirements of a high- p_T muon. Once the L1 muons are found, they passed to the CTP with only threshold

and detector location information.

The Level-1 trigger in the endcaps is slightly more complicated and has undergone multiple changes in the past few years due to the larger rate increase of activity in the endcaps with increasing luminosity and pileup. The L1 threshold used in this analysis, MU10, requires a three-station coincidence between TGC1, TGC2, and TGC3.

5.2.3 *Level-1 topological trigger*

Between the end of Run-1 and beginning of Run-2, the L1 trigger system was upgraded to allow for topological decisions, increasing the breadth of the trigger system.

Topological criteria allow for lowered L1 thresholds given a preferred event topology. For example, if two approximately back-to-back muons are desired in the final event, then the trigger at level-1 can now be configured to select events at L1 with back-to-back muon RoIs, allowing for a lower L1 rate given the same muon p_T thresholds.

The current application for displaced decays is for decays in the hadronic calorimeter. Traditionally, L1TAU items have been used because they select narrow jets. LLPs decaying in the calorimeters do not have time to spread out as much as prompt jets, and thus look narrower to an algorithm designed for prompt physics. However, L1TAU items also depend on the EM energy deposits, which is not desired for displaced HCal jets. To this effect, a new L1 item has been designed that selects L1TAU RoIs that have no more than 4 GeV deposited in the ECal. Additionally, an item that will compute the ratio between the HCal and ECal energy in L1TAU RoIs is being commissioned.

These new triggers will allow for the selection of much lower energy events, allowing for increased sensitivity to Higgs boson decays to light long-lived particles.

5.2.4 *Level-1 zero bias trigger*

Zero-bias triggers are a fancy way of saying that there is no bias to the events being selected: they are truly random. These triggers do not rely on there being a hard scatter, or even a minimal elastic collision with track activity. Instead, events are selected at L1 from bunch

crossings that occur one full turn of the accelerator *after* a high- p_T L1 trigger. This extra step, rather than selecting random bunch crossings, ensures that events are selected at a rate corresponding to the luminosity of the beams. The L1 zero-bias trigger is denoted L1_ZB.

The sample of events collected by L1_ZB is representative of what could be happening in the background during collision events, whether it is due to pileup, cosmic activity, or other non-collision backgrounds. Occasionally a collision event is selected by random, but because the cross-section for the hard scatter is incredibly low, these events are a minimal fraction of the overall number selected at random. The zero-bias trigger used in this analysis is called HLT_noalg_zb.L1ZB, which means that there are no additional requirements happening in after the initial L1 selection.

5.3 High level trigger

The high level trigger, or HLT, is a software-based trigger that processes the 100 kHz input from level-1 and outputs events at a manageable, average rate of 1 kHz [28]. The HLT has the difficult problem of rapidly processing and selecting a small fraction of events while simultaneously performing accurate reconstruction of physics objects. Many objects are reconstructed using the same algorithms that run in full, offline event reconstruction. Others, such as track-formation in the inner detector, are more computationally intensive processes and instead have dedicated algorithms running at trigger level [29].

Many of these algorithms are too computationally intensive to allow for the reconstruction of full events at HLT at the rate delivered by the L1 seed. Thus, instead of reconstructing objects in the full detector, only a fraction of the detector is used. The locations the HLT reconstructs are defined by the L1 RoIs seeding the HLT chain. The size of the area around the L1 seed is designated by the physics requirements of the objects being reconstructed. This procedure is called *partial scan* reconstruction, as opposed to the normal *full scan* reconstruction that uses the full detector.

The analysis in this thesis uses a number of different triggers: high- p_T jet triggers for control regions, zero-bias triggers for background determination, and custom triggers to select

candidate signal events.

The high- p_T jet trigger employed is called HLT_j400. This trigger runs the standard anti- k_T jet-finding algorithm [30] with jet radius 0.4, and selects events with at least one offline jet with $p_T > 400$ GeV. It is seeded from the L1 item L1_J100.

5.4 Specialized triggers

Despite the breadth of standard triggers ATLAS employs, they are nearly all designed to detect particles with microscopic lifetimes. Thus, instead of relying on the standard triggers to select potential long-lived particle decays, new triggers were developed to ensure that as many potential signal events as possible are recorded. Two sets of triggers were running in 2015 and 2016: CalRatio triggers for selecting decays in the hadronic calorimeter and the Muon RoI Cluster trigger for selecting decays in the muon spectrometer [31].

5.4.1 CalRatio triggers

A long-lived, neutral particle decaying at the end of the ECal or in the HCal will still deposit most of its energy in the calorimeters, but the balance of energy in the ECal relative to the HCal is unlike a jet produced in standard model processes. The CalRatio triggers exploit this energy difference using the variable $\log_{10}(E_{\text{HAD}}/E_{\text{EM}})$, where E_{EM} is the energy deposited in the ECal and E_{HAD} is the energy deposited in the HCal. At HLT, partial-scan ID tracking is performed using a $(\Delta\eta \times \Delta\phi) = (1.0 \times 1.0)$ square around any jets with $\log_{10}(E_{\text{HAD}}/E_{\text{EM}}) > 1.2$. Only jets with no ID tracks above 2 GeV within $\Delta R = 0.2$ of the jet axis will fire the trigger. Further detail is provided in Reference [31] (Run-1 implementation only).

5.4.2 Muon RoI Cluster trigger

The Muon RoI Cluster trigger (HLT_j30_mvtx) is specially designed to select events with hadronic decays in the ATLAS Muon Spectrometer (MS). Depending on the mass of the

decaying particles, there will be a significant amount of energy deposited in the MS. Unless there are long-lived particles decaying to muons, even standard “MS-only” muon triggers do not have an appreciable efficiency, since they require standalone *muons* reconstructed at HLT. The novel solution was thus not to rely on existing triggers but to define a specialized trigger that exploits the relatively wide shower of energy caused by a displaced hadronic decay in the MS. The trigger is carefully designed to accept as little background as possible while not compromising signal efficiency, since searching for a rare process means that any prescaling would immediately and directly reduce the power of the search.

In addition to the signal-driven trigger, a new trigger was developed for inclusion in the Run-2 trigger menu. This background trigger (HLT_j30_mvtx_noise) selects events with the characteristic MS activity, but has no requirements on the upstream isolation that are normally imposed such that the activity can be considered displaced. This trigger can be used to determine the amount of expected collision-related background. With high enough available bandwidth, this trigger could be the only trigger, since all “signal” trigger events are contained within the background trigger sample. However, at higher luminosity the rate of the trigger will exceed the allotted bandwidth and thus will be prescaled to a rate of 1 Hz. This rate will accept enough events such that the background trigger does not impose a statistical limitation on background estimations.

It should be noted that so far, through the end of 2016, the background trigger has remained unrescaled. The background trigger can therefore be used for signal analysis, with the only isolation criteria being those applied in the offline sample. This slightly boosts efficiencies and simplifies the overall analysis.

At the HLT, the first step is determining whether or not there are sufficient L1 muon RoIs present to indicate the possibility of a displaced hadronic decay. This is done by clustering using a crude cone algorithm that looks for clusters of at least three(four) RoIs in the barrel(endcaps). The RoIs must all be in a cone of radius $\Delta R = 0.4$.

If a cluster of RoIs is found, this is sufficient to pass the aforementioned background trigger. However, the signal trigger requires additional isolation criteria to be imposed.

When a long-lived particle decays outside the calorimeters, it does not deposit any energy in the inner detector or calorimeters. Thus, requiring that there be no reconstructed jets or tracks within some pre-defined cone below the muon RoI cluster will reduce backgrounds from standard model activity while maintaining a high signal efficiency.

Table 5.1: Isolation criteria for the Muon RoI Cluster trigger.

Criteria	Barrel value	Endcap value
Jet min ΔR	0.7	0.7
Jet max E_T	30 GeV	30 GeV
Jet max $\log_{10}(E_{\text{HAD}}/E_{\text{EM}})$	0.5	0.5
Track min ΔR	0.4	0.4
Track max p_T	5 GeV	5 GeV

Jet and track isolation are performed in the trigger with objects reconstructed in the HLT, and the values used are summarized in Table 5.1. If a track has both $p_T > 5$ GeV and $\Delta R(\text{track}, \text{cluster}) < 0.4$, the cluster is considered not isolated. Jets are only able to cause isolation to fail if they have both $E_T > 30$ GeV, $\log_{10}(E_{\text{HAD}}/E_{\text{EM}}) < 0.5$, and $\Delta R < 0.7$. The $\log_{10}(E_{\text{HAD}}/E_{\text{EM}})$ cut is present to ensure that if a LLP decay occurs near the outside of the calorimeters it will still be selected if it has left sufficient activity in the muon spectrometer.

The values for each isolation criterion are currently the same in the barrel and endcaps, though it is clear from Figure 5.1 that the behavior differs between the two regions and that future optimization should take this into account. Figure 5.1 shows the nearest ΔR for jets and tracks in the barrel and endcaps. Note that these plots were made by comparing offline reconstructed objects to muon RoI clusters, since the exact sets of HLT objects were not properly available in the samples for analysis. The results are not expected to be significantly different, but all values should be taken as motivation for a specific criterion rather than proof of an exact value.

Figure 5.2 shows the ΔR in simulated signal events between each muon RoI cluster and

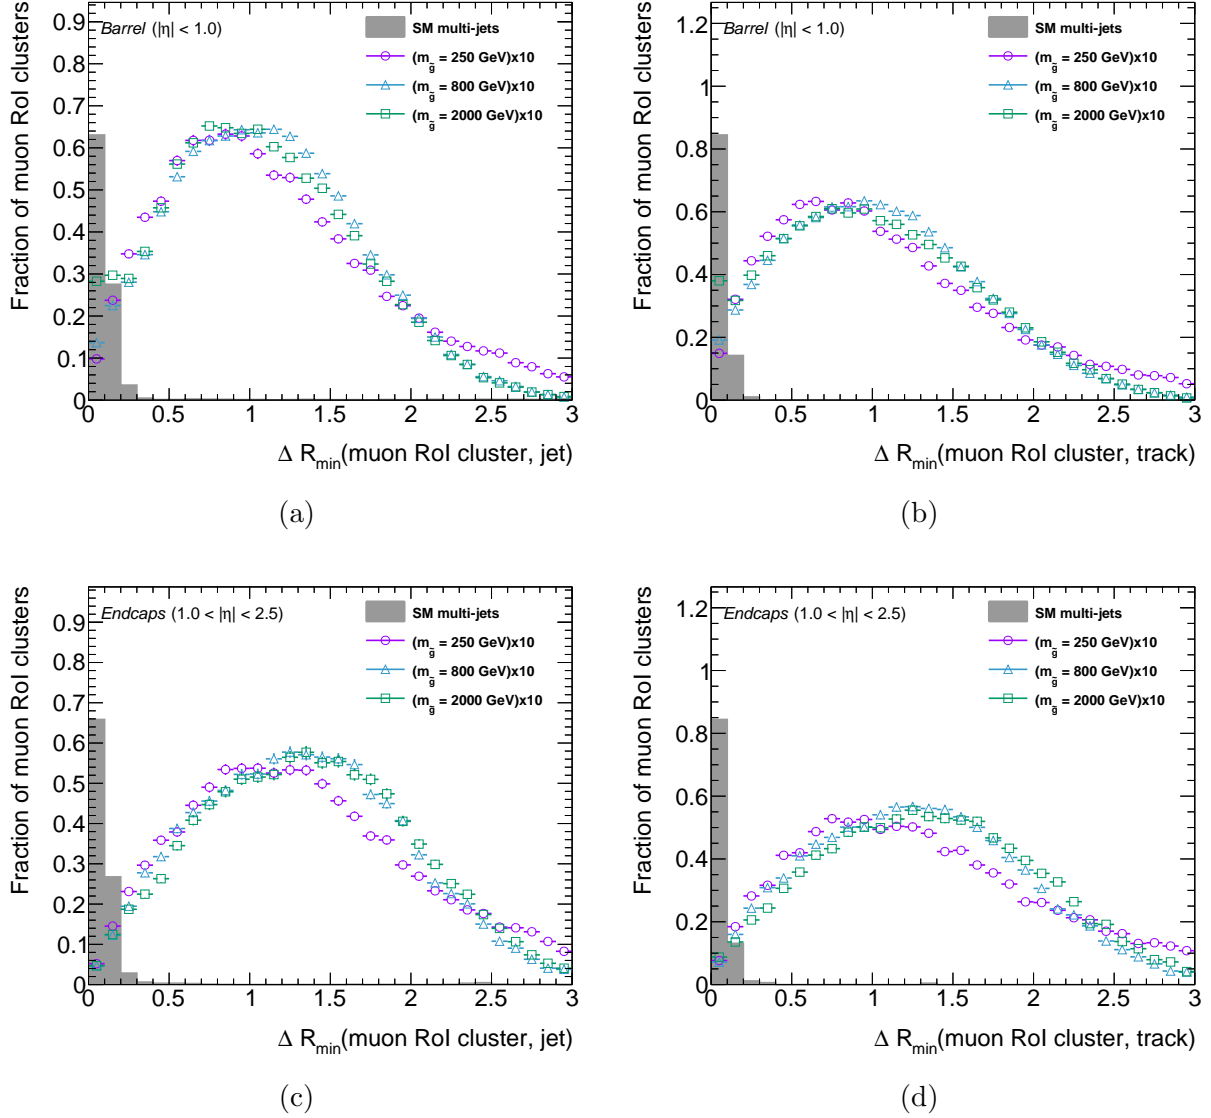


Figure 5.1: The minimum distance between charged particle deposits and a cluster of muon RoIs in SM multi-jet events compared to simulated signal events. “Track” here refers to well-reconstructed, offline inner detector tracks with $p_T > 5$ GeV, while “jet” refers to offline jets with $p_T > 30$ GeV and $\log_{10}(E_{\text{HAD}}/E_{\text{EM}}) < 0.5$. The nearest ΔR between a muon RoI cluster and jets and tracks in the barrel are shown in (a) and (b), while for the endcap they are shown in (c) and (d), respectively.

the nearest long-lived particle decay axis. A long-lived particle is considered *truth-matched* to a muon RoI cluster if the center of the cluster is within $\Delta R = 0.4$ of the LLP decay axis. The efficiency for a long-lived particle to fire the trigger is defined as the number of long-lived particles matched to a triggering muon RoI cluster divided by the total number of long-lived particles decaying. The efficiency is usually presented as a function of the long-lived particle decay position, as in Figure 5.3.

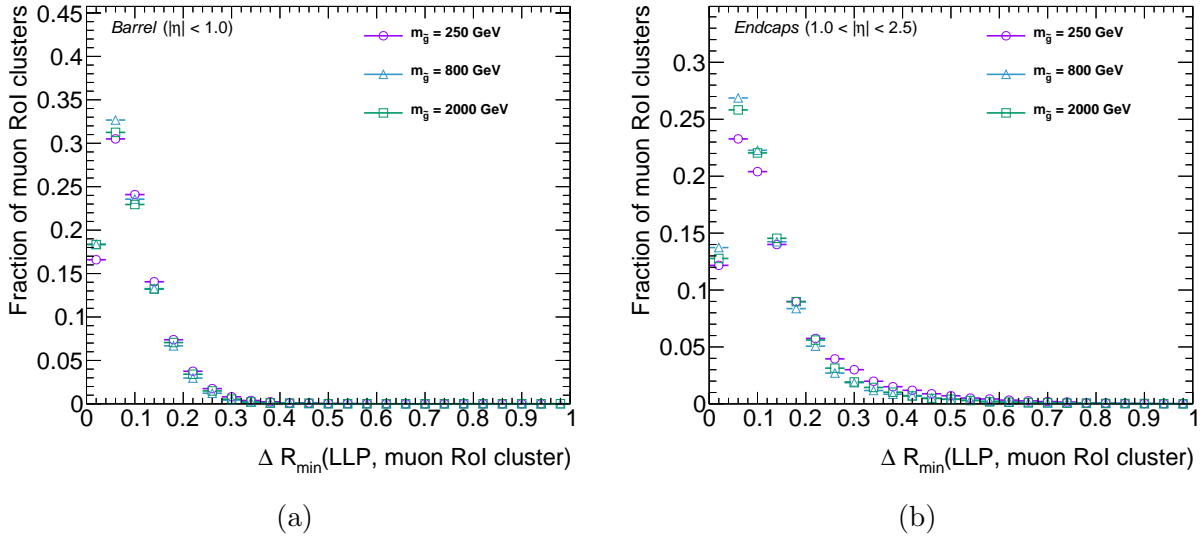


Figure 5.2: The minimum ΔR between the long-lived particle and a cluster of muon RoIs in the (a) barrel and (b) endcap.

A close analysis of Figures 5.3a and 5.3b shows how the detector geometry affects the trigger efficiency, for the lowest and highest mass Stealth SUSY samples. Note that the long-lived particle masses are the same in both cases, but as the gluino mass increases so does the LLP boost.

In the barrel, the efficiency starts turning on inside of the HCal. The higher energy LLPs can afford to pass through more of the calorimeters and still have enough energy to fire the trigger. To create a Level-1 MU10 RoI, there needs to be three out of four hits in the inner two RPC planes. This means that once a LLP decays past the first RPC plane it can no

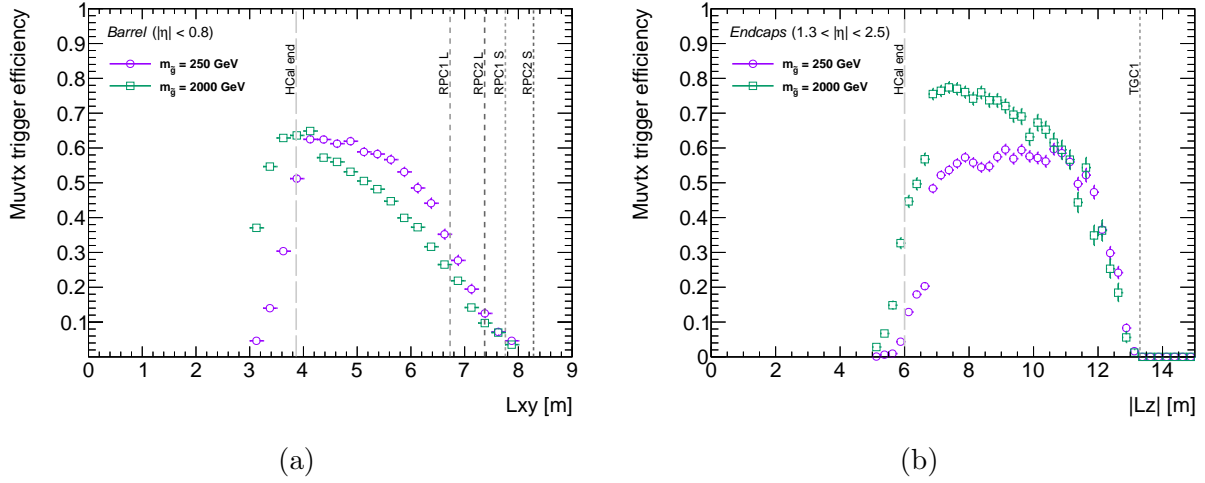


Figure 5.3: The efficiency for a single long-lived particle to fire the isolated Muon RoI Cluster trigger, HLT_j30_muvtx, for two extreme Stealth SUSY mass points. Dashed lines mark the relevant detector layers. Each RPC layer has alternating large (L) and small (S) sectors in ϕ , meaning the exact L_{xy} of each trigger plane depends on the ϕ location. Barrel details are shown in (a) and endcap details in (b).

longer fire the trigger. The RPC planes alternate small and large sectors in the barrel (see Figure 4.8b), with the small sectors being located over a meter beyond the large sectors. This causes the efficiency to rapidly decrease between RPC1 L and RPC1 S, but not drop to zero until the LLP decays past RPC1 S at 7.8 m. The efficiency decreases between the HCal end and RPC1 because the LLP decay products have less time to open up and cannot create enough distinct muon RoIs to form a good cluster.

A similar logic explains the structure of the efficiency in the endcaps. The efficiency increases from just before the HCal end until ~ 7 m. As the LLPs approach the first TGC plane, their decay products have less time to open up and create multiple RoIs, just as occurs in the barrel.

Non-collision background triggers

The Muon RoI Cluster trigger also selects events with significant MS activity that occur during empty or unpaired bunches. The triggers in empty bunch crossings allow for the study of detector noise, cosmic activity, or residual cavern backgrounds, while the trigger in unpaired bunches allows the study of beam halo. In 2015, these triggers were both seeded by a heavily prescaled L1_MU4 item, while in 2016 this was improved to L1_MU6.

Additionally, an overnight run with the full detector operational but no beams in the LHC was performed with the physics triggers running. This provides a large sample of cosmic activity and detector noise backgrounds.

Chapter 6

DATA AND MONTE CARLO SIMULATION SAMPLES

This chapter summaries the recorded datasets and simulated samples of events. All simulated samples have a centre of mass energy of 13 TeV to match the collision energy of the LHC in 2015 and 2016.

6.1 Monte Carlo simulation

Monte Carlo (MC) methods are used to simulate the manifestation of a given physics process in the ATLAS detector. This type of simulation relies on random numbers to provide an accurate sampling of final-state particles and their kinematic distributions. Both the signal process being searched for and standard model background events are generated in this fashion.

One important quantity to consider during event simulation is pileup. Pileup is generally discussed in terms of the quantity μ , a measure of the average number of inelastic interactions per bunch crossing, and $\langle\mu\rangle$, the average value over all proton bunches. The latter value, $\langle\mu\rangle$, gives the average number of expected proton-proton collisions per event. The distribution of $\langle\mu\rangle$ used in the simulation follows the predicted conditions for the 2016 LHC run, with enough statistics at lower $\langle\mu\rangle$ to accommodate for the inclusion of the dataset recorded in 2015.

In simulated events, pileup is added by *overlaying* a generated truth event with enough minimum-bias events to correctly model the desired number of interactions in the event. The number of overlaid minimum-bias events is varied to model how $\langle\mu\rangle$ varies throughout a given dataset. All the simulated events used in this thesis were simulated with the pileup profile shown in Figure 6.1. Events are further weighted at analysis time to ensure full agreement

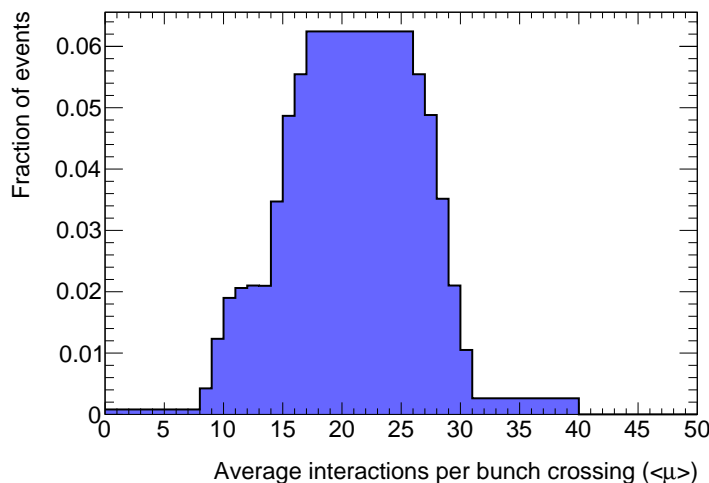


Figure 6.1: Distribution of the average number of interactions per bunch crossing in Monte Carlo simulation samples. This profile was generated to best represent the average interactions per bunch crossing expected in 2016, plus the existing 2015 data.

of the pileup distribution to that from the recorded dataset used in the analysis.

6.1.1 Signal samples

Monte Carlo simulated signal events are generated using `MadGraph5_aMC@NLO`, version 2.3.2 [32], with the proton parton distribution function `NNPDF2.3L0` [33] with QED corrections.

The `MADGRAPH` model used is not a full implementation of Stealth SUSY. Instead, a scalar S and fermion \tilde{S} are added to the default `mssm` model. The necessary vertices (interactions) are also added: $\tilde{g} \rightarrow g\tilde{S}$, $\tilde{S} \rightarrow S\tilde{g}$, and $S \rightarrow gg$. A working point with squarks decoupled is chosen such as not to overcomplicate the model.

Once the truth events are generated with `MADGRAPH`, `PYTHIA8` [34] with the `A14` [35] tune is used to generate showering and underlying event from proton fragmentation. Additionally, `EVTGEN` [36] is used after initial signal generation to ensure that the lifetimes and branching ratios of b - and c -hadrons are accurate and up-to-date.

The sample parameters are summarized in Table 6.1.

Table 6.1: Summary of the mass and lifetime parameters used in the Stealth SUSY signal sample generation. Every sample started with 400,000 events, but occasionally a job or two fail in one of the many simulation steps, causing the loss of a small fraction of events. Production cross-sections are 13 TeV gluino production cross-sections, with squarks decoupled [6].

$m_{\tilde{g}}$ [GeV]	$m_{\tilde{s}}$ [GeV]	m_S [GeV]	σ_{SUSY}^{prod} [pb]	$c\tau_{\tilde{s}}$ [m]	Events
250	100	90	1190.35	0.96	396,000
500	100	90	27.4171	0.76	391,000
800	100	90	1.4891	0.62	389,000
1200	100	90	0.0856481	0.50	396,000
1500	100	90	0.0141903	0.45	394,000
2000	100	90	0.000981077	0.37	390,000

For SM Higgs boson decays, a MADGRAPH implementation of the Hidden Abelian Higgs Model (HAHM) [15] was used to generate Higgs boson decays to scalars ($H \rightarrow ss$). The scalars are allowed to decay to muons; taus; and s -, c -, and b -quarks. The branching ratios to each fermion–antifermion pair are determined by the Yukawa coupling. Parameters for each of the generated mass points are summarized in Table 6.2. The mass of the Higgs boson is set to 125 GeV, and its production cross-section at 13 TeV is 48.48 pb, as computed at N3LO [37].

