

THE UNIVERSITY OF CHICAGO

SHADOWS IN THE DETECTOR: SEARCHING FOR AXINO DARK MATTER AS
MISSING ENERGY IN ATLAS AND COMMISSIONING MET AND JET TRIGGERS

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

DEPARTMENT OF PHYSICS

BY
KRISTIN DONA

CHICAGO, ILLINOIS
AUGUST 2025

Copyright © 2025 by Kristin Dona

All Rights Reserved

TABLE OF CONTENTS

LIST OF FIGURES	xii
LIST OF TABLES	xvii
LIST OF ACRONYMS	xx
ACKNOWLEDGMENTS	.xxviii
ABSTRACT	.xxxiv
1 INTRODUCTION	1
2 PHYSICS MOTIVATION	6
2.1 The Standard Model	6
2.1.1 Particle Content	8
2.1.1.1 Fermions	8
Spin and Statistics	8
Mass and Generations	9
Chirality	9
Color Charge and Confinement	10
Hadrons	10
Antiparticles and Conservation Laws	11
Flavor Mixing	12
2.1.1.2 Bosons	13
2.1.2 Mathematical Formalism	14
2.1.3 Electroweak Unification	15
2.1.4 Higgs Mechanism	17

2.1.4.1	Quantum Chromodynamics	19
2.2	Beyond the Standard Model	19
2.2.1	Fine-Tuning Problems	20
2.2.1.1	The Strong CP Problem	20
2.2.1.2	The Hierarchy Problem	21
2.2.2	Observational Phenomena Beyond the Standard Model	22
2.2.3	Motivation for New Physics	24
2.3	Axions as a Solution to the Strong CP Problem	25
2.3.1	Searching for Axions Through Photon Conversion	25
2.3.1.1	Axion Modified Maxwell's Equations	26
2.3.1.2	Axion Models and Coupling Predictions	27
2.3.1.3	The Axion Search Program	29
	BREAD	32
	ALP searches at the LHC	36
2.4	Supersymmetry	37
2.4.1	The Minimal Supersymmetric Standard Model	37
2.5	Axion + SUSY Theory: The PQMSSM	40
2.5.1	Axino as the LSP	41
2.5.2	Long-Lived NLSP	42
2.6	Searching for Axinos with Collider Experiments	44
3	THE LARGE HADRON COLLIDER AND ATLAS	46
3.1	CERN	46
3.2	The Large Hadron Collider (LHC)	47
3.2.1	Machine Overview	49
3.2.2	The LHC Experiments	51

3.2.3	Luminosity and Pile-Up	53
3.3	The ATLAS Experiment	55
3.4	ATLAS Geometry	58
3.5	Magnets	60
3.6	Inner Detector	61
3.7	Calorimeters	67
3.7.1	Electromagnetic Calorimeter	70
3.7.1.1	Electromagnetic Barrel Calorimeter	70
3.7.1.2	Electromagnetic Endcap Calorimeter (EMEC)	72
3.7.2	Hadronic Calorimeter System	74
3.7.2.1	Tile Calorimeter	74
3.7.2.2	LAr Hadronic Endcap Calorimeter	75
3.7.2.3	Forward Calorimeter	76
3.8	Muon Spectrometer	76
3.9	Trigger and Data Acquisition System	79
3.9.1	Run 2 Trigger System	81
3.9.2	Run 3 Trigger System Upgrades	83
4	GFEX - THE GLOBAL FEATURE EXTRACTOR	87
4.1	Hardware Design	88
4.2	Firmware Algorithms	93
4.2.1	gTower Builder	93
4.2.2	Large Radius Jet Algorithm: Jet Finder	94
4.2.3	MET Algorithm: Jets without Jets	95
4.2.4	gTower Calibration	97
4.2.4.1	Effects of gTower calibration	98

4.2.5	Pile-Up correction	101
4.3	gFEX Trigger Objects summary	102
4.4	Operating system and low-level monitoring	102
4.4.1	Custom operating system	104
4.4.2	I2C Bus Monitoring	107
4.4.3	IPMC and ATCA Monitoring	108
4.4.4	DCS Integration	110
4.5	Performance Monitoring	115
4.5.1	Trigger Efficiency Curves	115
4.5.2	Jet Trigger Efficiency Monitoring Algorithm	116
4.5.2.1	Quality Assessment and Web Display Integration	121
4.5.2.2	Issue Identification and Resolution	123
4.5.3	Trigger Rate Monitoring	125
4.6	Performance in Run 3	127
4.6.1	Outlook	128
5	EVENT RECONSTRUCTION AND OBJECT DEFINITION	131
5.1	Tracking	131
5.1.1	Standard Tracking	132
5.1.2	Large Radius Tracking	134
5.1.3	Vertexing	136
5.1.4	Secondary Vertexing Algorithms	137
5.1.4.1	VrtSecInclusive (VSI) Algorithm	137
5.1.4.2	Fuzzy Vertexing Algorithm	140
5.2	Object Definitions	142
5.2.1	Jets	143

5.2.1.1	Jet Clustering Algorithm	143
5.2.1.2	Jet Types	144
	EMTopo Jets	144
	Track Jets	145
	Particle Flow Jets	145
5.2.1.3	b -tagging	146
5.2.1.4	Jet Energy Scale and Resolution Uncertainties	146
5.2.2	Electrons and Photons	147
5.2.3	Muons	147
5.3	Overlap Removal	149
5.3.1	Missing Transverse Energy	149
6	DISPLACED VERTICES AND E_T^{miss} ANALYSIS	152
6.1	Motivation and Previous Searches	154
6.2	Analysis Strategy	155
6.3	Event and Object Selection	157
6.3.1	Data Collection and Trigger Strategy	157
6.3.1.1	Run 2 Data Sample	157
6.3.1.2	E_T^{miss} Trigger Selection	158
6.4	Non-Collision Background Veto	160
6.5	Displaced Vertex Algorithms and Track Selections	161
6.5.1	Vertex Reconstruction Algorithms	162
6.5.2	Track Cleaning Procedures	163
6.5.3	Baseline Vertex Selections	163
6.5.4	Standard Physics Objects	164
6.5.4.1	Overlap Removal	167

6.6	Material Map	167
6.7	Signal Models and Monte Carlo Simulation	169
6.7.1	Gluino R -hadron	169
6.7.2	Wino-Bino co-annihilation	171
6.7.3	Higgs portal	172
6.7.4	DFSZ Axino	173
6.7.5	Monte Carlo Generation	175
6.7.5.1	Gluino R -hadron MC Production	176
6.7.5.2	Bino-Wino Coannihilation MC Production	177
6.7.5.3	Higgs portal MC Production	177
6.7.5.4	DFSZ Axino MC Production	178
6.8	Signal Regions	179
6.8.1	Signal Region Selection Criteria	180
6.8.2	Signal Efficiencies	181
6.8.3	Validation Regions	184
6.9	Photon Triggered Control Region	186
6.10	Background Estimation	187
6.10.1	SDV Background Estimate	188
6.10.1.1	SDV Background Estimate Validation	189
6.10.1.2	Hadronic Interaction Background Estimate	190
6.10.1.3	Accidental Crossing Background Estimate	192
6.10.1.4	Merged Vertices Background Estimation	196
6.10.1.5	Summary of SDV SR Background Estimates	200
6.10.2	1FDV Background Estimate	201
6.10.2.1	1FDV Background Estimate Validation	202

6.10.3	2FDV Background Estimate	202
6.10.3.1	2FDV Background Estimate Validation	204
6.10.4	Summary of Background Estimates	205
6.11	Unblinding and Observed Events	206
6.11.1	Blinded Analysis Strategy	206
6.11.2	Unblinding Results	206
6.11.3	Event Characterization	207
6.11.4	Statistical Interpretation	209
6.12	Uncertainties	210
6.12.1	Signal Yield Statistical Uncertainties	210
6.12.2	Background Estimation Statistical Uncertainties	211
6.12.3	Background Estimate Systematic Uncertainties	212
6.12.3.1	Pile-Up Uncertainty	212
6.12.3.2	Non-linearity Uncertainty	213
6.12.3.3	Non-closure Uncertainty	213
6.12.4	Signal Yield Systematic Uncertainties	213
6.12.4.1	Track Reconstruction Uncertainties	213
6.12.4.2	Jet and MET Uncertainties	217
6.12.4.3	MET Trigger Uncertainties	217
6.12.4.4	Theoretical Uncertainties	224
6.12.4.5	Additional Uncertainties	224
6.12.5	Summary of Uncertainties	225
6.12.5.1	Background Uncertainty Summary	225
6.12.5.2	Signal Uncertainty Summary	225
6.12.6	Impact of Uncertainties on Results	226

6.13	Statistical Analysis and Limit Setting	226
6.13.1	Maximum Likelihood Fitting	227
6.13.2	Test Statistic and CLS Method	228
6.13.3	Upper Limit Setting	230
6.13.3.1	Procedure for Individual Grid Points	230
6.13.3.2	Constructing Exclusion Contours	231
6.13.4	Setting Limits on the DFSZ Axino	231
6.13.5	DFSZ Axino Model Results and Interpretation	232
6.13.6	Gluino R Hadron Model Results	235
6.13.7	Bino-Wino Coannihilation Model Results	236
6.13.8	Higgs portal model Results	238
6.13.9	Final Remarks	239
7	CONCLUSION	243
	APPENDICES	245
A	DIRECT DETECTION AXION SEARCHES	246
A.1	Solar Searches	246
A.1.1	CAST Experiment	247
A.1.2	Next-Generation Helioscopes	248
A.2	Laboratory Photon-Axion Conversion Searches	248
A.2.1	ALPS Experiments	249
A.2.2	Other LSW Experiments	249
A.3	Cavity Resonance Searches	250
A.3.1	ADMX	250
A.3.2	Other Cavity Experiments	251

A.4	Conducting Surface Experiments	251
A.4.1	MADMAX	252
B	THE BREAD EXPERIMENT	254
B.1	Experimental Design	254
B.2	The BREAD Multi-Stage Approach	256
B.3	GigaBREAD First Results	259
B.4	Detecting Single Photons	259
B.4.1	Terahertz Spectroscopy Development	261
B.4.2	Superconducting Nanowire Single Photon Detectors	263
B.4.3	SNSPD Characterization Studies	265
B.5	Future Plans	270
C	OBSERVED EVENT KINEMATICS IN THE DV + MET ANALYSIS	271
C.1	Event Properties and Signal Comparisons	271
C.1.1	Standard Displaced Vertex Signal Region	272
C.1.2	Single Fuzzy Displaced Vertex Signal Region	272
C.1.3	Two Fuzzy Displaced Vertex Signal Region	272
	BIBLIOGRAPHY	278

LIST OF FIGURES

2.1	Particles of the SM	7
2.2	Energy density composition of the universe	23
2.3	Axion-photon interactions	26
2.4	Axion-photon conversion at a conducting surface	28
2.5	Axion mass to signal photon wavelength scale	30
2.6	BREAD geometry and axion sensitivity projections	33
2.7	Extended axion parameter space including LHC searches	34
2.8	Axion parameter space with experimental limits	35
2.9	Cartoon inspired by the MSSM particle content	38
2.10	Decay length of the long-lived neutralino in the PQMSSM model	43
2.11	SM long-lived particles	45
3.1	CERN accelerator complex	48
3.2	LHC beam schematic	52
3.3	Integrated luminosity collected by ATLAS	54
3.4	Cutaway diagram of ATLAS	56
3.5	ATLAS detector particle signatures	57
3.6	ATLAS coordinate system	59
3.7	ATLAS magnet system	61
3.8	Cutaway of the ATLAS ID	63
3.9	Cross section of the ATLAS ID	65
3.10	Cross section diagram of the barrel region of the ID	67
3.11	A cutaway diagram of the ATLAS calorimeter system	69
3.12	Diagram of a LAr barrel module	73
3.13	The ATLAS muon system	77

