

SEARCH FOR NEW PHENOMENA IN DIJET EVENTS USING 13 TEV
PROTON-PROTON COLLISION DATA COLLECTED WITH THE ATLAS
DETECTOR

by

BENJAMIN WILLIAM ALLEN

A DISSERTATION

Presented to the Department of Physics
and the Graduate School of the University of Oregon
in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy

June 2018

DISSERTATION APPROVAL PAGE

Student: Benjamin William Allen

Title: Search for New Phenomena in Dijet Events Using 13 TeV Proton-Proton Collision Data Collected With the ATLAS Detector

This dissertation has been accepted and approved in partial fulfillment of the requirements for the Doctor of Philosophy degree in the Department of Physics by:

| | |
|---------------|------------------------------|
| Eric Torrence | Chair |
| David Strom | Advisor |
| Davison Soper | Core Member |
| Dev Sinha | Institutional Representative |

and

| | |
|----------------|---|
| Sara D. Hodges | Interim Vice Provost and Dean of the Graduate School |
|----------------|---|

Original approval signatures are on file with the University of Oregon Graduate School.

Degree awarded June 2018

© 2018 Benjamin William Allen
This work is licensed under a Creative Commons
Attribution-NonCommercial-NoDerivs (United States) License.

DISSERTATION ABSTRACT

Benjamin William Allen

Doctor of Philosophy

Department of Physics

June 2018

Title: Search for New Phenomena in Dijet Events Using 13 TeV Proton-Proton Collision Data Collected With the ATLAS Detector

Dijet events are studied in the proton–proton collision dataset recorded at $\sqrt{s} = 13$ TeV with the ATLAS detector at the Large Hadron Collider in 2015 and 2016, corresponding to an integrated luminosity of 37 fb^{-1} . Invariant mass distributions are compared to background predictions and no significant deviation is observed. A new method for fitting the background component of the invariant mass distribution is employed. The dataset is used to set upper limits at a 95% confidence level on a range of new physics scenarios. Excited quarks with masses below 6.0 TeV are excluded, and limits are set on quantum black holes, heavy W' bosons, W^* bosons, and a range of masses and couplings in a Z' dark matter mediator model. These results represent a substantial improvement over those obtained previously with lower integrated luminosity.

This dissertation includes previously published and unpublished co-authored material.

CURRICULUM VITAE

NAME OF AUTHOR: Benjamin William Allen

GRADUATE AND UNDERGRADUATE SCHOOLS ATTENDED:

University of Oregon, Eugene
University of California, Irvine

DEGREES AWARDED:

Doctor of Philosophy, Physics, 2018, University of Oregon, Eugene
Bachelor of Science, Physics, 2010, University of California, Irvine

PROFESSIONAL EXPERIENCE:

University of Oregon, Graduate Research Assistant, ATLAS Experiment,
June 2014-Present

University of Oregon, Graduate Teaching Assistant, September 2012-June
2014

GRANTS, AWARDS AND HONORS:

Fall Qualifying Exam Award – October 2014

PUBLICATIONS:

ATLAS Collaboration, "Search for new phenomena in dijet mass and angular distributions from pp collisions at with the ATLAS detector", Physics Letters B, Volume 754, 10 March 2016, Pages 302-322

ATLAS Collaboration, "Search for new phenomena in dijet events collected in 2015 and 2016 pp collisions with the ATLAS detector at $\sqrt{s} = 13$ TeV", ATLAS-CONF-2016-069, 2016

ACKNOWLEDGEMENTS

To my fellow undergrads, grads, and post-docs at Oregon: Walter, Kate, Aparajita, Chris, Liza, Max, Peter, Ian, Jason, John, Johan, Elliot, Bri, Amanda, Aaron, and many more. Thank you all for for encouragement, help, questions, and camaraderie.

My dijet partners-in-crime: James, Marco, Matteo, Karishma, Kate, Meghan, Caterina. Thank you for showing me the ropes and suffering along side me through each pressing deadline.

For Sarah, who kicked my butt into going to grad school, who thought we'd never end up in Oregon, much less Europe. Thank you for being my support and my travel buddy, and forcing me to go outside.

For January, who refused to nap and let this be written.

TABLE OF CONTENTS

| Chapter | | Page |
|---------|--|------|
| I. | INTRODUCTION | 1 |
| II. | THEORY | 3 |
| | 2.1. The Standard Model | 3 |
| | 2.2. Physics beyond the Standard Model | 5 |
| | 2.3. Benchmark Signals | 7 |
| | 2.3.1. Excited Quarks | 7 |
| | 2.3.2. W' Bosons | 8 |
| | 2.3.3. Excited W Bosons | 9 |
| | 2.3.4. Z' Dark Matter Mediator | 10 |
| | 2.3.5. Quantum Black Holes | 12 |
| III. | THE STRONG FORCE AND JETS | 15 |
| | 3.1. Parton Distribution Functions | 16 |
| | 3.2. Anatomy of a Collision | 18 |
| | 3.2.1. Hard Scatter and Parton Showering | 19 |
| | 3.2.2. Pileup and Underlying Event | 20 |
| | 3.2.3. Hadronization | 22 |
| | 3.3. Jets | 23 |

| Chapter | Page |
|--|------|
| 3.3.1. The Anti- k_t jet algorithm | 24 |
| IV. MACHINE AND DETECTOR | 28 |
| 4.1. The Large Hadron Collider | 28 |
| 4.1.1. Bunches and bunch trains | 29 |
| 4.2. The ATLAS Detector | 30 |
| 4.2.1. Coordinate System | 31 |
| 4.2.2. Magnet System | 32 |
| 4.2.3. Inner Detector | 32 |
| 4.2.4. Calorimeters | 34 |
| 4.2.5. Muon Spectrometers | 37 |
| 4.2.6. Forward Detectors | 38 |
| 4.2.7. Trigger and Data Acquisition | 38 |
| 4.2.8. Dead Time | 42 |
| 4.2.9. Data Processing | 43 |
| V. CALORIMETRY AND TRIGGERING | 45 |
| 5.1. The LAr Calorimeter Pulse | 45 |
| 5.2. Readout Path | 46 |
| 5.3. Energy Summing and Triggering Path | 47 |
| 5.4. The Level-1 Calorimeter System | 48 |
| 5.4.1. Receiver and PreProcessor Module | 50 |
| 5.4.2. Cluster and Jet/Energy-Sum Processors | 53 |

| Chapter | Page |
|--|------|
| 5.5. High-Level Trigger | 54 |
| VI. TRIGGER AND CALIBRATION | 57 |
| 6.1. Trigger Strategy | 57 |
| 6.2. Event Selection Criteria | 57 |
| 6.3. Jet Formation and Definition | 59 |
| 6.4. Jet Calibration | 60 |
| 6.5. Jet Cleaning | 62 |
| 6.6. Mistimed Events in 2015 and 2016 Data | 64 |
| 6.6.1. The L1Calo Saturated BCID Trigger | 64 |
| VII. SEARCH STRATEGY | 70 |
| 7.1. Background Fit | 71 |
| 7.1.1. The Wilks' Statistical Test | 72 |
| 7.1.2. The Sliding Window Fit (SWiFt) | 73 |
| 7.2. Search Phase | 77 |
| 7.2.1. BumpHunter | 77 |
| VIII. RESULTS | 79 |
| 8.1. Limit Setting and Bayes' Theorum | 83 |
| 8.2. Uncertainties | 84 |
| 8.2.1. Jet Energy Scale Uncertainty | 84 |

| Chapter | Page |
|---|------|
| 8.2.2. Luminosity Uncertainty | 86 |
| 8.2.3. Function Choice Uncertainty | 86 |
| 8.2.4. Statistical Uncertainty | 86 |
| 8.3. Limits on Benchmark Signals | 87 |
| | |
| IX. COMPLEMENTARY SEARCHES AND OUTLOOK | 93 |
| | |
| 9.1. Trigger-Level Analysis | 93 |
| 9.2. Dijet + Initial State Radiation Searches | 96 |
| 9.3. Combined limits on Dark Matter Models | 99 |
| 9.4. Dijet Angular Search | 103 |
| 9.5. Outlook | 106 |
| 9.6. Conclusion | 106 |
| | |
| REFERENCES CITED | 108 |

LIST OF FIGURES

| Figure | Page |
|--|------|
| 2.1. Predictions for the exclusion reach of dijet searches for excited quarks at the LHC for Run-1 and Run-2 energies and integrated luminosities.[9] Comparisons to published results for CDF and LHC Run 1 are shown in the filled shapes. | 6 |
| 2.2. Feynman diagrams for (a) s -channel processes which could produce a new resonance and (b) t -channel background processes. | 7 |
| 2.3. Truth level mass (a) and analysis acceptance (b) for q^* samples at various mass points. The acceptance is after the full analysis selection which is described in Chapter VII. | 8 |
| 2.4. Truth level mass (a) and analysis acceptance (b) for W' samples at various mass points. The acceptance is after the full analysis selection which is described in Chapter VII. The W' signal points have a much larger low-mass tail, especially at high mass, due to significant off-shell production. | 9 |
| 2.5. (a) The angular distribution for the W^* signal (red) compared to the QCD dijet background (blue). (b) Signal significance as a function of angular cut. The cut used in this analysis corresponds to a cut of 2.4 on the given plot. | 10 |
| 2.6. Feynman diagrams for the Z' decays searched for in the ATLAS dark matter mediator searches. The cross-sections and kinematics depend upon the mediator and dark matter masses (M_{med}, m_χ), and couplings to dark matter and quarks (g_{DM}, g_q). Diagrams similar to (a) are searched for using missing transverse energy along with initial state radiation, while (b) is the diagram of interest for the dijet search. | 11 |
| 2.7. Truth level mass (a) and analysis acceptance (b) for QBH samples at various mass points. The acceptance is after the full analysis selection which is described in Chapter VII. | 14 |
| 3.1. The measurement of parton distribution functions (PDFs) from the CTEQ collaboration. PDFs are shown at $Q = 2 \text{ GeV}$ and $Q = 100 \text{ GeV}$ for $u, \bar{u}, d, \bar{d}, s = \bar{s}$, and g . The full results can be seen in Ref. [23]. | 17 |

| Figure | Page |
|--|------|
| 3.2. Simulation of a sample proton-proton collision. The incoming partons are in blue, the underlying event interactions are in purple, the hard scattering event is in red, and the hadronization processes are shown in green. Image from [25]. | 19 |
| 3.3. Luminosity-weighted distribution of the mean number of interactions per crossing for the 2015 and 2016 datasets. | 21 |
| 3.4. Event display of a $Z \rightarrow \mu\mu$ event with 24 other reconstructed vertices from pileup interactions. The paths of the two muons are shown in red, while the blue lines are tracks with $p_T > 500$ MeV | 22 |
| 3.5. Examples of the behavior of infrared and collinear unsafe jet algorithms. In the top two diagrams, the addition of a soft gluon becomes the seed which a jet is centered on, causing the two jets to merge together into one jet. In the bottom diagram, the splitting of the most energetic seed into two collinear ones, shifting the jet location and lowering the jet energy. Diagram from [30]. | 24 |
| 3.6. An example of the anti- k_t recombination algorithm. a) The original input subjects to the algorithm. b) The highest p_T subjet combines with its nearest neighbor to produce a new combined subjet. c) This subjet again combines with its neighbor to produce a new subjet. d) The remaining subjects are separated by more than the radius parameter of the algorithm. The subjects become the final jets. | 26 |
| 3.7. A sample parton-level event together with many random soft particles, clustered together using the anti- k_t algorithm.[31]. | 27 |
| 4.1. Schematic of the CERN accelerator complex.[38] | 29 |
| 4.2. A cutaway view of the ATLAS detector showing the various components.[34] | 31 |
| 4.3. Average track region TRT occupancy as a function of track η using data taken during the 2015 and 2016 ATLAS runs for four different values of the average interactions per bunch crossing. | 34 |
| 5.1. LAr Calorimeter pulse shape as produced in the EM barrel before and after shaping. The dots represent the 25 ns spacing used in the pulse sampling and digitization.[44] | 46 |
| 5.2. The ATLAS LAr Calorimeter front-end system in Run 2, along with the connections to the L1 trigger and readout systems.[45] | 47 |

| Figure | Page |
|--|------|
| 5.3. The ATLAS trigger and data acquisition system in Run 2. L1Topo was being commissioned in 2015. [47] | 49 |
| 5.4. (a) The average pedestal correction per trigger tower as a function of the Bunch Crossing ID (BCID) in the Electromagnetic Barrel partition of the LAr calorimeter during a 2017 run. 1 ADC count is approximately equal to 0.25 GeV of transverse energy in a single tower.[48] (b) The per-bunch trigger rate for the L1 missing transverse momentum trigger with a threshold of 50 GeV (L1_XE50) as a function of the instantaneous luminosity per bunch. The rates are shown with and without pedestal correction applied. [47] | 53 |
| 5.5. Jet trigger algorithm windows using 0.2×0.2 ($\eta \times \phi$) jet elements with the local maxima cores shaded in. Only the 0.8×0.8 window is currently used for jet finding.[46] | 54 |
| 6.1. (a) The average energy response as a function of η_{det} for jets after origin and pile-up corrections are applied. (b) The signed difference between truth and reconstructed jet η due to biases in jet reconstruction.[53] | 61 |
| 6.2. Ratio of the EM+JES jet response in data to that in the nominal MC event generator as a function of jet p_T for Z -jet, γ -jet, and multi-jet in-situ calibrations. The final derived correction and its statistical and total uncertainty bands are shown. | 62 |
| 6.3. (a) Comparison of the effect on the m_{jj} uncertainty due to missing event information in BCID corrected events compared to the overall statistical uncertainty. (b) The dijet invariant mass spectrum showing the normal events in blue and the corrected events in red. | 69 |
| 7.1. Wilks' test between the 4-parameter fit function and the 3/5-parameter versions. For the full 2015+2016 dataset, the 4-parameter fit provides a much better description of the data than the 3-parameter fit, but the 5-parameter fit does not provide additional improvement. | 72 |
| 7.2. Statistical measures used to determine the optimal window size for SWiFt. (a) Wilks' p-value comparing the nominal SWiFt 3-parameter fit with an alternate 4-parameter version (b) KS p-value comparing the fit to pseudodata (c) χ^2 /NDF comparing fit to pseudodata. | 74 |
| 7.3. An example of the bin-by-bin construction of the SWiFt background as seen in Ref. [55]. The gray box shows the window considered at each point while the green line is the bin center for which the fit value is evaluated. In (b) and (c) only the value at the bin center is used for the total background | |

| Figure | Page |
|--|------|
| estimate, while in (a) and (d) the fit values before and after the center, respectively, are used as part of the overall background. | 75 |
| 7.4. Comparison of the global 4-parameter fit function to the SWiFt background using a 3-parameter fit in each window. No significant variations between the two fits are observed. | 76 |
| 8.1. The reconstructed dijet mass distribution m_{jj} (filled points) is shown for events with $p_T > 440$ (60) GeV for the leading (subleading) jet in the resonance $y^* < 0.6$ signal region. The solid line depicts the background prediction from the sliding-window fit. The vertical lines indicate the most discrepant interval identified by the BumpHunter algorithm, for which the p-value is 0.63. The middle panel shows the bin-by-bin significances of the data–fit differences, considering only statistical uncertainties. The lower panel shows the relative differences between the data and the prediction of Pythia 8 simulation of QCD processes, corrected for NLO and electroweak effects, and is shown purely for comparison. The shaded band denotes the experimental uncertainty in the jet energy scale calibration. | 81 |
| 8.2. The reconstructed dijet mass distribution m_{jj} (filled points) is shown for events with $p_T > 440$ (60) GeV for the leading (subleading) jet in the W^* resonance $y^* < 1.2$ signal region. The solid line depicts the background prediction from the sliding-window fit. The vertical lines indicate the most discrepant interval identified by the BumpHunter algorithm, for which the p-value is 0.83. The middle panel shows the bin-by-bin significances of the data–fit differences, considering only statistical uncertainties. The lower panel shows the relative differences between the data and the prediction of Pythia 8 simulation of QCD processes, corrected for NLO and electroweak effects, and is shown purely for comparison. The shaded band denotes the experimental uncertainty in the jet energy scale calibration. | 82 |
| 8.3. The highest mass dijet event observed in the 2015+2016 dataset. It has an invariant mass of 8.12 TeV, and the two jets each have transverse momenta of 3.79 TeV. | 83 |
| 8.4. Combined uncertainty in the JES of fully calibrated jets as a function of jet p_T at $\eta = 0$. Systematic uncertainty components include pileup, punch-through, and uncertainties propagated from Z/γ -jet and multijet balance, as well as the η -intercalibration. The flavor composition and response uncertainties assume a quark and gluon composition taken from Pythia dijet MC simulation (inclusive jets). | 85 |
| 8.5. The sliding window fit to the m_{jj} distribution with $ y^* < 0.6$ (red) with its two uncertainties. The dark blue line with long dashes indicates | |

| | |
|---|----|
| <p>the statistical uncertainty in the fit, corresponding to the variability in the fit results across a large collection of pseudoexperiments. The light blue line with short dashes indicates the uncertainty in the function choice, based on comparison to a function with a higher number of parameters. Statistical uncertainty is shown at the one-sigma level, while the function uncertainty shown is the mean difference between the nominal and alternate functions.</p> | 87 |
| <p>8.6. Limits on the cross-section times branching ratio times acceptance for the three resonance benchmark models (a) Quantum black holes (ADD, n=6 model) (b) excited quarks (c) W' bosons, and (d) W^* bosons. The first three limits use the nominal selection while the fourth uses the widened $y^* < 1.2$ selection. The theory cross-section is given by the blue dashed line.</p> | 88 |
| <p>8.7. 95% CL limits on the production of leptophobic Z', shown in the 2D plane of the mass of the Z' (m_Z) and the coupling to quarks (g_q). The excluded region is the region in the plane above the solid line, as indicated by the hatching. Exclusions are valid only up to $g_q = 0.5$, as coupling values above that create resonances which are too wide to be properly detected by this search.</p> | 90 |
| <p>8.8. The 95% CL upper limits obtained from the dijet invariant mass m_{jj} distribution on cross-section times acceptance times branching ratio to two jets, $\sigma \times A \times BR$, for a hypothetical signal with a cross-section σ_G that produces a Gaussian contribution to the particle-level m_{jj} distribution, as a function of the mean of the Gaussian mass distribution m_G. Observed limits are obtained for five different widths, from a narrow width to 15% of m_G. The expected limit and the corresponding $\pm 1\sigma$ and $\pm 2\sigma$ bands are also indicated for a narrow-width resonance.</p> | 92 |
| <p>9.1. In-situ calibration in the range of $85 \text{ GeV} < \text{jet } p_T < 2 \text{ TeV}$ for both the traditional spline method (dashed line) and new fitted (solid line) combination methods. Data points from the three input measurements are overlaid. The lower panel shows the ratio of the two calibration curves. . .</p> | 94 |
| <p>9.2. Reconstructed dijet mass distribution for events in the $y^* < 0.3$ (red) and $y^* < 0.6$ (blue) signal regions. The solid lines depict the background estimate obtained by a sliding window fit. The lower panel shows the bin-by-bin significances of the differences between the data and the background estimate, and the most discrepant region in both signal regions is indicated by the colored vertical lines.</p> | 95 |

| Figure | Page |
|---|------|
| 9.3. The 95% credibility-level observed and expected upper limits on g_q as a function of $m_{Z'}$. The lower-mass portion of the limits from the high-mass dijet result is also shown. Couplings above the solid lines are excluded. . . | 95 |
| 9.4. Reconstructed dijet mass distribution for events containing two jets with $p_T > 25$ GeV and either a jet with $p_T > 430$ GeV (left) or a photon with $p_T > 150$ GeV (right). The vertical lines indicate the most discrepant region as identified by the BumpHunter algorithm, and the bottom panel shows the bin-by-bin significances of the data with respect to the background fit. | 97 |
| 9.5. The 95% CL upper limits obtained from the X+jet (left) and X+photon (right) searches on coupling g_q as a function of the resonance mass $m_{Z'}$ for the leptophobic Z' dark matter mediator model. | 97 |
| 9.6. Distribution of Large-R jet mass in the ISR jet (left) and ISR photon (right) channels. The signal Z' shown has a mass of 160 GeV and coupling $g_q=0.5$. The bottom panel shows the ratio of the data to the estimated background; the estimate varies for different signal masses. | 98 |
| 9.7. Observed and expected limits at 95% CL on the coupling to quarks g_q for the combined ISR jet and photon channels for the boosted search. . . . | 98 |
| 9.8. Regions in a dark matter mass-mediator mass plane excluded at 95% CL by a selection of ATLAS dark matter searches, for one possible interaction between the Standard Model and dark matter, the leptophobic axial-vector mediator as described in [59]. The exclusions are computed for a dark matter coupling $g_{DM} = 1.0$, a quark coupling $g_q = 0.25$ universal to all flavors. The lepton coupling g_l in this model is set to zero. This choice for the couplings corresponds to the "A1" scenario in [59]. The results use 13 TeV data except for [7]. Dashed curves labeled "thermal relic" indicate combinations of dark matter and mediator mass that are consistent with a dark matter density of $\Omega_c = 0.12h^2$ and a standard thermal history, as computed in MadDM ([60], [59]). Between the two curves, annihilation processes described by the simplified model deplete Ω_c below $0.12 h^2$. A dotted curve indicates the kinematic threshold where the mediator can decay on-shell into dark matter. Excluded regions that are in tension with the perturbative unitary considerations of [61] are indicated by shading in the upper left. The exclusion regions, relic density contours, and unitarity curve are not applicable to other choices of coupling values or model. . . | 100 |
| 9.9. Regions in a dark matter mass-mediator mass plane excluded at 95% CL by a selection of ATLAS dark matter searches, for one possible interaction between the Standard Model and dark matter, the vector mediator as | |

- described in [59]. The exclusions are computed for a dark matter coupling $g_{DM} = 1.0$, a quark coupling $g_q = 0.1$ universal to all flavors, and lepton coupling $g_l = 0.01$, corresponding to the "V2" scenario in [59]. With this choice of couplings, the Z' decays to leptons are reduced with respect to the decays to quarks. The results use 13 TeV data except for [7]. The lepton constraints use 13 TeV data from 2016, reinterpreting the model independent limits in a way similar to what done for the dijet searches. The $E_T^{miss}+X$ exclusion regions are obtained by rescaling, using acceptance and cross-section information from samples simulated at truth-level, the exclusion contours published in the corresponding papers. Dashed curves labeled "thermal relic" indicate combinations of dark matter and mediator mass that are consistent with a dark matter density of $\Omega_c = 0.12h^2$ and a standard thermal history, as computed in MadDM ([60], [59]). To the left of the curve, annihilation processes described by the simplified model deplete Ω_c below $0.12 h^2$. A dotted curve indicates the kinematic threshold where the mediator can decay on-shell into dark matter. The exclusion regions, relic density contours, and unitarity curve are not applicable to other choices of coupling values or model. 101
- 9.10. Summary plot of ATLAS bounds in the coupling-mediator mass plane from dijet searches using 2015 and 2016 data. The 95% CL upper limits are on coupling g_q as a function of the resonance mass $m_{Z'}$ for the leptophobic Z' model described in [16]. The expected limits from each search are indicated by dotted curves. Coupling values above the solid curves are excluded, as long as the signals are narrow enough to be detected using these searches (10% signal width/mass for dijet+ISR and TLA, 15% for high-mass dijets, approximately corresponding to $g_q < 0.5$ and $g_q < 0.6$, respectively). 102
- 9.11. Reconstructed distributions of the dijet angular variable χ in different regions of the dijet invariant mass m_{jj} . The data (points), Pythia predictions with NLO and electroweak corrections applied (solid lines), and examples of the contact interaction (CI) signals(dashed lines) are shown. The theoretical uncertainties and the total theoretical and experimental uncertainties in the predictions are displayed as shaded bands around the SM prediction. The SM background prediction and corresponding systematic uncertainty bands are extracted from the best-fit to the data. Data and predictions are normalized to unity in each m_{jj} bin. 104
- 9.12. Results from the CMS Collaboration showing the 95% CL upper limits on the quark coupling g_q , as a function of mass, for an axial-vector or

vector DM mediator with $g_{DM} = 1.0$ and $m_{DM} = 1$ GeV. The observed limits (solid), expected limits (dashed) and the variation in the expected limit at the 1 and 2 standard deviation levels (shaded bands) are shown. The corresponding limits on the width of the mediators are shown on the vertical axis on the right-hand side of the figure.[62] 105

LIST OF TABLES

| Table | Page |
|---|------|
| 2.1. Particles of the Standard Model. Data from [6]. | 5 |
| 8.1. Comparison of limits on benchmark models set by ATLAS [1] and CMS [58]. | 89 |

CHAPTER I

INTRODUCTION

Over the last 50 years, the Standard Model of particle physics has proven itself time and time again, successfully describing phenomena covering several orders of magnitude in energy as well as describing new particles, such as the Higgs boson, which was finally discovered in 2012. However, this spectacular resilience is also a source of great frustration. The Standard Model is known to be incomplete, lacking answers for neutrino masses, the mass of the Higgs boson, dark matter...yet every precision search has yet to poke any new holes in it. With the Large Hadron Collider (LHC) reaching a new center of mass energy in 2015, the hope is to discover new particles which only become accessible at very high energies.