Table 6.2: Summary of the mass and lifetime parameters used in the $H \rightarrow ss$ signal sample generation.

m_s [GeV]	$c\tau_s$ [m]	Events
5	0.127	390,000
8	0.200	393,000
15	0.580	398,100
25	0.760	385,000
40	1.18	399,400
55	1.54	381,000

6.1.2 Background samples

Simulated QCD multi-jet events are used for data–MC comparison studies and to develop signal cuts for background rejection. These events are generated using `Pythia8` with PDF `NNPDF2.3L0` and tune `A14`, and `EvtGen`. The events are generated such that the number of events is uniform as a function of leading truth jet p_T , allowing for detailed MC studies of high-energy jets. The resulting 13 different samples referred to as *slices*, because they each provide a different range of leading truth jet p_T .


Table 6.3: Summary of the QCD multi-jet slice names, their leading jet (LJ) p_T range, and the number of events per slice.

Slice	LJ p_T range [GeV]	Events	Slice	LJ p_T range [GeV]	Events
JZ0W	0 – 20	1999400	JZ1W	20 – 60	1999000
JZ2W	60 – 160	1994600	JZ3W	160 – 400	7884500
JZ4W	400 – 800	7979800	JZ5W	800 – 1300	7977600
JZ6W	1300 – 1800	1893400	JZ7W	1800 – 2500	1770200
JZ8W	2500 – 3200	1743200	JZ9W	3200 – 3900	1813200
JZ10W	3900 – 4600	1996000	JZ11W	4600 – 5300	1993200
JZ12W	5300 – ∞	1974600			

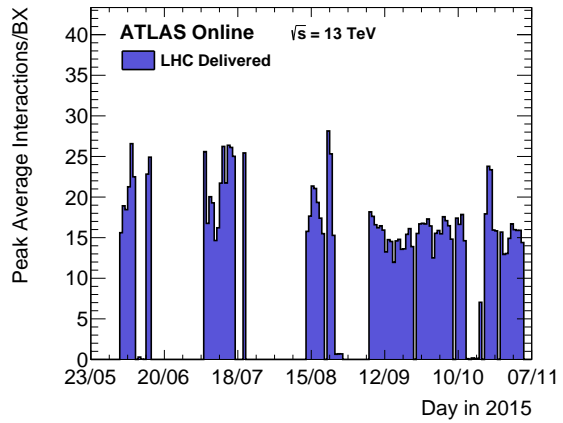
These slices are each labelled JZXW, where X corresponds to the slice number (JZ0W – JZ12W). Occasionally plots in the latter sections of this thesis will be labelled by a jet slice number. This provides a convenient mechanism for showing the difference between jets of different energies. The leading truth jet p_T ranges for each slice are shown in Table 6.3.

6.2 Data


This thesis uses a sample of data recorded in 2015 and 2016, up to 17 August 2016 (the start of machine development 2). The bunch spacing was constant at 25 ns, but the instantaneous luminosity and mean number of interactions per bunch crossing, $\langle\mu\rangle$, varied throughout data taking. Timelines for the amount of data collected and the peak number of interactions per bunch crossing are shown in Figure 6.2.




(a)



(b)



(c)



(d)

Figure 6.2: Integrated luminosity delivered and recorded by day in 2015 (a) and 2016 (c). The peak average interactions per bunch crossing, per day, are shown for 2015 (b) and 2016 (d). The 2015 dataset is both smaller and has a lower nominal pileup than 2016.

All recorded data is carefully analyzed to ensure that it meets data quality standards. These ensure that the detector was sufficiently operational, all calibrations were properly applied, alignment was correct, and there were no issues with the trigger or data acquisition systems. In addition, this analysis only uses data in which both the solenoid and toroids were turned on, since it requires the use of inner detector and muon spectrometer tracks.

After application of these criteria, the 2015 dataset totals 3.19 fb^{-1} , while the 2016 dataset totals 18.9 fb^{-1} . Thus, this analysis uses 22.1 fb^{-1} of recorded data. The systematic uncertainty on the luminosity, as determined in the luminosity calibration, is 2.9%.

Chapter 7

RECONSTRUCTING EVENTS

7.1 *Standard physics objects*

ATLAS has standard algorithms to reconstruct many types of physics objects, including photons, electrons, muons, taus, and jets from hadronic showers. The analysis presented in this thesis is a search for non-standard objects: displaced vertices from hadronic decays in the muon spectrometer. The algorithm for these vertices is described in detail in Section 7.2. Additionally, tracks from charged particles traveling through the inner detector and jets from hadronic showers are used in various contexts. The reconstruction of these objects is described in Sections 7.1.1 and 7.1.2, respectively.

7.1.1 *Inner detector track reconstruction*

Charged particles that travel through the inner detector pass through and potentially leave hits in the three inner detector subsystems: the pixel detector, SCT, and TRT. A cross-section of the inner detector is shown in Figure 4.3. A number of algorithms are used to reconstruct these hits as *tracks*. Track reconstruction starts by creating *space points* from hits in each of the pixel, SCT, and TRT layers.

Each pixel hit is individually a space point, since it provides a three-dimensional measurement. SCT modules are strips, and thus it takes a hit in a strip and its stereo strip partner to form a three-dimensional point. TRT hits are each individually space points, but since TRT straws are 37 mm long, aligned in z , they only provide fine resolution in $(R \times \phi)$ (in the endcaps, the straws are aligned vertically and provide good resolution in $(z \times \phi)$).

Tracks are formed by first looking for three-hit seeds (space point triplets). These can

either be three pixel hits (PPP), three SCT space points (SSS), or a combination: PPS or PSS. In Run-1, track finding proceeded with the set of three-hit seeds, but for Run-2 an additional step is taken [38]. Only three-hit seeds that are matched to one additional space-point to make a four-hit seed are used for track finding. Then, only seeds that are in approximate agreement with originating from the beam-line are used for track reconstruction.

The resulting four-hit seeds feed into an inside-out track-finding algorithm, which uses the seed for an initial trajectory. A Kalman fitter [39] is used to determine if hits along this trajectory form a real track. After pixel and SCT hits are considered, the track is matched to TRT-segments; if there are no matching TRT hits, the track is still kept.

Further restrictions on tracks are made before the track is deemed valid: offline cuts are made to ensure that only well-reconstructed tracks are used. For this analysis, reconstructed tracks must have at least 7 silicon hits (out of a possible 9), no more than 1 pixel *hole* or 2 total silicon holes, and no more than 1 shared space-point. Further kinematic and spatial cuts are made on the track momentum and impact parameters: track $p_T > 400$ MeV, $d_0 < 10$ mm, and $|z_0| < 320$ mm.

7.1.2 Jet reconstruction

When a quark or gluon is produced in a collision, it immediately begins to hadronize. Depending on the energy of the quark or gluon, the hadronization will produce many more quarks and gluons, resulting in the energy of the original particle being distributed into a shower of particles, mostly contained in some cone around the particle's direction (*collinear* with the original particle). A sketch of this process is shown in Figure 7.1.

The shower of particles resulting from a hadronic decay is generally referred to as a jet. Using the energy collected by the electromagnetic and hadronic calorimeters, this jet can be reconstructed and used to determine energy and position measurements. The first step in reconstructing a jet is forming topological clusters from nearby cells with energy deposition [40].

Topological clusters (*topo-clusters*) are formed by combining nearby calorimeter cells that




Figure 7.1: Formation of two jets from a $q\bar{q}$ pair. Straight lines indicate quarks, while the curly lines indicate gluons.

are above a pre-defined signal-to-noise (S/N) threshold. Each topo-cluster is seeded by a cell with $S/N > 4$, and any neighboring cells with $S/N > 2$ are added. The last step is to add in all cells neighboring the cluster, regardless of their S/N . The topo-clusters in this analysis are calibrated at the electromagnetic scale [40], meaning that their energy is calculated assuming it was all caused by electromagnetic showers.

There are a number of algorithms used to reconstruct a jets from topo-clusters. The analysis described in this thesis uses jets with a radius $R = 0.4$ reconstructed with the anti- k_t algorithm [30]. The anti- k_t algorithm starts by computing the “distance” between all topo-clusters i, j using the metric:

$$d_{i,j} = \min \left(\frac{1}{p_{T_i}^2}, \frac{1}{p_{T_j}^2} \right) \frac{\Delta R_{i,j}^2}{R}, \quad (7.1)$$

where $p_{T_{i,j}}$ are the transverse momenta of the two topo-clusters, and ΔR is the angular distance between them. The distance $d_{i,j}$ is compared against a boundary value, $d_{i,B} = 1/p_{T_i}^2$, which is the distance between the i -th topo-cluster the beam. The two topo-clusters with the smallest distance $d_{i,j}$ are combined if $d_{i,j} < d_{i,B}$. This procedure is repeated until no

distances are smaller than $d_{i,B}$, at which point a jet has been formed.

The anti- k_t algorithm produces very circular jets, compared to other flavors of sequential recombination. This is because of the $\min\left(\frac{1}{p_{T_i}^2}, \frac{1}{p_{T_j}^2}\right)$ component, which ensures that topo-clusters with very low- p_T are combined with topo-clusters with high- p_T before low- p_T topo-clusters are combined with each other.

The energy of the jets is calibrated to accurately match the energy of the original decaying particle [40]. This is done to compensate for energy not measured by the calorimeters. Causes for this include energy absorbed in the inner detector or by non-measuring detector material. Also, the ATLAS calorimeter is non-compensating, which means that the energy needs to be corrected based on the expected amount of energy deposited as electromagnetic or hadronic energy. The net response of all these effects is folded into a single jet energy scale (JES), which ensures that jets have a uniform response across all regions of the calorimeter.

Once jets are reconstructed, a set of cleaning criteria are applied to ensure the jets originating from true hadronic activity from the hard scatter are well-measured. ATLAS has different flavors of jet cleaning, including specialized criteria for the decays of long-lived particles. However, the jets used in this analysis are primarily from prompt activity, or promptly-decaying long-lived particles, so the standard cleaning criteria are used. The cleaning criteria are designed to reject jets caused by beam-induced backgrounds (“beam halo”), cosmic showers, and calorimeter noise. There are also specific criteria to identify and reject jets from pileup interactions. The variables used are as follows [41]:

- LAr calorimeter cell energy deposit quality, $Q_{\text{cell}}^{\text{LAr}}$. This is a measure of how well the energy deposit pulse shape agrees with the response from a signal-like pulse, as opposed to electronic noise. It is used for three separate quality criteria, in the LAr electromagnetic calorimeter and the LAr hadronic endcaps (HEC). In each calorimeter, the fraction of energy in the jet, f_Q , that comes from cells with noise-like $Q_{\text{cell}}^{\text{LAr}}$ is determined (f_Q^{LAr} and f_Q^{HEC} in the individual calorimeters). Additionally, the unity normalized average quality of all LAr cells in the jet, $\langle Q \rangle$, is used as a measure of

overall LAr signal quality.

- The electromagnetic energy fraction, f_{EM} . This is the fraction of the jet’s overall energy deposited within the electromagnetic calorimeters.
- The hadronic energy fraction, f_{HEC} is the fraction of the jet’s overall energy deposited in the LAr hadronic endcaps.
- The maximum energy deposited in any single calorimeter layer, f_{max} .
- The energy sum of all the cells in a jet with negative energy deposits, E_{neg} . Negative energy is caused by noisy calorimeter cells that fluctuate negatively instead of positively.
- The fraction of charged hadrons in the jet, as measured by reconstructed tracks in the inner detector, $f_{\text{ch}} = \sum(p_{\text{Ttracks}})/p_{\text{Tjet}}$. The numerator $\sum(p_{\text{Ttracks}})$ is the scalar p_{T} sum of all tracks associated to both the primary vertex and the jet, with $p_{\text{T}} > 500$ MeV, and p_{Tjet} is simply the measured p_{T} of the jet.

Additionally, a jet-vertex association variable, JVT (Jet Vertex Tagger), is computed for each jet with $p_{\text{T}} < 60$ GeV [42]. This provides a measure of how likely the jet is to have come from a given primary vertex, and is used to distinguish between jets from pileup collisions and jets from the hard scatter. It consists of two terms. The first is the jet vertex fraction (JVF), which provides a measure of the fractional p_{T} of tracks associated to the jet that come from the primary vertex. The JVF also has a correction factor that accounts for the increasing number of “fake-associated” tracks with increasing number of pileup vertices. The second term is the aforementioned variable f_{ch} (referred to as $R_{p_{\text{T}}}$ in the context of JVT). Details on how these two terms are combined are available in Reference [42].

A jet is considered bad and not used for analysis if it meets any of the following criteria:

1. $f_{\text{EM}} < 0.05$ and $f_{\text{ch}} < 0.05$ and $|\eta| < 2.0$

2. $f_{\text{EM}} < 0.05$ and $|\eta| \geq 2.0$
3. $f_{\text{max}} > 0.99$ and $|\eta| < 2.0$
4. $E_{\text{neg}} > 60$ GeV
5. $f_{\text{HEC}} > 0.5$ and $|f_Q^{\text{HEC}}|$ and $\langle Q \rangle > 0.8$
6. $f_{\text{EM}} > 0.95$ and $|f_Q^{\text{LAr}}|$ and $\langle Q \rangle > 0.8$
7. $|\eta| < 2.4$ and $JVT < 0.59$ and $p_T < 60$ GeV

Additionally, if *any* jet is rejected with the first six criteria **and** passes the seventh (the JVT cut), the event as a whole is rejected. If only jets above 60 GeV are being required in the event, then the event is only rejected if there is a jet with $p_T > 60$ GeV that fails the any of the selection criteria. The event-based jet cleaning is only applied when jets are required in an event.

7.1.3 Composite variables

Sometimes it is useful to be able to describe the overall properties of an event rather than just the individual objects. For example, in events with many jets, it is useful to describe the total transverse energy deposited in the calorimeters. However, this scalar value would not give any indication of the amount of energy escaping the detector or the imbalance of the jet energy, so additional variables are also computed.

The total jet energy in the event is usually discussed in terms of H_T , the scalar sum of all transverse jet energies:

$$H_T = \sum_{\text{all jets}} |\vec{E}_T|. \quad (7.2)$$

The vector sum can also be determined, and should equal approximately zero unless there is a substantial component of missing jet energy, denoted H_T^{miss} . Thus,

$$H_T^{miss} = - \sum_{all\ jets} \vec{E}_T. \quad (7.3)$$

Combining both the total jet energy and the missing jet energy it is possible to define the effective mass of an event, m_{eff} , where

$$m_{eff} = H_T + |H_T^{miss}|. \quad (7.4)$$




Figure 7.2: Transverse depiction of three event topologies with the same H_T but differing m_{eff} and H_T^{miss} .

Examples of event topologies with $H_T = 200$ GeV but different H_T^{miss} and m_{eff} are shown in Figure 7.2.

7.2 Muon Spectrometer vertices

The description of MS vertex reconstruction relies heavily on the detector description (see Section 4.3). Due to high noise rates and specialized configurations (*e.g.* only a single multilayer per chamber), not all MDT chambers are used for vertex reconstruction. The following chambers are skipped: BME, BEE, BIL7, BIL8, EEL, and EES. The remaining chambers all contain two multilayers, and each multilayer is made up of either three or four layers of MDT tubes.

Because some MDT chambers in the barrel are contained within the magnetic field, while the endcap chambers are not (see Figure 4.6), there are two separate MS Vertex reconstruction algorithms: one for each the barrel and the endcaps [43]. Both the algorithms follow the same general procedure:

1. MDT hits within a chamber multilayer are formed into 3- or 4-hit segments
2. Segments in the two multilayers of a chamber are joined to form tracklets
3. Clusters of tracklets are identified
4. Tracklets are back-extrapolated
5. Locations where multiple tracklets cross are identified as vertices

7.2.1 Multi-layer segment reconstruction

Multi-layer segment reconstruction begins by selecting 3-hit seeds. First, good ($\text{ADC} \geq 50$, drift radius is a physical value, and error > 0.001) MDT hits are ordered sequentially in tube number, along each layer, such that tube 1 is the leftmost tube from the bottom layer (layer 1), and the highest-numbered tube is the rightmost tube in the top layer. Figure 7.3 shows how the numbering works for an imaginary multi-layer with three layers of five tubes.




Figure 7.3: The numbering scheme for MDT tubes within a given multi-layer.

Segment seeds are formed independently in each multilayer of each chamber. All combinations of seeds meeting the following criteria are initially formed:

1. $N_{MDT3} > N_{MDT2} > N_{MDT1}$, where N_{MDTi} is the number of the i -th MDT tube following the scheme outlined in Figure 7.3.
2. $|z_{MDT1} - z_{MDT2}| < 100$ mm ($|R_{MDT1} - R_{MDT2}| < 100$ mm in endcaps)
3. $|z_{MDT1} - z_{MDT3}| < 160$ mm ($|R_{MDT1} - R_{MDT3}| < 160$ mm in endcaps)
4. Hits 1 and 2 must be in the same layers or sequential layers, and hit 3 must be in the same layer or one higher than hit 2, and not in the same layer as hit 1 (i.e. seed is formed from hits in at least two layers, with no skipped layer between hits)

Using the innermost and outermost hits of the segment seed, four possible fit lines are formed. These hits are assumed to always be hits 1 and 3. This is not *necessarily* always the case: see Appendix B for details. The four possible fit lines are shown in Figure 7.4a. The slope m and intercept b of each fit line are kept for future segment fitting if the residual of the middle MDT hit is less than 5 mm, where the residual for hit i is calculated as:

$$\text{Residual}_i = \left| \frac{r_{drift,i} - |(R_i - b - m \times z_i)/\sqrt{m^2 + 1}|}{\sigma_{drift,i}} \right|. \quad (7.5)$$

This equates to the shortest distance between the seed fit line and the outside of the drift circle for hit i , divided by the uncertainty on the drift radius for hit i ($\sigma_{drift,i}$).

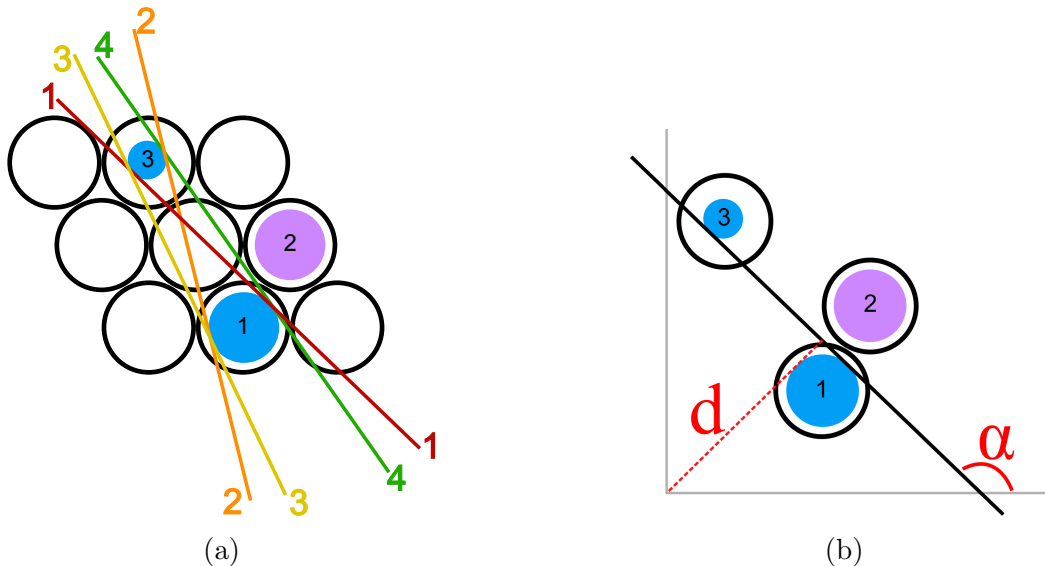


Figure 7.4: (a) The four possible fit lines for a three-hit segment seed, and (b) the parameters of the fit line used to perform a χ^2 -minimization fit.

Tracklet segments are then created using the following method for each seed fit line:

1. Find the average weighted position in R, z of all MDT hits contributing to the seed.
2. Ensure the seed fit line points to the second multilayer in the chamber (in the barrel, the angle from vertical must be ± 15 degrees, while in the endcaps the angle from horizontal must be $\pm 1.7\%$).
3. Determine if the 3-hit seed can actually form a coherent segment: a χ^2 fit is performed on the seed fit line angle, α and perpendicular shift from the origin, d (see Figure 7.4b). The exact fit procedure is outlined in Ref. [44].
4. Only store the segment if the χ^2 probability of fit is greater than 0.05.

5. Find the hit pattern for the segment, denoting whether the tubes are on the left or right side of the segment.
6. Calculate the 3D point segment position: the point of closest approach to the weighted centre of the MDT hits contributing to the segment. The ϕ coordinate of the segment (used to calculate x and y from R , since the fit is done in the $R - z$ frame) is the ϕ position of the chamber.
7. Ensure that the segment is actually unique, compared to to the previous segments. Non-unique segments are possible, since there are four seed fit lines for each set of 3-hit seeds. After minimisation, it is possible that the seed fit lines all return the exact same segment.

After the initial pass of segment finding, the resulting segments are cleaned to ensure that any overlapping segments are combined to form a single, better fit segment with all the involved hits:

1. Find all sets of segments that are in the same multilayer of the same chamber whose trajectories α differ by no more than 0.005 radians, and distance of the midpoints is less than 0.5 mm
2. If the set of unique hits on all the segments is larger than the initial segment, refit the tracklet segment using the previously outlined procedure.
3. If the refit fails and there are two segments involved, keep both segments as unique.
4. If the refit fails and there are more than two segments, remove the MDT hit used by the fewest involved segments and attempt to refit again. Iterate the remove hit–refit segment process until a good segment is found. If no new, combined segment is found and there are only three hits remaining, keep the original segment as clean.

5. Iterate ten times through the container, each time running the output of the previous cleaning as the input.

7.2.2 Tracklet reconstruction

The clean segments are then ready to use for tracklet formation. Chamber by chamber, the segments in each multilayer are collected. For each combination of ML1 and ML2 segments, the differences in angle $\Delta\alpha$ and parameter Δb are calculated. If both are below the acceptable threshold for the given chamber, the combination proceeds to tracklet fitting.

In the barrel, the momentum of a tracklet is calculated using the angle between multilayer segments, $\Delta\alpha$. The magnetic field is not the same in each chamber, nor is the spacing between multilayers, thus the angle between multilayer segments will not be the same for a given momentum. Thus, a set of constants k_i are calculated for barrel chambers. The momentum of a tracklet with angle $\Delta\alpha$ in the i -th chamber is then $p = k_i/\Delta\alpha$.

If the tracklet has a momentum of less than 800 MeV, the segment combination is not considered. If the tracklet has a measurable momentum, between 800 MeV and 10 GeV, there is no fitting of the segments, and the tracklet is stored as the two segment and their angular parameters.

If the chamber is in the endcaps, or the momentum is greater than 10 GeV, the tracklet is assumed straight, and is fit by collecting the MDT hits from both segments, and running the hit collection through the tracklet segment fitting algorithm. In this case the seed fit lines are formed with the innermost hit from ML1 and the outermost hit from ML2, with all remaining hits (at least four) in between.

7.2.3 MS vertex reconstruction

The presence of a magnetic field inside the barrel MDT chambers but not the endcap MDT chambers (recall the endcap toroid is located between the inner and middle layers, while the barrel toroid encompasses the inner two layers of the barrel muon system) necessitates different reconstruction algorithms in the two regions. Thus, tracklets in the barrel and

endcaps are treated completely separately, and vertices are constructed only with barrel tracklets or only with endcap tracklets. It should be noted that this does create an issue in the overlap region between the barrel and endcaps. This issue is discussed in Section 8.1.1.

Each vertex is required to be made from at least three tracklets, so in the event there are fewer than three total tracklets, vertex reconstruction is not run. Additionally, if there are more than 40 tracklets in the event, the vertex reconstruction will take too long to run and instead a dummy vertex is created at $(x, y, z) = (-9.99, -9.99, -9.99)$. Note that this is not an ideal position, since in η and ϕ this looks like a real vertex location, so care should always be taken to ensure these dummy vertices are not included in performance plots.

In general, events in data do not exceed the 40 tracklet threshold, although signal models with higher mass LLPs or mediators can produce far more than 40 tracklets, so this cut does impact the signal reconstruction efficiency of the algorithm (see Section B).

Details on the performance and implementation of this algorithm are detailed in Ref. [43]. The algorithm itself is documented in detail in the following sections.

Tracklet clustering

In both the barrel and endcaps, the first step in reconstruction is to identify clusters of tracklets. The clustering follows the same algorithm as for the Muon RoI Cluster trigger outlined in Section 5.4.2, though with different parameters. The tracklets are iterated through twice: once as cluster seeds, and a second for cluster formation.

For each seed, clusters are formed by adding any tracklets within $\Delta\eta = 0.7$ and $\Delta\phi = \pi/3$ of the cluster centre. The cluster center is recalculated every time a new tracklet is added. Once all tracklets are iterated through, the centre is recalculated as the average of all tracklets in the cluster. If a cluster with more tracklets than the previous iteration is found, it is saved and clustering is rerun. The clustering is reiterated up to five times, each time with the input seed location being the output centre of the previous cluster iteration.

After this procedure is complete, the cluster with the most tracklets is selected, and the unused tracklets are stored for reclustering. The clustering is rerun until there are less than

three tracklets remaining, or a the largest cluster found has fewer than three tracklets.

It should be noted that this clustering procedure is neither collinear safe nor infrared safe, which could lead to undesirable results in certain situations.

Barrel MS vertex reconstruction

Vertex reconstruction in the barrel MS takes into account the bend of tracklets due to the magnetic field. Because MDT tubes are aligned along ϕ , the ϕ coordinate of each tracklet is simply the centre of the MDT chamber it was reconstructed in. Thus, vertex reconstruction is performed in only two dimensions, in the $R - z$ plane. Note that vertices are not reconstructed from every tracklet in the muon system, but only those within a given cluster, so this procedure does not ever place tracklets from opposite sides of the detector in the same plane.

The vertex-finding proceeds as follows:

1. Calculate the average position of all tracklets in the cluster.
2. Refine the average ϕ position by calculating the average ϕ position of all RPC and TGC hits within $\Delta R = 0.6$ of the tracklet cluster centre.
3. Using the average position, calculate the line of flight (angle θ) of a particle that would have left such a vertex, in the $R - z$ plane.
4. Rotate all tracklets in ϕ , such that they are in the centre of the cluster.
5. Calculate a series of planes to search for a vertex along, starting at a radius of 3500 mm. These planes are intended to be spaced evenly in 200 mm increments along the line of flight, though the implementation of the algorithm used for this analysis miscalculated the number of planes and thus they are slightly farther apart. The exact distance varies from 208 mm at $|\eta| = 1$ to 220 mm at $\eta = 0$. See Section B for more details.