3.14	Cartoon of the trigger rate along the TDAQ path in ATLAS.	79
3.15	Overview of the Run 2 Trigger and Data Acquisition System for the ATLAS detector	81
3.16	Overview of the Run 3 Trigger and Data Acquisition System for ATLAS	83
3.17	Schematic of the granularity scheme used for LAr-L1Calo interface.	84
4.1	Example of large-radius jets as boosted objects	88
4.2	Comparison of current and large-radius jet triggers for boosted $t\bar{t}$ events	89
4.3	The gFEX board with main components labeled	90
4.4	Division of the η - ϕ space across the three pFPGAs in gFEX	91
4.5	gFEX pFPGA link connections	92
4.6	gFEX calibration factors as a function of η	99
4.7	gFEX performance with and without calibration	100
4.8	Overview of the gFEX monitoring systems	104
4.9	Workflow for setting up and building the gFEX OS	105
4.10	L1Calo DCS system overview	110
4.11	gFEX OPC-UA server class diagram	112
4.12	gFEX FSM minipod panel at Point 1	113
4.13	Main gFEX panel in the STF showing two boards present and labeled.	114
4.14	Numerator and denominator histograms that make up the TEfficiency curve for the gLJ100 trigger turn on curve for run 486315.	118
4.15	Resulting TEfficiency curve for the gLJ100 trigger turn on curve for run 486315 from Figure 4.14 components.	119
4.16	DQ Tier 0 Web Display of Jet Trigger Efficiency Monitoring Algorithm results for run 486315	122
4.17	Example event with unexpected gJ400 performance	124
4.18	FEX Trigger rates plot	126
4.19	gFEX trigger performance curves and rates	128

4.20	MET efficiency for JWoJ	129
5.1	Illustration of tracking parameters for object reconstruction	133
5.2	Schematic of LRT and standard tracks in the ATLAS detector	134
5.3	Illustration of the incompatibility graph method used in the VSI algorithm to form n -track vertices	139
5.4	Illustration of how an LLP ($\tilde{\chi}_2^0$) decaying via b -quarks would be reconstructed by the VSI and FV algorithms. Blue circles represent the reconstructed vertices. Figure courtesy of Jan Offerman	141
6.1	Material map for DV + MET analysis	168
6.2	Feynman diagram of the Gluino R -hadron model	170
6.3	Feynman diagram for the Wino-Bino co-annihilation model	171
6.4	Feynman diagram for the Higgs portal model	172
6.5	Feynman diagram for the DFSZ Axino model	174
6.6	Expected signal yields for the DFSZ Axino model in the SDV SR	184
6.7	Estimated and observed background count in the sideband VRs of the SDV SR.	190
6.8	Fits performed to estimate the HI contribution to the SDV SR background	192
6.9	Schematic overview of the AX background estimation method	193
6.10	Example mass template from the AX background estimate for the SDV SR	195
6.11	Sample data set including 3 events with one or two DVs each.	198
6.12	Schematics of how the distance significance distribution histograms are populated for same-event-pairs and mixed-event-pairs of DVs in the MV SDV SR background estimate	199
6.13	Distance significance histograms for the MV SDV SR background estimate	200
6.14	Combined background estimate for the SDV SR	201
6.15	Estimated and observed background count in the sideband VRs of the 1 FDV SR.	203
6.16	Estimated and observed background count in the sideband VRs of the 2 FDV SR.	205
6.17	Event display for the single event observed in the SDV SR	209

6.18	Comparison of the number of pile-up vertices in the MTR and PTR	212
6.19	Tracking uncertainty yield plot	215
6.20	Estimated tracking inefficiency as a function of DV radial position R_{xy}	216
6.21	Illustration of the track-killing study methodology	216
6.22	MET trigger efficiency turn-on curves as a function of METNoMu for data and MC	220
6.23	Percent difference in signal efficiency for the R-hadron model when using the complete DRAW_RPVLL trigger list versus using only the lowest unscaled triggers	221
6.24	MET trigger efficiency turn-on curves with SDV SR like selection	222
6.25	MET trigger efficiency as a function of offline MET for data and signal MC . . .	223
6.26	Exclusion limits on the DFSZ Axino model as a function of axion mass	233
6.27	Exclusion limits on the DFSZ Axino model as a function of axino mass	234
6.28	Exclusion limits on the Gluino R -hadron model	236
6.29	Exclusion limits on the Bino-Wino Coannihilation model	237
6.30	Exclusion limits on the Higgs portal model	238
6.31	Axion-photon coupling parameter space as a function of axion mass m_a and decay constant f_a , showing experimental constraints and theoretical predictions given the new axino limits	241
A.1	Schematic of the CAST experiment	247
A.2	Schematic of the LSW experimental method	249
A.3	Schematic of the MADMAX dielectric disk experiment	252
B.1	3D rendering of the GigaBREAD experimental setup	255
B.2	Projected sensitivity for BREAD	257
B.3	GigaBREAD exclusion results	260
B.4	Interferometer arrangement	261
B.5	Accuracy validation of THz FTS	262
B.6	Demonstration of THz FTS broad frequency coverage	263

B.7	Image of SNSPD structure	264
B.8	SNSPD detection mechanism	265
B.9	ADR system at Fermilab	266
B.10	SNSPD characterization workflow	267
B.11	Schematic of SNSPD measurement	268
B.12	SNSPD mounted on cryogenic stage	269
C.1	Kinematic distributions for axino hh signals in SDV region	273
C.2	Kinematic distributions for axino ZZ signals in SDV region	274
C.3	Kinematic distributions for gluino R-hadron signals in SDV region	275
C.4	Kinematic distributions for wino-bino coannihilation signals in 1FDV region . .	276
C.5	Kinematic distributions for Higgs portal signals in 2FDV region	277

LIST OF TABLES

2.1	Fermions in the SM organized by generation and approximate masses in MeV.	9
2.2	Properties of bosons in the SM.	13
2.3	Electric charge, weak isospin, and weak hypercharge assignments for SM fermions.	16
2.4	Overview of major axion and ALP detection experiments, their methods, target mass ranges, and operational status. Experiments targeting masses above $\sim 10^{-2}$ eV probe ALPs rather than QCD axions that solve the strong CP problem.	36
3.1	LHC runs with their operational years and proton-proton collision center-of-mass energies. Runs 4–6 represent planned future operations.	49
3.2	Proton acceleration stages to LHC	50
3.3	Summary of the ATLAS sub-detectors, with their radial and η coverage	62
3.4	Summary of the ATLAS calorimeter system.	69
4.1	Summary table of gFEX algorithm objects	103
4.2	Relevant Yocto versions compared to Xilinx Vivado versions	106
4.3	gFEX IPMC temperature sensor threshold values.	109
4.4	Classification thresholds used to evaluate the fitted maximum efficiency of each trigger.	122
4.5	DQ flag ranges for all triggers monitored using the jet trigger efficiency monitoring algorithm, used to quickly determine if a trigger is not turning on as expected. Values are hard-coded in the algorithm and were determined by eye. All thresholds are in GeV.	123
4.6	Summary of gFEX and legacy trigger items, including TOB thresholds, rates, and approximate turn-on behavior. All rate values are quoted at an instantaneous luminosity of $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$	129
5.1	Definitions of variables used to parameterize tracks. The reference point used to determine the track’s perigee during track reconstruction is the origin of the detector coordinate system.	132
5.2	Differences between track reconstruction criteria for standard tracking and large-radius tracking [127].	135

6.1	Lowest un-prescaled MET triggers included in the DRAW_RPVLL filter	159
6.2	Standard and strict NCB veto definitions.	161
6.3	Track cleaning selections for SDV and FDV algorithms.	163
6.4	Baseline displaced vertex selection requirements.	164
6.5	Electron selection requirements.	166
6.6	Muon selection requirements.	166
6.7	Photon selection requirements.	166
6.8	Jet selection requirements.	166
6.9	Parameter space explored for the Gluino R -hadron signal model.	170
6.10	Parameter space explored for the Wino-Bino co-annihilation signal model. . . .	172
6.11	Parameter space explored for the Higgs portal signal model.	172
6.12	Parameter space explored for the DFSZ Axino signal model.	173
6.13	Parameter space explored for the DFSZ Axino signal model.	174
6.14	Summary of MC generation details for signal and background samples.	176
6.15	Full vertex selection criteria for SDV SR and FDV SR (including the 1 FDV and 2+ FDV subregions).	181
6.16	Examples of SR selection efficiencies for representative grid points from each signal model. These show how efficiency varies with lifetime for typical parameter choices in each model's optimized signal region.	182
6.17	Vertex selection criteria for the various VRs for each of the SRs.	185
6.18	Background estimates for the SDV SR comparing the combined estimate and component contributions.	202
6.19	A summary of the background estimates for all SRs.	205
6.20	Observed event yields compared to predicted background in each SR.	207
6.21	Event-level properties for observed events in each SR.	207

6.22	Properties of each DV for observed events in each SR. Jet association is determined by $\Delta R < 0.4$ matching criterion with the nearest jet. For the 2 FDV SR events, “shared” indicates both FDVs in the same event are matched to the same jet.	208
6.23	Lowest unscaled MET triggers used for the MET turn-on curve study. For 2016, prescales were added and removed from different triggers throughout the year, requiring separate period-dependent analyses. For 2017 and 2018, single triggers remained unscaled throughout the data-taking years.	219
6.24	Summary of background estimate uncertainties by signal region	225
6.25	Ranges of systematic uncertainties (16th–84th percentile) per signal model . . .	225
B.1	Summary of the staged BREAD experiment implementations.	258
B.2	Comparison of candidate photosensor technologies and their typical performance characteristics relevant for BREAD implementations. Table originally from [24].	261

LIST OF ACRONYMS

- DRAW_RPVLL** Derivation of the RPVLL format for Analysis Workflow. 158, 159
- ACR** ATLAS Control Room. 109, 110, 114, 125
- ADC** Analog-to-Digital Converter. 124
- ADMX** Axion Dark Matter Experiment. 2, 31, 250, 251, 256
- ADR** Adiabatic Demagnetization Refrigerator. 265
- ALICE** A Large Ion Collider Experiment. 51, 53
- ALP** Axion-Like Particle. 2, 29, 36, 259
- ALPS** Any Light Particle Search. 31, 248, 249
- ATCA** Advanced Telecommunications Computing Architecture. 87, 88, 103, 108, 111, 112, 115
- ATLAS** particle detector at the Large Hadron Collider formerly known as A Toroidal LHC Apparatus. 4, 5, 42, 44, 46, 51–58, 60, 61, 63, 67, 72, 74, 75, 78–81, 83, 86, 87, 97, 103, 109–111, 114–118, 121, 124, 125, 127, 128, 130, 131, 136, 142, 143, 152, 154, 155, 157, 169, 175, 179, 214, 217, 228, 231, 232, 236, 237, 240, 243, 244, 271
- AX** Accidental Crossing. 156, 187, 189, 191, 192, 194, 196, 200, 206, 243
- BC** Bunch Crossing. 62, 79, 80, 83
- BCID** Bunch Crossing ID. 93
- BCR** Bunch Counter Reset. 92, 93
- BDT** Boosted Decision Trees. 141, 162
- BREAD** Broadband Reflector Experiment for Axion Detection. 3, 5, 32, 33, 251–254, 256, 259–265, 268, 270
- BSM** Beyond the Standard Model. 52, 53, 120, 153, 169, 175, 197, 206, 208–210, 227, 238, 239
- CAPP** Center for Axion and Precision Physics Research. 250, 251
- CAST** CERN Axion Solar Telescope. 32, 247, 248, 256
- CB** combined muon. 148