As a hadron collider, the LHC produces collisions which are dominated by strongly interacting particles which are detected in the form of jets. Given this very high cross-section, the dijet search is positioned to take advantage of the extraordinary statistics to look for evidence of strongly interacting new resonances across a very large mass range, and hopefully pointing the way for precision searches to explore any excesses that are seen.

This thesis presents the resonance portion of the dijet analysis performed on the combined 2015 and 2016 datasets taken by the ATLAS experiment.[1] The analysis looks at the dijet invariant mass spectrum for evidence of localized excesses which could point to new resonances or other forms of physics beyond the Standard Model. The analysis uses a model-agnostic methodology to be sensitive to as many possible Beyond the Standard Model (BSM) models as possible, and its results are presented

both in terms of limits on selected benchmark models as well as for generic Gaussians which can be used to interpret limits for other theoretical models.

This research paper result was by no means an individual effort, but was built on the efforts of dozens of other researchers, both past and present. The author's individual contributions included maintenance of the core analysis code, preparation, processing, and dissemination of real and simulated data samples, Wilks' statistical testing, year-to-year dataset validation, and serving as analysis contact and paper editor. Chapters VII and VIII contain work previously published in Physical Review D, Volume 96, on September 28th, 2017.[1] Chapter IX contains work which has been made public, some of which is awaiting publication.[2][3][4] These works are co-authored with the ATLAS Collaboration, comprising approximately 2900 authors, the full list of which are available in the full publications.

Chapter II introduces the Standard Model as well as the benchmark models used in the analysis. Chapters IV and V discuss the Large Hadron Collider and ATLAS Experiment, with a special focus on the calorimetry systems used to measure jets. Chapter III covers the strong force and how it gives rise to the jets measured in this analysis, while Chapter VI discusses how jets are reconstructed and calibrated into final physics objects. Chapter VII covers the search strategy used in this analysis and the sources of uncertainty in the result. Finally, Chapter VIII discusses the finding and limits set on models of new physics while Chapter IX puts these limits in context of other complementary searches as well as the future outlook.

CHAPTER II

THEORY

All of the observed particles and their interactions are described by what is known as the Standard Model (SM) of particle physics. The SM encompasses the three fundamental forces of electromagnetism, the weak force, and the strong force. Gravity is not included as part of the model, though its weakness compared to other forces means that it has a negligible impact (at least for energy scales currently probed at the LHC). Though the SM has several glaring issues it can not address, including the existence of dark matter and the non-zero masses of neutrinos, it has withstood decades of precision testing, much to the frustration of physicists looking for ways to address these outstanding questions.

2.1. The Standard Model

Particles in the Standard Model belong to one of two categories: fermions which have half-integer spin ($1/2$) and bosons which have integer spin (0 or 1).

The four vector (spin-1) bosons are the force carriers of the Standard Model and are the mediators of interactions between matter. The massless photon (γ) corresponds to electromagnetism and interacts with particles which carry electric charge, while the massless gluon (g) mediates the strong force and interacts with particles which carry color charge. The neutral Z and electrically-charged W bosons are both massive and are the carriers of the weak force. The Higgs boson is the only scalar (spin-0) boson in the SM and is not a force carrier. It corresponds to the Higgs field which is responsible for giving mass to all of the other particles in the SM. The particle content of the Standard Model is summarized in Table 2.1.

Gravity does not currently fit into the Standard Model, though if it were it could hypothetically be mediated by the *graviton*, which is as-of-yet unobserved. As the scale of gravity is many orders of magnitude weaker than the other forces, it can safely be neglected at the scale of the particle interactions under consideration.

There are 12 fermions in the SM, broken up into two types: quarks which carry color charge and interact under the strong force, and leptons which do not. The fermions can be divided into three generations, each with four particles: one quark with charge $+\frac{2}{3}$, one quark with charge $-\frac{1}{3}$, one massive lepton with charge -1 , and one nearly-massless lepton, called a neutrino, with zero charge. Each subsequent generation of particles is more massive and decays into the lighter generations, and only the lightest generation is stable with the exception of the neutrinos, all of which are stable. Neutrino non-zero masses and their observed oscillation between varieties[5] are not explained by the Standard Model.

All fermions also have an anti-particle: a particle with the same mass as the ordinary particle but with opposite quantum numbers such as charge and spin. While our universe is made up almost entirely of ordinary matter, particle collisions produce matter and anti-matter in pairs, such as pair production from a photon becoming an electron and a positron (anti-electron). When a particle and anti-particle collide, they annihilate each other in a burst of energy. The source of the asymmetry between the amount of matter and anti-matter in the present-day universe is a topic of active research.

While the lightest charged lepton, the electron, is a familiar part of matter, individual quarks are never observed. Instead, quarks form color-neutral bound states called hadrons. States are considered color-neutral if they either carry colors which cancel each other out, or have all three colors, in analogy to the colors of light

| | Particle | Spin | Charge | Color | Mass |
|----------|------------|---------------|----------------|----------|-----------|
| Fermions | | | | | |
| Quarks | u | $\frac{1}{2}$ | $+\frac{2}{3}$ | (r,g,b) | 2.2 MeV |
| | d | | $-\frac{1}{3}$ | (r,g,b) | 4.7 MeV |
| | s | | $-\frac{1}{3}$ | (r,g,b) | 96 MeV |
| | c | | $+\frac{2}{3}$ | (r,g,b) | 1.3 GeV |
| | b | | $-\frac{1}{3}$ | (r,g,b) | 4.2 GeV |
| | t | | $+\frac{2}{3}$ | (r,g,b) | 173.1 GeV |
| Leptons | e | $\frac{1}{2}$ | -1 | None | 0.51 MeV |
| | μ | | | | 105.7 MeV |
| | τ | | | | 1.78 GeV |
| | ν_e | $\frac{1}{2}$ | 0 | None | < 2 eV |
| | ν_μ | | | | < 2 eV |
| | ν_τ | | | | < 2 eV |
| Bosons | | | | | |
| Vector | g | 1 | 0 | 8 Colors | 0 |
| | γ | 1 | 0 | None | 0 |
| | W | 1 | ± 1 | None | 80.4 GeV |
| | Z | 1 | 0 | None | 91.2 GeV |
| Scalar | h | 0 | 0 | None | 125.1 GeV |

TABLE 2.1. Particles of the Standard Model. Data from [6].

producing white when overlapped. These hadrons are composed of either quark-antiquark pairs ($q\bar{q}$) called mesons, or sets of three (anti-)quarks (qqq or $\bar{q}\bar{q}\bar{q}$) called (anti-)baryons. Of the many possible combinations of quarks, only two states are stable and make up the matter in the universe: the proton (uud) and the neutron (udd), while the other hadronic states will decay after a short time into lighter hadrons or other particles.

2.2. Physics beyond the Standard Model

With the full particle content of the Standard Model now observed, the focus has shifted to Beyond the Standard Model (BSM) physics models which can hopefully

address some of the outstanding questions. With the LHC being a hadron collider, it makes sense to focus early searches on dijet final states, as any new heavy resonances produced directly from parton interactions (as opposed to processes such as vector boson fusion) will also have a large branching ratio to a dijet final state.

The dijet search is not strongly statistically limited, but instead gains much of its reach through the increase in center-of-mass energy. An example of the predicted reach for the excited quark benchmark model is shown in Figure 2.1; the exclusion limit for the entire Run 1 dataset was quickly surpassed within the first few months of Run 2 data. In the case of the quantum black hole benchmark, the 20 fb^{-1} 8 TeV result from Run 1 [7] was surpassed after just 80 pb^{-1} of Run-2 data at 13 TeV.[8]




FIGURE 2.1. Predictions for the exclusion reach of dijet searches for excited quarks at the LHC for Run-1 and Run-2 energies and integrated luminosities.[9] Comparisons to published results for CDF and LHC Run 1 are shown in the filled shapes.




FIGURE 2.2. Feynman diagrams for (a) s -channel processes which could produce a new resonance and (b) t -channel background processes.

2.3. Benchmark Signals

The dijet search is designed to be as model-agnostic as possible, searching for s -channel production of a new resonance which decays to two quarks, two gluons, or one quark and one gluon. Many possible models can describe this final state, of which a handful were chosen as benchmarks for this analysis. These models are introduced here. The goal of the analysis is not to discover or exclude certain models, but instead to provide results which can easily be interpreted in the context of other model-specific searches.

2.3.1. Excited Quarks

Excited quarks are a standard benchmark model for compositeness, where the quarks (and leptons) that we observe are not actually fundamental particles, but are made up of some smaller, unknown particles (typically called *preons*). Composite models aim to address the mass hierarchy of the quarks and the patterns of fermion generations in the same way that the quark model explains the behavior of baryons and mesons. Excited quark states would be produced around the scale of this new

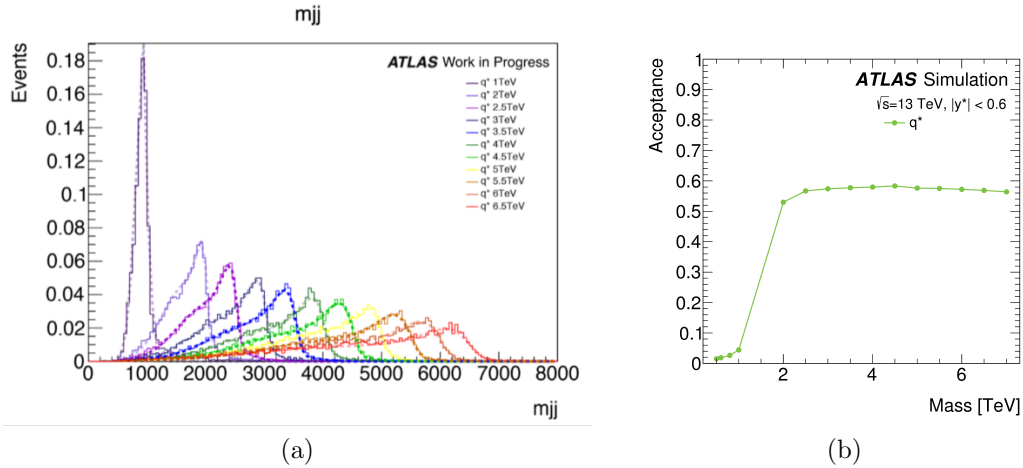


FIGURE 2.3. Truth level mass (a) and analysis acceptance (b) for q^* samples at various mass points. The acceptance is after the full analysis selection which is described in Chapter VII.

interaction. In pp collisions, a quark could be excited to a higher state through absorption of a gluon, and then radiate a boson as it returns to the ground state. This search looks for evidence of a $q^* \rightarrow qg$ decay which corresponds to $\sim 85\%$ of the total branching ratio. The other possible decay modes ($q^* \rightarrow \gamma, Z, W$) are not simulated and assumed to have zero acceptance in the search.

The model used for the benchmark signal is described in Refs. [10] and [11]. The parameters for the model are chosen such that the resonance has a relatively narrow width as shown in Figure 2.3. If excited quarks are detectable at LHC energies, they should produce peaks which are very pronounced compared to the background QCD prediction. The acceptance for the q^* model is roughly flat as a function of mass, with low acceptance points in Figure 2.3 being below the mass cut of the analysis.

2.3.2. W' Bosons

Many BSM theories which introduce additional W bosons are part of extended gauge theories. The benchmark model used for this analysis is the Sequential

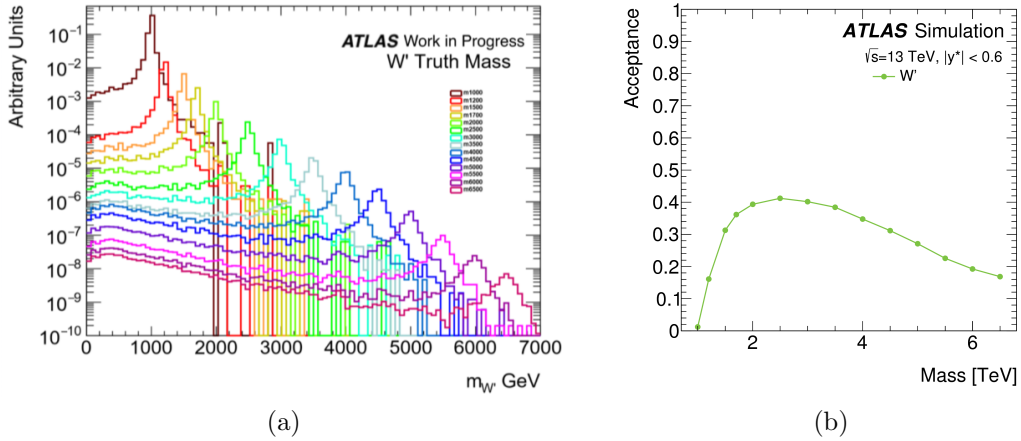


FIGURE 2.4. Truth level mass (a) and analysis acceptance (b) for W' samples at various mass points. The acceptance is after the full analysis selection which is described in Chapter VII. The W' signal points have a much larger low-mass tail, especially at high mass, due to significant off-shell production.

Standard Model (SSM), where the W' shares the same couplings as the SM W boson, varying only in mass. The corresponding branching ratio to quark final states of 75% from the process $W' \rightarrow qq$. No interference between the W' and the SM is simulated.

In comparison to the q^* signal, the W' has a much more pronounced low-mass tail, as shown in Figure 2.4. As such, the acceptance for the model decreases significantly at higher mass points.

2.3.3. Excited W Bosons

Similar to an excited quark, an excited state of the W boson could point to compositeness there, with decays to a qq final state.[12] Unlike all of the other models considered here, the angular distribution of W^* decays do not peak at $y^* = |\eta_1 - \eta_2|/2 = 0$, but rather at more forward angles, as shown in Figure 2.5. As such, the W^* search uses a separate second search region with a wider angular cut of $y^* < 1.2$ to maximize signal significance.[13]

The mixing angle $\sin(\theta_X)$ determines the coupling to the quarks and leptons, for this particular model it is set to 0, creating a leptophobic state with maximal coupling to quarks, and thus the highest branching fraction to the dijet final state. The coupling constants are selected in such a way to ensure that the W^* has the same width as the W' model used, helping to facilitate comparisons.

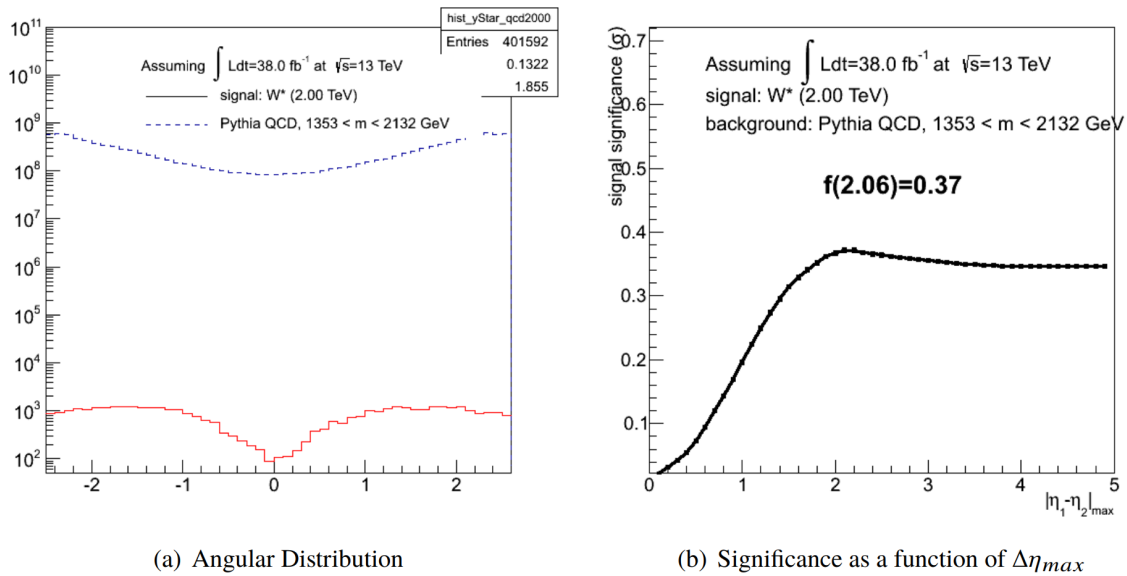


FIGURE 2.5. (a) The angular distribution for the W^* signal (red) compared to the QCD dijet background (blue). (b) Signal significance as a function of angular cut. The cut used in this analysis corresponds to a cut of 2.4 on the given plot.

2.3.4. Z' Dark Matter Mediator

The existence of dark matter is extremely well motivated, from measurements of the rotation speeds of distant galaxies [14] to analyses of the cosmic microwave background.[15] Current estimates put the amount of dark matter at more than five times the amount of normal matter, yet there is no particle in the Standard Model that can explain this density. A typical dark matter candidate interacts only with the weak and gravitational force; if it were charged under the EM or strong force it would




FIGURE 2.6. Feynman diagrams for the Z' decays searched for in the ATLAS dark matter mediator searches. The cross-sections and kinematics depend upon the mediator and dark matter masses (M_{med}, m_χ), and couplings to dark matter and quarks (g_{DM}, g_q). Diagrams similar to (a) are searched for using missing transverse energy along with initial state radiation, while (b) is the diagram of interest for the dijet search.

be visible to us, or interact with regular matter in a detectable way. These WIMPs (Weakly Interacting Massive Particles) are a feature of many BSM models, such as in R-parity conserving supersymmetry (SUSY) where the lightest stable SUSY particle does not decay and is uncharged. Other models postulate a "dark sector" of particles which do not directly interact with regular matter, but instead have some mediator particle which couples to both the SM and dark matter.

If there is some particle which couples to both Standard Model particles and the dark sector, then it could be produced at the LHC and be detected. This search takes two different forms depending on the decay mode of the mediator. If the mediator decays to the dark matter particle, it can be detected as missing transverse energy against a recoiling gluon or photon from initial state radiation. If the mediator instead decays back into Standard Model particles, then its products could directly observed. These two possible decays are shown in Figure 2.6.

The Z' model used is one of the benchmarks recommended by the ATLAS/CMS Dark Matter Forum [16], where the Z' couples to dark matter, has axial-vector couplings to all SM quarks, and does not couple to leptons in any way. The model then has four free parameters: M_{med} , the mediator mass, m_χ , the mass of the dark matter particle, g_q , the coupling to quarks, and g_{DM} , the coupling to dark matter. Both m_χ and g_{DM} have second order effects on the behavior of the Z' decays and are set to $M_{DM} = 10 \text{ TeV}$ and $g_{DM} = 1.5$ respectively. The relevant parameters are then the mediator mass, which sets the location of the excess, and g_q , which controls the decay cross-section and the width of the resonance. Signal samples are generated for a range of masses and g_q . Couplings beyond $g_q = 0.5$ are not considered as part of this analysis, as beyond this value the width of the resonance grows beyond 15% and can no longer be properly detected by this search. No interference with the SM is simulated.

2.3.5. Quantum Black Holes

In comparison to the other three fundamental forces, gravity is many, many orders of magnitude weaker. The fundamental scale of gravity, the Planck mass (m_P), is equal to

$$m_P = \sqrt{\frac{\hbar c}{G}} = 1.22 \times 10^{19} \text{ GeV}/c^2 \quad (2.1)$$

where G is the gravitational constant. This large gap between the scale of electroweak physics around 100 GeV and the scale of gravity is known as the *hierarchy problem*. One possible solution to this issue is to postulate that gravity is around the same scale as the other forces, but that it is diluted into some number of extra spacial dimensions. In the case of the ADD model (Arkani-Hamed, Dimopoulos, and Dvali) [17], the 4-

dimensional universe we observe is embedded in a higher-dimensional bulk spacetime where the extra dimensions are only accessible by gravitational fields. Since the inverse-square law for gravity has been tested down to the sub-millimeter scale, these extra dimensions must be compact such that they have thus far escaped detection.

If there are a number of such extra dimensions, the fundamental scale of gravity, M_D , could be low enough such that it could be accessible at LHC energies. This introduces the possibility that microscopic or Quantum Black Holes (QBH) could be produced in proton-proton collisions. If the black hole produced is well above M_D , it would quickly decay via Hawking radiation to a large number of particles, leaving behind a telltale high-multiplicity signature in the detector. In the more likely scenario of a black hole being produced close to M_D , quantum effects dominate and a low-multiplicity final state is dominant.[18]

The particular benchmark model used is an ADD-type QBH with $n = 6$ extra dimensions. In this scenario, QBH produced near M_D will decay to an average of two particles, and since it is created from the interaction of two strongly-charged partons, the decay to a two-parton final state is strongly favored. The model used has an estimated branching ratio of 96% to two partons. As such, the final state looks very much like a new resonance produced near M_D with a dijet final state, with a negligible low-mass tail and a signal width of approximately 10%, as shown in Figure 2.7.

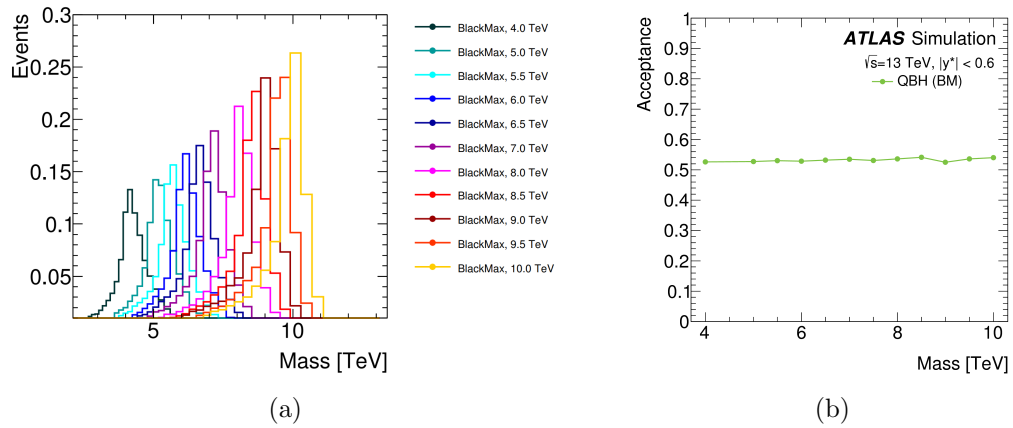


FIGURE 2.7. Truth level mass (a) and analysis acceptance (b) for QBH samples at various mass points. The acceptance is after the full analysis selection which is described in Chapter VII.