6. At each radial plane, the following steps are taken:
 - (a) Find tracklets whose radial position is beyond the plane location. Note that the tracklet radial position is calculated as centre of the chamber in which the tracklet was reconstructed.
 - (b) If there are at least three tracklets accessible by the plane, extrapolate the tracklets inwards to the point they would have crossed, given their charge and momentum. Straight-line tracklets (those with momentum > 10 GeV) are again treated as straight lines here.
 - (c) Calculate the longitudinal weighted position (in z) of the tracklets where they cross the plane.
 - (d) Perform a χ^2 minimisation to determine the best vertex: First, find the χ^2 contribution from each tracklet and the χ^2 probability of the vertex fit. If the probability of fit is < 0.05 , remove the tracklet with the largest χ^2 , and determine if the vertex χ^2 has improved. Continue this procedure until there is a vertex found with χ^2 fit probability > 0.05 or there are less than three tracklets remaining and thus a vertex can no longer be found.
7. Select the vertex that was found with the largest number of tracklets. If there is more than one vertex with equal numbers of tracklets, select the vertex with the largest χ^2 probability of fit.
8. Calculate the number of MDT hits within chambers whose centre is within $\Delta R = 0.6$ of the vertex centre, and the number of these MDT chambers with at least 25% occupancy.
9. Calculate the number of RPC + TGC hits that are within $\Delta R = 0.6$ of the vertex centre.
10. The final selection criteria for a vertex is a χ^2 probability > 0.05 , at least 250 MDT

hits and 200 RPC+TGC hits, and at least 2 chambers with hits in 25% of tubes (high occupancy chambers).

Endcap MS vertex reconstruction

In the endcaps, all tracklets are straight-line because the toroid is located between the inner and middle MDT layers, and thus the chambers themselves are not immersed in a magnetic field. It is likely that a long-lived particle that decays before or within the toroid will decay to charged particles that will be bent by the toroid. However, since these particles will leave straight-line tracklets in the MDT chambers themselves, there is a lack of charge and momentum information. Thus, the tracklets cannot be propagated as charged particles through the toroid, and are instead all treated as straight-line (infinite momentum) tracklets.

There are two separate straight-line vertex finding methods employed in this analysis. One is more precise, but takes longer and the CPU time suffers with large combinatorics. The other algorithm is faster, but does not take into account all combinatorial possibilities.

Before either algorithm runs, a list of all three-tracklet combinations with vertex-forming potential is compiled. Good three-tracklet combinations are those that combine with a least-squares fit to a line in slope-intercept space to form a vertex with $R < 10$ m and $7 < z < 15$ m, with no single tracklet's distance of closest approach to the vertex greater than 300 mm. These three criteria are referred to as the *preliminary vertex criteria*.

The slower algorithm runs when there are ≤ 20 tracklets in the endcap tracklet cluster:

1. For each of the three-tracklet combinations, attempt to add each of the tracklets that was not part of the original three-tracklet vertex.
2. For each additional tracklet, if the resulting vertex meets the preliminary vertex criteria, the vertex is added to the list of three-tracklet vertices.

If there are more than 20 tracklets in the endcap cluster, the following algorithm runs:

1. Find the average position of all the tracklets, which would correspond to the line-of-flight of a decaying long-lived particle.
2. Perform a least-squares fit in slope-intercept space to determine the best-fit vertex position in R and z .
3. Remove the tracklet with the farthest distance of closest approach to the vertex, if that distance is > 300 mm.
4. Iterate through this fit vertex - remove tracklets procedure until all tracklets are within 300 mm of the vertex position, and thus no tracklets are removed for the next iteration.
5. If the final vertex fit had less than three tracklets, exit.
6. The average ϕ position is refined by calculating the average ϕ position of all RPC and TGC hits within $\Delta R = 0.6$ of the average tracklet cluster centre calculated at the start of the algorithm.
7. Create a vertex with the R, z coordinates from the vertex fit and the ϕ position from the previous step.
8. Calculate the number of MDT hits within chambers whose centre is within $\Delta R = 0.6$ of the vertex centre, and the number of these MDT chambers with at least 25% occupancy.
9. Calculate the number of RPC + TGC hits that are within $\Delta R = 0.6$ of the vertex centre.
10. The final selection criteria for a vertex is a χ^2 probability > 0.05 , at least 250 MDT hits and 200 RPC+TGC hits, and at least 2 high occupancy chambers.

If no better vertices are found other than the original three-tracklet combinations, the first vertex in the list is chosen as the final vertex. Otherwise, if the slow algorithm runs

and finds a vertex, the first vertex with the largest number of tracklets is chosen. As with the other reconstruction algorithms, the MDT and trigger hits are analyzed, and the final selection criteria of χ^2 probability > 0.05 , at least 250 MDT hits and 200 RPC+TGC hits, at least 2 high occupancy chamber are applied.

Chapter 8

IDENTIFYING DISPLACED DECAYS

This thesis describes the search for events characteristic of one flavor of Stealth SUSY, as described in Section 2.3.2, and Standard Model Higgs boson decays to long-lived scalars, as described in Section 2.3. Two separate types of events are studied and analyzed. The first is a simple search for two displaced MS vertices, in events that pass the Muon RoI Cluster trigger. This search is very similar to that already performed with the 8 TeV dataset[45]. Although only plots for Stealth SUSY are shown in this section, care is taken to ensure that selection criteria are not overly tuned to the Stealth SUSY model. Since Stealth SUSY is only one of many possible models, it would be foolish to over restrict the search.

The second search takes advantage of a property particular to the Stealth SUSY model: the presence of two prompt gluons in addition to the long-lived particles (see Figure 2.4). The two prompt gluons are often produced with enough energy to leave a clear two-jet signature in the detector, in each event. Events must still pass the Muon RoI Cluster trigger, but then only a single vertex is required. This provides sensitivity to a much broader range of singlino proper lifetimes. A range of searches with this topology are discussed in Ref. [46].

Both search strategies will be discussed in detail in later chapters. In this chapter, the identification of displaced MS Vertices is discussed. Additionally, systematic uncertainties common to both searches are outlined.

8.1 Identifying muon spectrometer vertices

The muon spectrometer (MS) vertex reconstruction algorithm described in Section 7.2 reconstructs a signature consistent with a hadronic decay in the MS, but it does not require that the vertex is isolated or characteristic of a displaced decay. In order to ensure sufficient signal

acceptance and background rejection, a set of *good vertex criteria* (GVC) are established to assist in determining whether or not a vertex is consistent with a displaced hadronic decay.

The main goal is to distinguish true displaced vertices from those caused by jet activity, so the background sample considered here is MC QCD dijet events, as described in Section 6.1.2.

8.1.1 MS vertices in the barrel-endcap overlap region

Before the GVC are calculated, an η requirement that removes vertices from the overlap region between barrel and endcap chambers is made. Because vertex reconstruction is performed separately in the barrel and endcaps, hits from a single decay are split between the two algorithms, and vertices in these regions are thus constructed using fewer hits. This results in a dip in the reconstruction efficiency, but also occasionally results in two vertices being reconstructed from a single LLP decay.

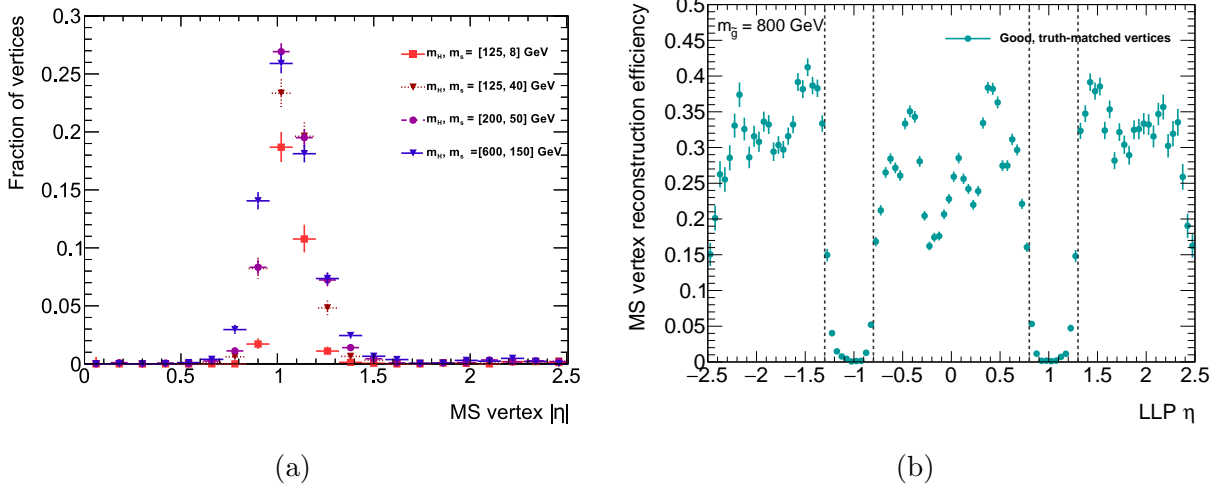


Figure 8.1: (a) Fraction of vertices coming from long-lived particles that match *two* otherwise good MS vertices. The simulated MC samples shown are heavy scalar decays (H) to lighter scalar pairs (s). This behavior also visible in Stealth SUSY decays. (b) The efficiency to reconstruct a MS vertex as function of η , for the Stealth SUSY $m_{\tilde{g}} = 800$ GeV sample. Dashed lines mark the excluded region. Note that the efficiency between each pair of lines is low relative to the rest of the detector, and thus removing the overlap region does not significantly harm the overall acceptance.

Figure 8.1a shows the fraction of LLP decays with two vertices within $\Delta R = 0.4$ of the decay axis. For the reconstruction efficiency in Figure 8.1b, if a LLP is associated to two vertices, only one is considered such that the two vertices are not double counted. The vertex reconstruction efficiency is relatively low in the overlap region. Thus, to eliminate the uncertainty caused by reconstructing two vertices (and the uncertainty over which region the vertex was reconstructed in), the region between $0.8 < |\eta| < 1.3$ is removed. In the future, the algorithm should be restructured to reconstruct vertices from both barrel and endcap chambers in the overlap region.

8.1.2 Good vertex criteria

Number of associated hits

The first good vertex quality criteria are based on the number of hits associated to a vertex. These cuts help eliminate background from both electronic noise bursts *and* punch-through jets. In general, a noise burst in one MS sub-detector will not be coherent with one in a different MS sub-detector. Therefore, a minimum number of hits is required for both MDTs and the relevant trigger system (RPCs in the barrel, TGCs in the endcaps). There are also minimum hit criteria applied at the time of reconstruction: $N_{\text{MDT}} > 250$ and $N_{\text{RPC+TGC}} > 200$. Vertices with less than these numbers of hits are not recorded, so plots showing the number of associated hits begin at 250 for MDTs and 200 for the trigger systems.

Additionally, in noise burst events, there are many MDT hits in a localized area in the MS. The sheer combinatorics from the number of hits provides enough combinations such that an MS vertex may be reconstructed. In order to protect against these events, a maximum of 3000 MDT hits are allowed in the vertex cone. The upper limit of 3000 MDT hits is not compared against a background, but simply selected high enough such that much of the signal falls below the cut. Unfortunately, for the highest mass gluino samples it is clear that

this cut is non-optimal and should be revisited for future searches¹

Acceptance plots for the number of hits in the vertex cone are shown in Figure 8.2. The minimum hit criteria were established with Run-1 data, and only vertices with at least 250 MDT hits and 200 trigger (RPC+TGC) hits are recorded to save computational time and space, causing the corresponding hard cuts in Figure 8.2.

Isolation criteria

In true collision data, many of the reconstructed vertices will be caused by high-energy jets punching through the calorimeter and depositing additional energy in the MS. However, such vertices will generally be found within a small angular distance of reconstructed calorimeter jets and inner detector tracks. Three separate criteria are established to ensure a maximal rejection of vertices from punch-through jets. The final values chosen are summarized in Table 8.1. The vertices considered for establishing these criteria must pass the η and N_{hits} cuts established in the previous sections.

The first is isolation from calorimeter jets. In order to not reject true vertices from displaced decays that occur near the end of the calorimeters, jets are only considered if $\log_{10}(E_{\text{HAD}}/E_{\text{EM}}) < 0.5$. Additionally, a minimum jet energy of 30 GeV is required, which prevents pileup or underlying event from spoiling a true displaced vertex. The maximum ΔR cone around the vertex for which there is no jet satisfying the aforementioned criteria is then calculated. This ΔR cone is referred to as an *isolation cone*. Figures 8.3a, 8.3b, 8.4a and 8.4a show how the signal and background acceptances change with varying isolation cone size and minimum jet E_{T} , in both the barrel and the endcaps.

Isolation from high-energy tracks and significant track activity is also considered. If there is significant prompt activity but poor calorimeter measurement (for instance, in the crack regions of the calorimeters), there is the chance that a sufficiently high quality jet will not be

1

This will involve a careful study of events with electronic noise in the muon spectrometer, and perhaps developing additional criteria to remove these events in a different way.

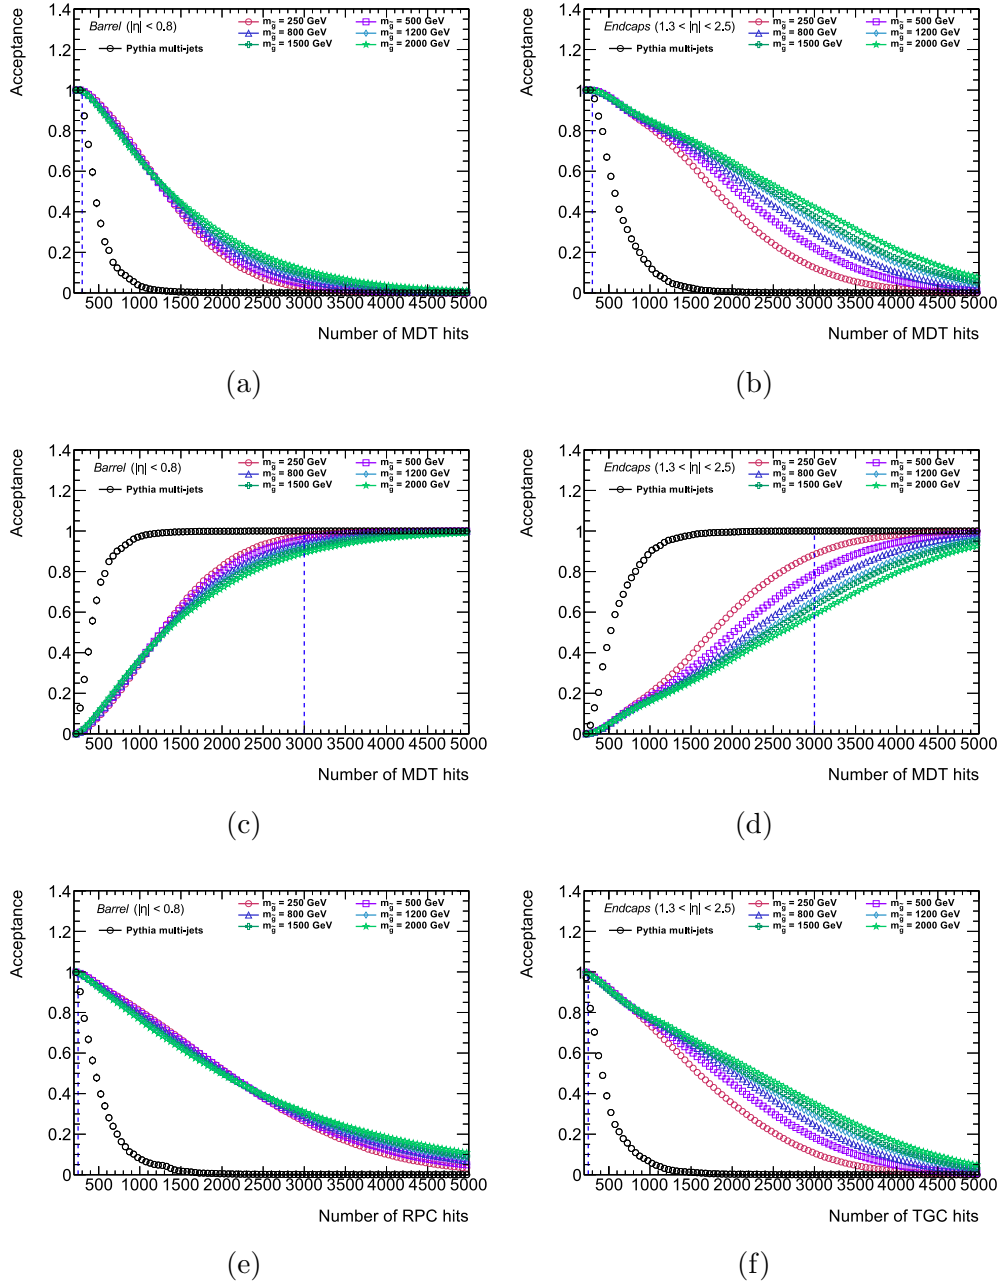


Figure 8.2: Acceptance for selecting barrel and endcap vertices passing hit-based good vertex criteria. Simulated QCD multi-jet events are compared against signal vertices to determine selection criteria that will remove most or all vertices from background QCD events and keep most signal vertices. In each case, dashed lines show the chosen cut value. Figures (a), (b), (e), and (d) select vertices with the number of hits *above* the cut line, while Figures (c) and (d) describe and upper cut, and vertices with the number of hits *below* the cut line are selected. The cut in Figures (c) and (d) is designed to reject electronic detector hits, and not QCD multi-jet events.

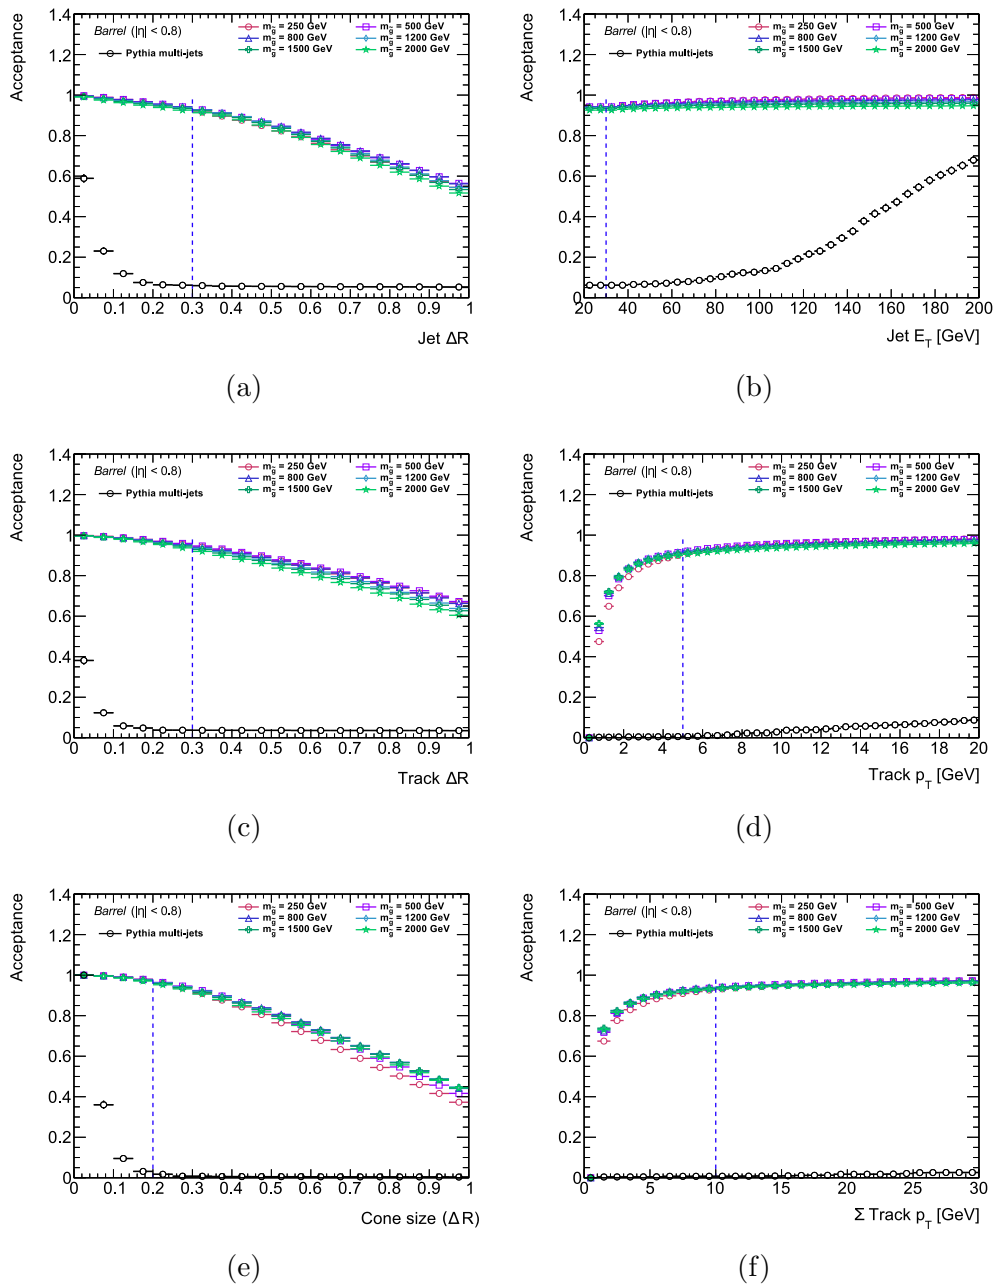


Figure 8.3: Acceptance for selecting barrel vertices passing good vertex criteria. Simulated QCD multi-jet events are compared against signal vertices to determine selection criteria that will remove most or all vertices from background QCD events and keep most signal vertices. In each case, dashed lines show the chosen cut value.

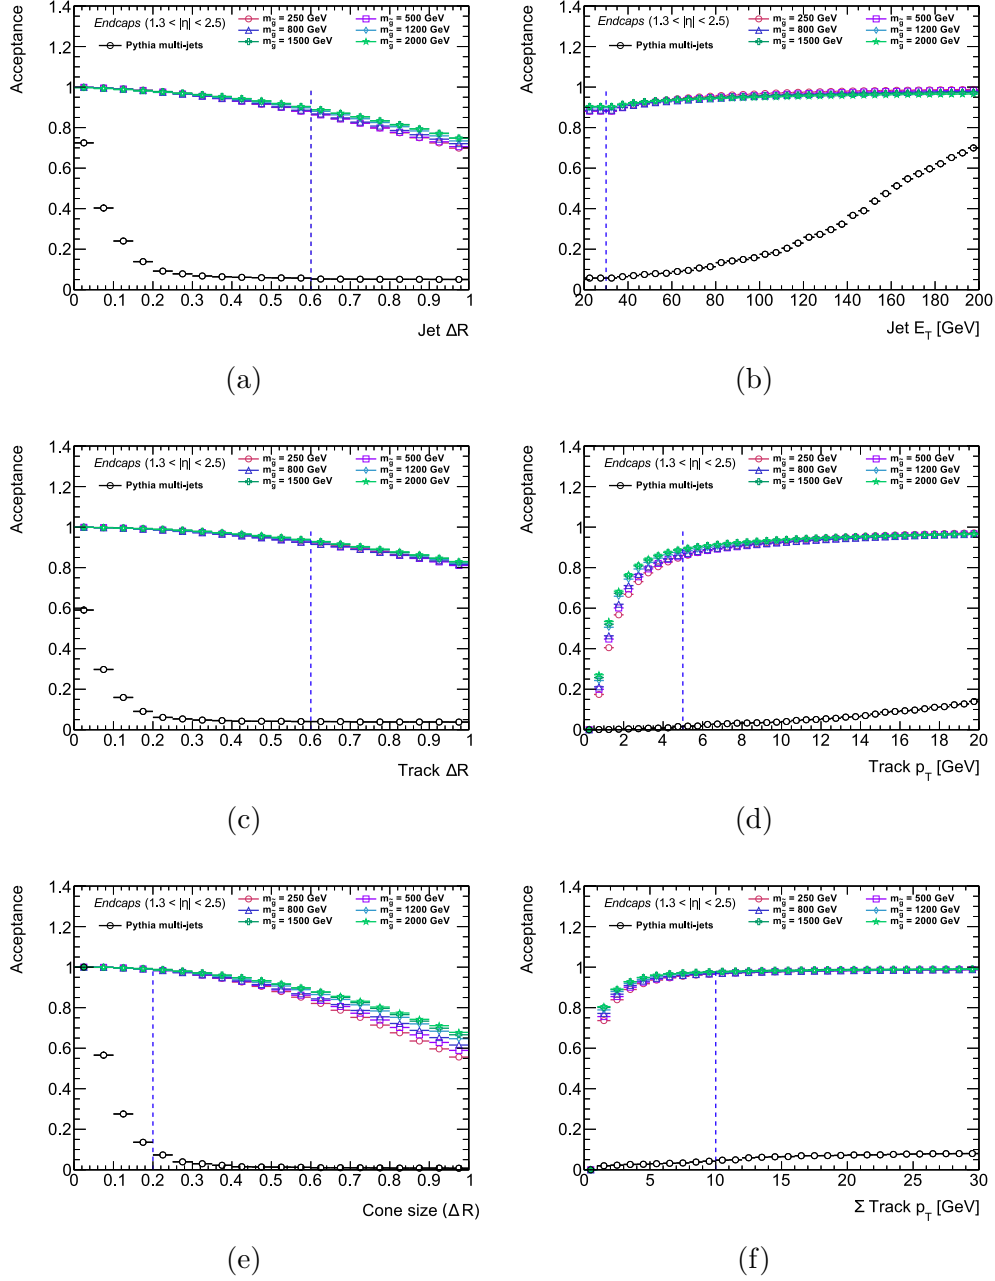


Figure 8.4: Acceptance for selecting endcap vertices passing good vertex criteria. Simulated QCD multi-jet events are compared against signal vertices to determine selection criteria that will remove most or all vertices from background QCD events and keep most signal vertices. In each case, dashed lines show the chosen cut value.

reconstructed, but there will be inner detector tracks pointing to the vertex. For high-energy track isolation, two parameters are considered: the transverse momentum, p_T , of the track, and the isolation cone size, ΔR . Figures 8.3c, 8.3d, 8.4c, and 8.4d show how both the signal and background acceptances vary with these parameters.