CERN European Organization for Nuclear Research, *Conseil Européen pour la Recherche Nucléaire*. 46, 47, 55, 80, 86, 158, 247, 249, 252

CKM Cabibbo–Kobayashi–Maskawa. 12, 18, 20

CMB Cosmic Microwave Background. 22

CMS Compact Muon Solenoid. 51, 53

CP Cluster Processor. 82, 85

CP Charge-Parity. 2, 3, 6, 9, 20, 21, 24, 25, 29, 36, 40, 41

CPU Central Processing Unit. 82

CR Control Region. 156, 160, 186, 187, 191, 201, 206, 243

CROWS CERN Resonant Weakly Interacting sub-eV Particle Search. 248, 249

CSC Cathode-Strip-Chamber. 78, 147

CSSK Constituent Subtraction and Soft Killer. 120

CV Collision Vertex. 188

DCR Dark Count Rate. 264

DCS Detector Control System. 103, 109–111, 114, 115

DESY Deutsches Elektronen-Synchrotron. 248, 249

DFSZ Dine–Fischler–Srednicki–Zhitnitsky model. 27, 29, 31, 40, 42, 173, 232, 243, 250, 251

DM Dark Matter. 1–5, 256

DQ data quality. 103, 117, 121, 125

DSP digital signal processing. 90

DV Displaced Vertex. 44, 140, 152, 154–156, 160, 161, 163, 167, 169, 179, 180, 183, 184, 186–188, 190, 196, 197, 207, 208, 210, 211, 213–216, 222, 231, 232, 239, 243, 271

DV + MET Displaced Vertices and E_T^{miss} . 152, 154, 240

ECAL Electromagnetic Calorimeter. 68, 70, 74–76, 84, 142, 147, 164

ECR Event Counter Reset. 92, 93

EDM Electric Dipole Moment. 21

EDP Event-DV probability. 186–189, 211, 213

eFEX Electromagnetic Feature Extractor. 85, 86, 88, 128

EM Electromagnetic. 3, 82, 144, 151

EMB Electromagnetic Barrel. 70–72, 74, 84

EMEC Electromagnetic Endcap. 70, 72, 74, 84

EMTopo jet electromagnetic topological jet. 144, 146, 151, 213

FCal Forward Calorimeter. 76

FDV fuzzy displaced vertex. 156, 161–163, 175, 179–187, 201–206, 213, 226, 272

FELIX Front End Link Exchange. 86, 92, 95, 96

FEX *Feature Extractor*. 85, 86

FFT Fast Fourier Transform. 259

FOX Fiber Optic Exchange. 86, 90

FPGA Field Programmable Gate Array. 85, 90, 95, 98, 108, 259

FSM Finite State Machine. 110–112, 115

FTS Fourier-Transform Spectroscopy. 33

FV Fuzzy Vertexing Algorithm. 137, 140–142, 152, 208, 236–239

gFEX Global Feature Extractor. 5, 46, 73, 84–88, 93–95, 97, 98, 101, 102, 104–108, 111, 112, 114, 115, 117, 120, 123–125, 127, 128, 130, 145, 151, 239, 243

GPIO General-Purpose Input/Output. 108, 109

GRL Good Runs List. 157

HAYSTAC Haloscope At Yale Sensitive To Axion CDM. 2, 31, 250, 251

HCAL Hadronic Calorimeter. 68, 74, 76, 142

HEC LAr Hadronic End-Cap Calorimeter. 75

HI Hadronic Interaction. 156, 163, 167, 185, 187, 189–191, 200, 202, 206, 243

HL-LHC High Luminosity–LHC. 84, 239, 244

HLT High-Level Trigger. 80, 82, 86, 125, 150, 218

I²C Inter-Integrated Circuit. 102, 107–109, 111, 114

IBL Insertable B-Layer. 64, 136

ID Inner Detector. 56, 60–65, 67, 71, 77, 80, 131, 136, 142, 143, 145, 147, 148, 153, 161, 164, 165, 167, 190, 192, 232, 271, 272

IO inside-out combined muon. 148

IP interaction point. 56, 61–64, 162

IPMC Intelligent Platform Management Controller. 103, 107–109, 111, 114

IS Information Service. 103, 108, 109, 114

JEP Jet Energy Processor. 82, 85

JER Jet Energy Resolution. 146, 217

JES Jet Energy Scale. 146, 151, 217

jFEX Jet Feature Extractor. 85, 86, 88, 117, 120, 127, 128

JVT Jet Vertex Tagger. 165

JwoJ “Jets without Jets”. 93, 95–97

KID Kinetic Inductance Detector. 260

KSVZ Kim–Shifman–Vainshtein–Zakharov model. 27, 29

L1 Level-1 Trigger. 80, 87, 90, 96, 101, 102, 125, 127, 218

L1A Level-1 Accept. 80, 92, 93

L1Calo Level-1 Calorimeter Trigger. 73, 80–85, 87, 88, 94, 109, 116, 124, 125, 128, 243

L1CTP Level-1 Central Trigger Processor. 80, 93

L1Muon Level-1 Muon Trigger. 80, 82

L1Topo Level-1 Topological Trigger. 82, 92, 93, 95, 96, 102

LAr Liquid Argon. 67, 70–72, 74–76, 84, 85, 94, 98

LArC LAr ATCA carrier blade. 85

LATOME LAr Trigger Object Mezzanine. 85

lcw local cell weighting. 151

LDPS LAr Digital Processing System. 85

LEP Large Electron-Positron Collider. 48

LHC Large Hadron Collider. 1, 3– 6, 14, 36, 39, 44, 46– 51, 53– 55, 79, 83, 88, 157, 243, 249

LHCb Large Hadron Collider beauty experiment. 51– 53

LHT LocHadTopo. 159

LLP Long Lived Particle. 10, 11, 44, 140, 152– 155, 158, 161, 162, 167, 169, 172, 175, 177– 179, 182, 186, 196, 237, 243

LOS Loss of Signal. 107, 114

LRT Large Radius Tracking. 134, 135, 137, 145, 152, 158, 213, 217

LS Long Shutdowns. 48

LSP Lightest Supersymmetric Particle. 2, 3, 39, 41, 42, 44, 155, 173, 243

LSW Light Shining through Walls. 248, 249

LTDB LAr Trigger Digitizer Board. 84

MADMAX Magnetized Disc and Mirror Axion experiment. 31, 251, 252

MC Monte Carlo. 42, 153, 154, 175, 177, 178, 181, 191, 205, 210, 211, 214– 217, 219, 222– 224, 239

MDT Monitored Drift Tube. 78, 147

MET Missing Transverse Energy. 3– 5, 76, 82, 83, 87, 93– 97, 100– 102, 127, 128, 130, 146, 149– 155, 157– 161, 169, 171, 172, 175, 177, 179, 186, 217– 219, 222, 223, 231, 235, 239, 243, 271, 272

MGT Multi-Gigabit Transceivers. 92

MHT Missing Hadronic Transverse energy. 95, 151

MHT MET Hard Term. 95, 96, 151

MPSoC Multi-Processor System-on-Chip. 90

MRI Magnetic Resonance Imaging. 32, 259

MS Muon Spectrometer. 61, 76, 77, 80, 147, 148, 165

MSSM Minimal Supersymmetric Standard Model. 37, 40, 41, 155, 173

MST MET Soft Term. 96

MTR MET-triggered dataset. 160, 186, 187, 201, 212

MuCTPI Muon-to-Central Trigger Processor Interface. 80

MV Merged Vertex. 156, 187, 189, 191, 196–198, 200, 206, 243

NCB Non-Collision Background. 160, 161

NLL Next-to-Leading Logarithmic. 177, 224

NLO Next-to-Leading Order. 177, 224

NLSP Next-to-Lightest Supersymmetric Particle. 39, 41, 42, 155

nMCM new Multi-Chip Module. 81

NNLL Next-to-Next-to-Leading Logarithmic. 176, 224

NNLO Next-to-Next-to-Leading Order. 176, 224

OPC-UA Open Platform Communications Unified Architecture. 103, 109, 111

ORGAN Oscillating Resonant Group AxioN. 250, 251

OS Operating System. 90, 102, 104–108, 111

OSQAR Optical Search for QED Vacuum Bifringence, Axions and Photon Regeneration. 248, 249

P1 Point 1. 55, 111

PDF Parton Distribution Function. 177, 224

PFlow jet Particle Flow jet. 115, 145

pFPGA Processor FPGA. 88, 90, 92–96, 98, 102

PMNS Pontecorvo-Maki-Nakagawa-Sakata. 12

PPM Pre-Processor Module. 81, 82

PQ Peccei-Quinn. 25, 27, 29, 42

PQMSSM Peccei-Quinn-Minimal Supersymmetric Standard Model. 40–42, 44, 240, 243

PS Proton Synchrotron. 49

PSB Proton Synchrotron Booster. 49

PTR Photon-Triggered Control Region. 160, 186–188, 201, 212

PUC Pile-Up Correction. 101

pufit pile-up fit. 151

PV Primary Vertex. 136, 137, 145, 157, 183, 192

QCD Quantum Chromodynamics. 2, 3, 9, 13, 14, 19, 20, 25, 27, 29, 32, 36, 53, 243, 250, 254, 258–260, 268, 270

QFT Quantum Field Theory. 6, 14, 21

QUAX QUaerere AXion. 250, 251

RF Radio Frequency. 32, 50, 258, 260

ROI Region of Interest. 80, 86

RPC Resistive Plate Chamber. 78, 147

SCT Semiconductor Tracker. 63, 64, 131–133, 138, 158, 167

SDV standard displaced vertex. 156, 161–164, 175, 179, 181–184, 187–192, 194, 196, 198, 200–202, 205–207, 213, 226, 231, 232, 235, 272

SM Standard Model. 1–4, 6, 8–11, 13–15, 17, 19–22, 24, 25, 29, 36, 37, 39, 40, 44, 52–55, 58, 102, 119, 152, 154, 155, 167, 169, 170, 172, 179, 183, 196, 206–210, 231, 243, 271

SNSPD Superconducting Nanowire Single Photon Detector. 3, 33, 258, 260, 263–265, 268

SoC System-on-Chip. 85, 93, 95, 96, 102–104, 107, 111, 112, 115

SPS Super Proton Synchrotron. 49

SR Signal Region. 152, 153, 155, 156, 160, 161, 179–192, 194, 196–198, 200–202, 204–208, 212, 213, 219, 225, 226, 229, 231, 232, 235–239, 272

STF L1Calo Surface Test Facility. 106, 111, 112

SUSY Supersymmetry. 3–5, 21, 24, 27, 37, 39–41, 44, 87, 128, 152, 154, 169, 171, 173, 232, 240

SV Secondary Vertex. 136, 137, 152

SWROD Software Readout Driver. 86

TDAQ Trigger and Data Acquisition. 56, 79–81, 83

TES Transition-Edge Sensor. 258, 260

TGC Thin Gap Chamber. 78, 147

TOB Trigger Object. 80, 92, 102, 117, 123, 124

topo-cluster topological cluster. 144, 147, 151, 159

TREX Tile Rear Extension Module. 85, 86

TRT Transition Radiation Tracker. 63, 65, 66, 77, 131, 134, 135, 138

UFO Unified Flow Object. 120

VR Validation Region. 156, 160, 184, 185, 189, 202, 204, 206, 213, 243

VSI VrtSecInclusiv Vertexing Algorithm. 137, 140–142, 152, 154, 179, 196, 197, 243

WinCCOA WinCC Open Architecture. 111, 115

XSA Xilinx Shell Archive. 104–106

ACKNOWLEDGMENTS

Completing this Ph.D. over the past six years has been both challenging and rewarding. I owe a great deal to the many people whose support, guidance, and friendship made this journey possible. Throughout these acknowledgments, names within research teams and friend groups are listed alphabetically with no intended order.