CHAPTER III

THE STRONG FORCE AND JETS

The cross-section for the collision of two hadrons can be written as

$$d\sigma_{AB \rightarrow J+X} = \sum_{a,b} \int dx_a dx_b f_{a/A}(x_a, \mu_F^2) f_{b/B}(x_b, \mu_F^2) \times d\hat{\sigma}_{ab \rightarrow J+X}(\alpha_s(\mu_R^2), Q^2) \quad (3.1)$$

where A and B are the incoming hadrons (in the case of the LHC, protons), $J + X$ are an observable plus any other particle, such as two jets, or a jet plus some invisible particle, a and b are the types of partons (the six quark flavors and the gluon), $d\hat{\sigma}_{ab \rightarrow X}(\alpha_s(\mu_R^2), \mu^2)$ is the cross-section for the collision of the two partons a and b to produce X , which can be calculated directly from theory and depends on the the scale of the interaction, μ , and strong QCD coupling, α_s , measured at renormalization scale μ_R^2 . This formula is approximate and has corrections of suppressed by a factor of $1 \text{ GeV}/Q$. For the energies probed in this analysis, with jets on the order of 1 TeV, these corrections are very small.

The renormalization scale is used to absorb the ultraviolet divergences which come from the higher-order diagrams of the process, imposing a cutoff on the energy in exchange for making α_s depend on this scale.[19] This is the origin of the running coupling of the strong interaction, where the strength of the coupling decreases as the energy of the process increases. Similarly, the parton distribution functions discussed in the next section depend on the choice of factorization scale.

While the choice of scale is arbitrary and should not have an effect on the cross-section, this is only true if one is able to expand out to infinite accuracy, as $\hat{\sigma}$ has an expansion in powers of α_s . With a good choice in scale, good accuracy to the full

expansion can be obtained with only a few terms in the series. With the choice of renormalization and factorization scales close to the scale of the given process such that $\mu_R^2 = \mu_F^2 \simeq Q^2$, the cross-section can be characterized by only the energy scale of the interaction and is well described by the expansion. This choice of scales is used for the remainder of this chapter.

The final factors, $f_{a/A}(x_a, Q^2)$, are the parton distribution functions, or PDFs. The PDFs represent the probability density for a parton of type a to carry a fraction x_a of the total momentum of the hadron. The total cross section for the process is then the sum over all of the types of partons from the two incoming hadrons.

3.1. Parton Distribution Functions

Since the proton is a composite particle, its overall momentum is equivalent to the sum of the momenta of its constituents. Given the large momenta of accelerated protons in the LHC for Run 2, this is equivalent to saying that the total energy of the proton (6.5 TeV for Run 2) is divided up between the partons that comprise the proton.

PDFs can not be calculated from theory but instead must be taken directly from data, much of which comes from deep inelastic scattering experiments from electron-proton colliders, in addition to Drell-Yan and inclusive jet production measurements from proton-proton colliders, including results from Run 1 of the LHC. Once the PDF at a given energy transfer Q is known, the PDF can be evolved to other energies using the Dokshitzer-Gribov-Lipatov-Altarelli-Parisi (DGLAP) equations.[20, 21, 22] Examples of the proton PDF at two different energy scales are shown in Figure 3.1.

While one would naively expect the momentum to be roughly split between the two up quarks and one down quark, this is not the case. Rather, a sizeable portion of

the momentum is carried by gluons which constantly radiate from and are reabsorbed by the valence quarks. These gluons can also fluctuate into quark-antiquark pairs, referred to as *sea quarks*. At higher Q , shorter distance scales are probed, meaning larger fluctuations are accessible, and thus larger amounts of momentum are carried by the gluons and sea quarks rather than just by the valence quarks.

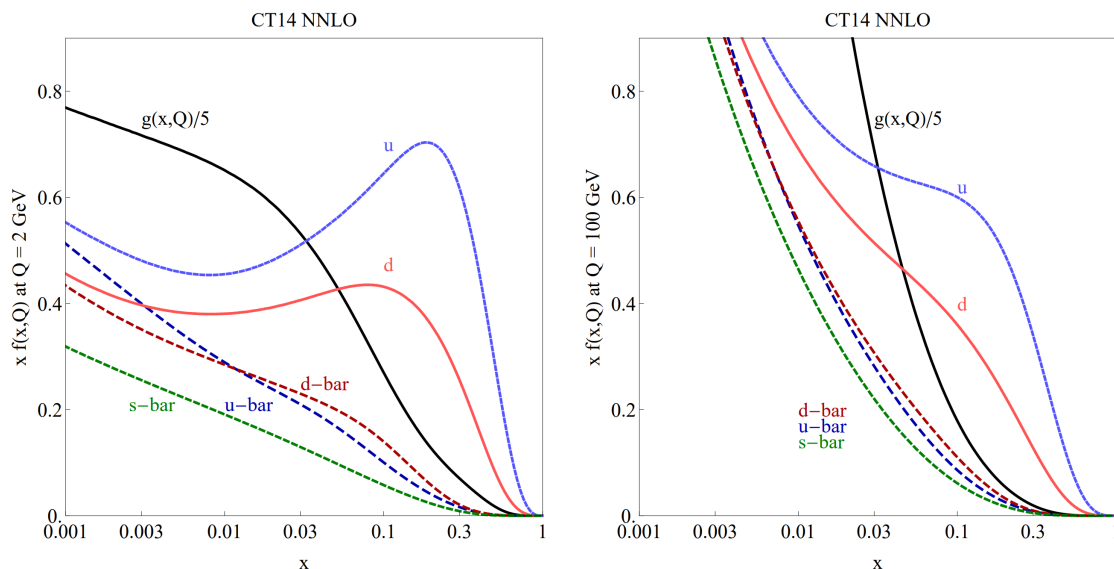


FIGURE 3.1. The measurement of parton distribution functions (PDFs) from the CTEQ collaboration. PDFs are shown at $Q = 2$ GeV and $Q = 100$ GeV for $u, \bar{u}, d, \bar{d}, s = \bar{s}$, and g . The full results can be seen in Ref. [23].

The high-mass dijet search relies on interactions at high- Q and partons at very high- x in the PDF, where x is the fraction of the total proton momentum carried by a given parton. However, while the valence quarks are the most common at high- x , there are no quark-quark vertices in the standard model, only quark-antiquark and (anti)quark-gluon ones. The t-channel quark-quark scattering is then the dominant background in a dijet search over s-channel processes from both Standard Model QCD and any possible new resonances. Figure 2.2 shows a comparison of these two channels. This background is suppressed in the analysis by imposing a maximum

difference in rapidities between the two leading jets, as t -channel processes peak at very high rapidity near the beam line. In the limit of massless quarks (a safe assumption given the energies probed by the high-mass dijet search), the Mandlestam variables can be written as:

$$\hat{s} = (p_1 + p_2)^2 = m_{jj}^2 \quad (3.2)$$

$$\hat{t} = -\frac{1}{2}\hat{s}(1 - \cos \theta^*) \quad (3.3)$$

where θ^* is the polar angle from the beam line to the outgoing parton in the center-of-mass frame, and m_{jj} is the invariant mass of the dijet system.[24] The matrix element for the s -channel process is proportional to \hat{s}^{-1} , while the t -channel process is proportional to \hat{t}^{-1} . Thus, the s -channel production is independent of angle, while the t -channel mode peaks at angles close to the beam line and is minimized in the central region of the detector.

3.2. Anatomy of a Collision

While the Feynman diagrams for dijet production look simple, proton-proton collisions are extremely convoluted, as evidenced by Figure 3.2. In addition to the hard scatter and decays (in red), the event also contains many gluon emissions from the strongly charged particles (blue) as well as secondary interactions from the other proton constituents (purple). These colored particles then group into colorless baryons and mesons, which in turn have their own decays (green).




FIGURE 3.2. Simulation of a sample proton-proton collision. The incoming partons are in blue, the underlying event interactions are in purple, the hard scattering event is in red, and the hadronization processes are shown in green. Image from [25].

3.2.1. Hard Scatter and Parton Showering

The hard scatter is the interaction of interest in an event, with two partons interacting and producing the high transverse energy final state of interest. The Feynman diagrams for a collision correspond to the hard scatter, such as the $gg \rightarrow tth$ process shown in Figure 3.2. Here, the Higgs decays into a pair of quarks, and the two top quarks decay to bW , with one W decaying hadronically and the other

leptonically, as shown in red. In addition, the free partons from the hard process undergo parton showering, whereby a quark or gluon can radiate a nearly-collinear gluon which carries away some portion of its momentum. In addition, a gluon can produce a quark-antiquark pair. This process continues until the splittings reach low enough virtuality that can no longer be described perturbatively, beyond which point confinement takes effect and the partons *hadronize* into the colorless particles which can be observed by the detector.

3.2.2. Pileup and Underlying Event

The underlying event, shown in purple in Figure 3.2, is the interaction of the remaining partons in the pp collision which did not participate in the hard scatter. These remnants can interact with each other and produce softer interactions which then overlay with the products from the hard scatter of interest.

ATLAS produces collisions by colliding bunches with $\sim 10^{11}$ protons each, and as such a variable number of collisions occur each bunch crossing. For a typical data event taken in 2016, a single hard scatter event was also accompanied by some two dozen other softer scatters, often referred to as pileup. The distribution of the number of pileup interactions can be seen in Figure 3.3. Pileup takes two forms: in-time pileup which produces particles which appear simultaneously in the detector with the hard scatter products, and out-of-time pileup where particles from previous bunch crossings are still traveling through the detector.

The effect of both the underlying event and a single pileup event is roughly the same: it adds additional calorimeter noise and physics objects which make it more difficult to pick out only the hard scatter products. An example event is shown in Figure 3.4 where a single $Z \rightarrow \mu\mu$ event with 24 pileup interactions is shown. All of the

energy deposited in the calorimeter, with the exception of the small amount along the path of the muons, comes from the underlying event and pileup interactions. While this can have a large effect on searches using lower-momentum objects or missing transverse energy, the high-mass dijet search is relatively insensitive to their effects owing to the large discrepancy in energy scales between typical pileup interactions and the hard scatters of interest.




FIGURE 3.3. Luminosity-weighted distribution of the mean number of interactions per crossing for the 2015 and 2016 datasets.




FIGURE 3.4. Event display of a $Z \rightarrow \mu\mu$ event with 24 other reconstructed vertices from pileup interactions. The paths of the two muons are shown in red, while the blue lines are tracks with $p_T > 500$ MeV

3.2.3. Hadronization

As the partons and their showering products in the collision move to longer distances, the process of hadronization occurs whereby the colored partons are formed into colorless baryons and mesons which will then decay into the final state particles

which are seen by the detector. An exact description of hadronization in QCD is not known, but there are several models used in Monte Carlo generators which work very well, including the Lund string model [26] used in the Pythia [27] generator, and the cluster model [28] used in the Herwig [29] generator. The particular showering process used has almost no effect on the sensitivity of this analysis.

3.3. Jets

The partons from the hard scattering process are not experimentally observable, only the wide range of hadrons produced in the hadronization step. However, momentum and energy are conserved throughout the showering process, and as such reflect the kinematics of the originating particle. Jets are the observable physics object for hadronic particles, roughly-conical sprays of particles which have their energies added together into a single object.

The *jet algorithm* defines the method by which observables, such as the final-state particles in simulated data or calorimeter cells in a detector, are combined to create jets. There is no one algorithm which is the correct one, and ATLAS uses several different jet definitions depending on the use case. For example, the Level-1 jet trigger uses the very crude method of a square around a given high-energy trigger tower seed. This is a very poor definition in general, but works very well given the constraints of the trigger.

A good jet definition will be immune to small changes in the treatment of the clustered particles, often called infrared and collinear (IRC) safe. The jet definition will return the same number of jets in both the case of an infinitesimally soft (infrared) particle being added to an event, and in the case of a hard single particle being split into two softer, collinear particles. Some examples of the behavior of jet algorithms

which are infrared or collinear unsafe are shown in Figure 3.5. IRC safety is important for linking experimental results and theoretical predictions, as these processes happen randomly through parton showering and should not change the interpretation between the original scattered particle and the jet it creates.




FIGURE 3.5. Examples of the behavior of infrared and collinear unsafe jet algorithms. In the top two diagrams, the addition of a soft gluon becomes the seed which a jet is centered on, causing the two jets to merge together into one jet. In the bottom diagram, the splitting of the most energetic seed into two collinear ones, shifting the jet location and lowering the jet energy. Diagram from [30].

3.3.1. The Anti- k_t jet algorithm

A jet algorithm begins with a list of subjects, the individual inputs which will be combined to make jets. In the case of ATLAS, the subjects are calorimeter topoclusters which are discussed in more detail in Section 5.5, while for truth-level jets individual final state particles are used as subjects. The ATLAS jet definition of choice is the anti-

k_t algorithm.[31] The anti- k_t algorithm belongs to a family of sequential recombination algorithms which take the form:

$$d_{ij} = \min(k_{ti}^{2p}, k_{tj}^{2p}) \frac{\Delta_{ij}^2}{R^2} \quad (3.4)$$

$$d_{iB} = k_{ti}^{2p} \quad (3.5)$$

where $\Delta_{ij}^2 = (y_i - y_j)^2 + (\phi_i - \phi_j)^2$ is the distance between two subjects with rapidities y_i and azimuths ϕ_i , k_{ti} is the transverse momentum of the subject, and R is the radius parameter which sets the size of the jet. The parameter p determines the behavior of the algorithm; $p = 1$ corresponds to the k_t algorithm [32], while $p = 0$ gives the Cambridge/Aachen algorithm.[33] Setting $p = -1$ gives the anti- k_t algorithm.

The algorithm begins by calculating d_{iB} for each subject and d_{ij} for each pair of subjects, where the subjects are the input type to the algorithm. The recombination proceeds from the smallest value of d . If a d_{ij} is smallest, the two subjects are combined, both their energy and their position, and all d are recalculated using the new set of subjects. If a d_{iB} is smallest, the subject is considered to be a full jet and is removed from the list of subjects. An illustration of how this method proceeds is shown in Figure 3.6.

The general behavior of the anti- k_t algorithm is that the subjects with the highest transverse energy will cluster first, gathering together all of the smaller radiation around them, before the softer subjects begin clustering together. The algorithm tends to produce roughly-conical jets, and gives preference to more energetic jets when two jets are close together. An example of the final result of the anti- k_t algorithm is shown in Figure 3.7. Here, a single simulated event is shown along with thousands of very soft "ghost" particles which are randomly distributed and do not contribute any




FIGURE 3.6. An example of the anti- k_t recombination algorithm. a) The original input subjects to the algorithm. b) The highest p_T subject combines with its nearest neighbor to produce a new combined subject. c) This subject again combines with its neighbor to produce a new subject. d) The remaining subjects are separated by more than the radius parameter of the algorithm. The subjects become the final jets.

meaningful energy to the event. These ghosts are used along with the real physics energy deposits in the jet algorithm but do not change the shape or location of the jets, allowing for the visualization and measure of the area encompassed by each jet. This area is used to correct for the amount of pileup energy which falls within a jet.

As mentioned previously, the model of hadronization used to simulate the processes studied in this analysis has no meaningful impact on the sensitivity. This is because the treatment of the hadronization step does not impact the location of the high- p_T particles, but rather the distribution of soft emissions between them, and as such will have little impact on the p_T as measured by the jet algorithm.

The anti- k_t algorithm is both collinear and infrared safe. In the case of a collinear splitting of a hard particle into two softer ones, their small angular separation Δ_{ij} ensures that they will be clustered together at the beginning of the recombination,




FIGURE 3.7. A sample parton-level event together with many random soft particles, clustered together using the anti- k_t algorithm.[31].

recovering the single hard particle. The algorithm is infrared safe because a soft emission will first cluster to a highly energetic one and will not shift the jet, preventing any change in the final jet. The choice of jet radius is somewhat arbitrary and depends on the physics process being investigated. The standard ATLAS jet uses $R=0.4$, but other radii are used in specialized contexts such as $R=0.2$ jets to investigate jet substructure, and $R=1.0$ jets to capture the energy of boosted objects such a top quarks.

CHAPTER IV

MACHINE AND DETECTOR

The particle physics program at CERN uses some of the most complex machines ever built, and have required the efforts of thousands of people to design, commission, and maintain. This section outlines the systems which are used to provide the collisions and measure their products.

4.1. The Large Hadron Collider

The Large Hadron Collider (LHC) is the worlds largest and most powerful particle accelerator, at over 27km in circumference and capable of accelerating protons a center of mass energy of 13 TeV. It is located 100 meters underground along the French-Swiss border at CERN, just outside Geneva, Switzerland. The beams are brought to collisions at different points along the ring to four major experiments: ATLAS [34], CMS [35], ALICE [36], and LHCb.[37]

Protons in the LHC start as hydrogen gas which is ionized, the protons of which are accelerated to 50 MeV by the LINAC2 linear accelerator. From here, particles are accelerated in a series of circular accelerators of increasing size, starting with the Proton Synchrotron Booster which raises their energy to 1.4 GeV. This is followed by the Proton Synchrotron (PS) (25 GeV) and Super Proton Synchrotron (SPS) (450 GeV), after which proton bunches are finally injected into the LHC where they are ramped up to their final collision energy. For Run 1, the LHC operated $\sqrt{s} = 7$ and 8 TeV, while in the ongoing Run 2 at $\sqrt{s} = 13$ TeV. The CERN accelerator complex is shown in Figure 4.1.

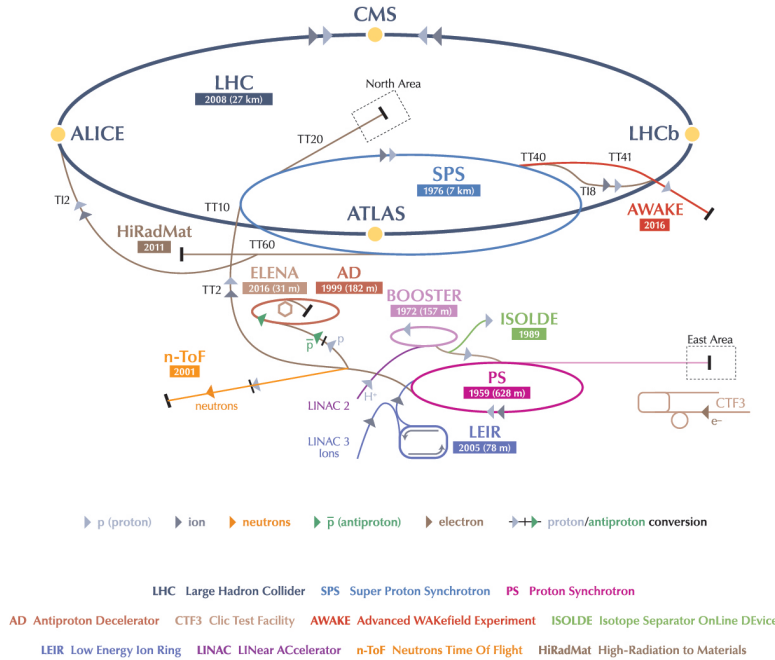


FIGURE 4.1. Schematic of the CERN accelerator complex.[38]

4.1.1. Bunches and bunch trains

The LHC operates at a radio frequency (RF) of 400 MHz which can be filled with bunches every 10 RF buckets for a separation of 25 ns between bunches.[39] During Run 1 and the early part of Run 2 the LHC operated with 50 ns spacing to reduce the effect of the electron cloud, but moved to the design spacing of 25 ns for the remainder of Run 2 data taking at the insistence of the detector collaborations. This spacing, coupled with the ~ 27 km circumference of the LHC, means there are 3564 bunches which could be filled in the ring at any time. In practice the LHC never fills all of these bunches at once, but instead the bunches are organized into *trains*.

Bunch trains are groups of sequentially filled bunches in the machine and are driven by the design of the PS and SPS rings. For a typical 2016 run, a single injection from the PS to the SPS was 72 sequential bunches, while the SPS transferred 288 bunches (4×72) at a time over to the LHC.[40] The gaps between each of these bunches are driven by the kicker magnet rise times for each of these systems. The minimum kicker times for the PS and SPS are 200 ns and 800 ns respectively, so each 72 bunch train is separated by 8 empty bunches, and each grouping of 288 is separated by 32 empty bunches. The filling scheme also includes an abort gap, a space of approximately 100 empty bunches which is used in the case that the beam needs to be redirected to the beam dump. In total, the number of filled bunches in the machine caps out around 2800 at a time.

While the overall luminosity can be directly increased by upping the number of bunches in the machine at one time, the LHC has also tested other filling schemes which reduce the number of filled bunches but increase the charge in each bunch, or decrease the beam emittance, leading to a higher instantaneous luminosity and overall integrated luminosity.

4.2. The ATLAS Detector

The ATLAS detector is a multipurpose physics detector, located at Point 1 on the LHC ring near the main CERN site. ATLAS weighs in at around 7000 tons, and is 44 m long and 25 m in diameter. The detector is composed of multiple layers which measure their own unique portions of the collision products, which are described in the following section. Figure 4.2 shows how these systems are laid out.




FIGURE 4.2. A cutaway view of the ATLAS detector showing the various components.[34]

4.2.1. Coordinate System

The ATLAS experiment uses a right-handed coordinate system, with the origin at the nominal interaction point (IP) in the center of the detector. The x-axis points from the IP towards the center of the LHC ring, and the y-axis points upward towards the surface. ATLAS uses a cylindrical coordinate system which defines a radial distance $r = \sqrt{x^2 + y^2}$ and azimuthal angle ϕ around the z-axis. Instead of using the polar angle from the beamline θ , pseudorapidity is typically used, defined as $\eta = -\ln(\tan \frac{\theta}{2})$.

Massive objects, such as jets, are measured to have a rapidity defined as $y = \frac{1}{2} \ln \frac{E+p_z}{E-p_z}$. In the limit where $p_T \gg m$, as is often the case in this search, $y \approx \eta$, meaning that the rapidity of a jet and its detector location given by pseudorapidity

are approximately interchangeable. Differences in rapidities are Lorentz-invariant under z-axis boosts, making them useful for measuring if two jets are back-to-back in their rest frame.

4.2.2. Magnet System

ATLAS is composed of four separate magnets: one central solenoid encompassing the ID, one barrel toroid, and two endcap toroids. The central solenoid provides a 2.0 T magnetic field and is responsible for bending charged particle tracks near the IP to measure their momentum. The namesake toroids provide a peak field of 4.0 T, bending muons in a different plane than the solenoid field does to improve their momentum measurement.

4.2.3. Inner Detector

The ATLAS Inner Detector (ID) is comprised of three different subsystems: the silicon pixel layers (including the Insertable B-Layer (IBL) new in Run 2), the semiconductor tracker (SCT), and transition radiation tracker (TRT). In total, the system covers the angular range $|\eta| < 2.5$. Information from all three subsystems is used to reconstruct tracks down to $p_T = 0.4 \text{ GeV}$. [41]

4.2.3.1. Pixel and Semiconductor Trackers

The pixel detector system is composed of four concentric layers of silicon pixels, providing multiple high-resolution hits for charged particles traveling through. The closest layer to the beampipe, the IBL, was added during Long Shutdown 1 (LS1) before the start of Run-2 operations. At only 3.3 cm away from the beamline, this additional tracking layer greatly improved the efficacy of tracking algorithms by an

additional hit for tracks passing through multiple layers, and allowing for better resolution of secondary decay vertices, such as those originating from b-quarks. The IBL also contains newer technologies than the rest of the pixel system which allow it to withstand the harsh radiation close to the beampipe. This will allow the IBL to continue performing well as the innermost layer of the old pixel system degrades due to radiation, allowing for the maintenance of the current tracking performance. The SCT uses silicon microstrips and sits just outside the pixel system. It is comprised of four barrel layers and 18 endcap disks (9 on each side of the detector), and range from 30 to 51 cm from the beamline, adding additional position information for track finding.