A second track isolation criteria aims to eliminate events where there is not a single high-energy track, but the sum of the track activity in a small cone around the MS vertex is exceptionally high. Note that this ensures that the GVC are collinear safe. This isolation criteria uses the vector-sum of all tracks within a ΔR cone around the vertex:

$$\left| \sum_{\Delta R^{cone}} \vec{p}_T \right|. \quad (8.1)$$

Figures 8.3f, 8.3e, 8.4e, and 8.4f show how varying both the track p_T sum and cone size affect the signal and background acceptances.

The final parameters selected to establish the good vertex criteria are summarized in Table 8.1. Figure 8.5 shows the vertex reconstruction efficiencies in the barrel and endcaps before and after application of the GVC, for the highest and lowest gluino mass samples. The largest effect comes from the application of the maximum number of MDT hits, which especially affects the $m_{\tilde{g}} = 2000$ GeV sample. Figure 8.2d shows that almost 50% of endcap vertices are lost in from the application of the maximum number of MDT hits cut. Despite the clear loss in efficiency, the GVC are necessary to improve the overall signal-to-background ratio.

8.2 Scale factors due to mismodelling in signal MC

In order to be sure the simulated MC signal samples accurately describe what would happen if a signal event occurred in data, a similar but orthogonal set of events are studied. Any large discrepancy between data and MC in the orthogonal selection will indicate that the MC is not sufficiently accurate and a scale factor needs to be applied. The uncertainty on the scale factor is applied as a systematic uncertainty.

Table 8.1: Summary of criteria for good MS vertices in the barrel and endcap regions.

Requirement	Barrel	Endcap
MDT hits	$300 \leq n_{\text{MDT}} < 3000$	$300 \leq n_{\text{MDT}} < 3000$
RPC/TGC hits	$n_{\text{RPC}} \geq 250$	$n_{\text{TGC}} \geq 250$
> 5 GeV track isolation	$\Delta R < 0.3$	$\Delta R < 0.6$
Track $ \Sigma p_{\vec{T}} $	$\Sigma p_{\text{T}} < 10 \text{ GeV}$	$\Sigma p_{\text{T}} < 10 \text{ GeV}$
> 30 GeV jet isolation	$\Delta R < 0.3$	$\Delta R < 0.6$

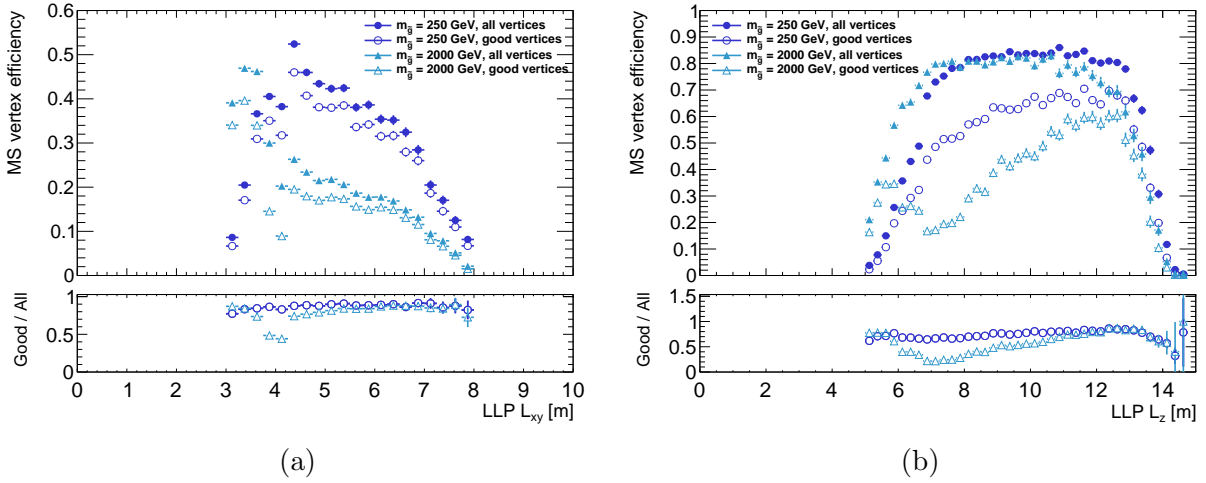


Figure 8.5: The MS vertex reconstruction efficiency in the barrel (a) and endcaps (b), with and without the scale factor applied, for the lowest and highest gluino mass Stealth SUSY samples.

This analysis uses two types of objects that need to be studied for data-MC agreement: muon RoI clusters, for the trigger, and MS vertices. However, since no Standard Model process has hadronic decays in the MS, there is not a reliable sample that can be used for the reconstructed final state. Thus, instead of clusters and vertices the previous reconstruction step is studied: level-1 muon RoIs and MS tracklets.

Systematic uncertainties and scale factors are determined by studying samples of multi-jet events in both data and MC. The details of these samples are outlined in Section 6.1.2. From the multi-jet events, leading jets with at least 20 ghost-associated muon segments are selected². This cut provides a sample of jets with a significant fraction of hadronic activity in the muon spectrometer. Because of the hadronic activity, they provide a good proxy for signal events.

A comparison of the number of MS tracklets in multi-jet event leading jet cones and signal MC events in the barrel and endcaps is shown in Figure 8.6. It is clear that LLP decays do leave more MS activity than punch-through jets. Figure 8.6 shows that higher- p_T jets leave increasing amounts of MS activity. Additionally, the scale factors are uniform as a function of increasing p_T (see Figure 8.7). Thus, one can conclude that it is acceptable to extrapolate the scale factors from the amount of MS activity in punch-through jets to the amount left by a hadronic LLP decay.

The scale factors and associated systematic uncertainties are summarized in Table 8.2.

The effect these scale factors and systematic uncertainties have on the final reconstructed objects (i.e. muon RoI clusters and MS vertices) are determined individually for each signal MC sample. The procedure for scale factors is as follows:

1. For each event, randomly remove a percentage of objects equal to $1 - \text{scale factor}$ (e.g. 13% of L1 muon RoIs in the barrel).

2

Ghost association simply selects particles that fall in the footprint of the jet shape. Since anti- k_T jets are used here, which are generally very circular, the footprint used for ghost association corresponds closely to a $\Delta R = 0.4$ cone around the jet axis.




Figure 8.6: (a) The number of MS tracklets in the leading jet cone in MC multi-jet events with leading jet $|\eta| < 0.8$ and $600 < p_T < 2000$ GeV or leading jet $p_T > 2000$ GeV. The number of MS tracklets in a cone around the LLP decay of the same size is shown for comparison. (b) The number of MS tracklets in the leading jet cone in MC multi-jet events with leading jet $1.3 < |\eta| < 2.5$ and $600 < p_T < 1000$ GeV or leading jet $p_T > 1000$ GeV. The lower p_T threshold is chosen for endcap jets because higher- p_T jets are preferentially at lower $|\eta|$ and thus statistics are poorer for TeV-scale jets.




Figure 8.7: The average number of MS tracklets in the leading jet cone in dijet events in both data (filled circles) and MC (open circles). The ratio between data and MC is shown below the plot and provides a scale factor and associated uncertainty on the number of tracklets reconstructed in MC versus data, in the (a) barrel and (b) endcaps.




Figure 8.8: The average number of level-1 muon RoIs in the leading jet cone in dijet events in both data (filled circles) and MC (open circles). The ratio between data and MC is shown below the plot and provides a scale factor and associated uncertainty on the number of level-1 muon RoIs reconstructed in MC versus data, in the (a) barrel and (b) endcaps.

Table 8.2: Scale factors and systematic uncertainties on data-MC agreement for MS objects.

Object	Scale factor	Systematic uncertainty
Barrel L1 muon RoIs	0.87	0.01
Barrel tracklets	0.55	0.01
Endcap L1 muon RoIs	0.97	0.01
Endcap tracklets	0.63	0.01

2. Reconstruct vertices and clusters given this reduced number of objects.
3. Re-evaluate efficiencies with the objects reconstructed after the scale factor is applied.

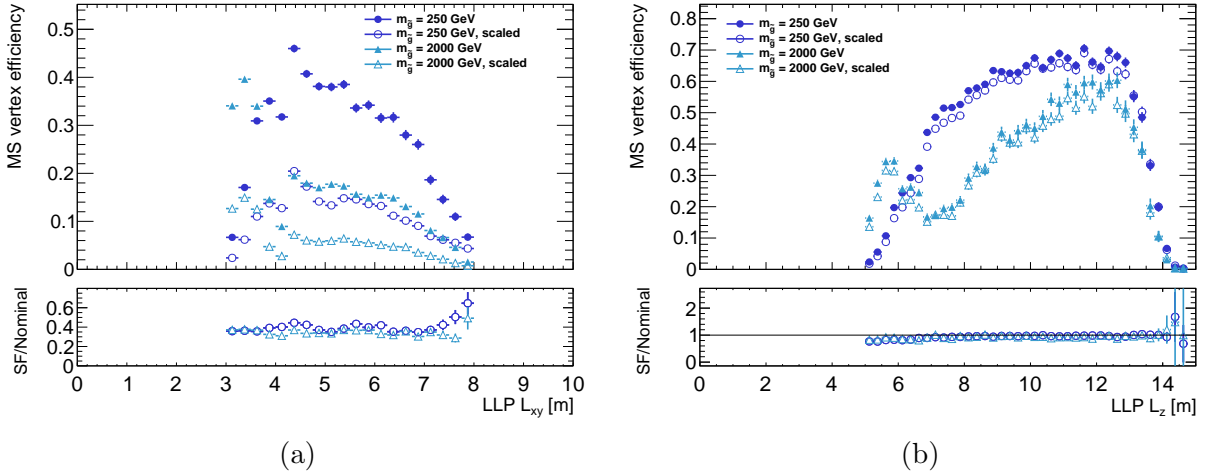


Figure 8.9: The MS vertex reconstruction efficiency in the barrel (a) and endcaps (b), with and without the scale factor applied, for the lowest and highest gluino mass Stealth SUSY samples.

The effect of applying the MS tracklet scale factors is shown in Figure 8.9. Because of the larger scale factor, the barrel vertex reconstruction is much more affected. The lack of a large effect on the endcap vertex reconstruction is due to the presence of many more tracklets than necessary, in most events (see Figure 8.6b). Thus, only keeping 63% of tracklets does not greatly affect the efficiency. Additionally, when the global number of tracklets is scaled down, *more* events have ≤ 40 tracklets. This means that more events are actually reconstructed.

The effects of these large scale factors are expected to be much greater in other signal models that do not have such high-energy LLP decays. For future analyses, the scale factors should be studied to fully understand their nature and thus be able to mitigate their effects on lighter, less boosted LLP decays.

8.3 *MS vertex reconstruction efficiencies*

MS vertex reconstruction efficiencies, for vertices after the application of all good vertex criteria and scale factors (see Section 8.2) are applied, are shown in Figure 8.10. Vertical dashed lines mark the physical locations of various sub-detectors that influence the efficiency shapes.

In the barrel, the vertex reconstruction efficiency increases as LLPs decay farther into the HCal. The highest gluino mass samples have a higher efficiency here because the LLPs are higher energy and thus their decay products can travel through more calorimeter before their energy is fully absorbed. As the LLP decay approaches the first MDT layer, the efficiency is at its peak, after which it slowly decreases. As the decay occurs farther out, the LLP decay products do not have time to separate before leaving tracks in MDT chambers, thus not creating enough distinct tracklets to reconstruct a vertex. Past the end of the second MDT plane, the LLP has also decayed past the last RPC layer. RPC hits are necessary to provide ϕ information for reconstruction, and thus at this point no more vertices can be reconstructed.

In the endcaps, the situation is much the same. Initially, decays occurring in the HCal are more likely to leave enough hits to reconstruct a vertex if they come from high mass gluino decays. As the LLP decays occur nearer the end of the HCal, more and more hits are distributed in the muon system, and thus many vertices fail the maximum MDT hit criteria. The effects of this are clear in Figure 8.5b. As the decays occur closer to the second MDT plane, the LLP decay products have a shorter distance to disperse and thus fewer hits are recorded. This has an understandably larger effect on the highest gluino mass samples, since they are the most affected by the maximum MDT hit cut. As in the barrel, after the second MDT plane there are no more trigger planes and thus no vertices can be reconstructed.

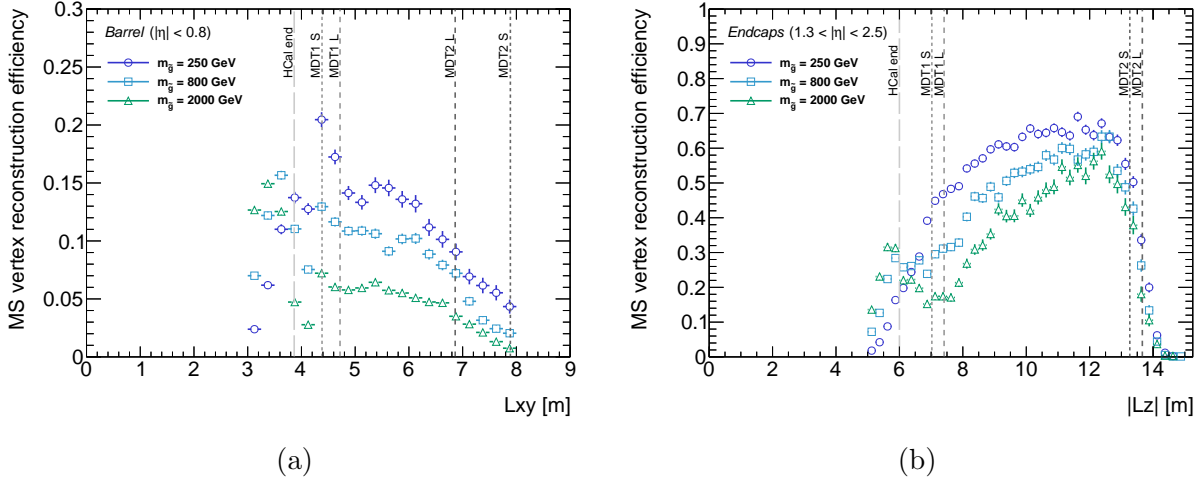


Figure 8.10: The MS vertex reconstruction efficiency in the barrel (a) and endcaps (b), with all good vertex criteria and scale factors applied, for the lowest and highest gluino mass Stealth SUSY samples. Vertical dashed lines show where relevant detector structures are located.

8.4 Systematic uncertainties associated to displaced MS objects

8.4.1 Systematic uncertainty due to mismodelling in signal MC

For systematic uncertainties, the scale factor is simply varied up and down by the systematic uncertainty. Three different sets of vertices and clusters are reconstructed for each sample: nominal, $+1\sigma$, and -1σ . The systematic uncertainty is uniform in LLP decay position, and thus before calculating the final systematic, the efficiencies are recalculated by rebinning the entire x -axis into a single bin. The overall systematic uncertainties on vertices and clusters are determined by calculating the ratio between the $+1\sigma$ and -1σ efficiencies with the nominal efficiency. Vertex reconstruction efficiencies with systematic uncertainties applied are shown in Figure 8.11 for the Stealth SUSY $m_{\tilde{g}} = 250$ GeV sample.




Figure 8.11: The MS vertex reconstruction efficiency at nominal scale factor, $+1\sigma$, and -1σ in the (a) barrel and (b) endcaps for $m_{\tilde{g}} = 250$ GeV Stealth SUSY vertices. The overall systematic is found by rebinning the entire x -axis into a single bin, and finding the net change in efficiency.

8.4.2 Systematic uncertainties due to pileup uncertainties

The distribution of $\langle\mu\rangle$ vs. the number of primary vertices in MC simulation does not exactly match that in data. Thus, even after reweighting the simulation to have a $\langle\mu\rangle$ distribution that matches data, there is an systematic uncertainty associated to the amount the analysis is affected by this disagreement.

Figure 8.12 shows the barrel vertex reconstruction efficiency for the Stealth SUSY $m_{\tilde{g}} = 250$ GeV sample, with and without the pileup systematic uncertainties applied. In Figure 8.12b, it is clear that any effect the pileup systematic uncertainty has on the vertex reconstruction efficiency does not depend on the LLP decay position. Thus, the final contribution to the vertex reconstruction efficiency is determined with a rebinned efficiency, just as with the data–MC scale factor systematic uncertainties addressed in the previous section. For this benchmark sample, pileup systematics contribute $+0.9$, -0.1% to the overall vertex reconstruction systematic uncertainty.

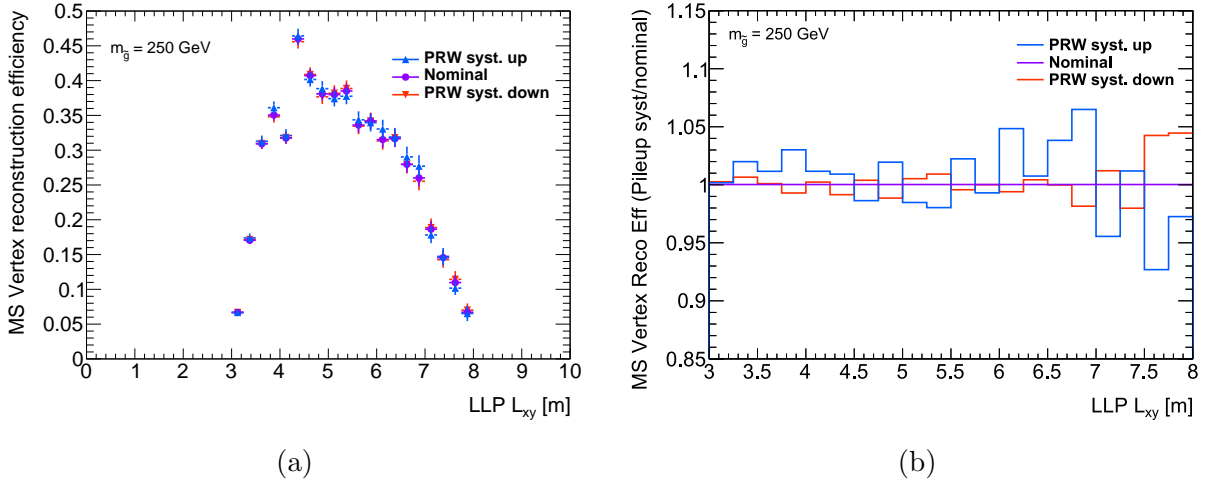


Figure 8.12: The effect of pileup systematic uncertainties on the MS vertex reconstruction efficiency in the barrel, for the Stealth SUSY $m_{\tilde{g}} = 250$ GeV sample.

8.4.3 Systematic uncertainties due to PDF uncertainties

The parton distribution function (PDF) used to generate signal MC events was NNPDF2.3 with QED corrections (NNPDF23_lo_as_0130_qed). The central value “best estimate” of the PDF is used to generate events, but this is formed from the 100 different PDF fits extracted from input data. In order to determine the uncertainty this ensemble of fits results in, each MC sample is reweighed to mimic each of the 100 PDF fits. Then, for each final-state efficiency (*e.g.* Muon RoI Cluster trigger or MS Vertex reconstruction), the ratio of the efficiency for the PDF fit weighted events to the central value PDF events is determined. An example of this is shown in Figure 8.13a.

The sample standard deviation of the ensemble is calculated for each L_{xy} bin, and the systematic uncertainty is the central value \pm the standard deviation. The central value $\pm 1\sigma$ is shown as black line in Figure 8.13a and the data points in Figure 8.13b. The per-bin variations are approximately uniform, and thus the net systematic uncertainty on the reconstruction efficiency is determined by rebinning the efficiency into a single bin, as with the pileup and data–MC scale factor systematic uncertainties. In this case, the systematic

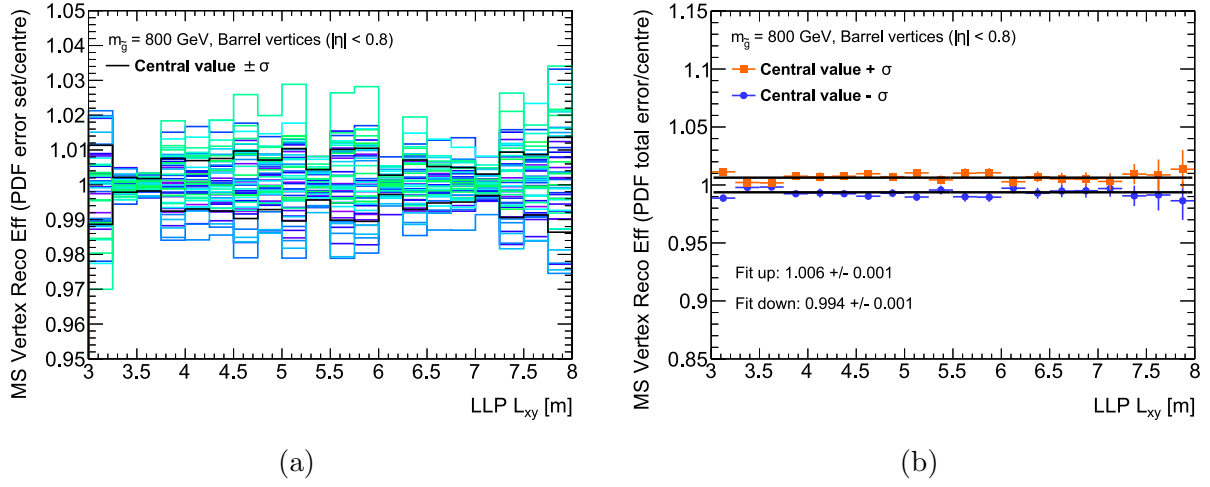


Figure 8.13: The effect of parton distribution function uncertainties on the MS vertex reconstruction efficiency in the barrel, for the Stealth SUSY $m_{\tilde{g}} = 800$ GeV sample.

uncertainty on barrel MS vertex reconstruction due to PDF uncertainties is 0.006, or 0.6%.

8.4.4 Summary of systematic uncertainties

The contributions to systematic uncertainties from each of the components, and the total systematic uncertainties, are given in Tables 8.3, 8.4, 8.5, and 8.6. Entries with only a systematic in one direction (e.g. +0.2%) mean that there was no downwards effect on the efficiency in question when the systematic was varied.

Table 8.3: Summary of systematic uncertainties on the muon RoI cluster trigger in the barrel for the Stealth SUSY benchmark samples.

$m_{\tilde{g}}$ (GeV)	Trigger reconstruction	Pileup	PDF	Total
250	+1.0, -1.2%	+0.4, -0.1%	$\pm 0.1\%$	+1.1, -1.2%
500	$\pm 1.1\%$	+0.04, -0.2%	$\pm 0.1\%$	$\pm 1.1\%$
800	$\pm 1.0\%$	+0.1, -0.02%	$\pm 0.1\%$	$\pm 1.0\%$
1200	+1.0, -1.1%	-0.1%	$\pm 0.04\%$	+1.0, -1.1%
1500	+1.0, -1.1%	+0.1, -0.3%	$\pm 0.1\%$	+1.0, -1.1%
2000	+1.0, -0.8%	-0.2%	$\pm 0.1\%$	+1.0, -0.9%

Table 8.4: Summary of systematic uncertainties on the muon RoI cluster trigger in the endcaps for the Stealth SUSY benchmark samples.

$m_{\tilde{g}}$ (GeV)	Trigger reconstruction	Pileup	PDF	Total
250	$\pm 1.8\%$	+0.1, -0.2%	$\pm 0.5\%$	$\pm 1.9\%$
500	+1.5, -1.6%	+0.1, -0.002%	$\pm 0.4\%$	+1.5, -1.6%
800	+1.4, -1.3%	+0.3, -0.1%	$\pm 0.2\%$	+1.4, -1.3%
1200	+1.2, -1.3%	+0.1%	$\pm 0.3\%$	+1.2, -1.3%
1500	+1.1, -1.2%	$\pm 0.1\%$	$\pm 0.2\%$	$\pm 1.2\%$
2000	$\pm 1.0\%$	$\pm 0.1\%$	$\pm 0.2\%$	+1.0, -1.1%

Table 8.5: Summary of MS vertex systematic uncertainty in the barrel for the Stealth SUSY benchmark samples.

$m_{\tilde{g}}$ (GeV)	Vertex reconstruction	Pileup	PDF	Total
250	+1.9, -2.5%	+0.9, -0.8%	$\pm 0.3\%$	+2.1, -2.6%
500	+2.5, -2.4%	+0.8, -0.4%	$\pm 0.7\%$	+2.7, -2.5%
800	+2.7, -2.5%	+0.7, -0.2%	$\pm 0.5\%$	+2.8, -2.5%
1200	$\pm 2.3\%$	+0.4%	$\pm 0.5\%$	+2.4, -2.3%
1500	$\pm 2.6\%$	+0.2%	$\pm 0.5\%$	+2.7, -2.6%
2000	+2.4, -2.0%	+0.3%	$\pm 0.4\%$	+2.5, -2.1%

Table 8.6: Summary of MS vertex systematic uncertainty in the endcaps for the Stealth SUSY benchmark samples.