First and foremost, I'd like to thank my advisor, David. Your patient mentorship and steady encouragement have not only shaped this thesis but also my identity as a physicist. Thank you for challenging me to be an independent scientist while also offering support when I needed it.

First, to the people I have met in this program, both in Chicago and at CERN.

To my lab mates (Ben, Cecilia, Emily S., Gabe, Jan, Jesse L., and Teresa), thank you for being such inspiring collaborators. Even though we weren't always in the same physical space or working on the same projects, I learned so much from our collaboration and conversations.

To the gFEX team, especially Anthony, Emily S., Greg, Rabia, and Simone, thank you for teaching me so much about gFEX, and for your commitment to making it work within the ATLAS system (no matter what time of night the errors occur). We're going to do great physics with gFEX!

To the DV+MET team (Ben, David, Emily F., Evan, Jan, Risa, and Zubair), it was a joy to collaborate with you, brainstorm, and problem-solve together. We were the analysis team that never slept, spread across the US, Canada, Australia, Japan, and Europe. Someone was always tired in our meetings (usually me), but we managed to bring this analysis to life together, and I'm thrilled that it's nearly complete.

As collaborators in my PhD, Emily S. and Jesse L. deserve special thanks. Emily was the

most wonderful office mate: my gym buddy, crepe-break partner, fellow “nail girly,” and bash alias genius who taught me so many useful shortcuts. I miss learning from her every day. Working on BREAD with Jesse in the Chicago lab remains one of my favorite physics experiences so far, because he made all elements of the work fun.

To my University of Chicago physics cohort (Casey, Paul, Tanvi, and Robby), thank you for making even the hardest coursework bearable. I’ll never forget our homework sessions, nights out around Hyde Park, brunches at Valois, and late-night Zoom calls. Your friendship brought light to some long, cold Chicago winters and hard problem sets.

Overall, living in France while working at CERN in Geneva was a once-in-a-lifetime experience. I feel incredibly lucky to have spent three years immersed in the culture, falling in love with new cheeses, taking bike rides in the French countryside, visiting the Ferney market on Saturday mornings, and, most of all, building strong friendships that defined this chapter of my life. CERN became home thanks to the incredible community that grew into my chosen family, the Soup-Squad. The friendships we built were the heart of my everyday life, full of laughter, comfort, and shared routines that made life at CERN beautiful and special, rather than isolating. These friends grounded me through the whirlwind of life abroad. To Anthony, Casey, Da Vinci, Emily S., Ezra, Greg, Hava, Jem, Jesse F., Meg, Mr. Mason, Nick, and Noah, thank you for making life in Geneva joyful. Our weekly cheese nights became a cherished tradition, often paired with wine and Overcooked. Around that, we really built a life together in France: playing board games, sharing dinners, spending sunny afternoons by the lake, and taking weekend getaways across Europe. Every winter, we would return to the Christmas market downtown. We had big outings for each birthday every year, we attended and later hosted beer tastings, and cooked American meals when we felt homesick. We even became loyal regulars at the English pub quiz (despite knowing very little about rugby or soccer). When you live over 4,000 miles from home, the friendships you form become

everything. Soup-squad, yours made Geneva feel like home.

Special shoutout to the CERN 2 Basketball team (2024 season champions of the GAB série C!) for providing entertainment and a community of sports enthusiasts. Cheering with Sara, Marija, and Justyna made the games even more fun. Drink the juice!

Next, the people who have been in my life long before grad school started.

To my friends from childhood — Clare, Eleanor, Jimmy, and Libby — thank you for growing up with me and always being there, no matter the distance or stage of life. We’ve stayed so close for so many years, even though we’ve spent so much of that time long-distance. I’m grateful to have you in my corner, and I’m excited to visit you more easily now that we’re on the same continent.

I also want to thank my childhood swim coaches: Katie, Sean, and the late John Drain. You taught me dedication, resilience, and the joy of swimming. These lessons shaped who I am. Even now, the pool is the place I return to when life feels overwhelming, and swimming has continued to bring new friendships into my life.

To my college friends from the University of Michigan, Becca, Charlotte, and Reid, thank you for being by my side during those formative years and helping me create my first home away from home. We’ve celebrated the highs (2023 National Champions!) and endured the lows together (like Trouble with the Snap). I’ll never forget dancing in the Big House to “Mr. Brightside,” “Pump It Up,” and “Hail to the Victors” with you. Thank you for visiting me and supporting me even when I was so far away. And to the wonderful friends I gained through Billy’s college roommates, I loved our adventures around Europe together: celebrating King’s Day in Amsterdam, Oktoberfest in Munich, St. Patrick’s Day in Dublin, and more. Go Blue forever!

To my friends from my semester at the University of Melbourne, Ashley, Grace, Irene, and

Jaimi, thank you for becoming such wonderful friends in just a few short months. That whirlwind semester gave me friendships for life and opened my world in ways that I'll always treasure. I'm grateful we stayed so close and that you visited me in France so many times. Irene, thank you for insisting we eat fondue when it was 100 degrees out; that was probably illegal in Switzerland, but it was delicious.

To my friends in Chicago, Alex, Ben, Irene, Joanna, Maddie, Nate, Robby, Sam, and Toka, thank you for turning this city into a home during the (far-too-brief) time I got to spend there. You're such awesome people, and the time we spent together lightened the intensity of grad school whenever I needed to get away from it all. I wish my time in Chicago had been longer because we had so much fun together, from Survivor watch parties to cooking nights to watching comedy shows at Second City. I'm sorry I ditched you all to move to France.

I'm also grateful to the many friends who made the journey to visit me in France, often more than once! Your visits filled my sleepy village with laughter, brought warmth to my routines, and reminded me how deeply I am loved. We kept a guest book to mark your stays and remember all the fun we had. RIP Saint-Genis Milk Machine.

To my mom and dad, Mary and Alan, you've inspired me my whole life, but your support while I lived so far away in France truly made this degree possible. Mom, you flew out just a week after I moved in to help me decorate my apartment and turn it from an empty space into a real *maison*. You always braved the summer heatwaves with no AC, and chaotic itineraries destined to be changed at the last minute thanks to EasyJet and Air Dolomiti. Dad, you made a French grocery store-compatible version of your famous chili that saved me on cold days when I missed home, and you turned every work trip to Europe into multiple visits to hang out with me. You were both so excited to see my life there, meeting my friends, visiting the CERN campus, even taking underground tours of ATLAS. I still can't believe

how much of Europe we explored together, from chasing the northern lights to swimming in the Mediterranean. I love you both so much. I did this all for you.

To my entire extended family, Billy's parents, and the people who became like second parents to me — Steve and Char, and the Khirallah family — thank you for being a steady source of love and pride throughout this journey and life. You've always made me feel like I could do anything.

To my undergraduate advisor, Myron, thank you for setting me on this path, sparking my enthusiasm for physics, and encouraging me to pursue this degree. To my undergraduate collaborators on the KOTO experiment at J-PARC, Brian, Josh, Melissa, Monica, Noah, and Stephanie, thank you for being so damn fun to do research with. The friendships we formed in Japan were the beginning of this story.

I also want to thank the teachers and mentors who believed in me along the way. Tom from the University of Michigan was a mentor who helped me navigate the world of physics research and graduate school. Todd was an excellent math teacher at Berkley High, whose enthusiasm for teaching made me love learning. Mrs. Kowalski, my fifth-grade teacher, who insisted that everyone in our class take the advanced math placement test (not just two students per class). This opportunity set me on a path in STEM that may not have come to light without you. Thank you. I'm also grateful to Olivier and Kyle, with whom I worked during my internship at CERN. You made research genuinely fun and inspired me to take on this whole Ph.D. nonsense in the first place. Thank you to all the educators who saw my potential and helped me reach it.

To my husband, Billy, thank you for everything. When we started graduate school, we chose long-distance because we were both pursuing our dream programs: I in Chicago and you in Austin. Through a twist of fate, we ended up together in France, and sharing this experience made everything more meaningful. You made life easier, funnier, and more joyful through

our weekend adventures across Europe, as well as fantastic home-cooked meals and date nights in the French countryside. We even bought a 2001 diesel minivan (the Blue Bus), and you learned to drive a stick so we could cross the French–Swiss border without waiting for the bus. You kept me grounded through the most challenging moments while finishing your own Ph.D. remotely. You are the smartest, kindest, and most thoughtful person I know, and I could not have done this without you.

Thank you to all the dogs who make the world a better place, especially my late best friend Jacobian. I wish you were here to celebrate. I know you’d be proud. Thank you to everyone who sent me pictures of my beloved long-distance dogs: Booker, Cami, Ernie, Martin Crane, Marzen, Stevie, and Violette. Your photos lit up many hard days.

A heartfelt merci beaucoup to Fromagerie Bouchet, located near my apartment in Saint-Genis-Pouilly. You fueled countless cheese nights that built friendships and made my time in France magical. Thank you for never judging my inability to speak French.

To everyone mentioned here (and to those whose names aren’t, but whose presence mattered all the same), thank you. This work is not mine alone; it is the result of an entire community of mentors, collaborators, friends, and loved ones who lifted me again and again. I am so deeply grateful.

ABSTRACT

This thesis presents a search for long-lived particles decaying to axinos using 137 fb^{-1} of proton-proton collision data collected by the ATLAS experiment at $\sqrt{s} = 13 \text{ TeV}$. The analysis focuses on events featuring displaced vertices with high invariant mass and track multiplicity, accompanied by substantial missing transverse energy. Results are primarily interpreted in the context of DFSZ axino models, where the axino serves as the lightest supersymmetric particle and a potential dark matter candidate. Background estimation is performed using data-driven techniques, with systematic uncertainties carefully evaluated through multiple validation regions. No significant excess over the estimated background is observed, allowing exclusion limits to be set on the model parameters at 95% confidence level. This work also details contributions to the gFEX system — a component of the ATLAS Level-1 trigger upgrade that enhances detection capabilities for missing energy and jet signatures. These developments strengthen future searches for axinos and other phenomena involving missing transverse energy, advancing our understanding of dark matter and beyond Standard Model physics.

CHAPTER 1

INTRODUCTION

The Standard Model (SM) of particle physics stands as one of the most precise and successful theoretical frameworks in science. It provides exceptional predictive power, describing the fundamental particles—quarks, leptons, and gauge bosons—and three of the four fundamental interactions that govern our universe: the electromagnetic, weak, and strong forces. The remarkable precision of the SM has been validated through numerous experiments, culminating in the discovery of the Higgs boson at the Large Hadron Collider (LHC) in 2012, which completed the particle spectrum predicted by the theory [1, 2]. Despite these extraordinary successes, significant gaps remain in our understanding of the universe. The SM fails to incorporate gravity, offers no explanation for neutrino masses, and crucially, provides no candidate particle for Dark Matter (DM)—a substance that astronomical observations indicate comprises approximately 85% of the matter content of the universe¹ [3, 4].

This introduction provides a broad overview of the theoretical landscape and experimental approaches that motivate the work presented in this thesis. The search for dark matter requires understanding connections between seemingly disparate areas of physics, from astrophysical observations to particle accelerators to sensitive laboratory detectors. While many technical terms and concepts are mentioned here, readers should not feel overwhelmed; each topic is introduced in detail in subsequent chapters with the necessary background and context.