4.2.3.2. Transition Radiation Tracker

The TRT operates on a very different principle than the other two tracking systems, taking advantage of the energy emitted when high energy particles transition between media with differing indices of refraction. It is comprised of nearly 300,000 polyimide straws with an external aluminum layer and central gold-plated, tungsten wire at the center of each 4mm diameter tube. The straws are filled with a gas mixture of Ar, CO₂, and O₂, and are ionized either by charged particles passing through or by emitted transition radiation. Each straw hit provides additional 2-D hit information for track finding, as well as acting as a discriminant in electron identification.

The TRT is scheduled to be removed as part of the ATLAS Phase-II upgrade starting in 2025, as the pileup conditions in Run 4 are expected to create TRT occupancies near 100%, significantly degrading the performance of the system. The average TRT occupancy as a function of the number of interactions per crossing is shown in Figure 4.3.




FIGURE 4.3. Average track region TRT occupancy as a function of track η using data taken during the 2015 and 2016 ATLAS runs for four different values of the average interactions per bunch crossing.

4.2.4. Calorimeters

ATLAS uses two different types of calorimeters to measure the energies of leptons, hadrons, and photons. First is the liquid argon (LAr) electromagnetic calorimeter, which captures most of the energy from EM showers. Second is the steel-scintillator Tile calorimeter which captures the energy from hadronic showers. Optimal performance of the calorimeter systems under challenging conditions is vital to nearly every physics signature ATLAS searches for. The system has nearly hermetic coverage in ϕ but is made up of several overlapping systems and varying thicknesses of material in η , and so special care must be taken to accurately calibrate the detector response in each region of the detector.

The journey from calorimeter hits to the eventual trigger is discussed in detail in Chapter V.

4.2.4.1. Electromagnetic Calorimeter

The EM calorimeter uses liquid argon as an active medium, a choice motivated by the need for radiation-hardness, especially over the 30+ year run time of the LHC. The calorimeter is composed of accordion shaped absorbers of lead with steel backing, with liquid argon flowing between and copper electrodes which measure the energy deposited.

The LAr calorimeter is a sampling calorimeter, meaning that only a small fraction of the total energy from a particle traversing it is actually measured. The power of an absorbing material is given by its radiation length (X_0), the distance an electron must travel in that material for its energy to be reduced by a factor of $\frac{1}{e}$. Liquid argon is not dense enough as a medium to fit in a reasonable size calorimeter, and so lead (with a radiation length $20\times$ shorter) is used to slow down the particles. The EM calorimeter contains enough material to fully contain all but the largest electromagnetic showers, typically 25-30 interaction lengths depending on the region of the detector. The actual pulse observed in the calorimeter is created when charged particles, showering from interactions in the absorber layers, cause the LAr to ionize and then release electrons which are accelerated by an electric field to the electrodes to become the calorimeter pulse. The amount of energy actually measured by the LAr is much less than the total shower energy, and so it must be scaled back up by comparing the behavior of particles using test beam data where the incoming particle energy is well known.

The EM Barrel covers the central region out to $|\eta| < 1.475$, and two endcaps cover the range $1.375 < |\eta| < 3.2$. The calorimeter is divided up into four layers of varying granularity, the first of which is the presampler. This layer is closest to the inner tracker and has no absorber material in it, but instead is used to measure

the amount of energy which has dissipated due to interactions with the material in the inner detector before it reaches the calorimeters. The other three layers vary in thickness and granularity, and the values from each of these layers help determine the shape evolution of the shower as it progresses through the calorimeter, an important tool for electron/photon identification. These pulses are read out and given to the level-1 trigger at coarse level, and if an event passes the trigger they are passed with full granularity to the high-level trigger.

4.2.4.2. Hadronic Calorimeter

Situated immediately outside the EM calorimeter, the hadronic calorimeter systems capture the remaining energy from hadrons which have already deposited a fraction of their energy in the EM layers. The depth of material is characterized by the number of nuclear interaction lengths, the distance a hadron travels before undergoing an inelastic collision. The depth varies as a function of η , but the hadronic calorimeters typically comprise 8-12 interaction lengths, enough to stop all but the most powerful showers. In the barrel region, it is comprised of steel absorbing layers and scintillating polystyrene tiles, and the deposited energy is measured by molecular excitations in the scintillator emitting UV light which is then read out by photomultiplier tubes (PMTs). The barrel systems covers the region $|\eta| < 1.7$, but a different, more radiation-hard technique is required for the more forward regions.

The Hadronic Endcap (HEC) covers the range $1.5 < |\eta| < 3.2$, and operates in largely the same manner as the EM systems. Instead of the lead-steel combination of the EM layer, the HEC uses much thicker copper layers as the absorbers, with LAr still acting as the active medium. The forward calorimeter (FCAL) covers the range $3.1 < |\eta| < 4.9$ and uses a combination of copper and tungsten plates as the absorbers.

Overall the hadronic calorimeters are much coarser than their EM counterparts, with a granularity of 0.1×0.1 ($\Delta\eta \times \Delta\phi$) in the region $|\eta| < 2.5$, and 0.2×0.2 in the forward regions.

4.2.5. Muon Spectrometers

Muons interact very little with the calorimeters in ATLAS, leaving behind only tracks in the inner detector before they reach the outermost systems in ATLAS, the muon spectrometers. In the muon systems, ATLAS's namesake toroids bend the tracks of passing muons to determine their momenta. The muon system uses different technologies in the barrel and endcap regions, as well as different systems for triggering and momentum measurement.

In the central region, monitored drift tubes (MDTs) are used to measure the tracks, and thus the momenta, of muons, while resistive plate chambers (RPCs) are used for triggering. This split is required because of the long drift time of the MDTs of ~ 700 ns, which is outside of the window the level-1 trigger system can use. To deal with the higher backgrounds in the forward regions, cathode strip chambers (CSCs) are used for momenta measurements and thin-gap chambers (TGCs) are used for triggering.

The muon system operates under the assumption that all other particles will be fully absorbed in the calorimeter layers, and thus any detected particles in the region must be muons. However, in the case of very high- p_T jets some of the energy can "punch through" the calorimeters and register hits in the muon system. This effect must be compensated for and is discussed later in the context of jet calibration.

4.2.6. Forward Detectors

Further down the beampipe from the IP are the various forward detectors for ATLAS which are primarily used for luminosity and cross-section measurements. LUCID (LUminosity measurement using Cerenkov Integrating Detector) sits 17 m down from the IP and consists of quartz fiber bundles connected to PMTs which detect Cherenkov radiation as charged particles pass through the quartz. The multiplicity is proportional to the amount of interactions taking place at the IP, and is therefore used to measure the instantaneous luminosity. ALFA (Absolute Luminosity For ATLAS) uses Roman Pots 240 m from the IP to detect protons at angles very close to the beamline; ALFA is used in dedicated runs to measure the total proton-proton cross-section. ZDC (Zero Degree Calorimeter) sits 140 m from the IP and is used to measure the amount of neutral particles (neutrons and photons) left as beam remnants once the rest of the beam has been bent away by the steering magnets of the LHC; this information is used as a minimum-bias trigger for ATLAS, and as a centrality measurement for heavy-ion collision data. Finally, AFP (ATLAS Forward Protons) is new for Run 2, primarily focusing on diffractive physics in special low-luminosity runs.

4.2.7. Trigger and Data Acquisition

During nominal operations, the LHC provides collisions at a rate of 40 MHz, or an event every 25 ns. As each event is slightly under 2 MB of data, this represents an extraordinary amount of information, and is far more than could be conceivably be written to disk, much less properly analyzed. In order to reduce the data rate to a manageable level, and to ensure that only the events with the most interesting physics are saved, a two-tiered trigger system is used.

4.2.7.1. Level-1 Trigger

The level-1 (L1) trigger is hardware based, making decisions on whether or not to keep an event within $2.5\ \mu\text{s}$. For Run-2 operation, the L1 trigger reduces the event rate to a maximum of 100 kHz. The L1 trigger is made up of several subsystems which bring together decisions from different areas of the detector to make a final decision on whether or not to keep an event.

The L1 calorimeter trigger (L1Calo) uses information from the EM and hadronic calorimeters in the form of rough 0.1×0.1 ($\eta \times \phi$) trigger towers. From this information, candidates for jets, electrons/photons, and taus are constructed and passed along as regions of interest (RoI) to the next stages in the trigger telling the type of object, its transverse momentum as measured at L1, and its location in the detector. The amount of global missing transverse energy (MET) is also calculated and used as part of the trigger signature. The L1Muon system uses a coincidence of hits in the RPCs and TGCs to identify muon candidates. Since precision momentum information from the MDTs and CSCs takes too long for the L1 decision, the muon momentum is roughly estimated using the hit pattern, and then marked as being above one of a set number of p_T thresholds.

Within each system, a list of pre-determined triggers are compared against the RoIs found in the event. For example, the muon system has triggers for muons above 20 GeV (L1_MU20) as well as for two muons above 6 GeV (L1_2MU6). The list of all triggers passed in each system is passed along to the Central Trigger Processor (CTP) where the final determination on whether or not to save an event is made.

RoIs from both L1Calo and L1Muon are then sent to the L1 Topological trigger. (L1Topo) Newly commissioned in 2016, L1Topo allows for more advanced combinations of trigger objects from the two systems, including operations such as

invariant mass, angular separation between objects, and transverse mass.[42] This additional processing step greatly reduces the background rates for certain processes, allowing for much lower energy thresholds when using objects triggered from L1Topo. Similar to L1Calo and L1Muon, the list of triggers passed by an event is passed along to the CTP.

The CTP is the final stage of the L1 trigger. It receives a list of all passed triggers from the other subsystems and compares it against an L1 trigger menu. The trigger menu determines which triggers should be turned on or off during a given run, as well as the prescale, or rate reduction, rates for triggers. Triggers with lower energy thresholds are often prescaled, meaning that only a fraction of events that pass a given trigger are kept. For example, a physics menu item with a prescale of 20 will have only 1 out of 20 events that pass the trigger saved, while the other 19 times the event is not saved unless some other trigger is also passed and marks the event for saving. The L1 trigger menu evolves over the course of time as well as within a single physics run, changing the prescales of physics items to keep the rate of saved data as close to 100 kHz as possible as the instantaneous luminosity decays. If the CTP determines that an event passes at least one trigger after prescaling, then the event is saved. At this point a level-1 accept signal is sent which triggers the readout of the full granularity detector data. The result of which triggers were passed at the CTP, along with the RoIs, are then passed along to the high-level trigger.

4.2.7.2. High Level Trigger

Events passing the L1 trigger are fed to the High-Level Trigger (HLT), which uses the full detector readout and granularity to measure physics quantities at close to offline precision. For Run 2, the HLT accepts events at nearly 100 kHz and reduces

this rate down to 1 kHz for the general physics stream. The HLT runs on a general computing cluster with thousands of cores and uses much more complex software algorithms for event analysis, such as b-tagging and anti- k_t jet clustering.

HLT trigger chains are seeded by an matching relevant L1 trigger. For example, the single-jet trigger HLT_j380 is seeded by the L1 trigger L1_J100. As events can pass multiple trigger thresholds and types, multiple HLT algorithms can be run on a single event. Each HLT item corresponds to a chain of processing steps which are optimized to reject events as early as possible, using the full readout and precision processing steps only on the most likely candidates. There are several different types of HLT chains which are used, including primary chains, the unrescaled standards used by most analyses, supporting chains, which are typically rescaled and are used mainly for determining backgrounds for other analyses in regions which can not support an unrescaled trigger, and monitoring chains, which readout only a small portion of the whole detector and are used to check the performance of specific subsystems.

The debug stream includes events which are unable to be processed in the HLT, either due to them timing out after taking too long to process, or because of some issue which caused a crash. All events which cause errors in the HLT are saved to the debug stream to ensure no interesting events are lost, and to troubleshoot which algorithms are problematic. The debug stream typically only contains a handful of events per run, but on occasion HLT farm crashes have caused hundreds of thousands of events to be saved to it. ATLAS also uses a data scouting stream which supports lower threshold physics searches by giving them drastically increased event rate in exchange for saving only a small portion of the event data. This stream is used by the ATLAS Trigger Level Analysis (TLA) to look at dijet events in the invariant mass range below the reach of the standard dijet search.

4.2.8. Dead Time

The ATLAS detector is limited in how many events it can process in a given amount of time, leading to "dead time" for the detector when collisions are occurring and the detector is operating nominally, but events are unable to be saved. The first type is referred to as simple dead time, which occurs after every L1Accept and takes about 100 ns, or 4 bunch crossings. During this time the detector is transferring the full readout to the front end buffers for processing in the HLT and can not save any other events. The second type is complex dead time which comes about from the limited size of these readout buffers for individual subsystems and the rate at which data can be sent off-detector. Complex dead time limits the number of accepted events in a given period of time; if too many events occur too quickly, the readout buffers fill up and events must be vetoed until there is space for another event in the buffer. The complex dead times are controlled by four "buckets" which are adjustable between runs depending on the requirements of the different hardware subsystems, characterized by a max number of L1Accepts in a given number of BCs. For example, 15/430 means that up to 15 events can be saved within a window of 430 BCs, and triggers are put on hold until the rate drops down below this level. For the dataset used in this analysis, ATLAS recorded data with 92% efficiency. The total amount of dead time varies based on the luminosity, but typically covers $\sim 3\%$ of collision time. The remaining inefficiency comes from reducible sources such as detector errors which send a busy signal and hold triggering while the system fixes an issue or resets a component.

4.2.9. Data Processing

Events which are accepted by the HLT are written to disk and are passed to the ATLAS Tier-0 computing facility at CERN, where the detector readout is transformed from raw data into the physics objects that will eventually be used for analyses. Immediately after a run is taken, the express stream (a small subset of the full physics data stream) is run through what is known as the calibration loop, a ~ 48 hour process in which the geometry and conditions for the run are calculated to properly process the whole dataset.

Once calibration is complete, the data enters bulk reconstruction where the full dataset is run through the standard ATLAS software to change the data from raw detector readout to containers which describe objects such as electrons and jets, or more granular objects such as track-candidates or calorimeter clusters. This data format is known as xAOD, or Analysis Object Data. This dataset is then duplicated around the world to the various computing facilities around the globe. From here, the data is run through the Derivation Framework.

To prevent analyses from needing to run over the entire ATLAS dataset, each group of analyses sharing a similar signature creates a specialized derivation which skims through the full dataset and only keeps events which are likely to be relevant to the analysis, and allows groups to remove unneeded data containers, or to create their own data containers which are not part of the standard set. For example, the derivation which the dijet analysis uses keeps only events which pass one of the single- or multi-jet HLT chains, and gets rid of muon and electron data as those objects are not used in any way in the analysis. The typical derivation is only $\sim 1 - 2\%$ of the full AOD in size, making it much more manageable for analyses to run their search over. These derived AODs (DxAODs) are then duplicated to multiple places in the CERN

computing grid and become available for analysis use. The main physics stream for ATLAS recorded 1.7 billion events in 2015 and 5.4 billion in 2016, comprising some 6.2 petabytes of raw data, making running over the full dataset exceedingly difficult!

CHAPTER V

CALORIMETRY AND TRIGGERING

The dijet analysis is highly dependent on the efficiency and accuracy of the jet triggering and reconstruction algorithms, as well as the performance of the calorimeters to hadrons. This section describes the long journey from when a particle first interacts with the ATLAS detector to when the decision is made to record the event. While the LAr side is described here (and in much greater detail in Ref. [43]), the method is very similar for hits originating in the tile calorimeter. An overview of the chain from calorimeter to data storage is shown in Figures 5.2 and 5.3.

5.1. The LAr Calorimeter Pulse

As particles from collisions traverse the LAr calorimeter, they interact with the absorber layers and produce showers of charged particles. As these particles travel through the liquid argon it ionizes it, releasing electrons in proportion to the deposited energy of the particle. The electrons are accelerated by the field created by the copper electrode, quickly reaching their terminal velocity in the fluid, and create the calorimeter pulse.

The created pulse is roughly triangular, with a very sharp rise time of a few nanoseconds and a very long tail of $\sim 450ns$ due to the long drift time for all of the electrons to cross the gap. This pulse is transmitted to the Front End Board (FEB) which is mounted on-detector and handles the shaping, processing, and digitization of the pulses before it is sent off-detector following a Level-1 accept. To begin, the input current pulses are converted to a voltage which is sent along four paths for two different purposes: one path goes through layer summing to be sent to the L1Calo

system (described in Section 5.3), while the other paths prepare the signal to be sent off detector.

5.2. Readout Path

The calorimeter pulse is sent through three paths which are each run through a bi-polar shaper, resulting in the zero-area shape seen in Figure 5.1. The pulses are then amplified using three different gains – 0.8, 8, and 80. This is done because of the energies of the detected particles encompass several orders of magnitude, and this ability to select a gain for a given channel allows for the precise measurement of smaller energy deposits without losing sensitivity in the case of the most energetic particles. The pulses are then sampled at 40 MHz and stored in switched-capacitor array (SCA) analog pipelines during the L1 trigger latency period. Each channel has an adjustable delay of 0-17.5 ns in steps of 2.5 ns, allowing of a delay to be added to ensure that the peak of the pulse is close to one of the samples.




FIGURE 5.1. LAr Calorimeter pulse shape as produced in the EM barrel before and after shaping. The dots represent the 25 ns spacing used in the pulse sampling and digitization.[44]

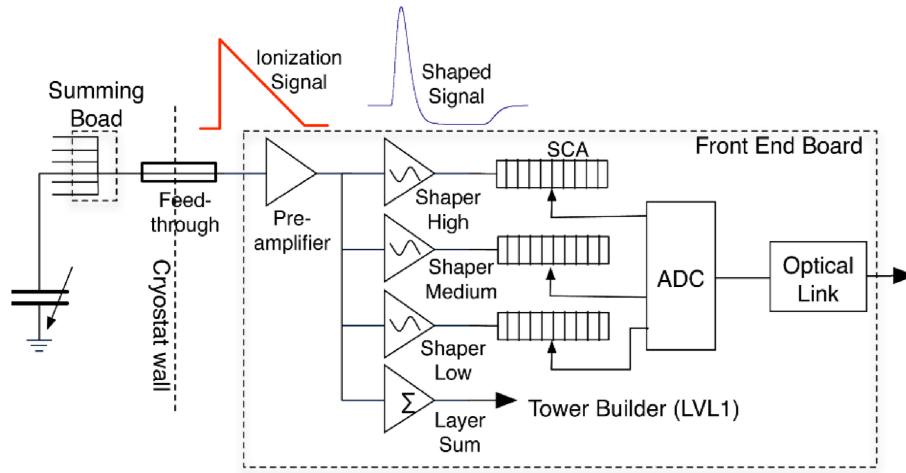


FIGURE 5.2. The ATLAS LAr Calorimeter front-end system in Run 2, along with the connections to the L1 trigger and readout systems.[45]

If the event is selected by the L1 trigger system, a Level-1 accept is sent which instructs the board to read out the samples for the corresponding bunch crossing from the SCA where they are digitized by a 12-bit analog-to-digital converter (ADC). The same gain is used for all channels to reduce any possible systematic effects, and the highest gain which does not saturate a channel is chosen, allowing for the highest possible granularity based on the energy deposited. The digitized samples are then temporarily stored in memory on the FEB before being sent across optical links to the read-out driver (ROD).

5.3. Energy Summing and Triggering Path

Parallel to this process, each set of four calorimeter channels on the FEB is passed to a linear mixer where they are summed together, and the resulting pulse is shaped in the same manner as the individual channels. Each FEB only processes channels from a single layer of the calorimeter, as different layers can have very different pulse shapes and delays. The resulting pulse is then sent to a neighboring Tower Builder

Board (TBB) which sums together the four layers of the calorimeter (presampler, front, middle, back) into a single pulse. To do this, it first equalizes the widths of the pulses from the different layers, adjusts the delay for each layer depending on cable length and time-of-flight in the calorimeter, and equalizes the relative gains of the different layers, before finally adding them together into the final pulse. Additionally, the energy scale is converted from raw energy to transverse energy, the quantity used for most trigger measurements.

These sums are the objects used by the L1 system and are known as trigger towers. In the central part of the calorimeter the trigger towers cover an area of 0.1×0.1 ($\eta \times \phi$), with the EM and hadronic layers each as their own distinct tower. In the forward regions of the calorimeter the towers become 0.2×0.2 ($2.5 < |\eta| < 3.2$) and 0.4×0.4 ($3.2 < |\eta|$).

5.4. The Level-1 Calorimeter System

Analog trigger tower signals from the calorimeter systems are sent off detector to the USA15 cavern adjacent to ATLAS to the Level-1 Calorimeter system for further processing. Signals are sent via twisted-pair cables which are routed through special holes between the detector cavern and the trigger system to minimize the latency of the system. The L1Calo system is described in detail in Ref. [46] and the changes for Run 2 in Ref. [47]; the processes most relevant for jet triggering are summarized here.

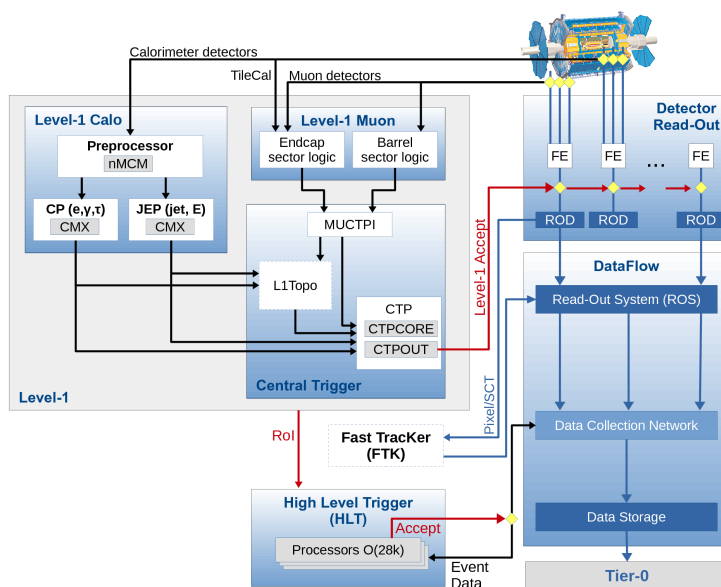


FIGURE 5.3. The ATLAS trigger and data acquisition system in Run 2. L1Topo was being commissioned in 2015. [47]

5.4.1. Receiver and PreProcessor Module

Their first stop are the L1Calo receivers which apply an $|\eta|$ -dependent gain adjustment to equalize the gain across the calorimeter, as well as to convert the hadronic towers from the tile calorimeters from energy to transverse energy. These signals are then sent to the PreProcessor Module (PPM) which handles the digitization of the calorimeter pulses as well as bunch crossing identification and the final calibration of the trigger towers before their use in the trigger identification algorithms.

Pulses are digitized by a flash-ADC with 10-bit precision, giving quarter-GeV steps and a saturation value of around 250 GeV for a single tower. For standard operation 5 samples, spaced by 25 ns, are read out, though this number can be configured for special runs. This information is used to understand the pulse shape and to apply the Bunch Crossing Identification(BCID) logic. To compensate for the different arrival times of each trigger tower due to varying cable lengths between the detector sections and the L1Calo system, the PPM also includes programmable delays in 1 ns steps to synchronize all of the trigger tower data. Trigger tower pulses have rise times of approximately two bunch crossings, and as such the pulses from consecutive bunch crossings overlap each other. To make sure that signals are associated with the correct bunch crossing, two different identification methods are used depending on the energy in a given tower.

For standard, unsaturated signals, the pulse is run through an auto-correlation filter to sharpen the pulse above the background noise level. This is done by multiplying the five samples by individual coefficients and summing result. Each section of the calorimeter has varying coefficients due to differences in pulse shape, but the general structure is to multiply by large coefficients for the central samples,

and have small or negative coefficients for the samples further away in time. This sum is then compared to the values obtained for the previous and following bunch crossings; if this value is a local maximum in time, then it is considered to be the peak of the pulse and it is assigned to that bunch crossing.