$m_{\tilde{g}}$ (GeV)	Vertex reconstruction	Pileup	PDF	Total
250	+0.4, -0.6%	+0.2, -0.7%	$\pm 0.6\%$	+0.7, -1.1%
500	$\pm 0.5\%$	+0.3, -0.7%	$\pm 0.6\%$	+0.8, -1.1%
800	+0.4, -0.3%	$\pm 0.3\%$	$\pm 0.5\%$	$\pm 0.7\%$
1200	+0.4, -0.5%	+0.2, -0.7%	$\pm 0.5\%$	+0.7, -1.0%
1500	+0.5, -0.6%	+0.06, -0.7%	$\pm 0.4\%$	+0.6, -1.0%
2000	+0.6, -0.001%	-0.2%	$\pm 0.5\%$	+0.8, -0.5%

Chapter 9

SEARCH FOR EVENTS WITH TWO DISPLACED VERTICES

Many models with long-lived particles predict that they are pair-produced. With two long-lived particles it is possible to require they both leave a displaced decay signature. This chapter outlines the search for events with two isolated, displaced MS vertices. While all the performance plots in the previous chapters were specific to Stealth SUSY, this search is in no way exclusive to that model. In particular, it has good sensitivity to the Standard Model Higgs boson decays to a pair of neutral, long-lived scalars (as described in Section 2.3.)

Requiring two displaced vertices dramatically reduces backgrounds, since while the probability for the existence of a single displaced vertex from any background source is very small, the existence of two in the same event is uncorrelated and thus proportional to the square of an already very small number.

Being able to accurately quantify the expected number of background events is crucial to begin able to make a discovery (or good exclusion limit). While there are no standard model backgrounds that leave isolated, displaced MS vertices under good operating conditions, various processes can contribute to a non-zero background. These can be from collision or non-collision processes:

- Jets can fluctuate to contain only neutral hadrons, be sufficiently mismeasured in the calorimeters, and then punch through to the muon system leaving a shower of tracks from non-absorbed hadrons.
- Beam-halo muons can shower in an endcap and fake an isolated MS vertex
- Cosmic muons can shower in the top or bottom of the detector, or a cosmic shower

could fake a vertex in barrel or endcaps

- Noise bursts in the muon system can create enough MDT and trigger hits that vertices could be reconstructed through random combinations.

None of these backgrounds can be measured with simulated events. In principle, a punch-through jet background could be determined with simulated QCD multi-jet events. However, the simulation will not accurately model such rare events, and even if it did the statistics necessary to calculate the background using simulation would be infeasibly large. Thus, data-driven methods are used to calculate the background in both the two vertex and single vertex searches.

9.1 Expected number of background events

The expected background is calculated based on the assumption that the probability of finding a single MS vertex is uncorrelated to the probability of finding a second. In the simplest of cases, this means that the number of events with two MS vertices could be presented as:

$$N[2 \text{ MS Vertex}] = N[1 \text{ MS Vertex}] \times P[\text{MS Vertex}]. \quad (9.1)$$

In practice, the calculation is somewhat more complicated because events are also required to pass the Muon RoI Cluster trigger, and once an event passes the trigger (*i.e.* has a cluster of muon RoIs), it is more likely to have a vertex. Thus, two types of vertices must be considered: those matched to a cluster, and those occurring without any trigger. Of course, at least one vertex must be associated to a cluster that passes the trigger, or the event would not pass the full cutflow.

Additionally, the barrel and endcaps must be treated separately. This is for two reasons. First, the trigger uses different technologies and has different requirements in the barrel and endcaps. Second, the MS endcap chambers are outside the endcap toroid and thus are all

straight-line segments, so the vertex reconstruction has less information to use.

Last, any events with two muon RoI clusters must be separated by at least $\Delta R = 1.0$. The muon RoI cluster algorithm can easily produce clusters that are nearby, since it is simply a naive cone algorithm with radius parameter 0.4. Since this radius is a hard cut, any *large* clusters with a sufficient number of muon RoIs will appear to be two small clusters. However, these two clusters would represent a *single* region of MS activity. Additionally, the vertex reconstruction algorithm (see Section 7.2) starts from tracklets clusters of size $(\Delta\eta \times \Delta\phi) = (0.7 \times 1.0)$. Therefore, even if there were two very nearby long-lived particle decays, it would be unlikely that the vertex reconstruction algorithm could resolve them.

Muon RoI clusters have a radius of $\Delta R = 0.4$, and thus clusters separated by more than 0.8 will be well-defined. Both Higgs boson decays and Stealth SUSY signal events tend to be separated by at least $\Delta R = 1.0$ (see Figure 9.1) and thus 1.0 is chosen as the minimum ΔR between two clusters or vertices.

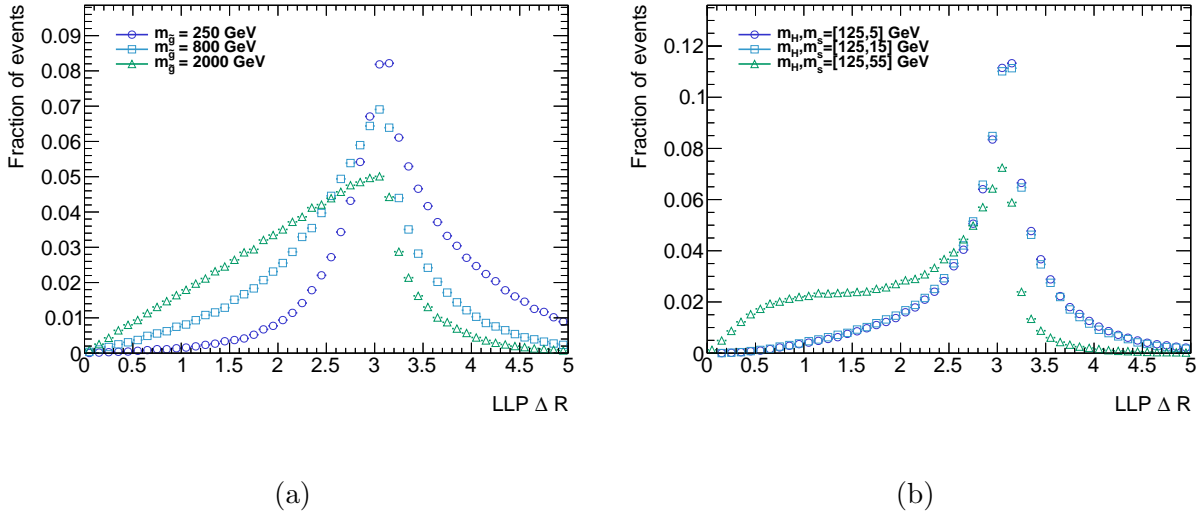


Figure 9.1: The angular separation, ΔR , between singlinos in Stealth SUSY events and scalars in Higgs boson events.

In the following paragraphs, cluster, or Clu, will be used to refer to a good, isolated Muon

Table 9.1: The numbers of events in the 13 TeV dataset passing the Muon RoI Cluster trigger without isolation criteria, and events passing said trigger also with a matched vertex. These values are necessary to calculate the background prediction. The notation is explained in the text. Uncertainties are statistical only.

Quantity	Value
Events passing the muvtx trigger	
N[1 BClu]	6894488
N[1 EClu]	2399957
N[2 BClu]	628
N[1 BClu, 1 EClu]	63
N[2 EClu]	140
Events passing muvtx, with $\Delta R(Vx, Clu) < 0.4$	
N[1 CluVx]	105247
N[1 BCluVx]	30833
N[1 ECluVx]	74414
N[1 BCluVx, 1 UM BClu]	0
N[1 ECluVx, 1 UM BClu]	0
N[1 BCluVx, 1 UM EClu]	0
N[1 ECluVx, 1 UM EClu]	0

Table 9.2: The probabilities needed to compute the background from events containing MS vertices from non-signal processes. $P[BVx | BClus]$ and $P[EVx | EClus]$ are the probabilities of finding one MS vertex in events passing the Muon RoI Cluster trigger that have a single cluster that is matched to a vertex in the barrel or endcaps, respectively. The probability of finding an MS vertex not matched to a cluster of muon RoIs is given by $P[Vx | \text{no MS trig}]$.

Quantity	Value
$P[Vx \text{no MS trig}]$	$4/27040061 = (1.5 \pm 0.7) \times 10^{-7}$
$P[BVx BClus]$	$(4.47 \pm 0.03) \times 10^{-3}$
$P[EVx EClus]$	$(3.10 \pm 0.01) \times 10^{-2}$

RoI Cluster that either fired the trigger, or would have fired the trigger but there was already another good cluster in the event. Good MS Vertices will be referred to simply as vertex,

or V_x . A vertex–cluster combination (*i.e.* a vertex within $\Delta R = 0.4$ of the cluster) will be denoted $\text{Clu}V_x$. The letters B or E before any of Clu , V_x , or $\text{Clu}V_x$ will indicate whether the object is in the fiducial regions of the barrel or the endcaps for reconstructing MS vertices. If a cluster is not matched to an MS vertex it is considered unmatched, indicated by UM Clu .

There are three types of background events that must be calculated. First, the number of events with one cluster and two vertices. This is determined from multiplying the number of events passing the Muon RoI Cluster trigger and having only one cluster with that cluster matched to a vertex ($N[1 \text{ Clu}V_x]$), with the probability that a vertex will randomly occur from *any* process, without the trigger firing ($P[V_x \mid \text{no MS trig}]$). The probability $P[V_x \mid \text{no MS trig}]$ is determined from zero-bias events by dividing the number of good MS vertices reconstructed by the total number of zero-bias triggered events that pass standard event quality criteria. The expected number of background events with one cluster and two vertices evaluates to:

$$N[1\text{Clu}V_x, 1V_x] = N[1 \text{ Clu}V_x] \times P[V_x \mid \text{no MS trig}] \quad (9.2)$$

$$= (105000 \pm 300) * (1.5 \pm 0.7) \times 10^{-7} \quad (9.3)$$

$$= (0.016 \pm 0.008). \quad (9.4)$$

The second and third terms encompass the number of events where there are two good vertices, both matched to a cluster. These are calculated in two parts, because the probability of an unmatched barrel cluster being associated to a vertex is much smaller than the probability of an unmatched endcap cluster being associated to a vertex. For an unmatched cluster in the barrel (endcaps), the value is determined from a number times a probability. The former is the number of events passing the Muon RoI Cluster trigger having *two* clusters with one them is matched to a vertex and the other unmatched in the barrel (endcaps). The latter is the probability of a barrel (endcaps) cluster being matched to a vertex.

In the dataset being analyzed there are zero events with two clusters separated by $\Delta R > 1.0$, so there is no evidence for the presence of a signal and 95% confidence level limits are set.

Chapter 10

SEARCH FOR EVENTS WITH ONE DISPLACED VERTEX AND TWO JETS

For a search with only one displaced object, a similar method to the two vertex search will not work: all one can do is *measure* the number of events with a single vertex. Instead, one can use the characteristics of “bad” vertices to estimate how many good vertices will be seen. This method relies on the assumption that the same physics process that gives bad vertices is the one creating fake good vertices, and thus one must be careful to construct the calculation in such a way that this is true. Before elaborating on the details, the methodology used for the background estimation is outlined in Section 10.1.

10.1 Description of the ABCD method for background determination

These two vertex-based criteria provide the backbone for an *ABCD method* to determine the background. The ABCD method relies on there being two, uncorrelated variables that distinguish signal-like events from background events. The two variables, denoted X and Y for simplicity, are plotted against each other. Values of X and Y that distinguish signal-like events from background events are chosen, such that the majority of signal events reside in region A, as shown in Figure 10.1. Regions B, C, and D thus describe three different regions that background events are split into.

If the two variables X and Y are uncorrelated, then the fractional distribution of events along X should not change as a function of Y , and vice-versa. This implies that for any given point $X = X_0$,




Figure 10.1: Sketch of the ABCD method for background estimation.

$$\frac{N_{X>X_0}}{N_{X<X_0}} \quad (10.1)$$

is constant as a function of Y .

Translating this back to the plane in Figure 10.1, the consequence of the variables being uncorrelated is that the expected number of events in region A can be determined with the following ratios:

$$\frac{N_A}{N_B} = \frac{N_C}{N_D}, \quad (10.2)$$

which implies that

$$N_A^{\text{expected}} = N_B \times \frac{N_C}{N_D}, \quad (10.3)$$

where N_A^{expected} is the number of expected events in region A for the physics process entering the ABCD method.

The ABCD method is applied to events passing all other *signal region* (SR) cuts, aside from the two cuts forming the X and Y axes. In order to validate that an ABCD method will correctly predict the background, it is necessary to identify a *validation region* (VR). This is developed by identifying event-level criteria that select a subset of events close to the signal region in characteristics, but not overlapping or containing a potentially significant fraction of signal. The validation region is used to verify that $N_A^{\text{expected}} = N_A^{\text{observed}}$ before the signal region is unblinded¹.

10.2 Variables distinguishing signal vertices from background vertices

The main criterion that is used to distinguish a vertex as good or bad is its degree of isolation. In the case of the two vertex search, the good vertex criteria outlined in Section 8.1.2 cover three types of isolation: jet, high- p_T track, and low- p_T track. These isolation criteria exist primarily to eliminate the overwhelmingly large number of MS vertices that appear in data due to punch-through jets. To characterise the *degree* of isolation with a single value, consider the variable:

$$\Delta R_{\text{min}} = \min(\Delta R(\text{vertex, closest jet}), \Delta R(\text{vertex, closest track})). \quad (10.4)$$

Here, $\Delta R(\text{vertex, closest jet})$ considers the same set of jets discussed for isolation in Section 8.1.2: $E_T > 30$ GeV, $\log_{10}(E_{\text{HAD}}/E_{\text{EM}}) > 0.5$, passes the standard good jet criteria, and passes JVT cuts. Similarly, $\Delta R(\text{vertex, closest track})$ considers good tracks with $p_T > 5$ GeV.

In most cases, a punch-through jet will have both a track *and* a jet within a small ΔR of the vertex axis. Combining the track and jet isolation into a single variable ensures that

1

Note that if the signal region were used to validate itself, one would have to operate under the assumption that there is no actual signal in the signal region, eliminating the usefulness of the analysis.




Figure 10.2: Isolation of barrel MS vertices from tracks and jets in events with two jets with $p_T > 50$ GeV.

any vertices caused by punch-through jets that happen to have no high- p_T ID tracks or have poorly-measured calorimeter activity will still be considered isolated. Figure 10.2 shows the relationship between jet and track isolation, for barrel MS vertices in data events with at least 2 jets with $p_T > 50$ GeV. Most vertices have both a jet and a track within $\Delta R = 0.3$. The spread of events along the x - and y -axes is due to vertices that have a nearby jet but no track, or a nearby track but no jet.

Another distinguishing characteristic of signal MS vertices is the number of hits they leave in the muon spectrometer. Even if a jet does manage to punch through and leave enough activity to reconstruct a vertex, much of the hadronic shower will have been absorbed by the calorimeters.

The power of the isolation variable and number of hits in distinguishing vertices from punch-through jets from signal-like vertices is shown in Figure 10.3.

Thus, the ABCD plane for background estimation is constructed with the following two

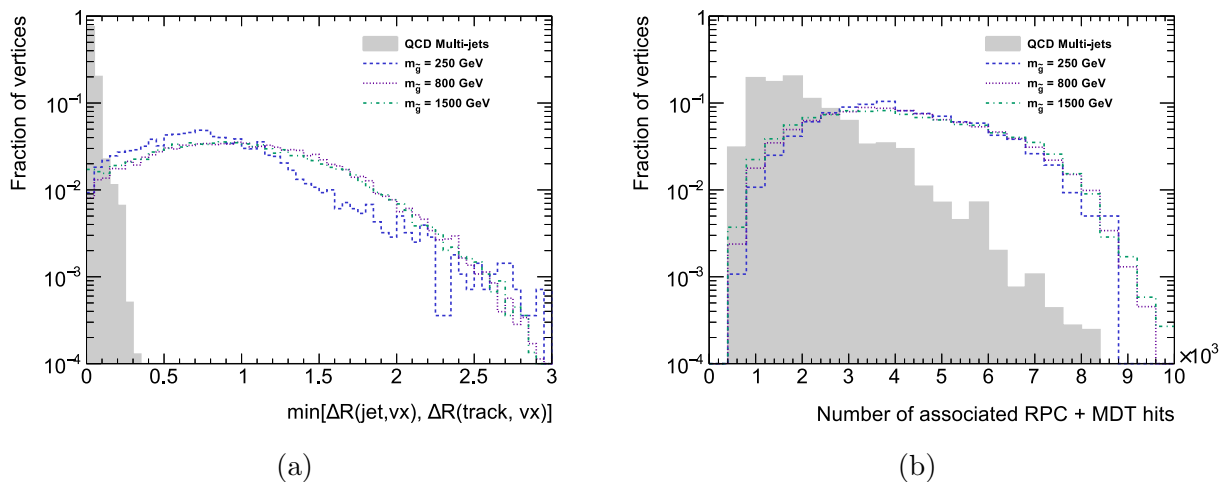


Figure 10.3: Two vertex-based criteria and their ability to distinguish signal vertices from punch-through jet vertices. (a) The degree of isolation of the vertex, where clearly punch-through jet vertices generally have a jet or track within $\Delta R = 0.3$ of the vertex axis. (b) The number of hits associated to a vertex. On average, most punch-through jet vertices have fewer hits than signal-like vertices, and the distribution for punch-through jet vertices falls off much faster than that for signal vertices.

vertex-based criteria: the ΔR between the closest jet or track and the vertex, and the number of associated MDT + RPC hits. These two variables will form the x - and y -axes of the ABCD plane, respectively.

10.3 Developing a signal region for Stealth SUSY events

Events in the Stealth SUSY model studied in this analysis have a very striking signature: in addition to two long-lived particles, there are two prompt gluons. With sufficient mass splitting between the gluino and singlino, these gluons create hard jets that can contribute to event selection. Figure 10.4 shows that the second highest (sub-leading) jet p_T is generally above 150 GeV, with some loss of signal efficiency in the lowest mass gluino sample.




Figure 10.4: Sub-leading jet p_T for Stealth SUSY signal MC samples, in events with a barrel MS vertex.

Since the chosen x -axis for the ABCD plane depends on the ΔR between a jet and the vertex, jets chosen for the VR and SR selections must have $\Delta R(\text{jet}, \text{vertex}) > 0.7$. This prevents the VR and SR cuts from selecting based on punch-through jets that leave vertices. In signal events this requirement has minimal effect, since jets and vertices originate from different particles and thus tend to be well-separated. In data events, the main background is QCD, and thus if a jet is going to be near a vertex it is generally well within $\Delta R = 0.7$.

This is clear from the good vertex criteria plots for jet isolation, notably Figure 8.3a.

Table 10.1 provides a full list of the criteria to select event from the signal and validation regions.

Table 10.1: Summary of the criteria used to select events for the two jets – one vertex search. All selection criteria are also applied to signal MC events when determining the amount of expected signal events in the dataset.

Region	Criteria
All	Event passes muvtx trigger
	Event passes cleaning criteria
	Event has a good primary vertex
	Event has no bad jets
	Event has one MS vertex
	MS Vertex has at least 300 MDT hits, 250 RPC/TGC hits
	MS Vertex is not in the overlap region ($ \eta_{vx} < 0.8$ or $1.3 < \eta_{vx} < 2.5$)
	MS Vertex matched to triggering muon RoI cluster ($\Delta R(\text{vertex}, \text{cluster}) < 0.4$)
Event has two good jets with $\Delta R(\text{vertex}, \text{jet}) > 0.7$	
Validation region	Two jets: $50 < E_{T,1} < 150$ GeV, $E_{T,2} > 150$ GeV
Signal region	Two jets: $E_T > 150$ GeV

10.3.1 Modelling the efficiency for Stealth SUSY events to have two jets

Determining the number of expected events at the proper lifetime the signal MC was generated at is trivial: the cutflow simply needs to be applied to signal events in the same way it is applied to data. However, in order to extrapolate the limits to the full range of proper lifetimes the analysis is sensitive to, an efficiency dependent on the truth parameters is necessary. To model the probability that an event has two jets, there are a number of variables to consider.

One long-lived particle in each event must decay in the muon spectrometer and leave an MS vertex. The other, however, may decay effectively promptly, displaced at some location inside the detector, or outside of the detector. If this non-vertex LLP decays before the end of the calorimeters, it has some probability of leaving a jet that can assist the event in passing

the two jet requirement. This probability will depend on both the LLP's decay position and momentum. If it decays in the barrel ($|\eta| < 1.5$: from the perspective of the calorimeters), the transverse decay position L_{xy} and transverse momentum p_T are used. Otherwise, the longitudinal decay position and momentum are used: L_z and p_z .

In addition to the long-lived particles, Stealth SUSY events also contain two prompt gluons, which also contribute to the efficiency for an event to contain two jets. Because these gluons are produced in the same decay as the long-lived particles ($\tilde{g} \rightarrow \tilde{S}g$), their momenta depend somewhat on the long-lived particle momenta (and thus decay position). However, the probability for an event to contain two jets with $E_T > 150$ GeV that do *not* come from either of the LLPs is relatively constant. Thus, the overall probability for an event to contain two jets is plotted in two dimensions, as a function of decay position and momentum. Since the analysis only requires jets in events that also have an isolated MS vertex and pass the Muon RoI Cluster trigger, the probability is determined from this subset of events.

In the Stealth SUSY $m_{\tilde{g}} = 250$ GeV sample, the probability that either prompt gluon leaves a high energy jet is fairly small, since each 250 GeV gluino decays into a gluon plus a 100 GeV singlino. Only when the gluinos are boosted are the gluons high enough energy to leave a 150 GeV jet. At the other end of the spectrum, the two-jet probability for the 2000 GeV gluino mass sample is nearly always exactly one. Here, a 2000 GeV gluino is decaying into a gluon plus a 100 GeV singlino, which means there is a significant amount of phase space where the prompt gluons are high energy.

When constructing the 2-jet probability plots, the final bin includes all events from the bin start out to infinity. This is because once the LLP decays outside of the calorimeters, it does not affect the jets in the event, and once it has sufficient p_T to always leave a 150 GeV jet with the same probability, there is no need to have finer binning. The cut values are motivated in Figure 10.5, where one can see that for barrel jets efficiency is uniform above LLP $p_T \sim 350$ GeV and zero when the LLP decays outside of the hadronic calorimeter, starting at ~ 3.8 m. Example global two-jet efficiency plots are shown in the following

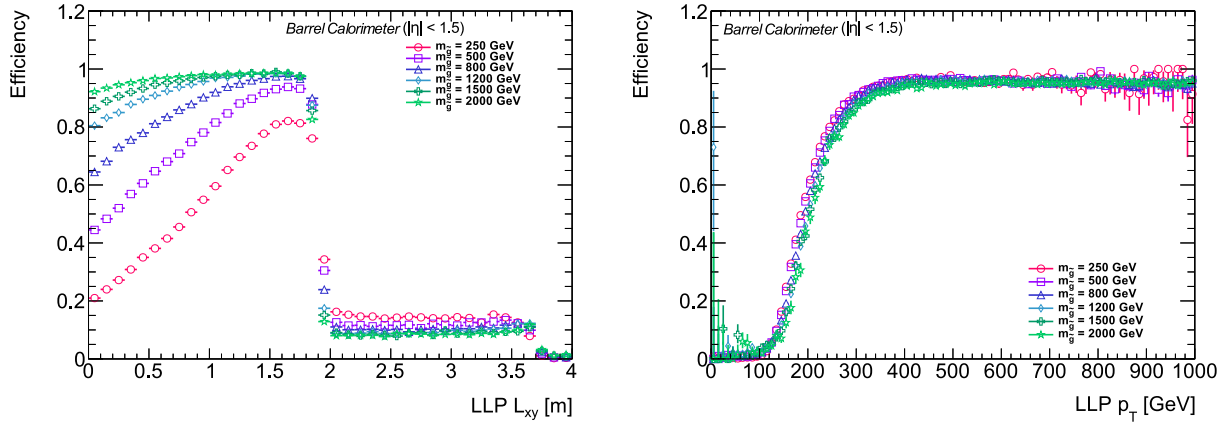


Figure 10.5: Efficiency for a long-lived particle (singlino) to leave a well-reconstructed jet in the barrel calorimeter ($|\eta| < 1.5$) as a function of (a) long-lived particle transverse decay position, and (b) long-lived particle transverse momentum.

chapter (Figure 11.7 for the $m_{\tilde{g}} = 800$ GeV sample).

Jet energy scale systematic uncertainties

Jets are carefully calibrated to ensure that the reported energy accurately represents what a particle at any given energy and location would deposit. Associated with this calibration is a set of systematic uncertainties. Each set of uncertainties is applied to the jets in an event, and the global efficiency is recalculated. The overall systematic uncertainty on the two-jet efficiency is determined by adding each component in quadrature.

A similar procedure is followed to determine the contribution of PDF and pileup systematic uncertainties to the overall jet efficiency systematic uncertainty. A summary of the contribution of each systematic uncertainty is given in Tables 10.2 and 10.3.

10.4 Applying the ABCD method to data

Applying the ABCD method in the validation region (Figure 10.6a) shows that 36 ± 4 events are expected in region A, while 31 ± 6 events are observed. The discrepancy between

Table 10.2: Summary of systematic uncertainties on the event-level probability for two jets with $E_T > 150$ GeV, in events with one MS vertex and a second LLP with $|\eta| < 1.5$, for the Stealth SUSY benchmark samples.

$m_{\tilde{g}}$ (GeV)	JES	Pileup	PDF	Total
250	+3.9, -3.0%	+0.6, -0.3%	$\pm 1.5\%$	+4.2, - 3.3%
500	+2.0, -2.2%	+0.5, -0.1%	$\pm 0.3\%$	+2.1, -2.2%
800	+0.7, -0.5%	+0.02, -0.1%	$\pm 0.1\%$	+0.8, -0.5%
1200	+0.2, -0.3%	-0.04%	$\pm 0.1\%$	+0.2 -0.3%
1500	+0.2, -0.1%	+0.03, -0.02%	$\pm 0.1\%$	+0.2, -0.1%
2000	$\pm 0.1\%$	+0.02, -0.1%	$\pm 0.1\%$	$\pm 0.1\%$

Table 10.3: Summary of systematic uncertainties on the event-level probability for two jets with $E_T > 150$ GeV, in events with one MS vertex and a second LLP with $|\eta| > 1.5$, for the Stealth SUSY benchmark samples.