Among the many theoretical candidates proposed to explain DM, axions are especially appealing because of their elegant origin. They were first introduced by Peccei and Quinn

1. This refers to approximately 85% of the matter content of the universe. When including dark energy, DM accounts for roughly 27% of the total energy density, according to cosmological measurements from the Planck collaboration [3].

as a solution to the “Strong Charge-Parity (CP) Problem” in Quantum Chromodynamics (QCD)² [5], and were later named by Frank Wilczek after a popular brand of dish soap—because the particle “cleaned up” a theoretical mess [6]³. Since then, axions have gained attention as promising DM candidates (Chapter 2.3) [7, 8, 9]. These hypothetical particles are expected to be extremely light—anywhere from 10^{-12} eV (peV) up to 10^{-3} eV (meV)—while being stable and interacting only very weakly with SM particles, making them a good fit for cold DM [10]. If axions make up the DM halo of our galaxy, they’d be floating all around us, with a local density of about 0.3 GeV/cm^3 [11].⁴

Direct detection experiments aim to measure axions by detecting the photons they convert into when interacting with a strong magnetic field, a process known as the Primakoff effect [17]. These experiments span an enormous mass range from 10^{-22} eV to 1 eV, where the photon frequency is directly proportional to axion mass: $f = m_a c^2 / h$. Traditional cavity haloscopes represent the most established approach, with Axion Dark Matter Experiment (ADMX) using a tunable resonant microwave cavity immersed in a strong magnetic field to search for axion-induced photon signals in the 1–10 μeV mass range, corresponding to photon frequencies of approximately 0.2–2.4 GHz [18, 19]. Haloscope At Yale Sensitive To Axion CDM (HAYSTAC) extends this technique to higher frequencies in the 15–50 μeV range by employing quantum-limited amplifiers to improve sensitivity [20, 21].

Despite ongoing improvements in sensitivity across cavity haloscope experiments, no axion

2. Discussed in detail in Chapter 2.2.1.1

3. Before being named the axion, the particle was simply referred to as the “PQ boson” in early literature [5]. Frank Wilczek introduced the name “axion” in 1978 but actually mislabeled the dish soap as a laundry detergent.

4. Throughout this thesis, the term “axion” refers specifically to the QCD axion arising from solutions to the strong CP problem, and “axino” refers to its supersymmetric partner. These particles appear in well-motivated models such as the KSVZ [12, 13] and DFSZ [14, 15] constructions, and should be distinguished from more general Axion-Like Particles (ALPs), which are not necessarily tied to QCD or the CP problem [16]. The collider-based search presented in this work is grounded in the QCD axion framework, with particular focus on scenarios where the axino serves as the Lightest Supersymmetric Particle (LSP) and a viable DM candidate.

signal has been observed at the writing of this thesis, prompting the development of next-generation detectors that overcome the fundamental limitation of narrowband frequency scanning. The newly established Broadband Reflector Experiment for Axion Detection (BREAD)(Chapter 2.3.1.3) [22] introduces a broadband dish antenna approach [23] that enables simultaneous sensitivity across multiple decades of axion mass, rather than requiring slow sequential frequency scanning like traditional haloscopes. BREAD uses a cylindrical metallic barrel with an internal Hershey’s Kiss-shaped parabolic reflector to focus potential axion-induced Electromagnetic (EM) radiation onto sensitive broadband detectors, as shown in Figure 2.6a. This design allows comprehensive coverage of theoretically motivated regions where QCD axion dark matter might exist, with the multi-stage program targeting masses from $\sim 20 \mu\text{eV}$ to 1 eV. *GigaBREAD* focuses on the 10 – 20 GHz range using cryogenic amplifiers, while *InfraBREAD* targets the terahertz regime near 1 eV using Superconducting Nanowire Single Photon Detectors (SNSPDs) [24], with first results already establishing new laboratory constraints in the 44–52 μeV range [25, 26].

While direct detection experiments continue their search for axions themselves, high-energy colliders offer complementary approaches to probe related particles within the axion sector. Specifically, in models where Supersymmetry (SUSY) (Chapter 2.4) [27, 28, 29, 30, 31, 32] and axions coexist, we can search for the axino - the fermionic supersymmetric partner of the axion. These SUSY-axion models provide a compelling extension to the SM, addressing both the strong CP problem and the DM puzzle. If the axino is the LSP and R-parity [33] is conserved, it would be stable and interact only weakly, appearing in collider experiments as Missing Transverse Energy (MET). Depending on the model parameters, the axino mass can lie within the energy reach of the LHC, typically ranging from a few GeV to several hundred GeV. This approach becomes especially valuable in scenarios where the axion-photon

coupling is significantly suppressed (Chapter 2.5)⁵, which would limit the sensitivity of direct detection experiments [34].

The potential presence of axinos in LHC collisions offers a unique opportunity to probe the axion sector indirectly, especially in scenarios where direct axion detection is suppressed. By searching for events with significant MET, we can potentially constrain the parameter space of SUSY-axion models and shed light on both the nature of DM and the possible existence of SUSY. The interplay between these two theoretical frameworks — axionic DM and SUSY—represents a rich area for exploration at the frontier of particle physics.

The interplay between these two theoretical frameworks—axionic DM and SUSY—represents a rich area for exploration at the frontier of particle physics. The search for DM is one of the most exciting open questions in modern physics, connecting ideas from astrophysics, cosmology, and particle physics. This thesis explores one piece of that puzzle by examining the axion sector from a collider perspective, specifically through searches for the axino. The methods and results presented here contribute to the broader effort to understand the nature of DM, whether through direct detection experiments seeking axions themselves or collider-based approaches targeting their supersymmetric partners.

This thesis presents a search for axinos using data collected by the particle detector at the Large Hadron Collider formerly known as A Toroidal LHC Apparatus (ATLAS) experiment [35] during LHC [36] Run 2 operations. Specifically, we focus on a scenario with a long-lived neutralino that decays to an axino plus either a pair of Z bosons (ZZ) or a Higgs boson pair (hh). The analysis targets collision signatures characterized by substantial MET, which occurs when the axino escapes detection, combined with the distinctive signatures from displaced decays of these long-lived neutralinos (Chapter 2.5.2). Given that few SM

5. The axion-photon coupling is not strictly zero, but can be suppressed by several orders of magnitude depending on the SUSY model [34].

processes mimic this combination of large MET and displaced decays, we employ a data-driven background estimation approach. The results are interpreted within the framework of SUSY-axion models, providing constraints on the relevant parameter space. The search methodology, event selection criteria, background estimation techniques, and statistical interpretation of the results are discussed in detail.

Beyond the present search for axinos, this work also contributes to the future of ATLAS and LHC physics through the development and support of the Global Feature Extractor (gFEX) triggering system. gFEX is specifically designed to select events with jets and MET — the latter being a key signature for axino searches. By improving the performance and monitoring of these triggers, we help ensure that ATLAS remains sensitive to new physics in future data-taking periods.

The remainder of this thesis is organized as follows. Chapter 2 provides the theoretical context for this search, introducing SUSY and axions in more detail and reviewing the direct detection experiments that aim to probe axion DM — including my contributions to the BREAD experiment. Chapter 3 introduces the LHC, followed by a description of the ATLAS detector in Chapter 3.3, as we fully transition to discussing the collider-based search for axinos. Chapter 4 focuses on the gFEX subsystem, including my work on its hardware and monitoring systems, which contribute to the long-term capabilities of ATLAS in detecting MET signatures. Chapter 5 describes how raw energy deposits and tracks observed in the detector are reconstructed into high-level physics objects—such as jets, leptons, and MET—which form the basis for the analysis presented later in this thesis. Finally, Chapter 6 presents the axino search itself, detailing the signal models, selection criteria, background estimation, and final results in the form of cross-section limits.

CHAPTER 2

PHYSICS MOTIVATION

The SM of particle physics describes the fundamental particles and three of the four known forces, but leaves several important questions unanswered. This chapter outlines the theoretical framework underlying the experimental search presented in this thesis. Chapter 2.1 reviews the SM structure and mathematical formalism. Chapter 2.2 discusses key limitations, including the hierarchy problem and the strong CP problem. Chapter 2.3 introduces axions as a solution to the strong CP problem and surveys current direct detection experiments. Chapter 2.4 presents supersymmetry as a framework for addressing the hierarchy problem. Chapter 2.5 demonstrates how combining axion and supersymmetry theories creates new collider search opportunities. Chapter 2.6 examines the expected LHC phenomenology for supersymmetric axion models, and establishes the motivation for the specific signature searched for in this thesis.

2.1 The Standard Model

The SM of particle physics describes elementary particles and three of the four fundamental forces. Fundamental particles are point-like objects with specific characteristics (mass, spin, charge, etc.) that distinguish them. The formalism for describing these particles and their interactions is Quantum Field Theory (QFT), where particles are described by quantum fields that respect certain symmetries. The theory successfully explains the electromagnetic, weak, and strong interactions through this framework, but excludes gravity. Built upon gauge symmetries, the SM has demonstrated remarkable predictive power, correctly forecasting the existence of the $W^\pm$ and Z bosons, the top quark, and the Higgs boson before their experimental discovery.




Figure 2.1: The fundamental particles of the SM organized by type and generation. The theory includes 12 fermions (6 quarks and 6 leptons) organized into three generations, four gauge bosons that mediate the fundamental forces, and the Higgs boson responsible for mass generation. Source: Wikipedia.

2.1.1 Particle Content

The SM contains 17 fundamental particles: 12 fermions, 4 gauge bosons, and 1 scalar boson. Fermions have half-integer spin and comprise all known matter, while bosons have integer spin and mediate interactions. Just as the periodic table organizes chemical elements by their properties, the SM organizes fundamental particles by their quantum numbers and interactions. However, while the periodic table describes composite atoms made of protons, neutrons, and electrons, the SM describes the truly fundamental constituents from which all matter is built.

2.1.1.1 Fermions

The 12 fermions are organized into three generations, with each generation containing two quarks and two leptons. The first generation consists of the up (u) and down (d) quarks, the electron (e), and the electron neutrino (ν_e). The second generation contains the charm (c) and strange (s) quarks, the muon (μ), and the muon neutrino (ν_μ). The third generation includes the top (t) and bottom (b) quarks, the tau (τ), and the tau neutrino (ν_τ).

Up-type quarks (u , c , and t) carry electric charge $+2/3$, while down-type quarks (d , s , and b) carry charge $-1/3$. Charged leptons (e , μ , and τ) have charge -1 , and neutrinos are electrically neutral.

Spin and Statistics All fermions have half-integer spin ($s = 1/2$) and obey Fermi-Dirac statistics. This means that no two fermions can occupy the same quantum state simultaneously, a principle known as the Pauli exclusion principle. This fundamental property is responsible for the structure of atoms and the stability of matter.

Mass and Generations The mass scale of fermions increases dramatically across generations. First-generation particles have masses in the MeV range, while third-generation quarks reach the GeV scale, with the top quark exceeding 170GeV. All fermions beyond the first generation are unstable and decay to lighter particles, which is why ordinary matter consists primarily of first-generation fermions: u and d quarks (forming protons and neutrons) and electrons. The fermion content of the SM and their masses are summarized in Table 2.1¹ The pattern of quark masses across generations is central to the Strong CP Problem discussed in Chapter 2.2.1.1. Together with QCD vacuum dynamics (Chapter 2.1.4.1), this structure determines the axion’s coupling strength and mass scale, and sets the possible decay channels that can produce axinos, as discussed in Chapter 6.

Generation	Quark	Quark Mass [MeV]	Lepton	Lepton Mass [MeV]
I	u	2.2	e	0.511
	d	4.7	ν_e	—
II	c	1.3×10^3	μ	106
	s	96	ν_μ	—
III	t	1.73×10^5	τ	1.78×10^3
	b	4.2×10^3	ν_τ	—

Table 2.1: Fermions in the SM organized by generation and approximate masses in MeV.