In the case of saturated signals, the peak is estimated by inspecting the behavior of the leading edge of the pulse and comparing it against predetermined low and high thresholds to determine where the peak is located. For pulses with only one saturated sample the peak is simple enough to determine; for pulses with multiple saturated samples the peak is either the first or second saturated sample.

It is important to note that the BCID logic does not select an algorithm based on whether or not the pulse is saturated, but rather both algorithms run in parallel for all pulses. In the case that the two algorithms point to different bunch crossings, the earlier of the two peaks is the one that ends up being used as the second peak falls into the simple dead time after the Level-1 Accept (L1A) from the saturated tower. The L1A is guaranteed because a saturated trigger tower is set to the maximum possible energy in each step of the jet algorithm, and will therefore trigger several different L1 jet algorithms.

If a pulse is found to peak in a given bunch crossing the result of the filter sum is sent to a look-up table to extract the final transverse energy which is used in the trigger algorithms. This information is sent as an 8-bit value from 0-255 GeV; saturated towers are assigned to the maximum possible energy, and towers which do not correspond to the peak of a pulse are set to zero. Following this step groups of 4 towers are summed to form 0.2×0.2 ($\eta \times \phi$) regions which are used in the Jet/Energy-sum Processor, with the sums now having a maximum value of 511 GeV.

If any of the towers were saturated in the previous sum, these towers are also set to their maximum value.

The PPM also handles the dynamic pedestal correction, a new addition in Run 2. For each bunch crossing in the machine, and for each channel, the PPM keeps a running value of the average filter output run over many bunch crossings. This gives an approximation of the average noise and pileup energy for a given tower, and this value is then subtracted from the peak finder filter output to give a corrected value which is translated to an energy via the look-up table. This bunch-dependent method is responsible for keeping L1 missing energy trigger rates in check with the increase in luminosity by cutting down on the amount of triggers that happen earlier in the bunch train; the pulse behavior in the beginning of a bunch train is different because it lacks the overlaid pulses and undershoots that bunches later in the train encounter. The magnitude and effect of this pedestal correction is shown in Figure 5.4. The correction spikes at the beginning of each bunch train and tapers off after the first few bunches.

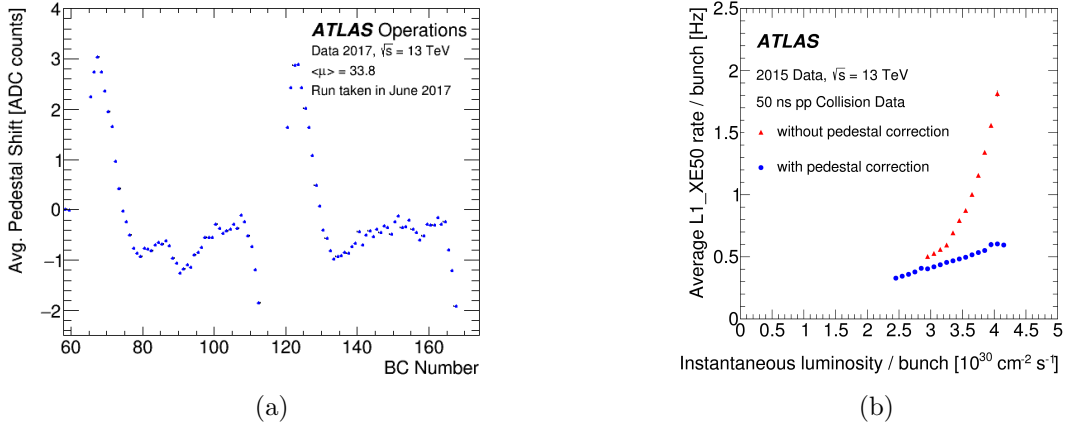


FIGURE 5.4. (a) The average pedestal correction per trigger tower as a function of the Bunch Crossing ID (BCID) in the Electromagnetic Barrel partition of the LAr calorimeter during a 2017 run. 1 ADC count is approximately equal to 0.25 GeV of transverse energy in a single tower.[48] (b) The per-bunch trigger rate for the L1 missing transverse momentum trigger with a threshold of 50 GeV (L1_XE50) as a function of the instantaneous luminosity per bunch. The rates are shown with and without pedestal correction applied. [47]

5.4.2. Cluster and Jet/Energy-Sum Processors

The Cluster Processors (CP) and Jet/Energy-Sum Processors (JEP) handle the calculations for detecting physics objects, with the CP handling electron/photon and tau algorithms while the JEP handles jets and missing transverse energy. The CP handles 0.1×0.1 calorimeter data out to $|\eta| < 2.5$ while the JEP uses the larger 0.2×0.2 towers out to $|\eta| < 3.2$ and the full calorimeter for the missing transverse energy algorithms. The CP uses the separate EM and hadronic tower information for hadronic isolation algorithms, but in the JEP this is unnecessary and so the EM and hadronic trigger towers are added together to create the 0.2×0.2 jet elements which are used in the trigger algorithms with transverse energy values up to 1023 GeV with 1 GeV granularity. If either the EM or hadronic towers are set to their maximum value indicating saturation, then the resulting jet element is also set to its maximum value.

The L1 jet algorithm uses a rectangular sliding window of 4×4 jet elements (0.8×0.8 in $\eta \times \phi$) which moves along the whole detector and looks for windows with a transverse energy above a given jet threshold. The algorithm first creates 2×2 core sums which are required to be local maxima; those that pass then have the transverse energy of the full window compared against a set of energy thresholds. The jet trigger algorithm window is shown in Figure 5.5; the original L1Calo designed allowed for algorithms which used the same sliding window and local maxima technique but with smaller jet areas of 2×2 or 3×3 jet elements, however these are not currently used for any jet triggers. All jet candidates are saved as Trigger Objects (TOBs) which carry the jet transverse energy as well as the $\eta - \phi$ position of the core, and are passed to the L1Topo system for use in topological algorithms. In the case of a Level-1 accept, the entire calorimeter is read out and used for jet reconstruction in the HLT.




FIGURE 5.5. Jet trigger algorithm windows using 0.2×0.2 ($\eta \times \phi$) jet elements with the local maxima cores shaded in. Only the 0.8×0.8 window is currently used for jet finding.[46]

5.5. High-Level Trigger

The High-Level Trigger has much longer time to process an event (500 ms compared to $2.5 \mu\text{s}$ for Level-1) as well as access to full-granularity information from the calorimeter. This allows the HLT to use much more precise, and computationally expensive, algorithms to calculate jet energies at levels very close to offline performance. Groups of calorimeter cells with significant deposits of energy are

grouped together using the technique of topological cell clustering, or topoclustering, which is described in Ref. [49] and summarized here.

Topoclustering attempts to create groups of topologically connected calorimeter cells which have signals above the background of the combined noise from electronics and the background pile-up in an event. To do this, the energy in a given cell is compared against an estimated average noise value assigned it using the ratio in Equation 5.1.

$$\zeta_{cell}^{EM} = \frac{E_{cell}^{EM}}{\sigma_{noise,cell}^{EM}} \quad (5.1)$$

Topoclustering proceeds using a 4-2-0 algorithm, corresponding to the values of $|\zeta_{cell}^{EM}|$ used for each step of the clustering algorithm. This particular configuration was found to provide the best energy resolution for pions in test beam data.[50] Topoclusters are seeded by cells which have energies at least four standard deviations larger than the noise. From these seeds, all neighboring cells are checked to see if they are more than two sigma above their noise value; if so they are added into the cluster. Neighboring cells are checked in three dimensions, so adjacent cells in the same layer are checked as well as cells overlapping the previous cell in the $\eta - \phi$ plane in an adjacent layer. This checking procedure is applied iteratively, so all neighbors to this first layer are checked as well, repeating until there are no more contiguous cells above two times their noise value. Finally, a perimeter layer using all cells regardless of energy is added to the cluster.

The goal of topocluster creation is to suppress the effects of noise and pileup by requiring highly-significant energy deposits to create and expand clusters, but it also allows for softer radiation to be captured in the perimeter step to maximize the amount of jet energy captured. The final cluster can grow to be very large

and encompass a significant portion of the detector depending on the distribution of particles in the event, and so a final splitting step is added to differentiate between close-by particles within the cluster. This is done by re-clustering the cells included in the current topocluster using all cells which are above a given energy threshold and are also local maxima.[51] The clusters are then created in the same manner as before, but only using the cells previously selected, not re-applying the threshold requirement, and not allowing clusters to overlap and merge. When two or more clusters are adjacent to a cell the energy is split between the two clusters based on their energies and distance from the shared cell. The final cluster is then kept if the total transverse energy is above a certain cutoff, with 400 MeV used as the threshold for the 2016 data reconstruction.

These topoclusters are what are used as inputs for the HLT jet reconstruction algorithm, which each cluster being treated as a massless four-vector with an energy equal to the sum of all of the cells in the cluster and a location based on the energy-weighted center of the cluster. Jets are reconstructed using the anti- k_t algorithm [31] with radius of $R = 0.4$ for single-jet triggers, and are calibrated using an approximation of the full offline calibration sequence.[47]. If the final jet energy exceeds the given energy threshold for its trigger chain then the event is saved to disk and processed for offline analysis.

CHAPTER VI

TRIGGER AND CALIBRATION

6.1. Trigger Strategy

The dijet analysis uses events that pass the lowest unrescaled single-jet trigger. For data taken in 2015 this was HLT_j360, meaning that the high-level trigger found at least one jet with 360 GeV of transverse momentum. In 2016, due to increased trigger rates from increasing luminosity and pileup, HLT_j380 was the lowest unrescaled single-jet trigger. For consistency, HLT_j380 is used as the final trigger selection for both datasets. This trigger becomes fully efficient ($>99.5\%$) around 420 GeV, and the final analysis cut of 440 GeV on the leading jet is well above this point.

6.2. Event Selection Criteria

In addition to the trigger requirement, physics events must pass checks related to the status of the detector at time of data taking, as well as the integrity of the event data. The standard method used by all analyses to determine which events are usable for analysis is the Good Runs List (GRL). GRLs are created centrally by the ATLAS Data Quality group and are based on data quality flags set by individual detector subsystems on a lumiblock-by-lumiblock basis for each data run. Lumiblocks (an ATLAS time unit of uniform detector conditions, typically ~ 1 minute) in which a subsystem sees some kind of error or degradation of performance can be flagged as having either tolerable or intolerable defects. For example, a single drawer of the tile calorimeter being offline is a tolerable defect, as it only covers a small portion of the total detector volume and the energy loss can be corrected for offline, whereas

a cooling failure knocking several drawers offline at once would be marked as an intolerable defect.

Several different GRLs are created depending on the needs of particular analyses. The most general (and widely used) of these is the All Good list, in which no systems report an error during a lumiblock. The dijet analysis uses two different varieties of GRL for the two different years of data taking. The 2015 data uses the `data15_13TeV.periodAllYear_DetStatus-v79-repro20-02.DQDefects-00-02-02_PHYS_StandardGRL_All_Good_tolerable_IBLSTANDBY-DISABLE` GRL. In addition to the standard All Good lumiblocks, this also includes those in which the innermost layer of the tracker was disabled. Extensive studies were conducted during the 2015 publication to compare data taken with the IBL off to the nominal dataset, and no significant differences were observed. The 2016 data uses the `data16_13TeV.periodAllYear_DetStatus-v84-pro20-16.DQDefects-00-02-04_PHYS_StandardGRL_All_Good_25ns_ignore_TOROID_STATUS.xml` GRL. This includes runs in which the toroid magnet was either off or ramping; similarly, this data was studied and found to have no significant differences compared to the nominal runs.

Individual events are also flagged as bad due to the presence of errors in systems that do not persist for the duration of a full lumiblock. Events with calorimeter noise or other issues can be marked with either warning or error flags for both LAr and Tile. Events with an error flag in either system are discarded. Additionally, events flagged as having incomplete event data are removed from consideration.

The total integrated luminosity for the dataset used in the analysis is 37.0 fb^{-1} . The total instantaneous luminosity \mathcal{L} is given by:

$$\mathcal{L} = n_b \frac{\langle \mu_{vis} \rangle f_r}{\sigma_{vis}} \quad (6.1)$$

where n_b is the number of colliding bunches, f_r is the bunch revolution frequency, μ_{vis} is the average visible number of inelastic collisions per bunch crossing, and σ_{vis} is the visible pp inelastic cross-section.[52] During data taking, μ_{vis} is directly measurable, while σ_{vis} is determined using special beam separation scans performed at a few points during the year. In these, the absolute luminosity is determined directly from measurements of the beam parameters, allowing for calculation of σ_{vis} . The instantaneous luminosity is averaged over an entire lumiblock and multiplied by the lumiblock duration to give the integrated luminosity. All of the lumiblocks which pass the GRLs used in this analysis combine to give the total luminosity.

6.3. Jet Formation and Definition

The jets used in the analysis are reconstructed from energy deposits in the calorimeter known as topoclusters.[49] Topoclusters are groups of neighboring cells in the calorimeter which contain deposits of energy which are well above the noise expected from the electronics and from pileup energy during events. A topocluster is seeded by a cell which is at least 4 standard deviations above the noise distribution, grows to add all contiguous neighboring cells which are above 2σ , and finally adds a perimeter around the cluster of all cells above 0σ . This process suppresses pileup and noise through its strong seed requirement, but also prevents softer jet radiation close to the noise threshold from being discarded.

Topoclusters are then clustered into jets with FastJet using the *anti* - k_t algorithm with a radius of $R=0.4$. This is the standard jet definition used by ATLAS; Run-1 searches used a wider jet radius of $R=0.6$, and a comparison between $R=0.4$

and $R=0.6$ using the Run-1 dataset showed no significant difference in sensitivity between the two jet sizes.

6.4. Jet Calibration

The multi-step process used by ATLAS to bring jets from the energy deposited in the calorimeter up to the final jet energy scale is described in Ref. [53], and is summarized here. The first step of the calibration is the origin correction, recalculating the jet four-momenta based on the primary vertex of the hard scatter as reconstructed by the tracking system, rather than the default center of the detector. This does not change the energy of the jets, but improves the angular resolution.

The next steps in the chain correct for the effects of pileup on the measured jet energy. First, an area-based p_t subtraction is applied to each jet based on the energy density in the central portion of the calorimeter for an event. Then, a residual correction is applied to jets based on the number of interactions per crossing, μ , and the number of primary vertices in the event, N_{pv} , based on coefficients derived from simulated data. In this way the jet energy is corrected to a level where it is the only interaction in an event.

The Jet Energy Scale (JES) calibration is derived using simulated data matching reconstructed jets to corresponding truth jets and energies. From this the average calorimeter response for jets is calculated as a function of the truth jet energy and η , as the calorimeter response varies significantly between regions of the detector. An additional correction is applied to the jet η (as a function of η_{det}) to remove biases created in jets which encompass regions with strongly varying responses. The values derived for jets in Run 2 is shown in Figure 6.1.

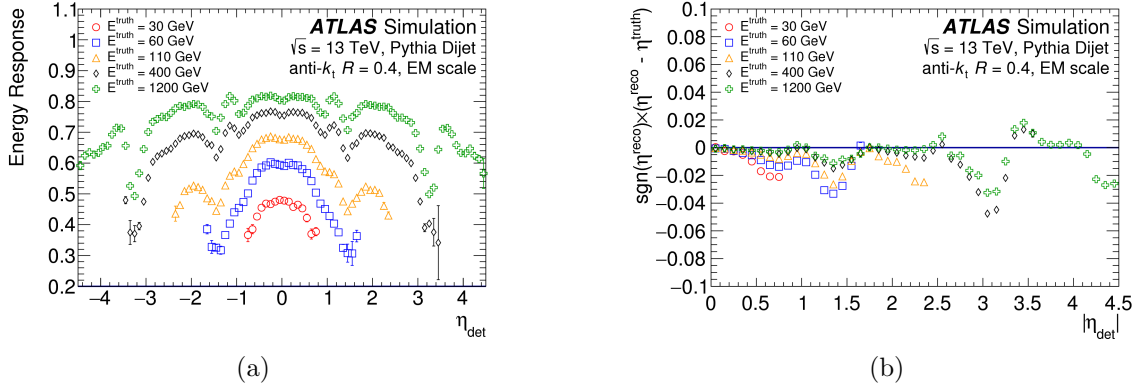


FIGURE 6.1. (a) The average energy response as a function of η_{det} for jets after origin and pile-up corrections are applied. (b) The signed difference between truth and reconstructed jet η due to biases in jet reconstruction.[53]

The Global Sequential Calibraton (GSC) is applied next and helps to account for the differences between gluon- and quark-initiated jets, as well as to correct for jets whose energy is not fully contained in the calorimeter. The GSC uses variables measuring the energy deposited in the first layer of the tile calorimeter, the third layer of the LAr calorimeter, the p_t weighted distance between the jet axis and associated tracks, the number of tracks associated with the jet, and the number of muon segments associated to a jet. The first four variables help to correct for the differences between low-multiplicity, highly penetrating quark jets versus high-multiplicity, softer radiation gluon jets. The number of muon segments is sometimes referred to as a punch-through correction, serving as a measurement of the amount of energy from a high- p_T jet that was not absorbed by the calorimeters and survived into the muon systems.

The final stage of the jet calibration is the in-situ calibration, accounting for differences between simulated and real data using well measured reference objects. In the lowest p_t region Z+jet events are used, γ +jet events are used in the intermediate p_t range, and multijet events are used in the high p_t range, balancing one high p_t jet

against several smaller jets calibrated using the other two regimes. The difference between simulation and data is then applied as a correction, shown in Figure 6.2. Additionally, η -intercalibration is used to derive corrections for jets in the more forward regions of the detector ($|\eta| \geq 0.8$) by balancing them against a jet in the better-modeled central region using dijet events.




FIGURE 6.2. Ratio of the EM+JES jet response in data to that in the nominal MC event generator as a function of jet p_T for Z -jet, γ -jet, and multi-jet in-situ calibrations. The final derived correction and its statistical and total uncertainty bands are shown.

6.5. Jet Cleaning

Jets in events considered by the analysis are subjected to cleaning criteria as defined in Ref. [54] to reject jets which do not originate from a hard scatter, but instead from calorimeter noise, cosmic-ray muons, or beam backgrounds resulting from interactions upstream of the detector. The variables used in the cleaning selection are:

- Q_{cell}^{LAr} , LAr cell quality. The quadratic difference between the expected and observed pulse shape.
- $\langle Q \rangle$, the average jet quality. An energy-weighted average of Q_{cell}^{LAr} for all cells in a jet.
- f_Q^{LAr} , the fraction of LAr calorimeter cells in a jet with poor quality. ($Q_{cell}^{LAr} > 4000$)
- f_Q^{HEC} , the fraction of HEC calorimeter cells in a jet with poor quality. ($Q_{cell}^{LAr} > 4000$)
- $|E_{neg}|$, the sum of all cells in a jet with negative energies.
- f_{EM} , the fraction of the total jet energy deposited in the EM calorimeter.
- f_{HEC} , the fraction of the total jet energy deposited in the HEC calorimeter.
- f_{max} , the maximum energy fraction of any single calorimeter layer.
- f_{ch} , the ratio of the scalar sum of the p_T of tracks from the primary vertex to the total jet p_T .

A jet is characterized as being bad if it fails any of the following criteria:

- $f_{HEC} > 0.5$ and $|f_Q^{HEC}| > 0.5$ and $\langle Q \rangle > 0.8$
- $|E_{neg}| > 60 \text{ GeV}$
- $f_{EM} > 0.95$ and $f_Q^{LAr} > 0.8$ and $\langle Q \rangle > 0.8$ and $|\eta| < 2.8$
- $f_{max} > 0.95$ and $|\eta| < 2$
- $f_{EM} < 0.05$ and $f_{ch} < 0.05$ and $|\eta| < 2$
- $f_{EM} < 0.05$ and $|\eta| \geq 2$

The first three selections are used to remove jets caused by noise bursts in the EM or HEC calorimeters, as indicated by a combination of poor cell quality and an overabundance of energy in a given layer. Calorimeter pulses from noise bursts tend to

be very unusual in shape and have large undershoots, and thus large negative energy. The last three selections target cosmic muons and beam-induced backgrounds, which are usually characterized by large amounts of energy in the hadronic portion of the calorimeter as well as little to no activity in the inner tracker or EM calorimeter.

6.6. Mistimed Events in 2015 and 2016 Data

The Run-1 jet cleaning criteria also included an additional variable t_{jet} , the energy-weighted average time of all calorimeter cells in the events. For Run 2, jet timing is not considered for cleaning. Jets originating from cosmic muons or beam backgrounds will tend to be out-of-time compared to hard scatter jets, and jets coming from prior or subsequent bunch crossings (BCs) can also be determined.

Early in the 2015 data taking a large deficit was observed in the number of high p_t jets. This was traced back to a misconfiguration in the Level-1 calorimeter trigger which was causing events with saturated trigger towers to fire the Level-1 Accept one BC too early. The issue was swiftly addressed, but all data taken prior to this point was unusable for physics (fortunately, only $\sim 80 \text{ pb}^{-1}$ of data was lost). Additionally, the problem persisted at a much reduced level for subsequent data taking. Thankfully, measures were put into place to detect and save mistimed events.

6.6.1. The L1Calo Saturated BCID Trigger

As was discussed in Section 5.4.1, to determine which BC an energy deposit is associated to, L1Calo uses two different algorithms which are run on each tower. The first algorithm is the standard peak finder and is what is used for the vast majority of triggers. The five readout slices are multiplied against filter coefficients which are determined for each individual tower, but follow the general pattern of having a

large multiplier for the central slice and small (or negative) values for the first and last slices. Between Runs 1 and 2, the system was changed to allow for negative coefficients. The slices are multiplied by their coefficients and then summed together, and this sum is then compared to the sums obtained from windows centered one BC before and after it; if it is found to be a local maximum then it is considered to have peaked in the nominal BC and is added in to the L1 trigger calculations for jets and other objects. This system is optimized for low energy pulses that are very close to the pedestal value; for larger pulses the peak is much more obvious and this algorithm points to the correct BC for the peak.

In situations where the tower readout saturates things become a bit more complicated; for pulses that saturate only one sample it is trivial to pick out the peak, but in cases where 2+ samples are saturated the peak is less obvious. A series of studies conducted during Run 1 determined that in these cases the peak always corresponded to the second saturated sample, at least for the energy regime accessible in Run 1. Thus, the saturated BCID algorithm is always set to point to the sample immediately after the first saturated sample.

The L1Calo system does not have the ability to turn on or off these algorithms based on the pulse in a given event, and so both algorithms independently make a determination on which BC a pulse belongs to. For small pulses peak finder will point correctly, for pulses with one saturated sample peak finder will point correctly and saturated BCID will point one BC later. For pulses that saturate many samples, saturated BCID will point to the correct peak while peak finder will point to the same or a later BC.

When a trigger accept is sent by the Central Trigger Processor (CTP), the ATLAS system must then read out the full state of the detector and pass it on

to the HLT. During this time ATLAS enters what is known as Simple Dead Time, a period of approximately 5 BCs where new triggers are not accepted because the system is busy saving the previous event. As such, any interesting physics events during this period are lost. So, if Level-1 Accepts are sent during two consecutive BCs, only the first event will be saved and the second will be lost. For the L1Calo trigger, this meant that even though the two algorithms could point to different BCs, only the event in the first would be saved while the second would be vetoed as it would fall during the simple dead time.

6.6.1.1. Identification and Reconstruction of Mistimed Events

In the case of the early 2015 data, the saturated BCID algorithm was incorrectly configured to point to the first saturated sample rather than the second. In events where more than one sample was saturated, this meant that the trigger would fire one BC too early and the actual event would fall into the simple dead time and be lost. Additionally, this error occurred during the 50 ns run configuration of the machine, and so a trigger that was one BC (or 25 ns) early pointed to an empty bunch in the machine. This meant that the event was not kept in the physics stream but rather placed in a calibration stream looking for noise bursts in the calorimeter. Additionally, this stream only saved a limited amount of data rather than the full event, so it was not possible to reconstruct the original event.