$m_{\tilde{g}}$ (GeV)	JES	Pileup	PDF	Total
250	+8.8, -3.7%	+5.4, -1.3%	$\pm 1.6\%$	+10.5, -4.3%
500	+2.7, -2.4%	+1.0, -2.6%	$\pm 0.5\%$	+2.9, -3.6%
800	+0.8, -1.0%	-0.2%	$\pm 0.2\%$	+0.8, -1.0%
1200	+0.3, -0.1%	+0.4, -0.1%	$\pm 0.1\%$	+0.5, -0.2%
1500	+0.2%	+0.1%	$\pm 0.1\%$	+0.2, -0.1%
2000	+0.1%	+0.04, -0.2%	$\pm 0.1\%$	+0.1, -0.2%

expected and observed events is less than one standard deviation. This shows that applying the ABCD method in the chosen plane is a valid method, and the signal region, which is closely related kinematically, should show similar agreement in the absence of any signal excess. The observed numbers of events in each of the regions A – D are shown in Table 10.4. Uncertainties are statistical only.

The ABCD plane for events in the signal region is shown in Figure 10.6b. Both the validation and signal regions show very low correlations between the x - and y -axes: 0.006 for the validation region, and 0.0003 for the signal region. The final background prediction in region A of the signal region is 10 ± 3 events (shown in Table 10.4). Upon unblinding, 10




Figure 10.6: ABCD plane in the validation and signal regions for events with a single barrel MS vertex.

Table 10.4: Event counts in each of the four regions of the ABCD plane, and the expected number in region A, for both the validation and signal regions.

Closest jet or track ΔR vs. Vertex nMDT+nRPC hits					
Region A cuts at $\min(\Delta R(\text{vertex}, \text{jet/trk})) > 0.3$ and $n\text{Hits} > 2000$					
Region	A	B*C/D	B	C	D
VR: 1j50,1j150	31 ± 6	36 ± 4	6580 ± 80	66 ± 8	12200 ± 100
SR: 2j150	10	10 ± 3	1740 ± 40	15 ± 4	2690 ± 50

events are observed so there is no evidence for a signal and limits are set.

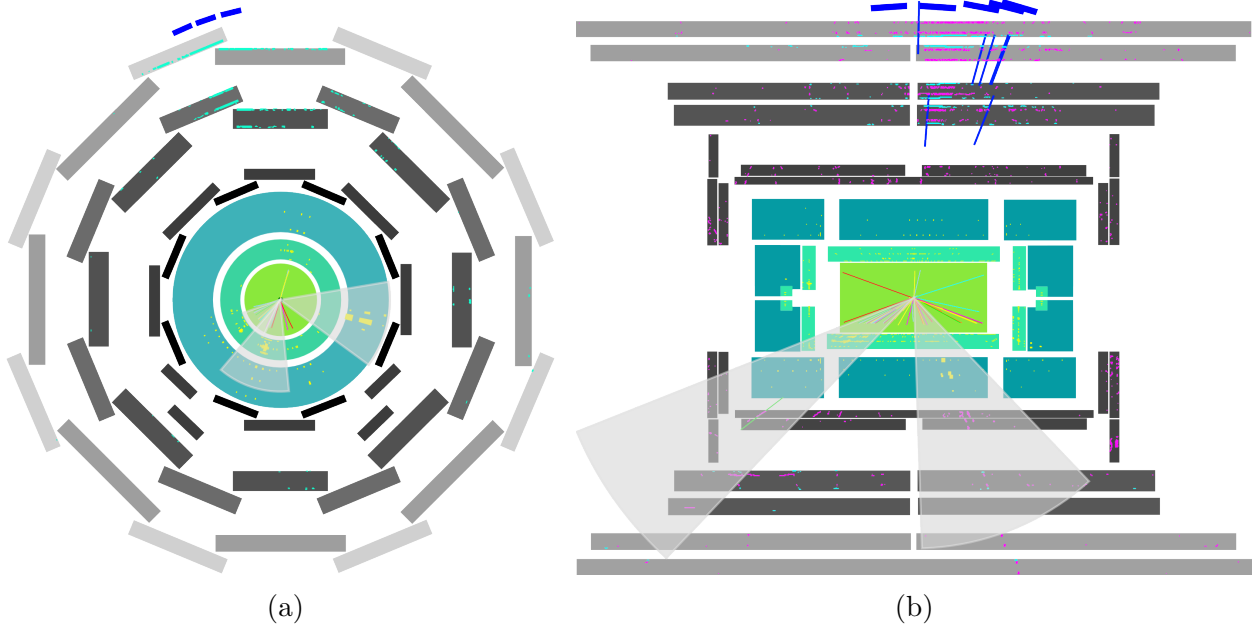


Figure 10.7: Event displays of a single, signal MC event from the $m_{\tilde{g}} = 800$ GeV sample. This event passes the full selection criteria for the single vertex plus two jets search. The displaced vertex is not shown, but a shower of hits shows the trajectory of a long-lived particle. (a) shows an xy -view, while (b) shows an Rz -view.

10.4.1 Systematic uncertainty on the ABCD method background estimation

The systematic uncertainty on the ABCD method is determined by varying the X - and Y -cuts in the validation region and determining the largest discrepancy between $N_A^{expected}$ and $N_A^{observed}$. The closest jet or track ΔR is varied by 0.05 up and down, producing a largest discrepancy of 35%. Additionally, the number of hits associated to the vertex is varied up and down by 100 hits, with a largest discrepancy of 23%.

It should be noted that while both these numbers are large, the statistical errors are such that none of the expected and observed events deviate from each other by more than 1σ . Thus, using 35% as a systematic uncertainty covers potential systematic issues, but is more

likely an overestimation due to statistical fluctuations. The full breakdown of events in the ABCD plane is given in Table 10.5.

Table 10.5: Event counts using varied x - and y -axis cuts in each of the four regions of the ABCD plane, and the expected number in region A, for both the validation and signal regions. The largest discrepancy between A and B*C/D is 35%, and is from the second variation, with cuts values $\min(\Delta R(\text{vertex}, \text{jet}/\text{trk})) > 0.35$ and $n\text{Hits} > 2000$.

2015+2016 data: varied X - and Y - cuts						
Closest jet or track ΔR vs. Vertex nMDT+nRPC hits						
Region	% difference	A	B*C/D	B	C	D
Region A cuts at $\min(\Delta R(\text{vertex}, \text{jet}/\text{trk})) > 0.25$ and $n\text{Hits} > 2000$						
VR: 1j50, 1j150	12%	44 ± 7	49 ± 5	6565 ± 81	90 ± 10	12200 ± 100
Region A cuts at $\min(\Delta R(\text{vertex}, \text{jet}/\text{trk})) > 0.35$ and $n\text{Hits} > 2000$						
VR: 1j50, 1j150	35%	24 ± 5	32 ± 4	6590 ± 80	60 ± 8	12200 ± 100
Region A cuts at $\min(\Delta R(\text{vertex}, \text{jet}/\text{trk})) > 0.3$ and $n\text{Hits} > 1900$						
VR: 1j50, 1j150	23%	33 ± 6	41 ± 5	7282 ± 85	64 ± 8	11500 ± 100
Region A cuts at $\min(\Delta R(\text{vertex}, \text{jet}/\text{trk})) > 0.3$ and $n\text{Hits} > 2100$						
VR: 1j50, 1j150	2%	30 ± 5	31 ± 4	5900 ± 80	67 ± 8	12900 ± 100

Chapter 11

RESULTS

11.1 *Expected number of signal events*

The simulated MC samples are generated only for a single proper lifetime (see Section 6.1.1). However, the analysis is sensitive to a wide range of proper lifetimes. Instead of generating MC events at many lifetimes, the results from the simulated MC sample are extrapolated to larger and smaller lifetimes. A toy MC using only the long-lived particle 4-momenta and reconstruction efficiencies is used to perform this extrapolation the expected number of signal events for each sample, for a range of proper lifetimes between 0 m and 300 m (this is the range between which a useful limit is set).

In order to determine if a criterion is passed based on an efficiency (*e.g.* will a LLP decay leave a good reconstructed vertex), a random probability between 0 and 1 is calculated. If that probability is less than the efficiency being considered, then the criterion is considered passed. If it is failed, the rest of the criteria are skipped and the toy moves onto the next event. A unique probability is determined for each criterion.

All efficiencies shown are for the Stealth SUSY sample $m_{\tilde{g}} = 800$ GeV, though this procedure applies to all Stealth SUSY samples. The steps corresponding to the two vertex topology also apply to the Higgs boson samples.

1. Starting with $c\tau = 0.0001$ m, generate a random decay position for each of the two LLPs, sampled from an exponential distribution, $f(t) = \exp -1/\beta\gamma c\tau$.
2. Calculate the physical decay position in the detector for each of the particles, using their stored 4-momenta.

3. Determine if the resulting decay position topology is a detectable one, thus worth keeping. These are classified into four categories:

- (a) Barrel MS - Barrel MS: Two particles decaying with $|\eta| < 0.8, 3 \text{ m} < R < 8 \text{ m}$.
- (b) Barrel MS - Endcap MS: One particle decaying with $|\eta| < 0.8, 3 \text{ m} < R < 8 \text{ m}$ and the other with $1.3 < |\eta| < 2.5, 5 \text{ m} < |z| < 15 \text{ m}, R < 10 \text{ m}$.
- (c) Endcap MS - Endcap MS: Two particles decaying with $1.3 < |\eta| < 2.5, 5 \text{ m} < |z| < 15 \text{ m}, R < 10 \text{ m}$.
- (d) Single barrel MS decay: One particle decaying with $|\eta| < 0.8, 3 \text{ m} < R < 8 \text{ m}$. This category includes both (a) and (b).

4. Each topology has a distinct set of reconstruction efficiencies and thus a slightly different procedure. They are outlined individually below:

- (a) Barrel MS - Barrel MS
 - i. Determine if the barrel decay with smaller Δt (not necessarily the decay with larger β - this also depends on how far the particle travels before decaying) is in time with respect to the level-1 barrel trigger response in data (see Appendix A). The probability of a decay passing based on its Δt delay with respect to a particle traveling at c is calculated from Equation 2 in Appendix A. The decay is considered “in time” if its randomly generated probability is greater than the efficiency for the RPC response at the particle’s Δt .
 - ii. If the event is in time, determine if it will pass the RoI cluster trigger by using the efficiency for the trigger dependent on both decay positions. Events in the numerator of this efficiency also must pass the global event criterion of ≤ 40 tracklets¹. For these events, the trigger efficiency is calculated from

1

This is computed *after* scale factors are applied




Figure 11.1: (a) Barrel RoI cluster trigger efficiency for barrel MS - barrel MS events. Since the RPC timing dependence is accounted for before this histogram is used, the timing dependence from MC is removed in the events in this histogram. The decays are ordered in terms of their radial decay position before entering the histogram - so the x -axis is the decay with the larger R , while the y -axis is the decay with the smaller R . (b) Barrel MS vertex reconstruction efficiency for triggered barrel MS - barrel MS events (*i.e* events in the numerator of (a)).

the subset of events where both decays are in the barrel MS. This efficiency is shown in Figure 11.1a.

- iii. If the event passes the trigger, the MS vertex reconstruction efficiency is considered for barrel MS decays in triggered barrel MS - barrel MS events (Fig. 11.1b). Each LLP is considered individually. If both decays “reconstruct” vertices, the event is counted as a good two vertex event.

(b) Barrel MS - Endcap MS

- i. Determine if the endcap decay is in time ($\Delta t < 25$ ns). If it is not, then see if the barrel decay is in time based on the efficiency for the RPC response.
- ii. If the event is in time, determine if it will pass the RoI cluster trigger and have ≤ 40 tracklets by using the efficiency for the trigger dependent on both decay positions. For these events, the trigger efficiency is calculated from the subset of events where one decay is in the barrel MS and the other is in the endcap MS. This efficiency is shown in Figure 11.2.




Figure 11.2: Barrel-Endcap RoI cluster trigger efficiency for barrel MS - endcap MS events. The x -axis is the decay position of the endcap decay, while the y -axis is the decay position of the barrel decay.

- iii. If the event passes the trigger, the MS vertex reconstruction efficiency is

considered for barrel MS decays in triggered barrel MS - endcap MS events (Fig. 11.3a) and for endcap MS decays in triggered barrel MS - endcap MS events (Fig. 11.3b). If both decays “reconstruct” vertices, the event is counted as a good two vertex event.




Figure 11.3: (a) Barrel MS vertex reconstruction efficiency for triggered barrel MS - endcap MS events. (b) Endcap MS vertex reconstruction efficiency for triggered barrel MS - endcap MS events.

(c) Endcap MS - Endcap MS

- i. Determine if the endcap decay with smaller Δt (not necessarily the decay with larger β - this also depends on how far the particle travels before decaying) is in time ($\Delta t < 25$ ns).
- ii. If the event is in time, determine if it will pass the RoI cluster trigger and have ≤ 40 tracklets by using the efficiency for the trigger dependent on both decay positions. For these events, the trigger efficiency is calculated from the subset of events where both decays are in the endcap MS. This efficiency is shown in Figure 11.4a.




Figure 11.4: (a) Endcap RoI cluster trigger efficiency for endcap MS - endcap MS events. The decays are ordered in terms of their longitudinal decay position before entering the histogram - so the x -axis is the decay with the larger $|z|$, while the y -axis is the decay with the smaller $|z|$. (b) Endcap MS vertex reconstruction efficiency for triggered endcap MS - endcap MS events.

- iii. If the event passes the trigger, the MS vertex reconstruction efficiency is considered for endcap MS decays in triggered endcap MS - endcap MS events (Fig. 11.4b). Each LLP is considered individually. If both decays “reconstruct” vertices, the event is counted as a good two vertex event.
- (d) One barrel MS vertex plus 2 Jets (ABCD method topology). This covers two types of truth decay topologies: events with a single decay in the muon spectrometer, with that decay in the barrel, and events with two MS decays, one in the MS barrel.
 - i. For events with a single decay in the muon spectrometer, first find if the LLP decays in time with the RPC trigger (as described in Step 4(a)).
 - ii. If the event is in time, determine if the event will pass the RoI cluster trigger and have ≤ 40 tracklets, using an efficiency made from the subset of events with a single MS decay in the barrel. This efficiency is shown in Figure 11.5a.






Figure 11.5: (a) Barrel Muon RoI Cluster trigger efficiency for events with one barrel MS decay. (b) Vertex reconstruction efficiency for triggered events with one barrel MS decay, for vertices matched to the trigger cluster.


- iii. If the event passes the trigger, the MS vertex reconstruction efficiency is considered for barrel MS decays in triggered barrel MS events (Fig. 11.5b).
- iv. For events with two MS decays (and at least one in the barrel), follow through the same procedure as the 2 vertex case to see if the event passes the trigger. The probability for reconstructing a vertex is also computed in the same way, but with efficiencies using the isolation criteria for vertices in the ABCD method, rather than the good vertex criteria. The vertex efficiencies are shown in Figure 11.6.
- v. If the event has a *single* good vertex, see if there are two jets with $E_T > 150$ GeV by examining the LLP that does not decay in the MS. The efficiency depends on whether this decay is in the barrel or endcap calorimeters. If LLP $|\eta| < 1.5$, the efficiency depends on L_{xy} and p_T . Otherwise, if the decay has $1.5 \leq |\eta| < 3.2$, the efficiency depends on L_z and p_z . The efficiencies are



(a)



(b)



(c)

Figure 11.6: (a) Barrel MS vertex reconstruction efficiency for triggered barrel MS - barrel MS events. (b) Barrel MS vertex reconstruction efficiency for triggered barrel MS - endcap MS events. (c) Endcap MS vertex reconstruction efficiency for triggered barrel MS - endcap MS events.

shown in Figure 11.7.

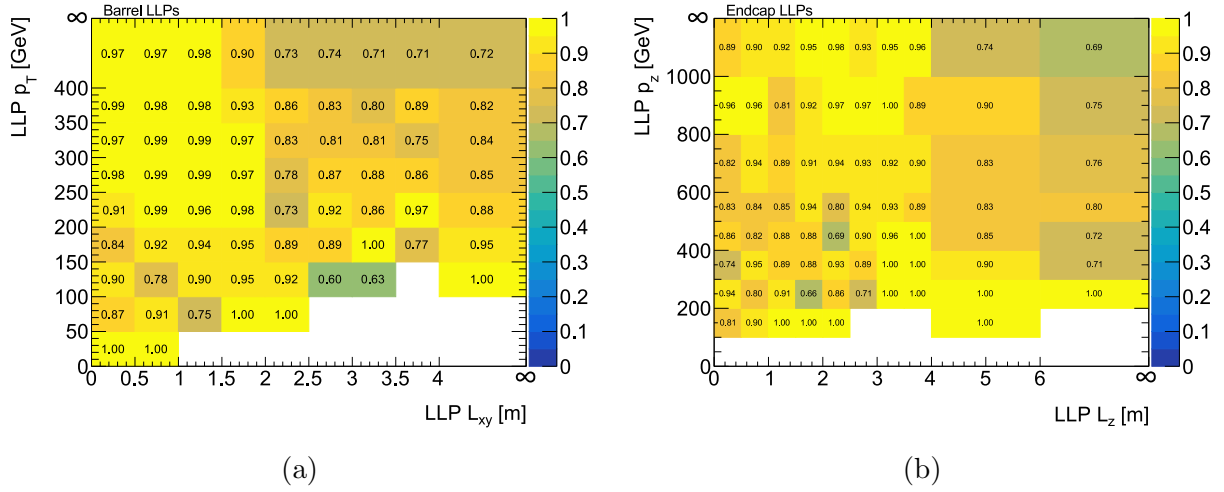


Figure 11.7: (a) Barrel Muon RoI Cluster trigger efficiency for events with one barrel MS decay. (b) Vertex reconstruction efficiency for triggered events with one barrel MS decay, for vertices matched to the trigger cluster.

- Steps 2 through 4 are repeated in proper lifetime increments of 0.05 m until $c\tau = 10.0001$ m, and then in 0.5 m until $c\tau = 300.0001$ m. The step size is changed because the scale on which the expected signal changes between $c\tau$ values becomes larger and a step size of 0.05 m is unnecessary.
- The final result is the expected signal events as a function of proper lifetime for each topology, for each search signature (two vertex, and one vertex plus two jets). These functions are then summed by search signature and rescaled by (production cross-section) \times (integrated luminosity)/(number of toy events) to obtain the number of expected signal events in $22.1fb^{-1}$ of data (for an example, see Figure 11.8).

The toy MC results are validated by comparing the expected signal at the generated lifetime of the official MC (in this case, 0.8 m) to the output of the official MC. The sigma

difference is then calculated for each topology. Individual differences are not required to be within some value, but it is ensured that the overall distribution between all topologies in all samples is normally distributed with a mean of approximately zero.

11.1.1 Inclusion of systematic uncertainties

The effect of all systematic and statistical uncertainties of the efficiency histograms are each considered individually. As an example, consider the barrel MS - endcap MS topology. There are four systematic uncertainties that must be taken into account: barrel RoI cluster trigger, barrel MS vertex, endcap RoI cluster trigger, and endcap MS vertex. For the $m_{\tilde{g}} = 800$ GeV sample these values are $\pm 1.0\%$, $+2.8\%$, -2.5% , $+1.4\%$, -1.3% , and $\pm 0.7\%$, respectively.

Each time a criterion is considered, the random probability is also compared against the criterion + systematic uncertainty, - systematic uncertainty, + statistical uncertainty, - statistical uncertainty. Only one uncertainty is applied at a time, so for example to count the number of expected BMS–EMS events with the barrel RoI cluster trigger systematic uncertainty up, all other efficiencies would be held at their central values.

However, if we were considering barrel MS - barrel MS events with the barrel MS vertex reconstruction efficiency systematic uncertainty up, that systematic uncertainty would be up for *both* decays.

At the end, the effect of the uncertainties on the total number of two vertex events are added in quadrature. There is a 2.9% systematic uncertainty on the integrated luminosity that is added at the limit setting step, and thus is not included at this stage [47].

The result of the toy MC is a point at each proper lifetime, which is subject to statistical uncertainties from the number of toy events. To present a smooth limit, the expected event histograms are all fit with a Novosibirsk function [48]. The Novosibirsk function is defined as follows:

$$f(x | x_0, \sigma, \tau) = \exp \left[-\frac{1}{2} \left\{ \frac{\ln^2 (1 + \Lambda \tau (x - x_0))}{2\tau^2} + \tau^2 \right\} \right], \quad (11.1)$$

$$\text{where } \Lambda = \frac{\sinh(\tau \sqrt{\ln 4})}{\sigma \tau \sqrt{\ln 4}}. \quad (11.2)$$

An example of a fit in the $m_{\tilde{g}} = 250$ GeV, two MS vertex event topology case is shown in Figure 11.8.




Figure 11.8: An example fit used in the signal extrapolation. The 2 MS vertex topology in the 250 GeV gluino sample is shown, with the smooth Novosibirsk fit on top of the raw output from the toy MC.

11.2 Expected number of signal events

The global efficiency as a function of proper lifetime for each Stealth SUSY benchmark sample are shown in Figure 11.10. Separate efficiencies are provided for the 2 MS vertex topology, the one MS vertex + 2 jets topology, and their sum. For any given lifetime, the global efficiency provides the fraction of events that are expected to pass the full analysis cutflow.

For the lower mass signal samples, both topologies provide a significant contribution to the overall efficiency. However, at higher masses, two processes interplay to make the 2 MS vertex topology less relevant. First, the higher gluino mass samples have a very high fraction of events passing the two jet requirement, since the two prompt gluons nearly always leave two jets with $E_T > 150$ GeV. Also, the global efficiency for the two vertex case decreases, since a larger fraction of events have > 40 tracklets and thus the vertex reconstruction algorithm is not run for computational reasons.

11.3 Expected and observed limits

No excess of events over the expected background is observed, and thus upper limits on the production of long-lived particles are derived. The CLs method [49] is used to derive upper limits for each channel. Additionally, the two channels are summed and upper limit on the combination of topologies are provided as well.

A profile likelihood is used as a test statistic and a frequentist calculator is used to generate toy data. The likelihood incorporates a Poisson probability term describing the total number of observed events. All the systematic uncertainties described in preceding sections are propagated to derive the corresponding uncertainties on the number of expected signal events. These uncertainties are then included as nuisance parameters through their effect on the mean of the Poisson functions and through convolution with their assumed Gaussian distributions. The number of expected events in signal MC, together with the expected background estimation and systematic uncertainties, are input for computing a CLs value. This CLs value represents the probability for the observed events to be compatible with the hypothesis of signal plus background.

For each signal benchmark, the limit is calculated for a chosen lifetime. It is then extrapolated to a range of lifetimes using the number of expected events as a function of the lifetime described in 11.1.

Limits are presented separately for each signal benchmark, for each of the three channels, in terms of the production cross-section for the studied Stealth SUSY process (see

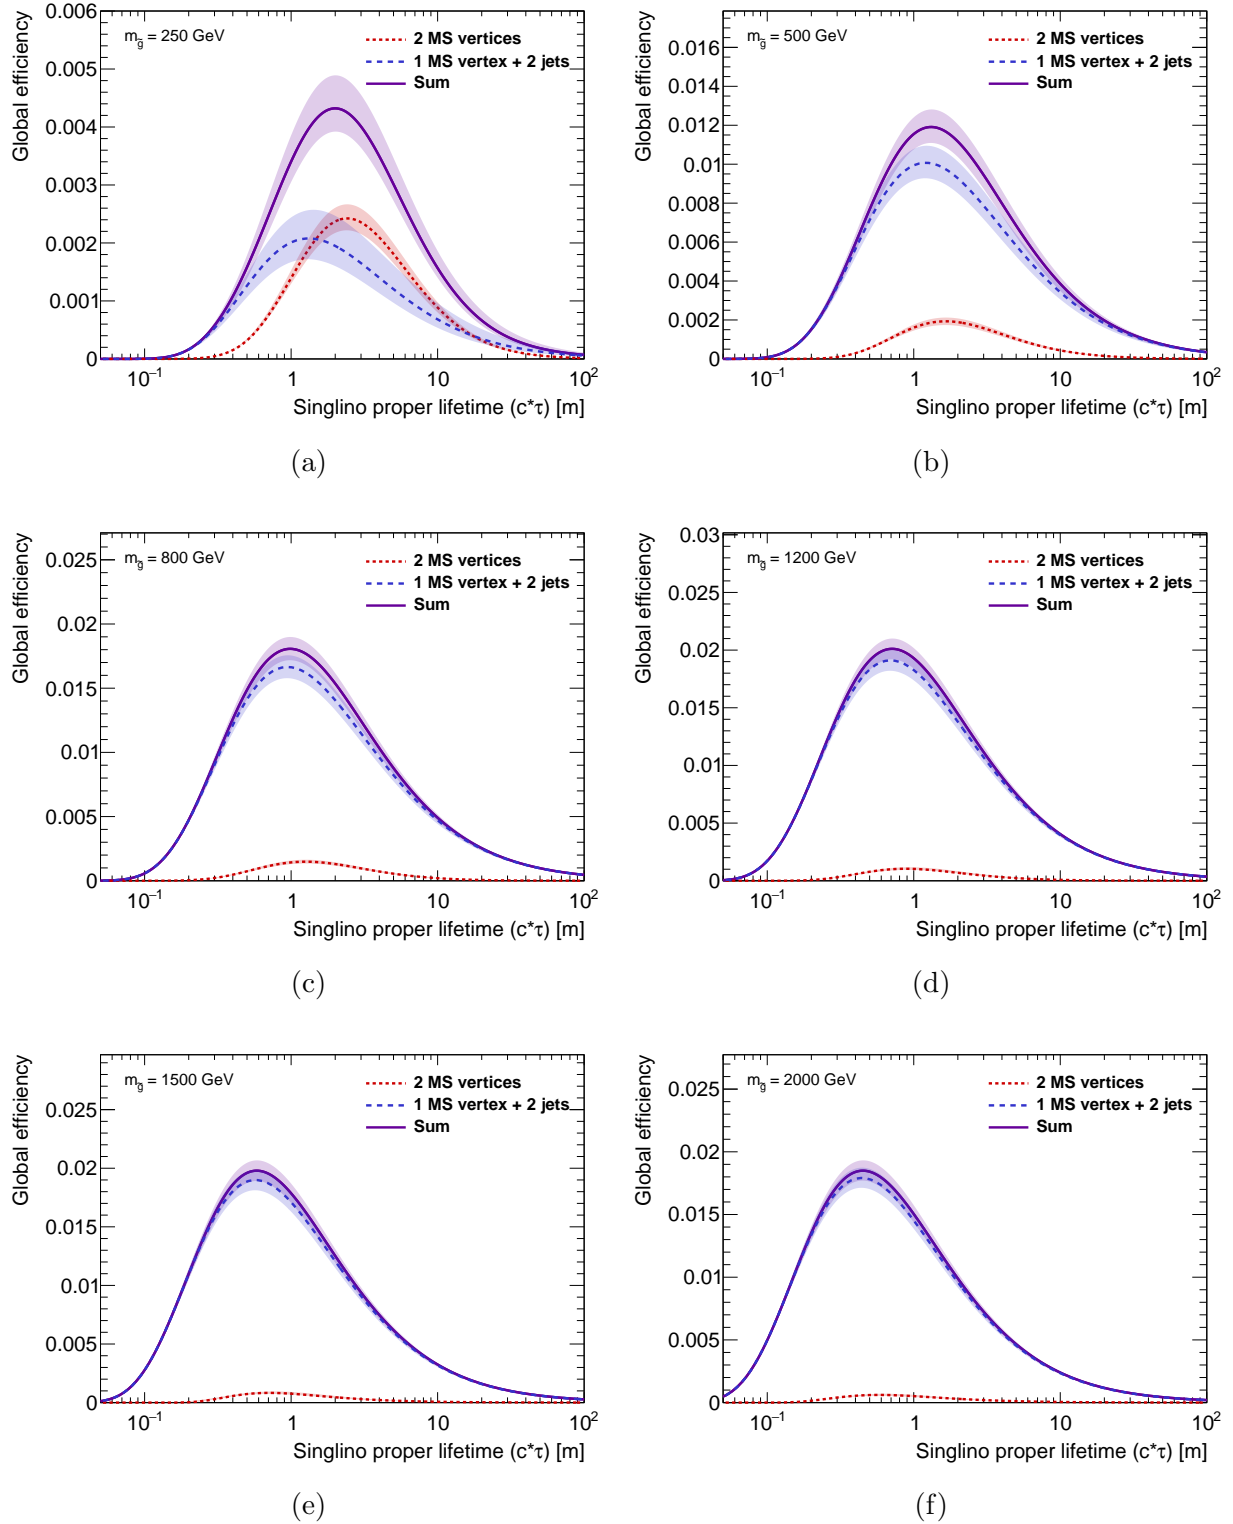


Figure 11.9: Expected number of signal events for each of the Stealth SUSY benchmark samples. At low masses, where a smaller fraction of events have two jets with $E_T > 150$ GeV, both the 2 MS Vertex topology and the 1 MS vertex + 2 jets topology have equal contributions to the global efficiency. At higher masses, the 2 MS vertex topology becomes nearly negligible.