Chirality An essential property of fermions in the mathematical structure of the SM is their chirality. The more familiar concept of helicity refers to the projection of a particle’s spin along its momentum: a particle has right-handed helicity if its spin and momentum are aligned, and left-handed helicity if they are anti-aligned. However, helicity is not Lorentz invariant for massive particles, since a boost can reverse the direction of momentum. Chirality, by contrast, is a Lorentz-invariant property that plays a central role in fermion interactions,

1. Neutrinos are experimentally known to have small but nonzero masses, as demonstrated by neutrino oscillation observations. The SM originally assumed massless neutrinos. While theoretical extensions exist to accommodate neutrino masses, they are not central to the physics discussed in this thesis.

particularly in the weak interaction, where only left-handed fermions and right-handed antifermions participate in charged-current processes. Chirality also constrains how axions couple to SM fermions and determines which SUSY decay channels can produce axinos, as discussed in Chapter 6.

Color Charge and Confinement Quarks carry a type of charge known as color charge, which comes in three types: red, green, and blue. Due to the strong interaction, quarks are never observed in isolation but are confined within color-neutral combinations called hadrons. This phenomenon, known as color confinement, arises because the strong force becomes stronger as the distance between quarks increases, unlike electromagnetic or gravitational forces, which weaken with separation. As a result, when quarks are pulled apart, the energy stored in the color field grows until it becomes energetically favorable to produce a new quark-antiquark pair from the vacuum, preventing the isolation of individual quarks. This will be relevant for the search for axinos described in Chapter 6 because axinos can be produced with particles that decay to quarks, which then hadronize due to color confinement.

Hadrons There are two main types of hadrons based on their quark content. Mesons consist of a quark-antiquark pair ($q\bar{q}$) and achieve color neutrality because the antiquark carries the anti-color of the quark (for example, red and anti-red). Examples include pions, π , kaons, K , and the J/ψ meson. All mesons are unstable and decay rapidly, with proper decay lengths ($c\tau$) ranging from femtometers to kilometers depending on the specific meson. The neutral kaon K_S^0 is particularly important for Long Lived Particle (LLP) searches, as its characteristic decay length and two-track decay signature make it an ideal benchmark particle for validating tracking efficiency and estimating systematic uncertainties in displaced vertex reconstruction algorithms. These K_S^0 studies are employed for background estimation (Chapter 6.10.1.3) and tracking reconstruction uncertainty evaluation (Chapter 6.12.4.1).

Baryons consist of three quarks (qqq) and achieve color neutrality by combining one quark of each color: red, green, and blue. The most familiar baryons are the proton (uud) and neutron (udd), which form atomic nuclei. The proton, composed of two up quarks and one down quark, has electric charge $+1$ and mass approximately 938 MeV. The neutron, composed of two down quarks and one up quark, has electric charge 0 and mass approximately 940 MeV. Other baryons include the Λ , Σ , Ξ , and Ω particles, which contain one or more strange quarks and are produced in high-energy interactions before decaying rapidly to lighter particles. Of particular relevance to LLP searches, the Λ (uds) baryon has a measurable lifetime and typically produces low-mass, low-multiplicity displaced vertices that serve as important background considerations, as discussed in Chapter 6.1.

Together with electrons, protons and neutrons constitute all ordinary matter in the universe. This form of matter is referred to as baryonic matter, since it is fundamentally composed of baryons bound with electrons to form atoms. The proton is stable, or at least has a lifetime exceeding 10^{34} years, while the free neutron decays via beta decay with a lifetime of about 15 minutes: $n \rightarrow p + e^- + \bar{\nu}_e$. Neutrons bound in stable nuclei do not decay.

Antiparticles and Conservation Laws Each fermion has a corresponding antiparticle with the same mass and spin but opposite quantum numbers. Common examples include the positron (e^+), the antiparticle of the electron, and the antiproton ($\bar{p}$). Antiparticles are typically denoted with a horizontal bar over their symbol, such as $\bar{b}$ for the bottom antiquark. When a particle and its antiparticle meet, they can annihilate and convert their mass into energy, often producing photons.

This symmetry between matter and antimatter is reflected in the conservation laws that govern all SM interactions. The theory assigns several conserved quantum numbers to fermions, including electric charge, lepton number, and baryon number. These additive

quantum numbers must be conserved in every interaction, constraining which processes can occur in nature.

Lepton number is separately conserved for each generation: electron lepton number L_e , muon lepton number L_μ , and tau lepton number L_τ . For example, the electron (e) and electron neutrino (ν_e) have $L_e = +1$, while the positron (e^+) and electron antineutrino ($\bar{\nu}_e$) have $L_e = -1$. The same pattern applies to the muon and tau families. All other particles carry zero lepton number. Similarly, quarks carry baryon number $B = +1/3$, and antiquarks carry $B = -1/3$. This assignment ensures that baryons have $B = +1$, antibaryons have $B = -1$, and mesons have $B = 0$, reflecting their quark content and maintaining the fundamental matter-antimatter symmetry. These conservation laws dictate which decays are allowed and therefore constrain the possible couplings and production modes of candidate new particles like the axino, as discussed in Chapter 6.

Flavor Mixing While quarks are organized into three generations based on their mass eigenstates, their interactions via the weak force occur in a different basis called flavor eigenstates. The relationship between these bases is described by the Cabibbo–Kobayashi–Maskawa (CKM) matrix for quarks and the Pontecorvo–Maki–Nakagawa–Sakata (PMNS) matrix for neutrinos. This mixing allows for transitions between different generations in weak interactions and is responsible for phenomena such as neutrino oscillations. The mixing explains why only first-generation particles are stable; heavier generations decay rapidly to lighter ones through weak interactions, ensuring that ordinary matter consists entirely of first-generation fermions.

2.1.1.2 Bosons

The fundamental forces in the SM are mediated by gauge bosons, whose properties are summarized in Table 2.2. The massless photon mediates electromagnetic interactions between electrically charged particles and gives rise to a long-range force. At low energies, its coupling strength is given by the fine-structure constant, $\alpha \approx 1/137$.

Boson	Mass (GeV)	Associated Force / Field	Range
Photon (γ)	0	Electromagnetic	Long-range
$W^\pm$	80.4	Weak (charged)	Short-range
Z	91.2	Weak (neutral)	Short-range
Gluon (g)	0	Strong	Effectively short-range
Higgs (H)	125	Higgs field / EWSB	Not a force mediator

Table 2.2: Properties of bosons in the SM.

The weak interaction is mediated by the massive $W^\pm$ and Z bosons, with masses of approximately 80 GeV and 91 GeV, respectively. Their large masses limit the range of the weak force to subatomic distances, making it significant only when particles are in close proximity. The weak coupling constant at low energies is $\alpha_W \approx 1/30$. The $W^\pm$ bosons mediate charged-current processes such as beta decay, while the Z boson is responsible for neutral-current interactions.

The strong force is mediated by eight massless gluons, which carry color charge and can interact with each other. This self-interaction occurs because QCD is a non-Abelian gauge theory based on the symmetry group $SU(3)$, unlike electromagnetism, which is an Abelian $U(1)$ gauge theory where photons cannot interact with each other. The gluons' ability to carry color charge and couple to other gluons leads to unique properties of the strong force, most notably color confinement. Unlike other fundamental forces that weaken with distance, the strong force actually increases as quarks are pulled apart, making it energetically favorable to create new quark-antiquark pairs rather than separate existing quarks. This confines

the strong force to distances smaller than atomic nuclei (typically less than 1 fm) and ensures that free quarks are never observed in nature.

The strong coupling constant changes dramatically with energy. At high energies, it becomes small enough to allow precise calculations, while at low energies, it grows so large that quarks and gluons cannot exist as free particles and must form jets of hadrons. At LHC energies, $\alpha_s \approx 0.1$, enabling the perturbative QCD calculations used to predict collider processes, while the low-energy regime governs the jet formation observed in detectors, discussed in Chapter 5.2.1.

The Higgs boson is the quantum excitation of the Higgs field, which underlies the mechanism that gives mass to the $W^\pm$ and Z bosons, as well as to the fermions in the SM.

2.1.2 *Mathematical Formalism*

The SM is a QFT, in which each particle corresponds to a quantized field, and their interactions are governed by the symmetries of the model. These symmetries are central to the structure of the SM: according to Noether's theorem, every continuous symmetry gives rise to a conserved physical quantity [37], such as electric charge or energy.

The SM is defined by a mathematical expression called a Lagrangian, which specifies how the fields behave and interact. This Lagrangian includes all terms allowed by the symmetries, and nothing more. Predictions about physical processes are derived by systematically expanding this expression and calculating probabilities for different outcomes using Feynman diagrams. In essence, the symmetries determine which interactions are possible, the Lagrangian encodes them mathematically, and Feynman diagrams provide the computational tools to make testable predictions.

The SM is also a gauge theory, meaning the equations remain unchanged under transfor-

mations that can vary from place to place in space and time. These symmetries correspond to the fundamental forces and require force-carrying particles (gauge bosons) to maintain consistency. The SM is based on three symmetry groups:

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

where each group corresponds to a fundamental interaction:

- $SU(3)_C$ describes color charge and the strong force, with eight generators corresponding to the eight gluons
- $SU(2)_L$ describes weak interactions acting on left-handed particles, with three generators for the original weak gauge bosons
- $U(1)_Y$ governs weak hypercharge, a property related to electric charge, with one generator and one gauge boson

At energies below the electroweak scale (~ 100 GeV), the weak and electromagnetic forces appear unified through spontaneous symmetry breaking via the Higgs mechanism. This gives mass to the $W^\pm$ and Z bosons while keeping the photon massless. The structure reflects the three forces in the SM: $SU(3)_C$ for the strong interaction, and $SU(2)_L \times U(1)_Y$ for the electroweak interaction. The model contains 26 parameters, including particle masses and interaction strengths, that must be measured experimentally.

2.1.3 Electroweak Unification

The electromagnetic and weak forces are unified under the $SU(2)_L \times U(1)_Y$ gauge symmetry. Electromagnetic interactions between charged particles are mediated by the massless photon

and are described by the $U(1)_{\text{EM}}$ gauge theory known as quantum electrodynamics (QED). At energies much higher than the masses of the W and Z bosons (~ 100 GeV), the massive gauge bosons behave effectively as massless particles, revealing the underlying electroweak unification. At lower energies, however, the forces appear very different: electromagnetism is a long-range force that couples universally to electric charge, while the weak force is a short-range force due to its massive mediators and couples only to left-handed fermions.

In electroweak theory, particles carry two types of weak quantum numbers. The first is weak isospin, analogous to spin but associated with the weak force. Left-handed fermions have weak isospin $I = 1/2$ and appear in doublets (like the electron and its neutrino), where the third component $T_3 = \pm 1/2$ distinguishes the two members. Right-handed fermions have $I = 0$ and are singlets under $SU(2)_L$. The second quantum number is weak hypercharge Y , associated with the $U(1)_Y$ gauge group.

A particle's electric charge Q is related to its weak quantum numbers via the Gell-Mann–Nishijima relation:

$$Q = T_3 + \frac{Y}{2}. \quad (2.2)$$

For example, the left-handed electron has $T_3 = -1/2$ and $Y = -1$, giving $Q = -1/2 + (-1)/2 = -1$, which matches its known electric charge. The complete set of quantum number assignments can be found in Table 2.3.

Particles	Q	T_3	Y_L (Left-Handed)	Y_R (Right-Handed)
Neutrinos (ν_e, ν_μ, ν_τ)	0	+1/2	-1	-
Leptons (e, μ, τ)	-1	-1/2	-1	-2
Up-Type Quarks (u, c, t)	+2/3	+1/2	+1/3	+4/3
Down-Type Quarks (d, s, b)	-1/3	-1/2	+1/3	-2/3

Table 2.3: Electric charge, weak isospin, and weak hypercharge assignments for SM fermions.