Once this error was fixed, measures were put into place to guard against future data loss. A special physics stream was set aside for L1Calo to automatically save all events which were at or near saturation regardless of whether they occurred in empty or filled bunches in the machine. Additionally, a timing analysis was set up to check all events in this stream for signs of early triggers. Near the end of 2015 data taking

it was discovered that a small number of events were still triggering early, this time due to an unforeseen side effect of the peak finder algorithm.

The typical filter coefficients for the three central readout slices were positive, but the first and last coefficient were negative. In very highly saturated events where 3 to 4 readout slices were saturated, this first negative coefficient was multiplying a large number in the correct timing window, leading it to give a lower value than a window centered one BC earlier. This meant that for these very high energy towers the peak finder algorithm would fire one BC early rather than the intended design of firing after the saturated BCID algorithm. Thankfully, these events were saved by the special L1Calo physics stream, but also tended to be inadvertently saved in the main physics stream as the LHC had switched to 25 ns.

To identify events which were triggered early by this error, an analysis was set up on the mistimed stream by identifying events with pulses 25 ns late. This was done by performing a parabolic fit on all trigger towers above 10 GeV but not at saturation, where the parabola was centered on the largest value of the 5 slices of readout. The mean value of these L1 tower timings was used to create a timing value for the event, which led to the following selection criteria:

- At least one trigger tower with samples 3-5 saturated, sample 2 below saturation
- Event flagged as having been triggered by peak finder algorithm
- Event timing > 20 ns
- Tower timing RMS < 3 ns
- At least one tower above 10 GeV but below saturation in both the EM and hadronic calorimeters.

- Vetoing certain lumiblocks with multiple bad events, and events where a large number of EM towers with the same ϕ all saturate, indicative of unflagged noise burst events.

After all cuts were applied, 110 events were identified as possible early triggers. These events were then run through the ATLAS reconstruction framework with a special configuration to shift the calorimeter timing to the nominal bunch crossing. This process corrected the energies of the jets to the proper values, with a few pieces of information missing from the events. First, a few of the calorimeter channels underestimated their energies due to their readouts saturating. This is because the gain on the calorimeter channel is chosen based on the value during the nominal bunch crossing, and an early trigger could cause the wrong gain to be chosen. In studied events, this affected less than 10% of channels, and the energy was underestimated by $\sim 20\%$ in such channels, leading to a negligible overall effect.

Second, the tracking information in the event could not be corrected to the proper BC, and as such the tracks in these corrected events had no correlation to the true event. Even with this complication nearly all events still passed the primary vertex and charge fraction cleaning criteria, and events that did not were discarded even though they may have been perfectly good. Finally, muon information was also missing from corrected events, and as such the punch through correction of the global sequential calibration could not be performed correctly. Testing revealed that this step only affected the energy by a few percent, well below the total JES uncertainties, as shown in Figure 6.3. After selection, around 90 of these rehabilitated events passed into the final analysis selection, but they comprised some of the highest invariant mass events in the 2015 dataset.

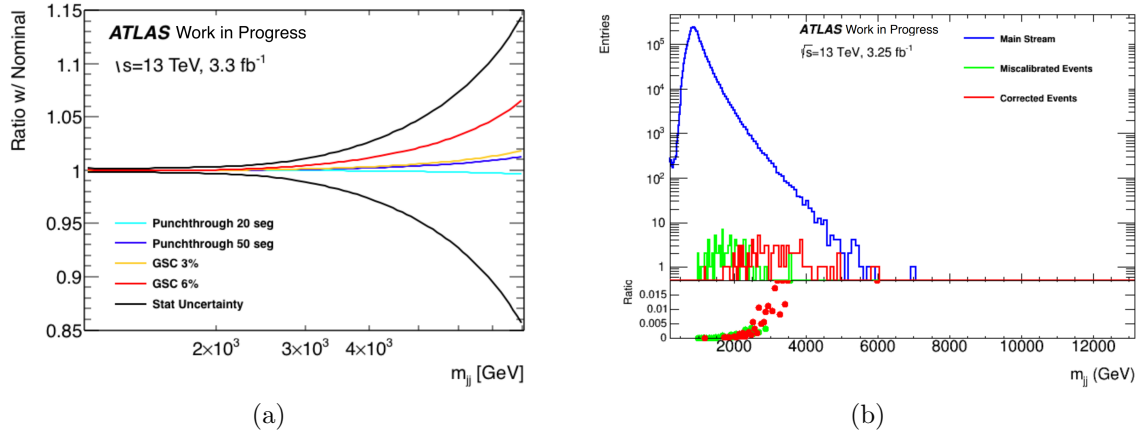


FIGURE 6.3. (a) Comparison of the effect on the m_{jj} uncertainty due to missing event information in BCID corrected events compared to the overall statistical uncertainty. (b) The dijet invariant mass spectrum showing the normal events in blue and the corrected events in red.

For the 2016 dataset additional measures were put into place to reduce the number of events by forcing peak finder to turn off in towers with four saturated samples; however, a small number still triggered early with three saturated samples. The cut could not be tightened further without risking late triggers, and as such the few events that snuck through had to be rehabilitated in the same manner. The early trigger rate for the 2016 dataset was approximately one event/ fb^{-1} , and the affected energies were much lower than those in the 2015 events. The issue of early triggers was fully resolved near the end of 2016 data taking with the introduction of the 80 MHz sampling algorithm which uses readout with twice the granularity to further refine the BCID identification.

CHAPTER VII

SEARCH STRATEGY

This chapter contains co-authored material from Ref. [1], written as part of the ATLAS Collaboration.

Compared to most ATLAS analyses, the search strategy is very straightforward:

- Event passes the lowest unrescaled trigger, HLT_j380
- Events are part of the Good Run List
- Events are not flagged as having a calorimeter noise burst, or as having incomplete event data
- Leading jet $p_T > 440$ GeV
- Subleading jet $p_T > 60$ GeV
- All jets with $p_T > 60$ GeV pass jet cleaning cuts

To help discriminate between the t -channel dominated QCD background and any possible s -channel resonances from new physics processes, a cut is placed on the angular variable

$$y^* = \frac{y_1 + y_2}{2} \tag{7.1}$$

where y_1 and y_2 are the rapidities of the leading and subleading jets, respectively. The cut is chosen as it has the best discrimination between signal and background for the ensemble of benchmark models used, or for the single benchmark in the case of the W^* selection. Finally, a cut is placed on the invariant mass of the two leading jets to ensure that the trigger is fully efficient when paired with the y^* cut. The cuts for the resonance and W^* resonance selections are:

- $|y^*| < 0.6$ (1.2)
- $m_{jj} > 1.1$ (1.7) TeV

7.1. Background Fit

Previous versions of the dijet resonance analysis, including those from CDF, ATLAS, and CMS, have fit the dijet invariant mass spectrum with a function of the form

$$f(z) = p_1(1 - z)^{p_2} z^{p_3 + p_4 \ln(x) + p_5 \ln(x)^2} \quad (7.2)$$

where $z = m_{jj}/\sqrt{s}$ and the behavior is governed by the free parameters p_i . Some past searches have required fewer terms in Eq. 7.2, such as by setting $p_4 = p_5 = 0$, but more parameters are required to properly describe the distribution as integrated luminosity increases. Eq. 7.2 was found to fit the observed spectrum in searches at CDF, and in ATLAS and CMS searches at both $\sqrt{s} = 8$ and $\sqrt{s} = 13$ TeV. The parameterization also fits the distribution obtained from simulated QCD samples.

For the analysis of the full 2015+2016 ATLAS dataset, there was a worry that the aforementioned function would no longer be sufficient to properly describe the dataset. While additional parameters could be added to Eq. 7.2, the function is already ad-hoc in nature, and any additional terms could possibly see diminishing returns in usefulness. As such, this analysis was used to test and implement a new method for fitting the background, resulting in the use of the sliding-window fit, or SWiFt as the specific implementation was named. The goal was to create a fitting method that was more robust against increases in luminosity while maintaining the same results in the regime where the global fit still held.

7.1.1. The Wilks' Statistical Test

To ensure that enough parameters are included in the fit function used to properly describe the observed spectrum, it must be determined that adding additional terms does not significantly improve the fit, and that any bumps in the observed spectrum do not come from under-fitting the data. To determine the number of parameters needed, the Wilks' statistic is used. When comparing lower and higher parameter versions of Eq. 7.2, the statistic reveals how much of an improvement an additional parameter in the fit function provides. In the case of a $p\text{-value} < 0.05$, the $n - 1$ parameter function is discarded, and the statistic for n vs. $n + 1$ is calculated. The evolution of the test statistics over time is shown in Figure 7.1. For the full 37 fb^{-1} dataset the 4-parameter function is preferred. The decline in the viability of the 3-parameter function is clearly visible, becoming unviable once around 25 fb^{-1} of data is used.

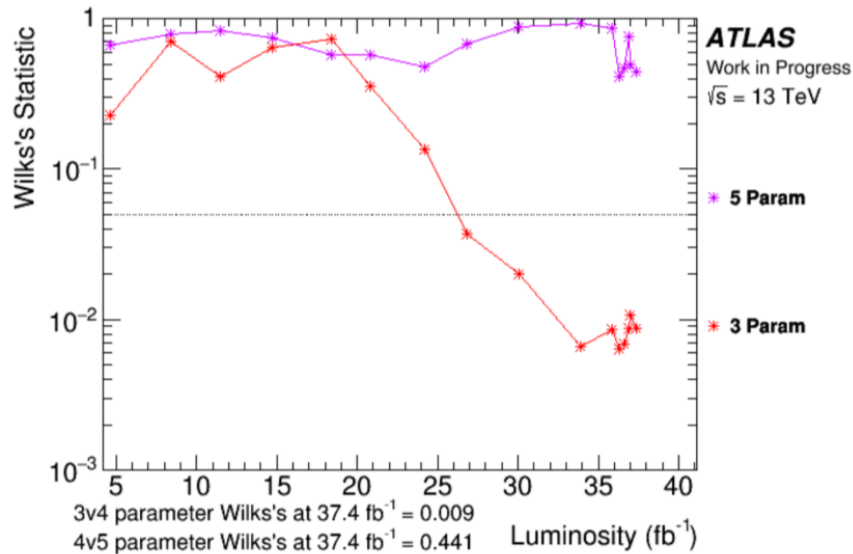


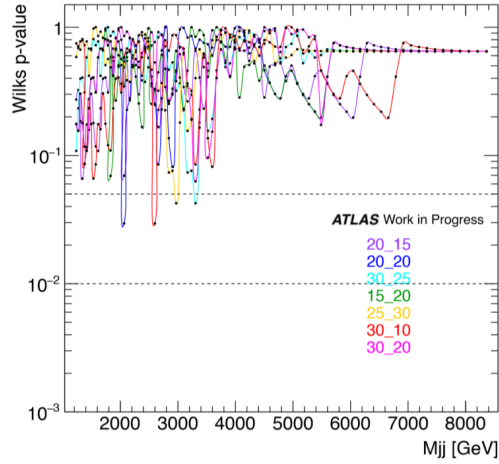
FIGURE 7.1. Wilks' test between the 4-parameter fit function and the 3/5-parameter versions. For the full 2015+2016 dataset, the 4-parameter fit provides a much better description of the data than the 3-parameter fit, but the 5-parameter fit does not provide additional improvement.

For the full dataset, a global fit function with four parameters is still a viable fit. (If a 5-parameter version was preferred, a 6th parameter would be needed to assess the goodness of fit, but such a term has not yet been determined) This means that a direct comparison can be made between the sliding window fit and the global fit to ensure that the sliding window gives the same results.

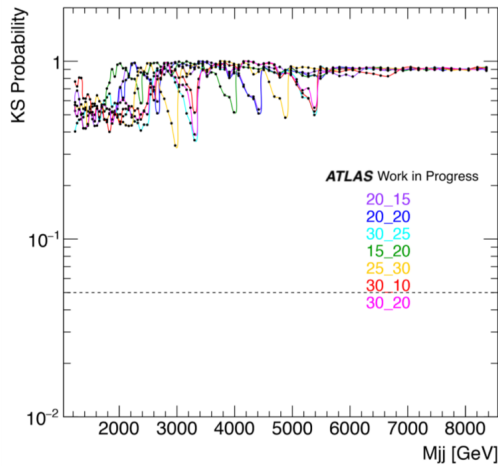
7.1.2. The Sliding Window Fit (SWiFt)

The nominal size for the sliding window was chosen as the largest window which used one less parameter than the global fit while still proving a good fit to the function for all possible windows. The three metrics used to measure this were the Wilks' statistic, chi-squared divided by the number of degrees of freedom, and the Kolmogorov-Smirnov (KS) test. The results for various window sizes are show in Figure 7.2. All of the tested windows performed very well with the data, but the final chosen window uses 30 bins to the left of the center and 20 bins to the right, comprising approximately half of the total distribution, and larger than any of the tested signal shapes. The asymmetry in the window allows for a stronger contribution from the higher-statistics, low- m_{jj} bins in a given window to better anchor the fit as the window slides to the tail of the distribution.


For each window, a fit is performed to the bins in the window using the three-parameter form of Eq. 7.2. The fit is used to give a background value at the bin "center" (keeping in mind that the window is asymmetric), and the center then slides one bin to the right and creates the next fit and returns a value for the next bin center. This is done across the entire spectrum, including beyond the final filled bin in the data distribution. The values for each center are then stitched together to create the full background distribution.



(a)



(b)



(c)

FIGURE 7.2. Statistical measures used to determine the optimal window size for SWiFt. (a) Wilks' p-value comparing the nominal SWiFt 3-parameter fit with an alternate 4-parameter version (b) KS p-value comparing the fit to pseudodata (c) χ^2/NDF comparing fit to pseudodata.

To ensure the best fit across the whole spectrum, the sliding window changes shape at the two extremes of the distribution. As it approaches the end of the distribution at low m_{jj} , the left side of the window is compressed down, shrinking to a smallest size of 10 bins to the left, 20 bins to the right. This does not cause any loss in efficacy as the window fit is relatively insensitive to window size, as demonstrated

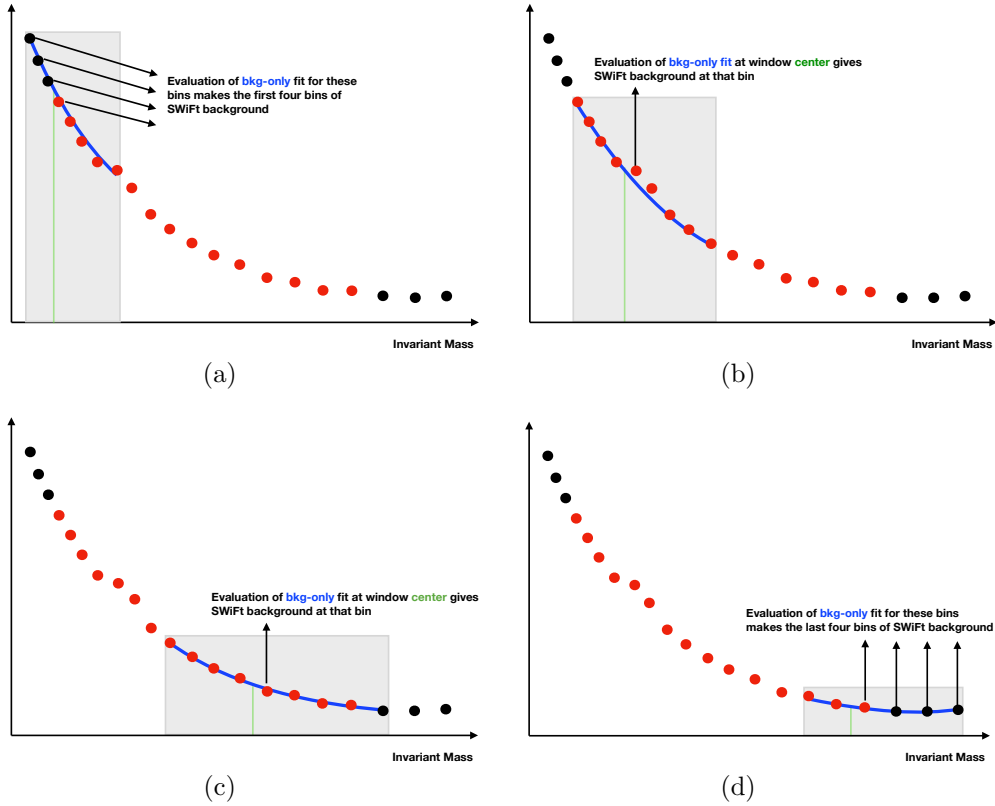


FIGURE 7.3. An example of the bin-by-bin construction of the SWiFt background as seen in Ref. [55]. The gray box shows the window considered at each point while the green line is the bin center for which the fit value is evaluated. In (b) and (c) only the value at the bin center is used for the total background estimate, while in (a) and (d) the fit values before and after the center, respectively, are used as part of the overall background.

by the 15_20 window line in Figure 7.2. Below this point in the spectrum, the fit value for all bins below the window center is determined by their fit in this smallest window. At the high- m_{jj} end, the left side of the window is fixed once the window center reaches the 5.5 TeV bin, and past that point the window expands in size. This is done to maintain adequate statistics in the fit, especially as the window slides beyond the end of the data spectrum. This system is demonstrated in Figure 7.3.

Figure 7.4 shows a comparison of the background obtained from SWiFt with the 4-parameter global fit function. Across the full spectrum no significant deviations

are observed, and the limits obtained using the two different background fits are essentially identical. For the paper results, only the SWiFt background results are shown to prevent any confusion between the two methods, and to encourage future use of the SWiFt method for similar searches.

The full SWiFt method is now outlined in Ref. [55], though it has evolved somewhat from its original implementation in the dijet search to its current form for use in future searches.

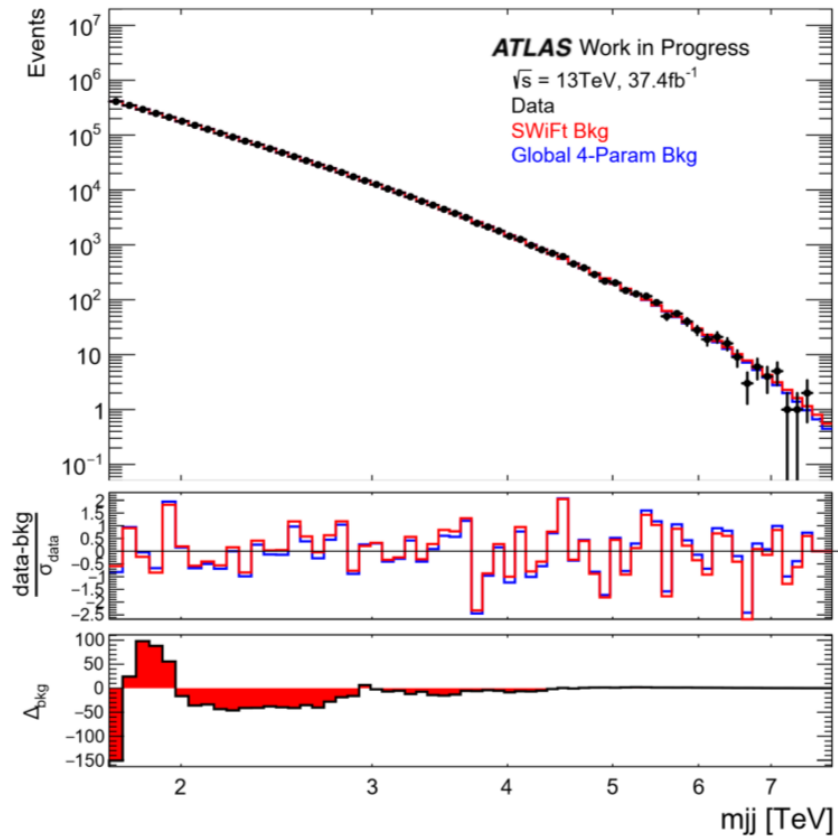


FIGURE 7.4. Comparison of the global 4-parameter fit function to the SWiFt background using a 3-parameter fit in each window. No significant variations between the two fits are observed.

7.2. Search Phase

The search phase of the analysis deals with determining if there are any signs of new physics in the data set. In contrast to the previous fitting section which sought to address whether or not a better description of the data was available through a more complex fit function, the search phase asks whether or not the obtained fit and the actual data spectrum are consistent with each other. The dijet analysis uses BumpHunter to create its test statistic as it is insensitive to the shape of any possible excess.

7.2.1. BumpHunter

Statistical tests such as the χ^2 test measure the discrepancy between data and expectation in a bin-by-bin basis, but do not take into account the relation between bins. For example, five consecutive bins with excesses over background is much more interesting for a physics search than if those bins had the same significance but alternated signs between bins. To look for the "bumpiness" of data compared to prediction, BumpHunter scans over sets of adjacent bins and measures the significance if those bins were combined together into one. The dijet analysis scans for all possible combinations of bins from 2 bins to half the full spectrum, much wider than any of the signal templates considered in the analysis and certain to cover any possible cases of interest. (Since the bin widths are chosen to approximate the detector resolution, a resonance creating an excess in only one bin is exceedingly unlikely) From this scan BumpHunter returns the most discrepant region, identified in this search as the region 4326–4595 GeV.

The global significance of the region is calculated by running BumpHunter over a set of 10,000 pseudo-experiments. In each pseudo-experiment, a toy spectrum is

created by Poisson-fluctuating each bin of the spectrum and then calculating the BumpHunter test statistic for that toy. From this set it is possible to calculate the portion of pseudo-experiments in which the observed data has an excess that is less significant.

For this particular search an excess was not observed, but it is worth noting the procedure that was in place for such an eventuality. In the case that BumpHunter returned a most discrepant region with a p-value < 0.01 , representing a very significant excess, the background fit would be re-derived using the sliding window method with all bins in the discrepant region removed from the fit, in turn removing any bias in the fit caused by the large fluctuation. This new fit would then be used as the background estimate for the limit setting phase.

CHAPTER VIII

RESULTS

This chapter contains co-authored material from Ref. [1], written as part of the ATLAS Collaboration.

Figure 8.1 shows the observed m_{jj} distribution for events passing the resonance selection. The red line overlaying the data shows result of the sliding window fit across the entire spectrum, while the second panel shows the significances of the data-fit discrepancies in each bin. Two sample mass points of the q^* benchmark model are shown at $10\times$ their predicted cross-section to illustrate the shape of a possible excess, and to make the points visible above the data distribution.

The BumpHunter algorithm quantifies the statistical significance of any localized excess in the m_{jj} distribution. The algorithm compares the binned m_{jj} distribution of the data to the fitted background estimate, considering contiguous mass intervals in all possible locations, from a width of two bins to a width of half of the distribution. For each interval in the scan, it computes the significance of any excess found. The algorithm identifies the interval 4326–4595 GeV, indicated by the two vertical lines in Figure 8.1, as the most discrepant interval in the signal region. The global significance of this outcome is evaluated using the ensemble of possible outcomes across all intervals scanned, by applying the algorithm to pseudo-data samples drawn randomly from the background fit. Without including systematic uncertainties, the probability that fluctuations of the background model would produce an excess at least as significant as the one observed in the data anywhere in the distribution

(the BumpHunter probability) is 0.63. Thus, there is no evidence of a localized contribution to the mass distribution from BSM phenomena.

The 3rd panel of Figure 8.1 shows a comparison of the data spectrum to the prediction of Pythia8 simulated QCD processes, with next-to-leading order and electroweak correction factors applied, along with the band of uncertainty due to the jet energy scale calibration. This comparison is not used for any of the fits or limits, but is only to demonstrate that the spectrum is in agreement with theory expectations.