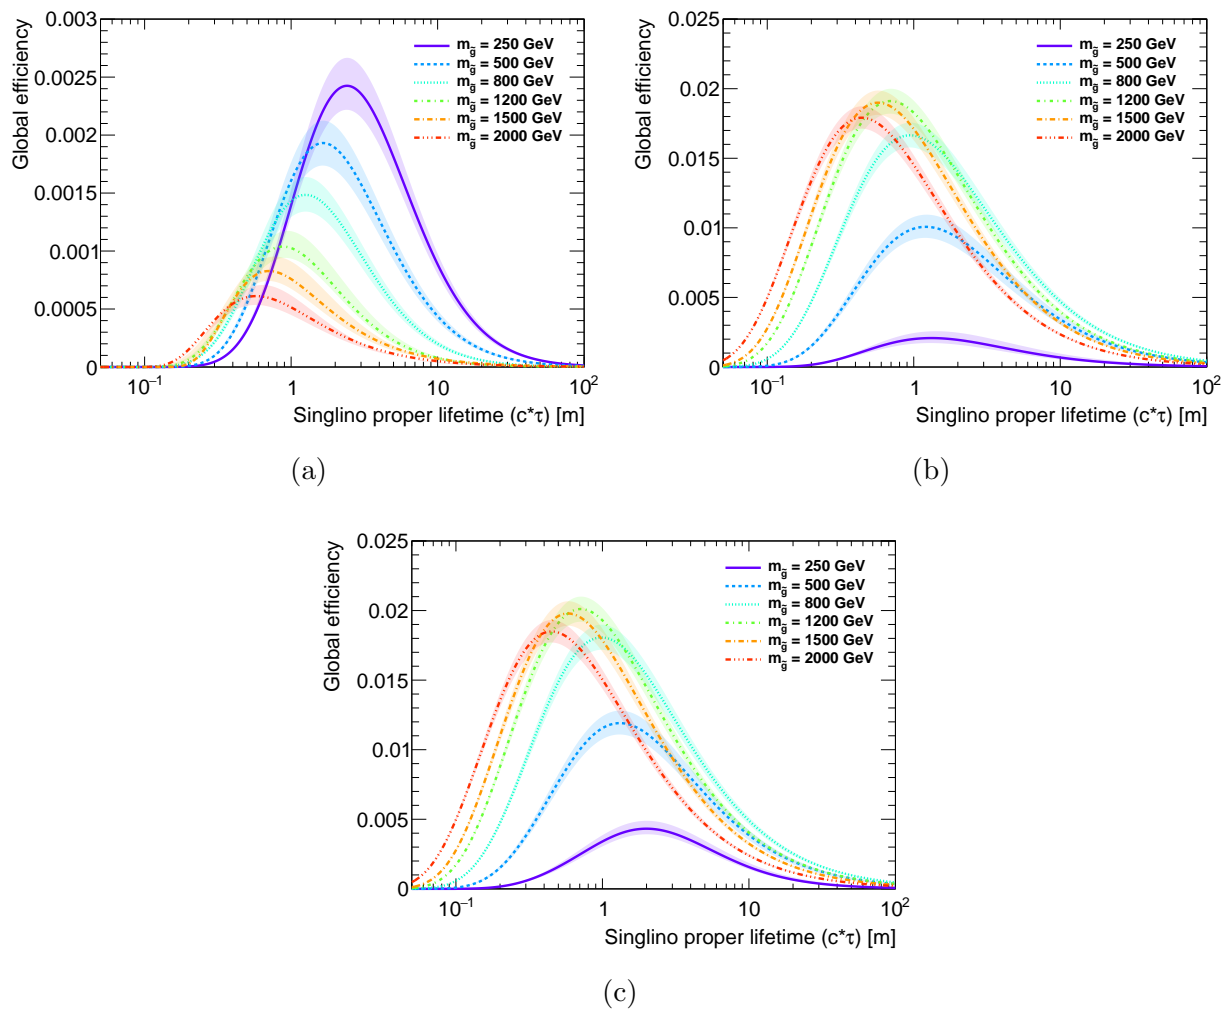


Figure 11.10: Expected number of signal events for each of the Stealth SUSY benchmark samples, for the two vertex channel, the one vertex plus two jets channel, and the sum of the two channels.

Section 2.3.2) divided by the theoretical production cross-section for pair-produced gluinos with squarks decoupled [6]. For these plots, in Figures 11.11 – 11.16, horizontal lines are provided at branching ratios for the Stealth SUSY process. When the observed limit falls below a branching ratio line, those proper lifetimes are excluded at 95% confidence level.




Figure 11.11: Expected and observed 95% CL exclusion limits on the production cross section for Stealth SUSY events with $m_{\tilde{g}} = 250$ GeV, as a function of the singlino proper lifetime. Limits for the two vertex search are shown in (a), (b) shows the one vertex plus two jets search, and (c) shows the combination of both topologies.

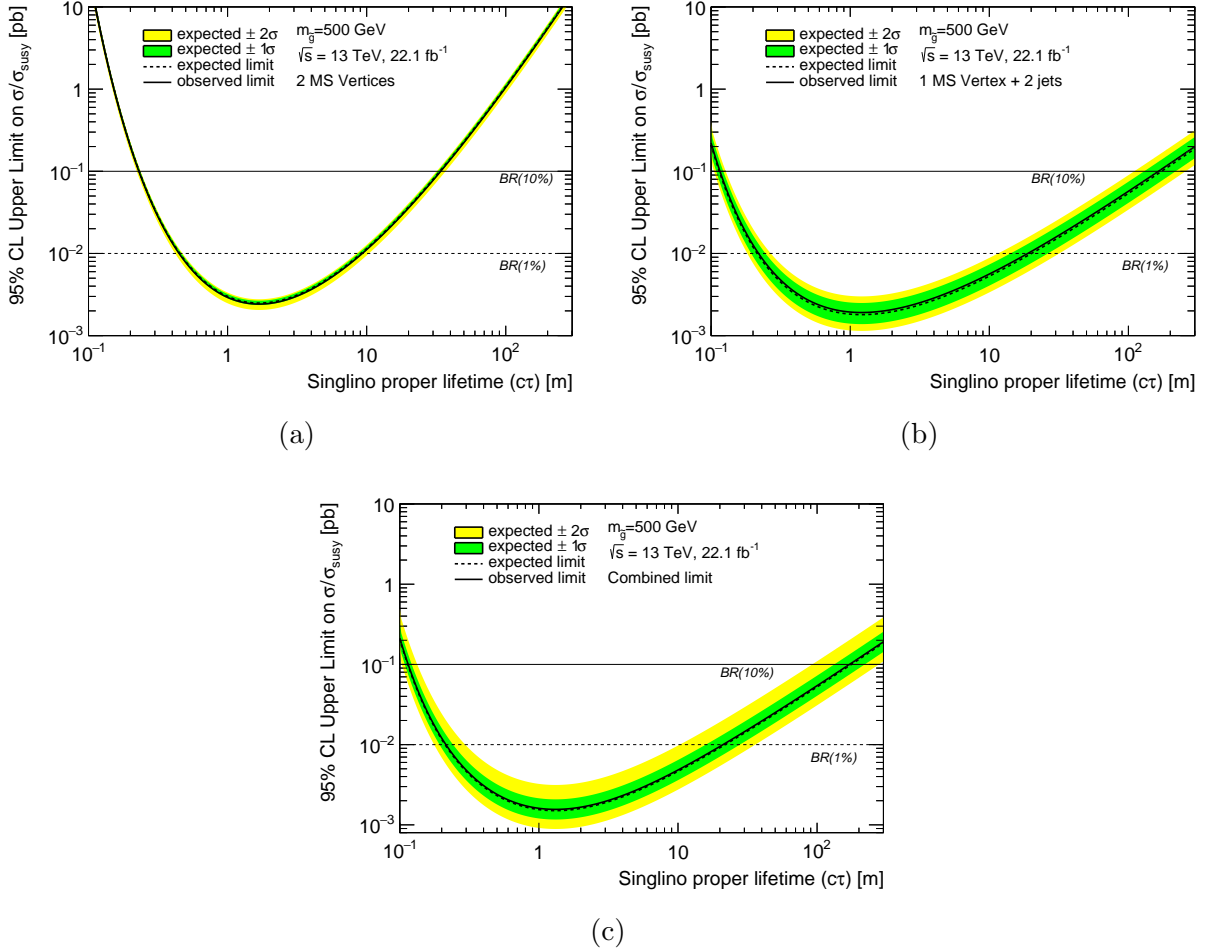


Figure 11.12: Expected and observed 95% CL exclusion limits on the production cross section for Stealth SUSY events with $m_{\tilde{g}} = 500$ GeV, as a function of the singlino proper lifetime. Limits for the two vertex search are shown in (a), (b) shows the one vertex plus two jets search, and (c) shows the combination of both topologies.

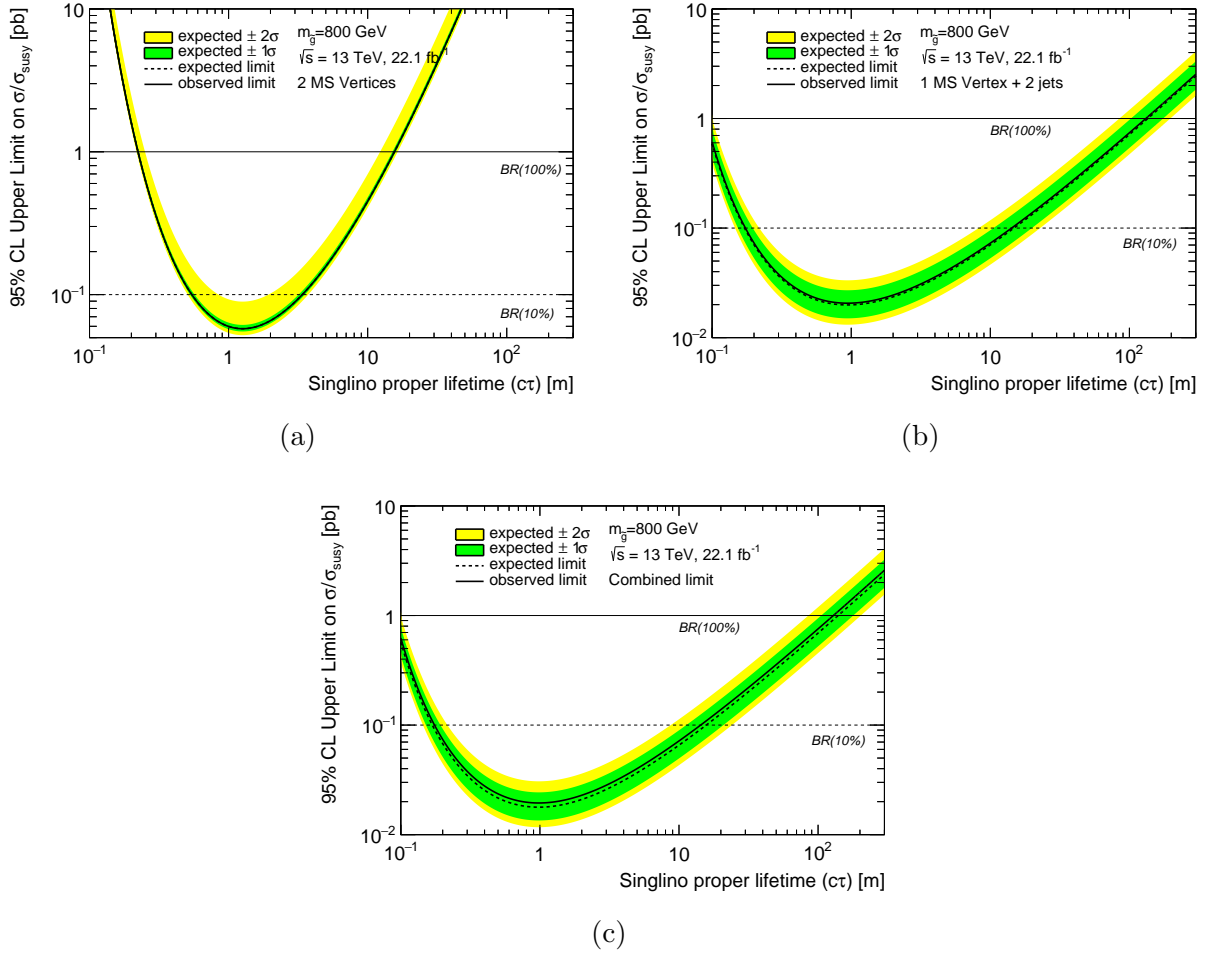


Figure 11.13: Expected and observed 95% CL exclusion limits on the production cross section for Stealth SUSY events with $m_{\tilde{g}} = 800$ GeV, as a function of the singlino proper lifetime. Limits for the two vertex search are shown in (a), (b) shows the one vertex plus two jets search, and (c) shows the combination of both topologies.




Figure 11.14: Expected and observed 95% CL exclusion limits on the production cross section for Stealth SUSY events with $m_{\tilde{g}} = 1200$ GeV, as a function of the singlino proper lifetime. Limits for the two vertex search are shown in (a), (b) shows the one vertex plus two jets search, and (c) shows the combination of both topologies.

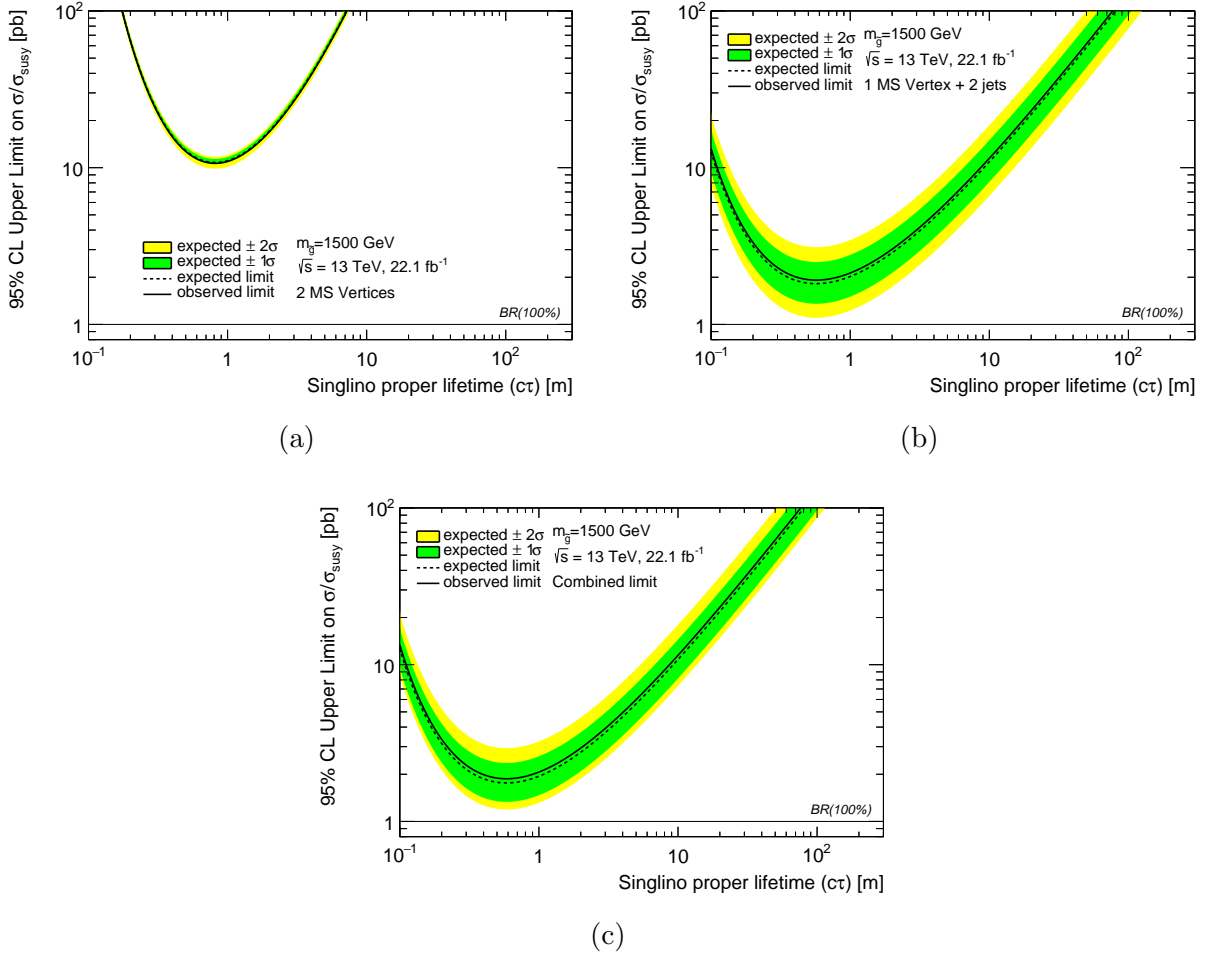




Figure 11.15: Expected and observed 95% CL exclusion limits on the production cross section for Stealth SUSY events with $m_{\tilde{g}} = 1500$ GeV, as a function of the singlino proper lifetime. Limits for the two vertex search are shown in (a), (b) shows the one vertex plus two jets search, and (c) shows the combination of both topologies.



(a)

(b)



(c)

Figure 11.16: Expected and observed 95% CL exclusion limits on the production cross section for Stealth SUSY events with $m_{\tilde{g}} = 2000$ GeV, as a function of the singlino proper lifetime. Limits for the two vertex search are shown in (a), (b) shows the one vertex plus two jets search, and (c) shows the combination of both topologies.

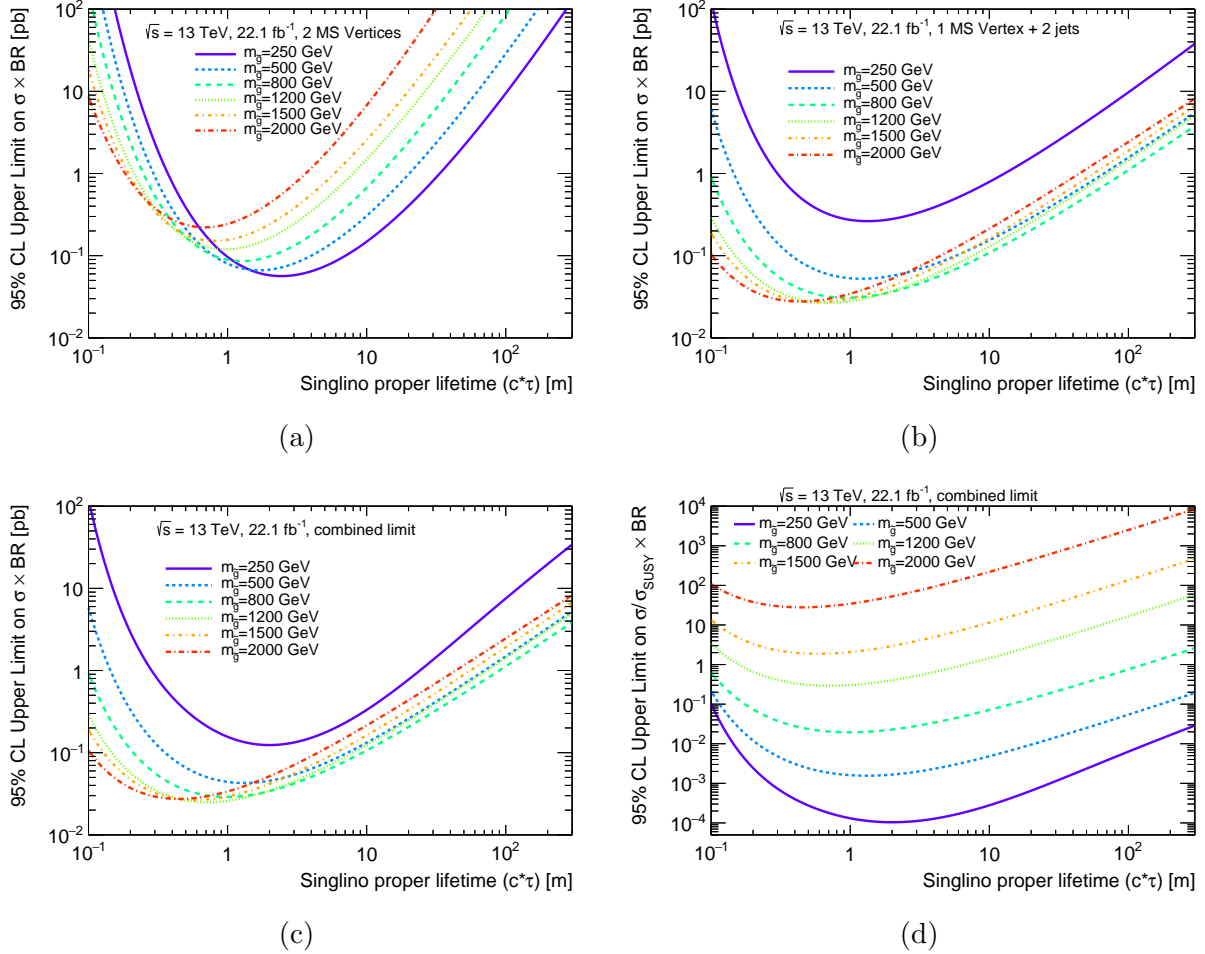


Figure 11.17: (a) – (c) Observed 95% CL exclusion limits on the production cross section for Stealth SUSY events, as a function of the singlino proper lifetime, for six gluino masses. Limits only considering the two vertex topology are shown in (a), the one vertex plus two jets topology in (b), and the combination of the two channels in (c). The same limits as (c) are shown in (d), but each limit is divided by the gluino production cross section at 13TeV for the corresponding gluino mass, and thus the y -axis shows the observed 95% CL exclusion limit on the branching ratio for $\tilde{g} \rightarrow \tilde{S}g, \tilde{S} \rightarrow \tilde{G}S, S \rightarrow gg$.

11.4 Application of two vertex search results to Higgs boson decays

In Section 2.3, it was mentioned that the results of a search for long-lived, neutral particles from Higgs boson decays would be presented. Unlike the Stealth SUSY model, SM Higgs boson decays are not generally produced along with two, high-energy jets², except in the rare case of initial- or final-state radiation. Thus, only the two-vertex search topology is relevant for Higgs boson decays.




Figure 11.18: Expected number of $H \rightarrow ss$ signal events produced in 22.1 fb^{-1} , for six scalar masses.

The full search is carried out in an identical manner to the Stealth SUSY search, using the six scalar masses listed in Table 6.2. Results, however, are presented slightly differently, since the Higgs boson has a well-established production cross-section of 48.58 pb . Figure 11.18 shows the number of expected events in the full 22.1 fb^{-1} dataset, divided by the branching

The model studied here only addresses Higgs bosons produced with the gluon-gluon fusion process, which has the highest production cross-section. Other production mechanisms, such as vector-boson fusion, could produce high enough energy jets for events to reach the ABCD search signal region.

ratio for the Higgs boson decaying to two long-lived, neutral scalars. Effectively, this means that if every single Higgs boson decayed to scalar pairs, the number of reconstructed events in 22.1 fb^{-1} would be given by the curves in Figure 11.18. A more realistic example is that the branching ratio for $H \rightarrow ss$ is 5%. In this case, the expected number of reconstructed events for each proper lifetime would be found by multiplying the value of the curve by 0.05.