Left-handed fermions transform as doublets under $SU(2)_L$ and participate in both charged and neutral weak interactions. Right-handed fermions are singlets and do not couple to the

$W^\pm$ bosons, participating only in neutral current interactions and electromagnetism.

At high energies above the electroweak scale, the full $SU(2)_L \times U(1)_Y$ symmetry is restored, and the theory contains four massless gauge bosons: W^1 , W^2 , W^3 from $SU(2)_L$, and B from $U(1)_Y$. At lower energies, this symmetry is spontaneously broken to $U(1)_{\text{EM}}$ via the Higgs mechanism, giving mass to the $W^\pm$ and Z bosons while leaving the photon massless. The physical gauge bosons are linear combinations of the original fields:

$$W^\pm = \frac{1}{\sqrt{2}}(W^1 \mp iW^2), \quad (2.3)$$

$$Z = \cos \theta_W W^3 - \sin \theta_W B, \quad (2.4)$$

$$A = \sin \theta_W W^3 + \cos \theta_W B, \quad (2.5)$$

where A is the photon field, and θ_W is the weak mixing angle (Weinberg angle).

2.1.4 Higgs Mechanism

The Higgs mechanism provides the theoretical framework for mass generation in the SM. It involves a scalar field known as the Higgs field, which permeates all of space. This field has a potential energy function:

$$V(\phi) = \mu^2 \phi^\dagger \phi + \lambda (\phi^\dagger \phi)^2 \quad (2.6)$$

For $\mu^2 < 0$ and $\lambda > 0$, this potential has a distinctive shape often called the “wine bottle potential.” The key feature is that when $\mu^2 < 0$, the minimum energy does not occur at zero field strength ($\phi = 0$). Instead, the potential has a maximum at the origin and a circular ring of minima at $|\phi| = v = \sqrt{-\mu^2/\lambda} \approx 246 \text{ GeV}$.

This creates a fundamental choice: the field must settle somewhere on this ring of equivalent minima. Like a ball rolling to the bottom of a wine bottle and settling in the circular groove

around the base, the Higgs field spontaneously chooses one particular direction and settles into this nonzero value everywhere in space. This value is called the vacuum expectation value. The choice of which minimum to settle into is arbitrary, but once chosen, it breaks the original symmetry of the system. This is spontaneous symmetry breaking: the underlying equations (the potential) respect the full $SU(2)_L \times U(1)_Y$ symmetry, but the actual vacuum state does not, reducing the symmetry to $U(1)_{\text{EM}}$.

The key consequence is that this nonzero field value permeates all of space, meaning that particles constantly interact with the Higgs field as they move. These interactions give mass to the W and Z bosons through their coupling to the field, while the photon remains massless because it doesn't couple to the Higgs field. The quantum excitations of the Higgs field around this vacuum state appear as the physical Higgs boson.

Fermions acquire their masses through this same mechanism, but each type interacts with a different strength. This interaction strength directly determines how massive the particle is. The relationship is $m_f = y_f v / \sqrt{2}$, where y_f measures the strength of coupling of that particular fermion to the Higgs field. This explains why different particles have such different masses: the top quark interacts very strongly with the Higgs field and is therefore very heavy, while the electron interacts much more weakly and is much lighter.

For quarks, there's an additional complication. The quarks that have definite masses (like the up, down, and strange quarks we observe) are not exactly the same as the quarks that participate in weak interactions. Instead, they're mixed together in a way described by the CKM matrix:

$$V_{\text{CKM}} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \quad (2.7)$$

This matrix tells us how much each mass eigenstate (the quarks we actually observe) contains

of each weak eigenstate (the combinations that participate in weak decays). For example, when a W boson decays to what we call a “down quark,” it’s actually producing a mixture that’s mostly down quark but with small amounts of strange and bottom quarks mixed in. For leptons, this mixing doesn’t occur in the SM. The electron that participates in weak interactions is precisely the same as the electron with a definite mass.

2.1.4.1 Quantum Chromodynamics

Quantum Chromodynamics (QCD) describes the strong interaction through an $SU(3)_C$ gauge theory. The subscript C denotes color charge. The strong force operates through three types of color charge: red, green, and blue. Quarks carry color charge, antiquarks carry anti-color charge, and gluons carry both color and anti-color charge. This results in eight independent gluon states corresponding to the eight generators of the $SU(3)$ group.

Rather than weakening with distance like other forces, the strong force actually becomes stronger as colored particles are pulled apart. This leads to color confinement: quarks and gluons cannot exist as free particles but are always bound within color-neutral hadrons. When colored particles are separated, the energy stored in the color field increases linearly with the distance between them. Eventually, it becomes energetically favorable to create new quark-antiquark pairs from the vacuum rather than continue the separation. This hadronization process produces the jets observed in high-energy collisions, consisting of sprays of color-neutral hadrons such as mesons and baryons.

2.2 Beyond the Standard Model

The SM achieves remarkable agreement with experimental data, but several fundamental problems demonstrate that it cannot be the complete theory of nature. These problems

fall into two categories: theoretical fine-tuning issues within the model and observational phenomena that the SM cannot accommodate.

2.2.1 *Fine-Tuning Problems*

The SM suffers from several fine-tuning problems where parameters must be set to extremely precise values without any natural explanation. Two stand out as particularly severe and directly relevant to extensions of the SM: the strong CP problem and the hierarchy problem. These problems suggest that the SM is an effective theory that emerges from a more fundamental framework at higher energy scales.

2.2.1.1 The Strong CP Problem

While CP violation is well-established in weak interactions through the CKM matrix², the situation in the strong force is puzzling. The QCD Lagrangian allows for CP violation through an additional term:

$$\mathcal{L}_\theta = \frac{\theta g_s^2}{32\pi^2} G_{\mu\nu}^a \tilde{G}^{a\mu\nu} \quad (2.8)$$

where $G_{\mu\nu}^a$ is the gluon field strength tensor and $\tilde{G}^{a\mu\nu}$ is its dual,

$$\tilde{G}^{a\mu\nu} = \frac{1}{2} \epsilon^{\mu\nu\rho\sigma} G_{\rho\sigma}^a, \quad (2.9)$$

with $\epsilon^{\mu\nu\rho\sigma}$ the totally antisymmetric Levi-Civita tensor. The dual tensor mixes space and time components, and the θ term violates CP symmetry when $\theta \neq 0$.

2. First observed in neutral kaon decays [38] and later confirmed in B-mesons [39, 40], CP violation in weak interactions arises from the irreducible complex phase in the CKM matrix. This phase cannot be removed by redefining quark fields when there are at least three generations, as shown by Kobayashi and Maskawa [41].

The presence of this term would induce a permanent Electric Dipole Moment (EDM) for the neutron, as CP violation can create an asymmetry between the distribution of positive and negative charges within the neutron. This effect is estimated as

$$d_n \approx 5.2 \times 10^{-16} \theta \text{ e} \cdot \text{cm}. \quad (2.10)$$

However, experiments constrain the neutron EDM to be less than $1.8 \times 10^{-26} \text{ e} \cdot \text{cm}$, implying $\theta < 3 \times 10^{-10}$ [42]. This extremely small value has no explanation within the SM: no known symmetry forces θ to vanish, and no known mechanism dynamically suppresses it.

This unnatural fine-tuning defines the *strong CP problem*. Unlike the weak force, where CP violation is observed and understood, the strong force appears to conserve CP to an extraordinary degree, despite no experimentally verified theoretical reason for this behavior. Elegant theoretical solutions exist to this puzzle through the introduction of axions (Chapter 2.3) and their SUSY partners, axinos (Chapter 2.5).

2.2.1.2 The Hierarchy Problem

In the SM, the Higgs boson mass receives contributions from all the particles to which it couples. This requires a high level of fine-tuning of model parameters, so that divergences in contributions from higher-order corrections cancel out across 30 orders of magnitude to make the SM consistent with observations.

In QFT, virtual particle interactions continually modify particle properties through quantum corrections. For the Higgs boson, these corrections grow quadratically with the energy cutoff scale Λ —the maximum energy up to which we expect the SM to remain valid. For instance,

the one-loop correction from a fermion of mass m_f and Yukawa coupling y_f contributes

$$\Delta m_h^2 \sim -\frac{y_f^2}{8\pi^2} \Lambda^2. \quad (2.11)$$

If the SM is valid up to the Planck scale ($\Lambda \sim 10^{19}$ GeV), these corrections would naturally push the Higgs mass toward that scale. However, the measured Higgs mass is only 125 GeV. To reconcile this, the Higgs mass parameter must be fine-tuned to cancel the quantum corrections with extraordinary precision, better than one part in 10^{32} .

While such fine-tuning is mathematically possible, it seems implausible without a deeper explanation. As physicist Gerardus 't Hooft emphasized, it is unlikely that nature would rely on delicate cancellations of quantum corrections without an underlying mechanism [43]. This motivates extensions of the SM that stabilize the Higgs mass through new symmetries or particles.

2.2.2 Observational Phenomena Beyond the Standard Model

Multiple astronomical observations provide compelling evidence for physics beyond the SM. Most notably, these observations indicate that the majority of matter in the universe is invisible and does not interact electromagnetically, referred to as dark matter.

Galaxy rotation curves first revealed this puzzle in the 1970s when Vera Rubin found that stars in spiral galaxies orbit much faster than expected from visible matter alone [44]. Gravitational lensing studies, including observations of the Bullet Cluster, directly measure total matter distributions and confirm the presence of non-luminous mass [45].

The strongest constraints come from Cosmic Microwave Background (CMB) measurements. The Planck satellite's observations of temperature fluctuations precisely determine the uni-

Energy Content of the Universe




Figure 2.2: Energy density composition of the universe from cosmological observations [3]. Dark matter accounts for approximately 26.8% of the total energy budget, while ordinary matter makes up only 4.9%. The SM does not contain a viable candidate for dark matter. Figure courtesy of David Miller.

verse’s composition, constraining the cold dark matter density to $\Omega_{\text{cdm}}h^2 = 0.120 \pm 0.001$ [3]. This implies dark matter accounts for approximately 26.8% of the universe, as shown in Figure 2.2.

Despite its success in describing particle interactions, the SM offers no viable dark matter candidate. Neutrinos are the only stable, weakly interacting particles in the model, but their small masses make them inconsistent with the cold dark matter required by structure formation. This provides strong motivation for searching for new fundamental particles beyond the SM.

2.2.3 Motivation for New Physics

The problems we’ve outlined in Sections 2.2.1 and 2.2.2 make it clear that the SM, despite its successes, isn’t the whole story. The fine-tuning issues suggest there’s some underlying structure we’re missing, while the dark matter evidence directly shows us there are particles out there that the SM doesn’t include. This naturally leads physicists to look for theories that can fill in these gaps. We want extensions that keep all the good parts of the SM while fixing its problems. The next sections introduce two popular candidates: axions and SUSY. Both offer elegant solutions to some of the issues we’ve discussed, and both make predictions we can test with experiments.

Most particle physics searches tend to focus on only one theory at a time, hunting for the specific particles that each model predicts. But there’s also the intriguing possibility of combining these frameworks. When axions and SUSY coexist in the same model, we can search for particles from one sector (like SUSY particles) that might reveal information about the other (like the axion sector). This combined approach is less common but offers a unique window into the strong CP problem, the hierarchy problem, and the dark matter puzzle

simultaneously.