Figure 8.2 shows the observed m_{jj} distribution for events passing the wider W^* resonance selection. The same interval is identified as the most discrepant, and the BumpHunter probability is 0.83, meaning no evidence of an excess in the invariant mass distribution.

The highest-mass event observed in the dataset is 8.12 TeV and is shown in figure 8.3. Energy deposits in the EM (green) and hadronic (red) calorimeters are indicated by the bars radiating outward from the detector. The white panels in the outside of the detector denote hits in the muon system, indicating that the energy for this event punched through the calorimeter system.




FIGURE 8.1. The reconstructed dijet mass distribution m_{jj} (filled points) is shown for events with $p_T > 440$ (60) GeV for the leading (subleading) jet in the resonance $y^* < 0.6$ signal region. The solid line depicts the background prediction from the sliding-window fit. The vertical lines indicate the most discrepant interval identified by the BumpHunter algorithm, for which the p -value is 0.63. The middle panel shows the bin-by-bin significances of the data-fit differences, considering only statistical uncertainties. The lower panel shows the relative differences between the data and the prediction of Pythia 8 simulation of QCD processes, corrected for NLO and electroweak effects, and is shown purely for comparison. The shaded band denotes the experimental uncertainty in the jet energy scale calibration.




FIGURE 8.2. The reconstructed dijet mass distribution m_{jj} (filled points) is shown for events with $p_T > 440$ (60) GeV for the leading (subleading) jet in the W^* resonance $y^* < 1.2$ signal region. The solid line depicts the background prediction from the sliding-window fit. The vertical lines indicate the most discrepant interval identified by the BumpHunter algorithm, for which the p-value is 0.83. The middle panel shows the bin-by-bin significances of the data–fit differences, considering only statistical uncertainties. The lower panel shows the relative differences between the data and the prediction of Pythia 8 simulation of QCD processes, corrected for NLO and electroweak effects, and is shown purely for comparison. The shaded band denotes the experimental uncertainty in the jet energy scale calibration.




FIGURE 8.3. The highest mass dijet event observed in the 2015+2016 dataset. It has an invariant mass of 8.12 TeV, and the two jets each have transverse momenta of 3.79 TeV.

8.1. Limit Setting and Bayes' Theorum

With no excess seen in the data, we can then use the observed spectrum to set limits on the presence of possible new physics. This analysis uses several benchmark models and sets limits on their cross-sections. Put another way, for each new physics model, at a given mass point, the largest number of events which would still be compatible with the background estimation is calculated. A Bayesian statistical approach using a flat, non-informative prior is used to set limits [56], with all nuisance parameters accounted for in the marginalization step of the procedure. To calculate the final limits, the analysis uses a Markov Chain Monte Carlo (MCMC) method to perform the numerical integration needed to combine the nuisance parameters.

This is implemented using the Bayesian Analysis Toolkit [57], a package specifically designed for physics use.

In addition to reporting an upper limit on the presence of a signal in the observed data, an expected limit is created to provide context for determining the deviation between observation and the expected null hypothesis of no signal. This is done by creating an ensemble of pseudo-datasets generated by Poisson-fluctuating the background prediction from the analysis, then calculating the upper limit for each pseudoexperiment using the same marginalization procedure as used for the data. The resulting distribution is used to define a central expected value, along with 1σ and 2σ error bands.

8.2. Uncertainties

The dijet analysis is relatively insensitive to many sources of uncertainty, including the negligible jet energy resolution uncertainty and no theoretical uncertainties on the background parameterization due to the direct fit to data. Each source of systematic uncertainty is assigned is treated as a unique nuisance parameter, the treatment of which is described in the next section.

8.2.1. Jet Energy Scale Uncertainty

The full treatment of jet energy scale uncertainty involves 84 nuisance parameters, the bulk of which come from the bin-to-bin correlations of the various in-situ analyses (Z+jet, γ +jet, and multijet balance), along with terms covering the eta intercalibration, pile-up, and the behavior of high- p_T jets. Taking into account all of these variations requires a considerable amount of computing power and is infeasible for this analysis. Instead, the analysis uses the strongly reduced uncertainties

provided by the JetETMiss combined performance group. These reduced sets consist of a single eta intercalibration term and three additional nuisance parameters which are combinations of the other 83 parameters. These parameters are reduced down in four different configurations, and the dataset is tested against all four configurations to determine if the final analysis observables are sensitive to jet correlations. The final dijet spectrum showed negligible differences between strong reduction sets, and thus any one of them can be used without loss of sensitivity compared to the full set of nuisance parameters. For the dijet analysis, the uncertainty is dominated by whichever parameter contains the high- p_T term. This uncertainty ranges from 1.5% at low invariant mass to 3% for masses above 4.5 TeV.




FIGURE 8.4. Combined uncertainty in the JES of fully calibrated jets as a function of jet p_T at $\eta = 0$. Systematic uncertainty components include pileup, punch-through, and uncertainties propagated from Z/γ -jet and mulijet balance, as well as the η -intercalibration. The flavor composition and response uncertainties assume a quark and gluon composition taken from Pythia dijet MC simulation (inclusive jets).

8.2.2. Luminosity Uncertainty

A luminosity uncertainty is applied as a scale factor to the normalization of the various signal samples used in the resonance analysis. For the combined 2015+2016 dataset, the uncertainty on the luminosity is 3.2%. This value was derived from a preliminary calibration of the luminosity scale using the results of the two van der Meer scans performed in August 2015 and May 2016. In these scans, the x-y beam separation of the two low intensity beams is scanned over, allowing for a measurement of the visible cross-section. The method used is similar to the Run-1 method documented in [52].

8.2.3. Function Choice Uncertainty

To estimate an uncertainty in the choice of fit function, an additional sliding window fit was performed using Eq. 7.2 with $p_4 \neq 0$. The result of the fit is compared to the result of the nominal three-parameter fit, and the average difference between the two fit results across a set of pseudodata drawn from Poisson fluctuations of the nominal background fit is taken as an uncertainty. The fit uncertainty is shown by the light blue line in Figure 8.5; the uncertainty is only in one direction unlike most other sources of uncertainty.

8.2.4. Statistical Uncertainty

The statistical uncertainty in the fit is calculated in a similar method to the function choice uncertainty, using pseudodata drawn from the nominal fit to data. The uncertainty in each m_{jj} bin is taken to be the root mean square of the fit results for that bin across all pseudo-experiments. The statistical uncertainty is shown by the dark blue line in Figure 8.5.




FIGURE 8.5. The sliding window fit to the m_{jj} distribution with $|y^*| < 0.6$ (red) with its two uncertainties. The dark blue line with long dashes indicates the statistical uncertainty in the fit, corresponding to the variability in the fit results across a large collection of pseudoexperiments. The light blue line with short dashes indicates the uncertainty in the function choice, based on comparison to a function with a higher number of parameters. Statistical uncertainty is shown at the one-sigma level, while the function uncertainty shown is the mean difference between the nominal and alternate functions.

8.3. Limits on Benchmark Signals

Starting from the invariant mass distribution obtained from the search selections, the previously described Bayesian method is applied to the data and simulated signal samples to set 95% credibility-level (CL) upper limits on the cross-section times acceptance. The limit is interpolated between the discrete simulated mass points to create a continuous exclusion curve.

Figure 8.6 shows the observed and expected limits for the benchmark signals with the data shown as black points and the theory curve in dashed blue. Limits are set on the product of cross-section times acceptance times branching ratio ($\sigma \times A \times BR$); the efficiency for all signal models is approximately 1. Expected limits are shown

by the black dashed line, and the 1σ and 2σ bands which describe the sensitivity of the analysis to random fluctuations of the data are denoted by the yellow and green curves. The observed data points are the black points, and the theory curve for each benchmark model is shown in blue. No uncertainties are assessed for the theoretical cross-sections.

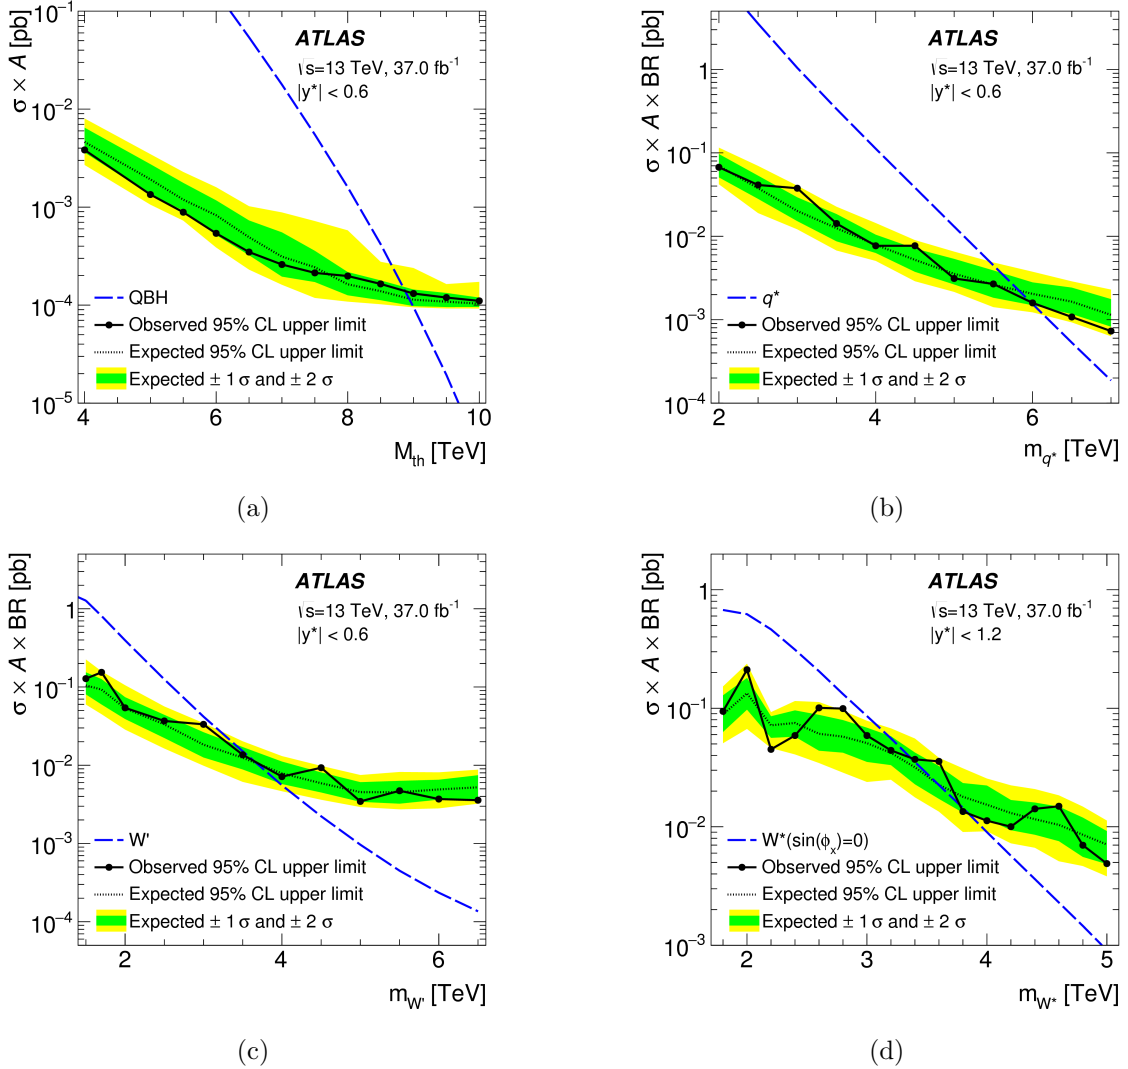


FIGURE 8.6. Limits on the cross-section times branching ratio times acceptance for the three resonance benchmark models (a) Quantum black holes (ADD, $n=6$ model) (b) excited quarks (c) W' bosons, and (d) W^* bosons. The first three limits use the nominal selection while the fourth uses the widened $y^* < 1.2$ selection. The theory cross-section is given by the blue dashed line.

| Observed (Expected) 95% CL Exclusion Limits | | |
|---|------------------------------|----------------------------|
| Benchmark Model | ATLAS (37 fb ⁻¹) | CMS (36 fb ⁻¹) |
| Excited Quark | 6.0(5.8) TeV | 6.0(5.8) TeV |
| W' | 3.6(3.7) TeV | 3.3(3.6) TeV |

TABLE 8.1. Comparison of limits on benchmark models set by ATLAS [1] and CMS [58].

This analysis sets observed (expected) 95% CL exclusion limits on these benchmarks: 6.0 (5.8) TeV for excited quarks, 3.6 (3.7) TeV for W' bosons, 3.4 (3.6) TeV for W^* bosons, and 8.9 (8.9) TeV for quantum black holes. These limits compare favorably to the limits obtained by CMS using an equivalent amount of data, as shown in Table 8.1.

Limits for the dark matter mediator Z' model are shown in the 2D plane of the Z' mass and the coupling to quarks, g_q . Limits are set in the region between 1.5 and 3.5 TeV, and all couplings above the curve (indicated by the hatching) are excluded up to $g_q = 0.5$, as beyond that the resonance width grows beyond 15% and becomes too wide to be properly detected by this analysis. Signal points are simulated with spacing of 0.5 TeV in mass and spacing between 0.05 and 0.1 in g_q ; a smooth curve is created by interpolating between points in g_q^2 and $m_{Z'}$.

Limits are also set on generic Gaussian signals, meant as a tool for the theory community to help interpret the results by providing a point of comparison for any signal model with an approximately-Gaussian peak in the dijet invariant mass distribution. The Gaussians used are trimmed to the central 95% of events to eliminate long tails, and limits are set on varying widths from 0% to 15%. The ATLAS results are given at particle level; this is achieved by creating a Monte Carlo-based transfer matrix which relates the particle level and reconstructed observables.




FIGURE 8.7. 95% CL limits on the production of leptophobic Z' , shown in the 2D plane of the mass of the Z' ($m_{Z'}$) and the coupling to quarks (g_q). The excluded region is the region in the plane above the solid line, as indicated by the hatching. Exclusions are valid only up to $g_q = 0.5$, as coupling values above that create resonances which are too wide to be properly detected by this search.

This allows for limits to be set on a given signal model, in this case a Gaussian shape, without the need for any additional information about the detector response. For sufficiently narrow resonances, these results can be used to set limits on other BSM models by comparing them to a signal shape at particle level after applying the event selection used in this analysis. The limits on Gaussian signal shapes are shown in Figure 8.8.




FIGURE 8.8. The 95% CL upper limits obtained from the dijet invariant mass m_{jj} distribution on cross-section times acceptance times branching ratio to two jets, $\sigma \times A \times BR$, for a hypothetical signal with a cross-section σ_G that produces a Gaussian contribution to the particle-level m_{jj} distribution, as a function of the mean of the Gaussian mass distribution m_G . Observed limits are obtained for five different widths, from a narrow width to 15% of m_G . The expected limit and the corresponding $\pm 1\sigma$ and $\pm 2\sigma$ bands are also indicated for a narrow-width resonance.

CHAPTER IX

COMPLEMENTARY SEARCHES AND OUTLOOK

This chapter contains co-authored material from Refs. [2][3][4], written as part of the ATLAS Collaboration.

The high-mass dijet search is just one piece of the strategy to search for new resonances which decay to dijet final states. ATLAS has several other recently-published searches for dijet events which probe the lower-mass regime, an area inaccessible in this search due to the kinematic limitation of the trigger used. These searches combine to set limits on new resonances down to 100 GeV.

9.1. Trigger-Level Analysis

The ATLAS trigger-level analysis uses the technique of data scouting to collect dijet events with lower p_T jets than can be normally saved. Data scouting gets around the trigger rate limitation by only saving a small portion of the event, usually just the jets used in the high-level trigger and some global event quantities, rather than the full set of physics objects and detector information. This creates a much smaller event size and allows for a much larger dataset to be collected than is normally possible, at the expense of not having the standard fully-calibrated jets to use in the analysis. A series of calibrations is used to bring these trigger-level jets close to the quality of offline jets. The very high statistics provided by this dataset has also revealed limitations in parts of the standard jet calibration sequence, such as the need for a smoother jet energy scale in-situ calibration to prevent the calibration from introducing bumps into the invariant mass spectrum. The comparison between

the the standard JES and the new smoothed fit it shown in Figure 9.1. This new, smoother in-situ calibration will be needed in the high-mass dijet analysis for future results as the amount of data increases even further, increasing the sensitivity of the search to such tiny fluctuations.




FIGURE 9.1. In-situ calibration in the range of $85 \text{ GeV} < \text{jet } p_T < 2 \text{ TeV}$ for both the traditional spline method (dashed line) and new fitted (solid line) combination methods. Data points from the three input measurements are overlaid. The lower panel shows the ratio of the two calibration curves.

The trigger-level analysis uses a similar series of cuts as the standard dijet analysis, analyzing events with a leading jet with $p_T > 220 \text{ GeV}$ and a second jet with $p_T > 85 \text{ GeV}$, with $|y^*| < 0.6$, and in the invariant mass range $700 \text{ GeV} < m_{jj} < 1800 \text{ GeV}$. The analysis also has a second search region which uses data taken in early 2016 with a lower-threshold trigger. This region lowers the leading jet cut to $p_T > 185 \text{ GeV}$, with $|y^*| < 0.3$ and probes the region $m_{jj} > 450 \text{ GeV}$. The results of the search are shown in Figure 9.2, and limits obtained for the Z' dark matter mediator in the two search regions are shown in Figure 9.3, along with a comparison to the result obtained from the high-mass dijet search.[3]




FIGURE 9.2. Reconstructed dijet mass distribution for events in the $|y^*| < 0.3$ (red) and $|y^*| < 0.6$ (blue) signal regions. The solid lines depict the background estimate obtained by a sliding window fit. The lower panel shows the bin-by-bin significances of the differences between the data and the background estimate, and the most discrepant region in both signal regions is indicated by the colored vertical lines.






FIGURE 9.3. The 95% credibility-level observed and expected upper limits on g_q as a function of $m_{Z'}$. The lower-mass portion of the limits from the high-mass dijet result is also shown. Couplings above the solid lines are excluded.

9.2. Dijet + Initial State Radiation Searches

To probe an even lower energy regime, ATLAS has performed two other searches which look for resonances in association with initial state radiation. This circumvents the trigger limitation by triggering on a high- p_T photon or jet rather than on the much softer jets from the new resonance. The first is the resolved dijet+ISR search which looks for two $R=0.4$ jets in association with a high- p_T jet or photon.[2] The results of that search are shown in Figures 9.4 and 9.5. The other is the boosted search which looks for one large radius ($R=1.0$) jet rather than for two distinct $R=0.4$ jets.[4] Background processes are suppressed by cutting on jet substructure variables which measure the likelihood that the energy deposition within the large jet is from two sub-jets rather than from the single jet expected from QCD background. The results of the search are shown in Figure 9.6. This search is used to set limits on the same Z' dark matter mediator model between 100 and 220 GeV, and the results are shown in Figure 9.7.




(a)




(b)

FIGURE 9.4. Reconstructed dijet mass distribution for events containing two jets with $p_T > 25$ GeV and either a jet with $p_T > 430$ GeV (left) or a photon with $p_T > 150$ GeV (right). The vertical lines indicate the most discrepant region as identified by the BumpHunter algorithm, and the bottom panel shows the bin-by-bin significances of the data with respect to the background fit.



(a)



(b)

FIGURE 9.5. The 95% CL upper limits obtained from the X+jet (left) and X+photon (right) searches on coupling g_q as a function of the resonance mass $m_{Z'}$ for the leptophobic Z' dark matter mediator model.

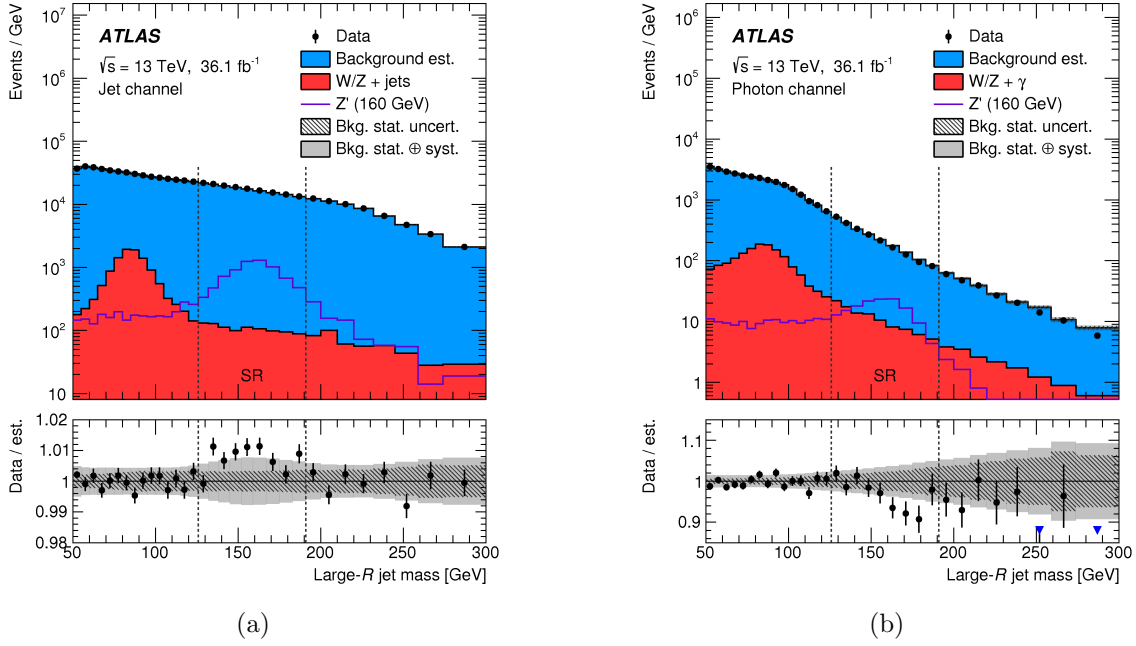


FIGURE 9.6. Distribution of Large-R jet mass in the ISR jet (left) and ISR photon (right) channels. The signal Z' shown has a mass of 160 GeV and coupling $g_q=0.5$. The bottom panel shows the ratio of the data to the estimated background; the estimate varies for different signal masses.




FIGURE 9.7. Observed and expected limits at 95% CL on the coupling to quarks g_q for the combined ISR jet and photon channels for the boosted search.

9.3. Combined limits on Dark Matter Models

The limits set on the Z' dark matter mediator model by the dijet analysis, combined with the results from the aforementioned complimentary analyses, provide strong constraints on the model across a wide range of mediator masses. Figure 9.8 shows the state of the analyses as of July 2017 following the release of the high-mass dijet result, but does not include the latest results from the dijet+ISR and TLA searches. This plot shows the limits for the Z' model with coupling $g_q=0.25$, shown in the plane of mediator (Z') mass and the dark matter mass. For this particular benchmark model, mediator masses between 200 GeV and 2.6 TeV have been excluded for all shown dark matter masses. The various colors represent the regions excluded by the different results and help demonstrate the complementarity between the various ATLAS searches. The very low mediator mass region is probed by a different class of searches which look for decays of the mediator to the dark matter particle in association with some kind of initial state radiation, such as the diagram in Figure 2.6a.

Figure 9.9 shows the same exclusions for a different model using a vector mediator as well as non-zero coupling to leptons, along with the limits obtained from the dilepton search looking for mediator decays to leptons. Here the dijet limits are much less stringent, but are able to exclude some higher masses than the dilepton search alone. For this particular model values corresponding to the thermal relic density are not excluded above a mediator mass of 1 TeV, offering a region ripe for exploration and exclusion in future dijet searches with greater integrated luminosity. Similar to Figure 9.8, this does not include the latest results from the dijet+ISR and TLA searches.