Table 11.1: Ranges of long-lived scalar proper decay lengths excluded at 95% CL assuming a 10%, or 1% BR, for $m_H = 125 \text{ GeV}$.

m_s [GeV]	Excluded $c\tau$ range [m]	
	1% BR	10% BR
5	0.14–0.76	0.059–3.93
8	0.23–1.19	0.092–6.04
15	0.30–4.68	0.15–18.6
25	0.47–10.2	0.25–37.7
40	0.96–15.1	0.50–59.5
55	2.75–21.0	1.30–92.6

Since there was no excess of events over the background expectation in the two-vertex topology, limits are set on $BR[H \rightarrow ss] = \sigma[H \rightarrow ss]/\sigma_{SM}$, for each of the six scalar mass samples. Figure 11.19 shows the observed limits on each of the six samples in a comparable way. From this figure it is clear that the boost of the scalars has a large effect on the range of excluded proper lifetimes at a given branching ratio. However, the *size* of the excluded range is much less variable, which shows the robustness and power of the two-vertex search over a large kinematic range. Individual exclusion limit plots, with 1σ and 2σ error bands, are shown in Figures 11.20a–11.20f. The range of excluded proper lifetimes for a 1% and 10% branching ratio for $H \rightarrow ss$ are shown in Table 11.1.

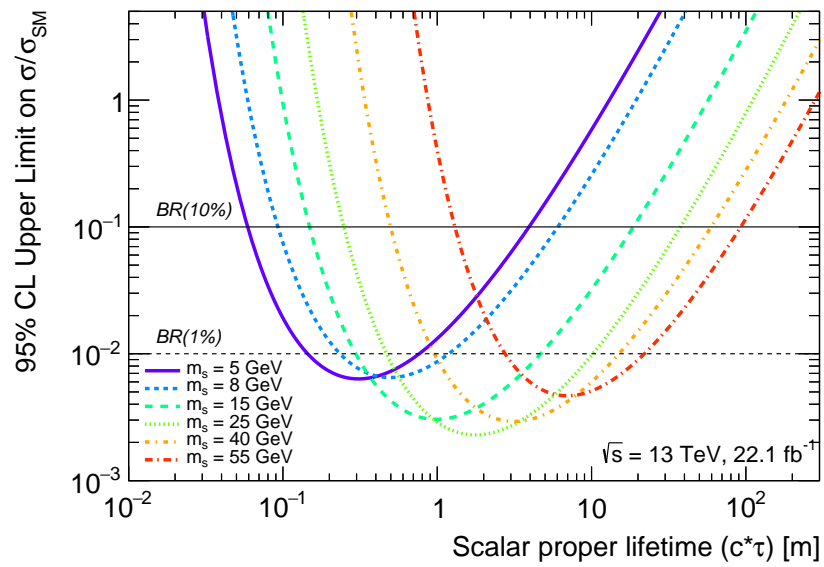


Figure 11.19: Observed 95% CL exclusion limits on the branching ratio for $H \rightarrow ss$, as a function of the scalar proper lifetime, for six scalar masses.




Figure 11.20: Expected and observed 95% CL exclusion limits on the branching ratio for $H \rightarrow ss$, as a function of the scalar proper lifetime, for six scalar masses.

Chapter 12

SUMMARY

This thesis presented the results of two searches using displaced muon spectrometer vertices: one for events with a single displaced MS vertex and two, high-energy jets, and the other for events with two displaced MS vertices. Data from proton–proton collisions, collected with the ATLAS detector at the LHC, was used for these searches. The data totaled 22.1 fb^{-1} and included the full set of data collected in 2015 in addition to 2016 data collected until 17 August 2016. The search results were presented in the context of two models: Stealth Supersymmetry, and Higgs boson decays to long-lived scalars.

The two search topologies were not applied to both models. Although the sensitivity of a search with a single displaced vertex can have greatly increased sensitivity for all lifetimes, such a search is only useful if a model exhibits additional characteristics that distinguish it from background events. The single vertex analysis presented in this thesis used the two prompt jets generally present in Stealth Supersymmetry events to provide the additional handle needed for background discrimination. However, since Higgs boson decays to long-lived particles are not generally associated with prompt activity this model was only studied in the context of two-vertex events.

No excess of events was found in either topology, and thus 95% confidence limits were set for both each topology separately, and their combination (for Stealth Supersymmetry only). While the sensitivity to lower gluino mass Stealth Supersymmetry benchmark models was very high and the full range of proper lifetimes studied was excluded, there was no sensitivity to masses above 1200 GeV. Future studies with the full 2016 dataset are necessary to extend sensitivity to higher masses. Additionally, for the powerful single vertex plus two jets topology, only events with an MS vertex in the barrel were studied. With careful

examination of background sources, the search could be extended to vertices in the endcaps as well.

For Standard Model Higgs boson decay, the path for future analyses is clear: if a suitable topology requiring only a single MS vertex can be found, the search sensitivity can be greatly increased. For these decays, additional incorporation of searches for invisible Higgs boson decays or prompt decays to four b -quarks ($H \rightarrow ss, s \rightarrow b\bar{b}$) should also be investigated to understand the full ATLAS sensitivity.

In addition to broadening the search breadth, there are some changes that should be considered to improve the overall analysis:

- In its current state, the MS vertex reconstruction algorithm does not run if there are more than 40 reconstructed tracklets in an event. While this limit was rarely reached for Higgs boson decays at $\sqrt{s} = 8$ TeV, both the increase in center of mass energy and the addition of models with higher-multiplicity decays means that the 40 tracklet limit is reached on a regular basis. The algorithm should be tuned such that it can handle high-multiplicity decays within the ATLAS reconstruction chain and not affect the search sensitivity.
- The maximum allowable number of MDT hits associated to an MS vertex, as a part of the good vertex criteria, is currently 3000 hits. In a similar manner to the previous item, this cut was originally established at 3000 hits because it eliminated little to no signal. In principle, the cut could simply be increased to not reduce the efficiency of any signal model. However, it has an important function in preventing the acceptance of noise bursts as good MS vertices. The number of hits generally seen in such events in 13 TeV data needs to be carefully studied, such that either this 3000 hit limit can be increased on its own, or alongside the introduction of an additional criteria that removes noise bursts.
- As described in Section 8.1.1, the barrel–endcap overlap region of the muon spectrom-

eter is currently excluded from the analysis. Recovering efficiency in this large angular region would boost efficiency. However, this would only be feasible if the algorithm were to properly incorporate tracklets from both barrel and endcap chambers. Otherwise, the efficiency is so low that losses from eliminating the region are minimal. A proper treatment of vertices in this region would require rewriting the vertex reconstruction such that clusters could contain both endcap and barrel tracklets, and vertex reconstruction could use both sets of tracklets. This would not be a simple change, but should be a longer-term goal.

- The data–MC scale factors associated to MS tracklet reconstruction, as determined with QCD multi-jet events, are 0.55 in the barrel and 0.63 in the endcaps. For decays with low-tracklet multiplicity (e.g. Higgs boson decays to 5 GeV LLPs), this scale factor hugely affects the global signal acceptance. While on one hand, the scale factors are legitimate (they are derived from equivalent sets of events in data and MC), the underlying reason for it is currently not understood. With careful study, the cause could be determined and its effects mitigated from within the tracklet reconstruction algorithm. The most likely scenario is that the tracklet reconstruction is being affected by cavern background, but even if this is verified to be the case, understanding where tracklet reconstruction fails could provide large increases in tracklet reconstruction efficiency.

BIBLIOGRAPHY

- [1] Standard model of elementary particles. https://commons.wikimedia.org/wiki/File:Standard_Model_of_Elementary_Particles.svg.
- [2] Sayipjamal Dulat, Tie-Jiun Hou, Jun Gao, Marco Guzzi, Joey Huston, Pavel Nadolsky, Jon Pumplin, Carl Schmidt, Daniel Stump, and C. P. Yuan. New parton distribution functions from a global analysis of quantum chromodynamics. *Phys. Rev.*, D93(3):033006, 2016.
- [3] Cinzia De Melis. The CERN accelerator complex. Complexe des accrateurs du CERN. Jul 2016. General Photo.
- [4] ATLAS Collaboration. The ATLAS Experiment at the CERN Large Hadron Collider. *JINST*, 3:S08003, 2008.
- [5] ATLAS Collaboration. Track Reconstruction Performance of the ATLAS Inner Detector at $\sqrt{s} = 13$ TeV. 2015.
- [6] Christoph Borschensky, Michael Krmer, Anna Kulesza, Michelangelo Mangano, Sanjay Padhi, Tilman Plehn, and Xavier Portell. Squark and gluino production cross sections in pp collisions at $\sqrt{s} = 13, 14, 33$ and 100 TeV. *Eur. Phys. J.*, C74(12):3174, 2014.
- [7] ATLAS Collaboration. Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Phys. Lett.*, B716:1–29, 2012.
- [8] CMS Collaboration. Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Phys. Lett.*, B716:30–61, 2012.
- [9] Peter W. Higgs. Broken symmetries and the masses of gauge bosons. *Phys. Rev. Lett.*, 13:508–509, Oct 1964.
- [10] LHCb Collaboration. Observation of $J/\psi\phi$ structures consistent with exotic states from amplitude analysis of $B^+ \rightarrow J/\psi\phi K^+$ decays. 2016.
- [11] LHCb Collaboration. Amplitude analysis of $B^+ \rightarrow J/\psi\phi K^+$ decays. 2016.

- [12] BESIII Collaboration. Observation of a charged charmoniumlike structure in $e^+e^- \rightarrow \pi^+\pi^-j/\psi$ at $\sqrt{s}=4.26$ GeV. *Phys. Rev. Lett.*, 110:252001, Jun 2013.
- [13] Belle Collaboration. Study of $e^+e^- \rightarrow \pi^+\pi^-j/\psi$ and observation of a charged charmoniumlike state at belle. *Phys. Rev. Lett.*, 110:252002, Jun 2013.
- [14] LHCb Collaboration. Observation of $J/\psi p$ Resonances Consistent with Pentaquark States in $\Lambda_b^0 \rightarrow J/\psi K^- p$ Decays. *Phys. Rev. Lett.*, 115:072001, 2015.
- [15] David Curtin et al. Exotic decays of the 125 GeV Higgs boson. *Phys. Rev. D*, 90(7):075004, 2014.
- [16] Matthew J. Strassler and Kathryn M. Zurek. Echoes of a hidden valley at hadron colliders. *Phys. Lett. B*, 651:374, 2007.
- [17] Matthew J. Strassler and Kathryn M. Zurek. Discovering the Higgs through highly-displaced vertices. *Phys. Lett. B*, 661:263, 2008.
- [18] Matthew J. Strassler. Possible effects of a hidden valley on supersymmetric phenomenology. 2006.
- [19] David Curtin and Christopher B. Verhaaren. Discovering Uncolored Naturalness in Exotic Higgs Decays. *JHEP*, 12:072, 2015.
- [20] Zackaria Chacko, David Curtin, and Christopher B. Verhaaren. A Quirky Probe of Neutral Naturalness. *Phys. Rev.*, D94(1):011504, 2016.
- [21] Yanou Cui and Brian Shuve. Probing Baryogenesis with Displaced Vertices at the LHC. *J. High Energy Phys.*, 02:049, 2015.
- [22] ATLAS Collaboration. Summary of the ATLAS experiment's sensitivity to supersymmetry after LHC Run 1, interpreted in the phenomenological MSSM. *JHEP*, 10:134, 2015.
- [23] Stephen P. Martin. A Supersymmetry primer. 1997. [Adv. Ser. Direct. High Energy Phys.18,1(1998)].
- [24] JiJi Fan, Matthew Reece, and Joshua T. Ruderman. Stealth Supersymmetry. *J. High Energy Phys.*, 11:012, 2011.
- [25] JiJi Fan, Matthew Reece, and Joshua T. Ruderman. A Stealth Supersymmetry Sampler. *J. High Energy Phys.*, 07:196, 2012.

- [26] Lyndon Evans and Philip Bryant. LHC Machine. *JINST*, 3:S08001, 2008.
- [27] R Achenbach et. al. The atlas level-1 calorimeter trigger. *Journal of Instrumentation*, 3(03):P03001, 2008.
- [28] Yu Nakahama. The atlas trigger system: Ready for run-2. *Journal of Physics: Conference Series*, 664(8):082037, 2015.
- [29] Carlo Schiavi. ATLAS High-Level Trigger algorithms for Run-2 data-taking. Technical Report ATL-DAQ-PROC-2015-018, CERN, Geneva, May 2015. Not published in the proceedings.
- [30] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. The Anti-k(t) jet clustering algorithm. *J. High Energy Phys.*, 04:063, 2008.
- [31] ATLAS Collaboration. Triggers for displaced decays of long-lived neutral particles in the ATLAS detector. *JINST*, 8:P07015, 2013.
- [32] Johan Alwall, Michel Herquet, Fabio Maltoni, Olivier Mattelaer, and Tim Stelzer. MadGraph 5 : Going Beyond. *J. High Energy Phys.*, 06:128, 2011.
- [33] Richard D. Ball et al. Parton distributions with LHC data. *Nucl. Phys.*, B867:244–289, 2013.
- [34] Torbjorn Sjostrand, Stephen Mrenna, and Peter Z. Skands. A Brief Introduction to PYTHIA 8.1. *Comput. Phys. Commun.*, 178:852, 2008.
- [35] ATLAS Run 1 Pythia8 tunes. Technical Report ATL-PHYS-PUB-2014-021, CERN, Geneva, Nov 2014.
- [36] D. J. Lange. The EvtGen particle decay simulation package. *Nucl. Instrum. Meth.*, A462:152–155, 2001.
- [37] LHC Higgs Cross Section Working Group, S. Heinemeyer *et al.* Handbook of LHC Higgs Cross Sections: 3. Higgs Properties. 2013.
- [38] Andreas Salzburger. Optimisation of the ATLAS Track Reconstruction Software for Run-2. *J. Phys. Conf. Ser.*, 664(7):072042, 2015.
- [39] R. Fruhwirth. Application of Kalman filtering to track and vertex fitting. *Nucl. Instrum. Meth.*, A262:444–450, 1987.

- [40] ATLAS Collaboration. Jet energy measurement with the ATLAS detector in proton-proton collisions at $\sqrt{s} = 7$ TeV. *Eur. Phys. J.*, C73(3):2304, 2013.
- [41] ATLAS collaboration. Selection of jets produced in 13TeV proton-proton collisions with the ATLAS detector. 2015.
- [42] ATLAS collaboration. Tagging and suppression of pileup jets. 2014.
- [43] ATLAS Collaboration. Standalone vertex finding in the ATLAS muon spectrometer. *JINST*, 9:P02001, 2014.
- [44] M J Woudstra and Frank L Linde. *Precision of the ATLAS muon spectrometer*. PhD thesis, Amsterdam U., Geneva, 2002. Presented on 4 Dec 2002.
- [45] ATLAS Collaboration. Search for long-lived, weakly interacting particles that decay to displaced hadronic jets in proton-proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector. *Phys. Rev.*, D92(1):012010, 2015.
- [46] Andrea Coccaro, David Curtin, H. J. Lubatti, Heather Russell, and Jessie Shelton. Data-driven Model-independent Searches for Long-lived Particles at the LHC. *Accepted to Phys. Rev. D*, 2016.
- [47] ATLAS Collaboration. Improved luminosity determination in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector at the LHC. *Eur. Phys. J. C*, 73(8):2518, 2013.
- [48] H. Ikeda et al. A detailed test of the CsI(Tl) calorimeter for BELLE with photon beams of energy between 20-MeV and 5.4-GeV. *Nucl. Instrum. Meth.*, A441:401–426, 2000.
- [49] Alexander L. Read. Presentation of search results: The CL(s) technique. *J. Phys. G*, 28:2693, 2002.

Appendix A

TIMING OF THE LEVEL-1 BARREL MUON TRIGGER

The barrel timing response of RPCs is not modelled correctly in the ATLAS MC. Slow particles ($\beta < 1$) can arrive inside the MC timing window but outside of the actual window in data, which results in artificially high efficiencies. To correct for this effect, a scale factor is applied to events that pass the trigger. The scale function is determined by fitting a Gaussian to the window of acceptance for the L1 RPC trigger from both data and MC, as show in Figure A.1. The parameters of the corresponding gaussian fits, rescaled to ns, are:

$$\mu_{data} = (94.079 \pm 0.001) \text{ ns}, \quad (\text{A.1})$$

$$\sigma_{data} = (2.9548 \pm 0.0004) \text{ ns}, \quad (\text{A.2})$$

$$\mu_{MC} = (92.998 \pm 0.004) \text{ ns}, \quad (\text{A.3})$$

$$\sigma_{MC} = (2.084 \pm 0.002) \text{ ns}. \quad (\text{A.4})$$

The actual efficiency of the trigger, based on arrival time difference alone (i.e. if we assume the trigger were 100% efficient if the hits arrived with a 0 ns time delay), is given by the following equation:

$$\epsilon_{\text{trigger}}^{\text{data,MC}} = \frac{1}{2} \left(1 - \text{erf} \frac{\Delta t - \Delta_{\text{data,MC}}}{\sqrt{2}\sigma_{\text{data,MC}}} \right). \quad (\text{A.5})$$

The efficiency depends on Δt , or how late a particle arrives at the last RPC trigger plane ($L_{RPC} = 10 \text{ m}$) as compared to a particle travelling at c . A long-lived particle will travel at speed β for a distance d at angle θ before decaying into particles with a speed negligibly




Figure A.1: Arrival time of hits accepted by the L1 RPC trigger in (a) data and (b) MC. 1 tick corresponds to 3.125 ns. The distributions are fitted with a gaussian distribution. Mean and standard deviation values, and their respective uncertainties, are employed for scaling the trigger efficiency on simulation and for assigning a systematic uncertainty on the scaling procedure.

different from c . Thus,

$$\Delta t = t_{\text{LLP}} - t_0 = \frac{d}{\beta c} + \frac{1}{c} \left(\frac{L_{\text{RPC}}}{\sin \theta} - d \right) - \frac{L_{\text{RPC}}}{c \sin \theta} \quad (\text{A.6})$$

$$= \frac{d}{c} \left(\frac{1}{\beta} - 1 \right). \quad (\text{A.7})$$

The case where $d/\sin \theta > 10$ m does not come into play, since we only consider particles decaying radially before 8 m in the barrel.

The parameter Δ is a measure of the width of the readout window in (data,MC):

$$\Delta_{\text{data,MC}} = t_{\text{end of data,MC readout window}} - \mu_{\text{data,MC}} \quad (\text{A.8})$$

$$\Delta_{\text{data}} = (34 * 3.125 - 94.079) \text{ ns} = 12.171 \text{ ns} \quad (\text{A.9})$$

$$\Delta_{\text{MC}} = (34 * 3.125 - 92.998) \text{ ns} = 13.262 \text{ ns}. \quad (\text{A.10})$$

Each particle that passes the RoI cluster trigger in MC is weighted by the ratio $\rho = \epsilon^{\text{data}}/\epsilon^{\text{MC}}$ when the trigger efficiency is determined. Figure A.2 shows the two efficiencies for data and MC.




Figure A.2: The scale factor applied as a weight to events passing the RoI cluster trigger, as a function of the time delay as compared to a particle travelling at c .

In principle, this scaling should be computed on a run-by-run basis and only runs with the same timing profile in data should be added. Then, MC events would be scaled differently proportional to the amount of data collected with each timing profile. However, the timing scaling had negligible effect on the overall efficiencies for the signal MC samples studied in this thesis and small run-to-run deviations would not impact the final result. Any signal MC samples with a slower β profile would need to be more carefully treated.

Appendix B

THE QUIRKS OF MS VERTEX RECONSTRUCTION

The MS tracklet and vertex reconstruction algorithms were scrutinised very closely during the course of writing Section 7.2. Consequently, several bugs were found, and are documented in this appendix.

B.1 Three-hit segment seed spacing

The first step in tracklet formation is to create three-hit seeds from MDT hits in individual multilayers. This is done by determining all combinations of hits that satisfy the following criteria (as described in Section 7.2.1):

1. $N_{MDT3} > N_{MDT2} > N_{MDT1}$, where N_{MDTi} is the number of the i -th MDT tube following the scheme outlined in Figure 7.3.
2. $|z_{MDT1} - z_{MDT2}| < 100$ mm ($|R_{MDT1} - R_{MDT2}| < 100$ mm in endcaps)
3. $|z_{MDT1} - z_{MDT3}| < 160$ mm ($|R_{MDT1} - R_{MDT3}| < 160$ mm in endcaps)
4. Hits 1 and 2 must be in the same layers or sequential layers, and hit 3 must be in the same layer or one higher than hit 2, and not in the same layer as hit 1 (*i.e.* seed is formed from hits in at least two layers, with no skipped layer between hits).

The goal of these criteria is to create seeds from neighboring hits in adjacent or subsequent layers, such as those shown in Figure B.1.

However, there are two problems with the seed formation. The first is the distance requirements between MDT hits. Based on descriptions and comments in the code, good




Figure B.1: Examples of three-hit seeds that the algorithm should select.

seeds should have hits in mostly adjacent tubes. But when the hits are allowed to be 100 mm or 160 mm apart, they can be separated by *many* empty tubes. Some of these valid combinations are shown in Figure B.2.




Figure B.2: Examples of three-hit seeds that the algorithm selects but should reject as bad.

The most likely explanation for the selection of these seeds is that the original author mistook the diameter for the radius of an MDT tube. MDT tubes have a diameter of 30 mm, which means that many tubes can fit between hits 160 mm apart. The desired outcome is achieved by cutting the values in half, such that:

- $|z_{MDT1} - z_{MDT2}| < 50$ mm ($|R_{MDT1} - R_{MDT2}| < 50$ mm in endcaps)
- $|z_{MDT1} - z_{MDT3}| < 80$ mm ($|R_{MDT1} - R_{MDT3}| < 80$ mm in endcaps).

The remaining issue originates with the interplay between the distance restrictions and the hit numbering requirements. Since the numbering is computed left to right, but the allowable distance $|z_{MDT1} - z_{MDT2}| < |z_{MDT1} - z_{MDT3}|$, there are certain combinations

that are allowed while their reflection is not. Examples of such combinations are shown in Figure B.3. Reflections of combinations shown in Figure B.3 are not allowed under the reduced hit distances. Under the original distances they would be permitted, but reflections of “stretched” combinations would not be.




Figure B.3: Examples of three-hit seeds that the algorithm selects, but whose reflections are not allowed combinations.

Fixing the hit distance requirements is simple and will reduce the number of three-hit seeds that never lead to reconstructed tracklets, thus reducing the computational time. The reflection issue is slightly more complicated, but since this could impact the reconstruction of real tracklets once the distance requirements are fixed, this should be carefully studied and altered such that all combinations are allowed in both orientations.

B.2 Vertical plane formation

The barrel MS vertex reconstruction algorithm computes the χ^2 probability that tracklets fit to a common vertex position along a series of radial planes. The intention is that these planes are spaced evenly in 200 mm increments along the line of flight. However, this is not what the algorithm actually does, due to a simple calculation error.

To calculate the location of each radial plane, the algorithm performs the following steps, where `MAXPLANES = 100`, `m_VertexMaxRadialPlane = 7000 mm`, `m_VertexMinRadialPlane = 3500 mm`, and `m_VxPlaneDist = 200 mm`:

```
float Rpos[MAXPLANES];
```

```

float RadialDist = m_VertexMaxRadialPlane - m_VertexMinRadialPlane;
float LoFdist = fabs(RadialDist/sin(LoF));
int nplanes = LoFdist/m_VxPlaneDist;
float PlaneSpacing = 3500/(nplanes-1.0);
for(int k=0; k<nplanes; ++k){
    Rpos[k] = m_VertexMinRadialPlane + PlaneSpacing*k;
}

```

The important calculation is `nplanes = LoFdist/m_VxPlaneDist`. This calculates how many increments of 200 mm fit along the line-of-fire between the radial distances of 3500 mm and 7000 mm, rounded down to the nearest integer value. However, the *radial* spacing of planes is then calculated by dividing the total radial distance, 3500 mm, by `nplanes-1.0`. This has two issues. First, if the line-of-fire distance is not divisible by 200, the radial planes will not fully span the 3500 mm radial distance, and thus calculating the radial plane spacing over the full distance will place them in a different location. Second, dividing by `nplanes-1.0` unnecessarily removes one of the planes.




Figure B.4: Horizontal planes, both as intended and calculated, for a line-of-fire corresponding to $\eta = 0.5$.

Figure B.4 shows the correct plane values and what the algorithm computes, for a line-of-fire of $\eta = 0.5$. The issue can be fixed by modifying the code as follows:

```
float Rpos[MAXPLANES];
float RadialDist = m_VortexMaxRadialPlane - m_VortexMinRadialPlane;
float LoFdist = fabs(RadialDist/sin(LoF));
int nplanes = LoFdist/m_VxPlaneDist +1;
float PlaneSpacing = fabs(200./cos(LoF));
for(int k=0; k<nplanes; ++k){
    Rpos[k] = m_VortexMinRadialPlane + PlaneSpacing*k;
}
```

B.3 Fitting of parallel lines to a vertex location

The endcap MS vertex reconstruction uses a straight-line fit, which computes the point of closest approach for the input tracklets. However, if the input lines are all parallel, the point of closest approach is infinity. While this is not explicitly incorrect, it creates many warnings in the algorithm execution. Because dividing by zero can cause unreliable behavior, cases with parallel lines should instead be treated more carefully. The issue occurs within the VxMinQuad function, and can be avoided by changing the following lines:

```
d = s*sxx - sq(sx);
float Rpos = (sxx*sy - sx*sxy)/d;
float Zpos = (sx*sy - s*sxy)/d;

Amm::Vector3D MyVx(Rpos,0,Zpos);
return MyVx;
```

to

```
d = s*sxx - sq(sx);
```

```
if(d == 0){
    Amg::Vector3D MyVx(100000,0,100000);
    return MyVx;
}
float Rpos = (sxx*sy - sx*sxy)/d;
float Zpos = (sx*sy - s*sxy)/d;

Amg::Vector3D MyVx(Rpos,0,Zpos);
return MyVx;
```

The vertex reconstruction only considers potential vertices within the detector volume, so here the value of 100000 is just some large number outside of the detector. However, this change should be studied before it is implemented to ensure that there are no unintended consequences.