FIGURE 9.8. Regions in a dark matter mass-mediator mass plane excluded at 95% CL by a selection of ATLAS dark matter searches, for one possible interaction between the Standard Model and dark matter, the leptophobic axial-vector mediator as described in [59]. The exclusions are computed for a dark matter coupling $g_{DM} = 1.0$, a quark coupling $g_q = 0.25$ universal to all flavors. The lepton coupling g_l in this model is set to zero. This choice for the couplings corresponds to the "A1" scenario in [59]. The results use 13 TeV data except for [7]. Dashed curves labeled "thermal relic" indicate combinations of dark matter and mediator mass that are consistent with a dark matter density of $\Omega_c = 0.12 h^2$ and a standard thermal history, as computed in MadDM ([60], [59]). Between the two curves, annihilation processes described by the simplified model deplete Ω_c below $0.12 h^2$. A dotted curve indicates the kinematic threshold where the mediator can decay on-shell into dark matter. Excluded regions that are in tension with the perturbative unitarity considerations of [61] are indicated by shading in the upper left. The exclusion regions, relic density contours, and unitarity curve are not applicable to other choices of coupling values or model.




FIGURE 9.9. Regions in a dark matter mass-mediator mass plane excluded at 95% CL by a selection of ATLAS dark matter searches, for one possible interaction between the Standard Model and dark matter, the vector mediator as described in [59]. The exclusions are computed for a dark matter coupling $g_{DM} = 1.0$, a quark coupling $g_q = 0.1$ universal to all flavors, and lepton coupling $g_l = 0.01$, corresponding to the "V2" scenario in [59]. With this choice of couplings, the Z' decays to leptons are reduced with respect to the decays to quarks. The results use 13 TeV data except for [7]. The lepton constraints use 13 TeV data from 2016, reinterpreting the model independent limits in a way similar to what done for the dijet searches. The $E_T^{\text{miss}} + X$ exclusion regions are obtained by rescaling, using acceptance and cross-section information from samples simulated at truth-level, the exclusion contours published in the corresponding papers. Dashed curves labeled "thermal relic" indicate combinations of dark matter and mediator mass that are consistent with a dark matter density of $\Omega_c = 0.12h^2$ and a standard thermal history, as computed in MadDM ([60], [59]). To the left of the curve, annihilation processes described by the simplified model deplete Ω_c below $0.12 h^2$. A dotted curve indicates the kinematic threshold where the mediator can decay on-shell into dark matter. The exclusion regions, relic density contours, and unitarity curve are not applicable to other choices of coupling values or model.




FIGURE 9.10. Summary plot of ATLAS bounds in the coupling-mediator mass plane from dijet searches using 2015 and 2016 data. The 95% CL upper limits are on coupling g_q as a function of the resonance mass $m_{Z'}$ for the leptophobic Z' model described in [16]. The expected limits from each search are indicated by dotted curves. Coupling values above the solid curves are excluded, as long as the signals are narrow enough to be detected using these searches (10% signal width/mass for dijet+ISR and TLA, 15% for high-mass dijets, approximately corresponding to $g_q < 0.5$ and $g_q < 0.6$, respectively).

Figure 9.10 shows a summary of the Z' mediator limits in the plane of the resonance mass $m_{Z'}$ and coupling to quarks g_q as of March 2018, reflecting the most recent results in the trigger-level analysis and boosted dijet + ISR searches. These new results have greatly extended the lower mass reach and have bridged the gap between the masses covered by the high-mass search and the previous TLA result. The Run-2 searches now cover a range from 100 GeV to 3.5 TeV, with exclusions ranging from 0.27 down to 0.04 in g_q .

9.4. Dijet Angular Search

In addition to the resonance bump hunt previously shown, ATLAS has also performed an angular shape analysis on dijet events, using the same event selection criteria but using events out to $|y^*| < 1.7$, and cutting on the variable $y_B = |y_1 + y_2|/2 < 1.1$. Data is binned in both m_{jj} and in $\chi = e^{2|y^*|}$, and the χ distribution is compared to the shape expected from QCD as determined from simulated data. Within each mass bin the shape is obtained from simulation and is scaled up to the amount of data seen. The results of the angular search are shown in Figure 9.11.

While the ATLAS result is only interpreted in the context of limits on the scale of new quark contact interactions, they can also be extended to the same Z' model, but can set limits for much higher values of g_q as it is sensitive to much wider widths than the resonance search. CMS has recently used their angular analysis [62] to set limits on the Z' DM mediator model for very high mediator masses, the results of which are shown in Figure 9.12. The results show a very large discrepancy from the expected limit for masses above 4.5 TeV; this may hint at a region of interest to be followed up on in the next resonance search, or may be due to a mis-modeling of the high invariant mass region by the simulated data which is used in this type of analysis.

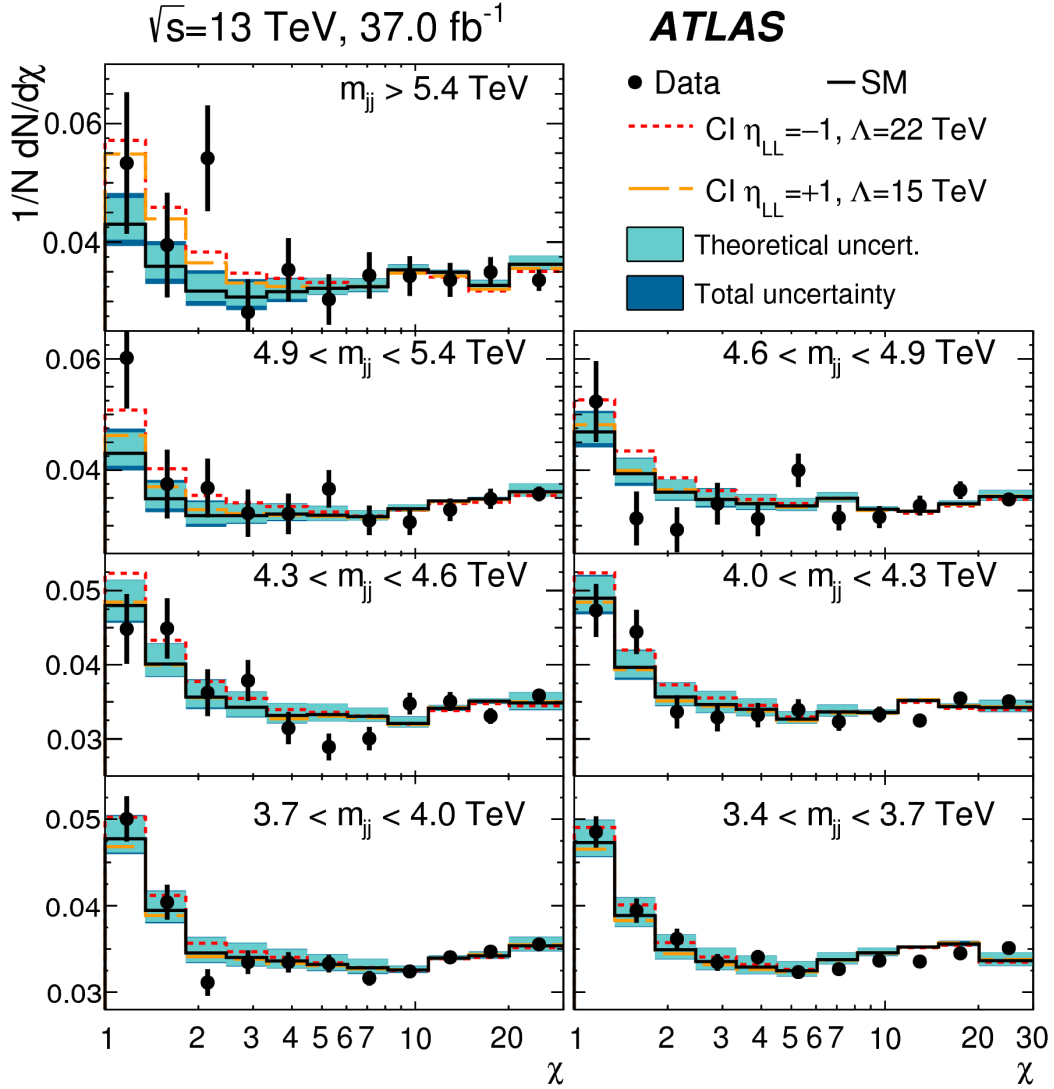


FIGURE 9.11. Reconstructed distributions of the dijet angular variable χ in different regions of the dijet invariant mass m_{jj} . The data (points), Pythia predictions with NLO and electroweak corrections applied (solid lines), and examples of the contact interaction (CI) signals (dashed lines) are shown. The theoretical uncertainties and the total theoretical and experimental uncertainties in the predictions are displayed as shaded bands around the SM prediction. The SM background prediction and corresponding systematic uncertainty bands are extracted from the best-fit to the data. Data and predictions are normalized to unity in each m_{jj} bin.




FIGURE 9.12. Results from the CMS Collaboration showing the 95% CL upper limits on the quark coupling g_q , as a function of mass, for an axial-vector or vector DM mediator with $g_{DM} = 1.0$ and $m_{DM} = 1 \text{ GeV}$. The observed limits (solid), expected limits (dashed) and the variation in the expected limit at the 1 and 2 standard deviation levels (shaded bands) are shown. The corresponding limits on the width of the mediators are shown on the vertical axis on the right-hand side of the figure.[62]

9.5. Outlook

Looking forward, the next ATLAS publication will be on the complete Run-2 dataset, comprising some $100+\text{fb}^{-1}$ of 13 TeV data. The LHC has already exceeded its design luminosity and expects to deliver even higher pileup in 2018, meaning that the ATLAS trigger menu must evolve to keep rates under control. However, this evolution will have little impact on the dijet analysis, only requiring the loss of the lower one or two bins in the invariant mass spectrum. The sensitivity of the dijet analysis is much more limited by the center-of-mass energy rather than by statistics, and as such minimal increases in limits are to be expected with the next paper result, with most benchmark models having their limits improve by 10% or less. For comparison, the jump between the 2015 and 2015+2016 papers had $10\times$ the integrated luminosity, but only improved limits between 9% and 40%.

For the full HL-LHC dataset, comprising some 3000fb^{-1} at a center-of-mass energy of 14 TeV, the expected limits on excited quarks and quantum black holes are estimated to be 8.0 TeV and 10.1 TeV, respectively. [63] Any increases beyond the design energy of the machine will lead to considerable improvements in these limits.

9.6. Conclusion

A search was performed analyzing dijet events in 37fb^{-1} of collision data with center-of-mass energy of 13 TeV, taken with the ATLAS detector during the 2015 and 2016 runs. The dijet invariant mass spectrum does not exhibit any excesses or deficits from the smoothly-falling background estimate which was obtained by a sliding window fit to the data. The spectrum is also consistent with the spectrum derived from a simulation of QCD processes. The most significant excess in the signal region has a p-value of 0.63, consistent with the background prediction. The analysis

sets limits on several benchmark signals, with an improvement of limits of 5-40% in mass reach over the previous published limits by ATLAS in 2015 with 3.2 fb^{-1} of data.

REFERENCES CITED

- [1] Morad Aaboud et al. Search for new phenomena in dijet events using 37 fb^{-1} of pp collision data collected at $\sqrt{s} = 13 \text{ TeV}$ with the ATLAS detector. *Phys. Rev.*, D96(5):052004, 2017.
- [2] Search for new light resonances decaying to jet pairs and produced in association with a photon or a jet in proton-proton collisions at $\sqrt{s} = 13 \text{ TeV}$ with the ATLAS detector. Technical Report ATLAS-CONF-2016-070, CERN, Geneva, Aug 2016.
- [3] Morad Aaboud et al. Search for low-mass dijet resonances using trigger-level jets with the ATLAS detector in pp collisions at $\sqrt{s} = 13 \text{ TeV}$. 2018.
- [4] Morad Aaboud et al. Search for light resonances decaying to boosted quark pairs and produced in association with a photon or a jet in proton-proton collisions at $\sqrt{s} = 13 \text{ TeV}$ with the ATLAS detector. 2018.
- [5] Y. Fukuda et al. Evidence for oscillation of atmospheric neutrinos. *Phys. Rev. Lett.*, 81:1562–1567, 1998.
- [6] C. Patrignani et al. Review of Particle Physics. *Chin. Phys.*, C40(10):100001, 2016.
- [7] Search for new phenomena in the dijet mass distribution using pp collision data at $\sqrt{s} = 8 \text{ TeV}$ with the atlas detector. *Phys. Rev. D*, 91:052007, Mar 2015.
- [8] The ATLAS collaboration. Search for New Phenomena in Dijet Mass and Angular Distributions with the ATLAS Detector at $\sqrt{s} = 13 \text{ TeV}$. 2015.
- [9] Gavin Salam. General aspects of qcd at future colliders, 2014.
- [10] U. BAUR, I. HINCHLIFFE, and D. ZEPPENFELD. Excited quark production at hadron colliders. *International Journal of Modern Physics A*, 02(04):1285–1297, 1987.
- [11] U. Baur, M. Spira, and P. M. Zerwas. Excited-quark and -lepton production at hadron colliders. *Phys. Rev. D*, 42:815–824, Aug 1990.
- [12] M. V. Chizhov and Gia Dvali. Origin and Phenomenology of Weak-Doublet Spin-1 Bosons. *Phys. Lett.*, B703:593–598, 2011.
- [13] M. V. Chizhov, V. A. Bednyakov, and J. A. Budagov. A unique signal of excited bosons in dijet data from pp-collisions. *Phys. Atom. Nucl.*, 75:90–96, 2012.

- [14] V. C. Rubin, W. K. Ford, Jr., and N. Thonnard. Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 $/R = 4\text{kpc}/$ to UGC 2885 $/R = 122\text{ kpc}/$. *The Astrophysical Journal*, 238:471–487, June 1980.
- [15] Tracy R. Slatyer and Chih-Liang Wu. General Constraints on Dark Matter Decay from the Cosmic Microwave Background. *Phys. Rev.*, D95(2):023010, 2017.
- [16] Daniel Abercrombie et al. Dark Matter Benchmark Models for Early LHC Run-2 Searches: Report of the ATLAS/CMS Dark Matter Forum. 2015.
- [17] Nima Arkani-Hamed, Savas Dimopoulos, and G. R. Dvali. The Hierarchy problem and new dimensions at a millimeter. *Phys. Lett.*, B429:263–272, 1998.
- [18] Patrick Meade and Lisa Randall. Black Holes and Quantum Gravity at the LHC. *JHEP*, 05:003, 2008.
- [19] G. Dissertori. The Determination of the Strong Coupling Constant. *Adv. Ser. Direct. High Energy Phys.*, 26:113–128, 2016.
- [20] Yuri L. Dokshitzer. Calculation of the Structure Functions for Deep Inelastic Scattering and $e^+ e^-$ Annihilation by Perturbation Theory in Quantum Chromodynamics. *Sov. Phys. JETP*, 46:641–653, 1977. [*Zh. Eksp. Teor. Fiz.*73,1216(1977)].
- [21] V. N. Gribov and L. N. Lipatov. Deep inelastic $e p$ scattering in perturbation theory. *Sov. J. Nucl. Phys.*, 15:438–450, 1972. [*Yad. Fiz.*15,781(1972)].
- [22] Guido Altarelli and G. Parisi. Asymptotic Freedom in Parton Language. *Nucl. Phys.*, B126:298–318, 1977.
- [23] 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.
- [24] R.K. Ellis, W.J. Stirling, and B.R. Webber. *QCD and Collider Physics*. Cambridge Monographs on Particle Physics, Nuclear Physics and Cosmology. Cambridge University Press, 2003.
- [25] T. Gleisberg et al. Event generation with SHERPA 1.1. *JHEP*, 02:007, 2009.
- [26] Bo Andersson, G. Gustafson, G. Ingelman, and T. Sjostrand. Parton Fragmentation and String Dynamics. *Phys. Rept.*, 97:31–145, 1983.

- [27] Torbjörn Sjöstrand, Stefan Ask, Jesper R. Christiansen, Richard Corke, Nishita Desai, Philip Ilten, Stephen Mrenna, Stefan Prestel, Christine O. Rasmussen, and Peter Z. Skands. An Introduction to PYTHIA 8.2. *Comput. Phys. Commun.*, 191:159–177, 2015.
- [28] B. R. Webber. A QCD Model for Jet Fragmentation Including Soft Gluon Interference. *Nucl. Phys.*, B238:492–528, 1984.
- [29] Johannes Bellm et al. Herwig 7.0/Herwig++ 3.0 release note. *Eur. Phys. J.*, C76(4):196, 2016.
- [30] Bora Isildak. *Measurement of the differential dijet production cross section in proton-proton collisions at $\sqrt{s} = 7$ tev*. PhD thesis, Bogazici U., 2011.
- [31] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. The Anti-k(t) jet clustering algorithm. *JHEP*, 04:063, 2008.
- [32] Stephen D. Ellis and Davison E. Soper. Successive combination jet algorithm for hadron collisions. *Phys. Rev.*, D48:3160–3166, 1993.
- [33] M. Wobisch and T. Wengler. Hadronization corrections to jet cross-sections in deep inelastic scattering. In *Monte Carlo generators for HERA physics. Proceedings, Workshop, Hamburg, Germany, 1998-1999*, pages 270–279, 1998.
- [34] The ATLAS Experiment at the CERN Large Hadron Collider. *J. Instrum.*, 3:S08003. 437 p, 2008.
- [35] The CMS experiment at the CERN LHC. *J. Instrum.*, 3:S08004. 361 p, 2008.
- [36] The ALICE experiment at the CERN LHC. *J. Instrum.*, 3:S08002. 259 p, 2008.
- [37] The LHCb Detector at the LHC. *J. Instrum.*, 3(LHCb-DP-2008-001. CERN-LHCb-DP-2008-001):S08005, 2008.
- [38] Cinzia De Melis. The CERN accelerator complex. Complexe des accélérateurs du CERN. Jan 2016. General Photo.
- [39] Daniel Boussard and Trevor Paul R Linnecar. The LHC Superconducting RF System. Technical Report LHC-Project-Report-316. CERN-LHC-Project-Report-316, CERN, Geneva, Dec 1999.
- [40] W. Bartmann et al. Impact of LHC and SPS Injection Kicker Rise Times on Lhc Filling Schemes and Luminosity Reach. In *Proc. of International Particle Accelerator Conference (IPAC17), Copenhagen, Denmark, 14-19 May, 2017*, number 8 in International Particle Accelerator Conference, pages 2043–2046, Geneva, Switzerland, May 2017. JACoW. <https://doi.org/10.18429/JACoW-IPAC2017-TUPVA007>.

- [41] Early Inner Detector Tracking Performance in the 2015 data at $\sqrt{s} = 13$ TeV. Technical Report ATL-PHYS-PUB-2015-051, CERN, Geneva, Dec 2015.
- [42] Imma Riu. The ATLAS Level-1 Topological Trigger performance in Run 2. Technical Report ATL-DAQ-PROC-2017-002. 3, CERN, Geneva, Jan 2017.
- [43] *ATLAS liquid-argon calorimeter: Technical Design Report*. Technical Design Report ATLAS. CERN, Geneva, 1996.
- [44] Georges Aad et al. Monitoring and data quality assessment of the ATLAS liquid argon calorimeter. *JINST*, 9:P07024, 2014.
- [45] Nikiforos Nikiforou. Performance of the ATLAS Liquid Argon Calorimeter after three years of LHC operation and plans for a future upgrade. In *Proceedings, 3rd International Conference on Advancements in Nuclear Instrumentation Measurement Methods and their Applications (ANIMMA 2013): Marseille, France, June 23-27, 2013*, 2013.
- [46] R Achenbach, P Adragna, V Andrei, P Apostologlou, B Åsman, C Ay, B M Barnett, B Bauss, M Bendel, C Bohm, J R A Booth, I P Brawn, P Bright Thomas, D G Charlton, N J Collins, C J Curtis, A Dahloff, A O Davis, S Eckweiler, J P Edwards, E Eisenhandler, P J W Faulkner, J Fleckner, F Föhlich, J Garvey, C N P Gee, A R Gillman, P Hanke, R P Hatley, S Hellman, A Hidvégi, S J Hillier, K Jakobs, M Johansen, E E Kluge, M Landon, V Lendermann, J N Lilley, K Mahboubi, G Mahout, A Mass, K Meier, T Moa, E Moyses, F Müller, A Neusiedl, C Nöding, B Oltmann, J M Pentney, V J O Perera, U Pfeiffer, D P F Prieur, W Qian, D L Rees, S Rieke, F Rühr, D P C Sankey, U Schäfer, K Schmitt, H C Schultz-Coulon, C Schumacher, S Silverstein, R J Staley, R Stamen, M C Stockton, S Tapprogge, J P Thomas, T Trefzger, P M Watkins, A Watson, P Weber, and E E Woehrling. The atlas level-1 calorimeter trigger. *Journal of Instrumentation*, 3(03):P03001, 2008.
- [47] Morad Aaboud et al. Performance of the ATLAS Trigger System in 2015. *Eur. Phys. J.*, C77(5):317, 2017.
- [48] Sebastian Mario Weber and Martin Wessels. L1Calo Pedestal Correction Plots. Technical Report ATL-COM-DAQ-2017-064, CERN, Geneva, Jun 2017.
- [49] Georges Aad et al. Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1. *Eur. Phys. J.*, C77:490, 2017.
- [50] Peter Speckmayer, Tancredi Carli, and Christian Wolfgang Fabjan. Energy Measurement of Hadrons with the CERN ATLAS Calorimeter, 2008. Presented on 18 Jun 2008.

- [51] W. Lampl, S. Laplace, D. Lelas, P. Loch, H. Ma, S. Menke, S. Rajagopalan, D. Rousseau, S. Snyder, and G. Unal. Calorimeter clustering algorithms: Description and performance. 2008.
- [52] Morad Aaboud et al. Luminosity determination in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector at the LHC. *Eur. Phys. J.*, C76(12):653, 2016.
- [53] Morad Aaboud et al. Jet energy scale measurements and their systematic uncertainties in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. 2017.
- [54] Selection of jets produced in 13TeV proton-proton collisions with the ATLAS detector. Technical Report ATLAS-CONF-2015-029, CERN, Geneva, Jul 2015.
- [55] Karishma Sekhon, Ryan Christopher Edgar, and Dante Amidei. SWiFt: Sliding Window Fit Method for Resonance Searches. Technical Report ATL-COM-PHYS-2018-161, CERN, Geneva, Feb 2018.
- [56] B.P. Carlin and T.A. Louis. *Bayesian Methods for Data Analysis, Third Edition*. Chapman & Hall/CRC Texts in Statistical Science. CRC Press, 2008.
- [57] A. Caldwell, D. Kollár, and K. Kröninger. BAT - The Bayesian analysis toolkit. *Computer Physics Communications*, 180:2197–2209, November 2009.
- [58] Searches for dijet resonances in pp collisions at $\sqrt{s} = 13$ TeV using data collected in 2016. Technical Report CMS-PAS-EXO-16-056, CERN, Geneva, 2017.
- [59] Andreas Albert et al. Recommendations of the LHC Dark Matter Working Group: Comparing LHC searches for heavy mediators of dark matter production in visible and invisible decay channels. 2017.
- [60] Mihailo Backovic, Antony Martini, Kyoungchul Kong, Olivier Mattelaer, and Gopolang Mohlabeng. MadDM: New dark matter tool in the LHC era. *AIP Conf. Proc.*, 1743:060001, 2016.
- [61] Felix Kahlhoefer, Kai Schmidt-Hoberg, Thomas Schwetz, and Stefan Vogl. Implications of unitarity and gauge invariance for simplified dark matter models. *JHEP*, 02:016, 2016.
- [62] Albert M Sirunyan et al. Search for new physics in dijet angular distributions using proton-proton collisions at $\sqrt{s} = 13$ TeV and constraints on dark matter and other models. 2018.
- [63] Dijet resonance searches with the ATLAS detector at 14 TeV LHC. Technical Report ATL-PHYS-PUB-2015-004, CERN, Geneva, Mar 2015.