2.3 Axions as a Solution to the Strong CP Problem

The Peccei-Quinn (PQ) mechanism [5] solves the strong CP problem by promoting the QCD θ parameter to a dynamical field. This is done by extending the SM with a global chiral symmetry $U(1)_{\text{PQ}}$, which is spontaneously broken at an energy scale f_a . When this symmetry breaks, it produces the axion, a pseudo-Nambu-Goldstone boson whose field couples to QCD in such a way that non-perturbative effects [46, 13] drive the effective θ parameter to zero. This mechanism removes the need for fine-tuning [47] and also gives the axion a small mass.

Because axions are stable, extremely weakly interacting, and can be abundantly produced in the early universe, they are also a compelling dark matter candidate [7, 8, 9]. Their unique connection to the strong CP problem and dark matter has motivated extensive experimental searches across a wide range of techniques.

2.3.1 Searching for Axions Through Photon Conversion

Most direct axion detection experiments rely on the axion-photon coupling, which allows axions to convert into photons in the presence of electromagnetic fields. Although the axion does not couple directly to photons, an effective interaction arises through loop diagrams mediated by QCD effects, as illustrated in Figure 2.3a. This loop-induced coupling, denoted $g_{a\gamma\gamma}$, is central to most experimental searches.




Figure 2.3: Axion-photon interactions. (a) The triangle loop process generates the coupling $g_{a\gamma\gamma}$, which allows axion decay into two photons. (b) The inverse Primakoff effect, where an axion converts into a single photon in an external magnetic field $\vec{B}_0$. Figures courtesy of [48].

2.3.1.1 Axion Modified Maxwell's Equations

The reason axions can be detected through photon conversion lies in how they modify the fundamental equations governing electromagnetic fields [49]. The interaction between axions and photons is described by adding a term to the electromagnetic Lagrangian:

$$\mathcal{L}_{a\gamma\gamma} = -\frac{g_{a\gamma\gamma}}{4}a(x)F_{\mu\nu}\tilde{F}^{\mu\nu} \quad (2.12)$$

where $a(x)$ is the axion field, $F_{\mu\nu}$ is the electromagnetic field tensor, $\tilde{F}^{\mu\nu}$ is its dual, and $g_{a\gamma\gamma}$ is the axion-photon coupling strength. In terms of electric and magnetic fields:

$$\mathcal{L}_{a\gamma\gamma} = g_{a\gamma\gamma}a(x)\vec{E} \cdot \vec{B} \quad (2.13)$$

The axion-photon coupling strength is given by [50]:

$$g_{a\gamma\gamma} = \frac{\alpha}{2\pi f_a} \left(\frac{E}{N} - 1.92 \right) \quad (2.14)$$

where α is the fine structure constant, f_a is symmetry breaking energy scale of the axion, E and N are the electromagnetic and color anomaly coefficients respectively. The ratio E/N is model-dependent: for the Kim–Shifman–Vainshtein–Zakharov model (KSVZ) model $E/N = 0$ [12, 13], for the Dine–Fischler–Srednicki–Zhitnitsky model (DFSZ) model $E/N = \frac{8}{3}$ [14, 15], and the constant 1.92 arises from axion-meson mixing at low energies [50]. The specific values of the ratio E/N in SUSY models will be discussed in Chapter 2.5.

This coupling modifies Maxwell’s equations, enabling axion-photon conversion in the presence of external magnetic fields through the Primakoff effect [17]. For axion dark matter, which behaves as a coherent, oscillating field with frequency $\omega = m_a c^2 / \hbar$, an external magnetic field $\vec{B}_{\text{ext}}$ generates an oscillating electric field that can be converted into detectable electromagnetic waves using appropriate boundary conditions.

2.3.1.2 Axion Models and Coupling Predictions

For the QCD axion, non-perturbative QCD effects give the axion a small mass that depends on the PQ symmetry-breaking energy scale f_a :

$$m_a = (1 \text{ meV}) \left(\frac{5.7 \times 10^9 \text{ GeV}}{f_a} \right) \quad (2.15)$$

This relationship shows that higher symmetry-breaking scales f_a correspond to lighter axions. Combined with the axion-photon coupling strength $g_{a\gamma\gamma} \propto f_a^{-1}$ from Equation 2.14, this creates a direct relationship between axion mass and coupling strength that is central to




Figure 2.4: Illustration of axion-photon conversion at a conducting surface immersed in a background magnetic field $\vec{B}_0$. The axion-photon coupling $g_{a\gamma\gamma}$ induces an effective current at the surface due to the oscillating axion field, generating a small electric field $\vec{E}_a$. To satisfy the boundary condition $\vec{E}_{\parallel} = 0$, an outgoing electromagnetic wave is emitted. This process, based on the inverse Primakoff effect, forms the basis for many direct detection experiments searching for axion dark matter. Figure courtesy of Andrew Sonnenschein.

experimental searches.

Two primary theoretical models describe how the PQ symmetry integrates with the SM, each making different predictions for axion-photon coupling strength depending on which particles carry PQ charges. The key difference lies in their electromagnetic anomaly contributions, characterized by the ratio E/N in Equation 2.14. The KSVZ model [12, 13] introduces new heavy quarks that carry PQ charges while leaving SM fermions unchanged. Assuming electrically neutral heavy quarks, this gives $E/N = 0$, leading to:

$$g_{a\gamma\gamma}^{\text{KSVZ}} = -3.9 \times 10^{-13} \left(\frac{m_a}{\text{meV}} \right) \text{ GeV}^{-1} \quad (2.16)$$

The DFSZ model [14, 15] assigns PQ charges directly to SM fermions and extends the Higgs sector with multiple Higgs doublets. This approach gives $E/N = \frac{8}{3}$, resulting in:

$$g_{a\gamma\gamma}^{\text{DFSZ}} = 1.5 \times 10^{-13} \left(\frac{m_a}{\text{meV}} \right) \text{ GeV}^{-1} \quad (2.17)$$

2.3.1.3 The Axion Search Program

The axion’s solution to the strong CP problem and its viability as a dark matter candidate have motivated an extensive experimental search program. Searches scan across an enormous mass range, from ultra-light particles near 10^{-22} eV to masses approaching the eV scale. While the QCD axion was initially proposed to solve the strong CP problem [5, 6], with theoretical models predicting masses typically in the μeV to meV range depending on the PQ symmetry breaking scale [51], the broader search program also targets ALPs that could constitute dark matter without addressing the strong CP problem.

The photon frequency from axion conversion is directly proportional to the axion mass:

$$f = \frac{m_a c^2}{h} \quad (2.18)$$

A $1 \mu\text{eV}$ axion converts to a microwave photon at 240 MHz, while a 1 eV axion generates infrared photons at 240 THz. This wide frequency range necessitates the use of completely different experimental approaches for different mass scales.




Figure 2.5: The axion mass (bottom) determines the frequency and wavelength of converted photons (middle), which sets the equivalent blackbody temperature that defines the minimum cryogenic operating temperature for thermal noise suppression. The figure also shows the corresponding DFSZ axion-photon coupling strength (top) and the PQ symmetry-breaking scale. Figure courtesy of David W. Miller.

Figure 2.5 illustrates the relationships that determine experimental requirements for different axion masses. The axion mass sets the converted photon frequency from Equation 2.18,

which corresponds to an equivalent blackbody temperature: $T = hf/k_B$, where h is Planck’s constant and k_B is the Boltzmann constant. This temperature scale is crucial for experimental design; for example, a $1\ \mu\text{eV}$ axion corresponds to 12 mK thermal energy, requiring cryogenic systems that operate below this temperature to detect photons above the thermal noise level. The figure also shows the DFSZ model prediction for the axion-photon coupling strength as a function of mass, illustrating how both thermal and coupling considerations determine experimental sensitivity requirements.

A variety of experimental methods target different axion mass ranges using specialized experimental designs:

- **Cavity haloscopes** use superconducting resonant cavities to enhance axion-photon conversion. Axion Dark Matter Experiment (ADMX) [19] has probed 2.7–4.2 μeV masses with sensitivity approaching theoretical predictions. Haloscope At Yale Sensitive To Axion CDM (HAYSTAC) [21] extends to 15–50 μeV using quantum-limited amplifiers. These experiments place a tunable microwave cavity in a strong magnetic field, where dark matter axions convert to detectable photons when the cavity resonance matches the axion mass.
- **Dielectric haloscopes** like Magnetized Disc and Mirror Axion experiment (MAD-MAX) [52] aim to probe 40–400 μeV masses using stacks of dielectric disks to enhance axion-photon conversion efficiency at higher frequencies and masses. The alternating dielectric interfaces create coherent enhancement of the converted photon signal, with disk spacing optimized for specific frequency ranges.
- **Light-shining-through-wall experiments** such as Any Light Particle Search (ALPS)-II [53] search for axions up to 0.1 meV by attempting to convert photons to axions, then back to photons through an opaque barrier. Laser light is converted to axions in

a magnetic field on one side of an opaque wall, with any reconverted photons detected by sensitive receivers on the other side.

- **Helioscopes** like CERN Axion Solar Telescope (CAST) [54] search for solar axions by looking for their conversion to X-rays in laboratory magnetic fields. These experiments point a strong dipole magnet toward the Sun and search for X-ray photons that would be produced when solar axions convert back to photons in the magnetic field.

BREAD Complementing this broad search program, new experiments such as the Broadband Reflector Experiment for Axion Detection (BREAD)³ have recently been developed to search for axions using novel broadband techniques [24]. The BREAD experiment follows a staged approach, with multiple frequency-targeted implementations designed to span different regions of the QCD axion parameter space. BREAD employs a unique reflector design, consisting of a cylindrical metallic barrel where axion-to-photon conversion occurs via the inverse Primakoff effect, with a Hershey’s Kiss-shaped parabolic reflector positioned inside the cylinder (Figure 2.6a). The entire apparatus sits within an Magnetic Resonance Imaging (MRI) magnet, providing a strong static magnetic field parallel to the metallic cylinder. As illustrated in Figure 2.4, axions traversing the strong static magnetic field inside the MRI magnet induce an effective surface current that radiates photons normal to the conducting surface. These photons then reflect off the internal Hershey’s Kiss-shaped reflector (Figure 2.6a) and focus onto a broadband single-photon sensor positioned above the tip of the reflector. Unlike traditional cavity haloscopes that must scan slowly in frequency, BREAD enables simultaneous sensitivity across multiple decades of axion mass, limited only by the sensor bandwidth and the reflectivity of the material. The first implementation, GigaBREAD, targets microwave frequencies around 44–52 μeV (10.7–12.5 GHz) using room-temperature Radio Frequency (RF) receivers and has already produced competitive axion

3. Proudly named by myself and Jesse Liu.




Figure 2.6: (a) shows the BREAD reflector geometry, showing the Hershey's Kiss-shaped parabolic reflector inside a metallic cylinder housed within an MRI magnet. Rays (yellow lines) are emitted from the cylindrical conversion surface and focused at the vertex by the parabolic surface. (b) shows the projected sensitivity of BREAD in the coupling vs mass plane for different photosensing technologies and NEP assumptions. These assume no backgrounds beyond the intrinsic detector noise. The upper axis corresponds to the detected photon frequency. Shaded blue regions denote existing astrophysical and haloscope exclusions. The shaded green region shows the viable QCD axion models, and the dark green line denotes the KSVZ scenario. Figures from [24].

constraints.

The full target mass range for the BREAD experiment, from roughly 1,000 to 1,000,000 μeV (1 meV to 1 eV), will ultimately require a suite of detectors using different sensor technologies to span the corresponding frequency ranges, demonstrated in Figure 2.6b. During the early development of BREAD, I contributed to establishing key experimental infrastructure for future high-frequency searches, including building a terahertz Fourier-Transform Spectroscopy (FTS) and testing facilities for SNSPDs. More details on my work with BREAD are provided in Appendix B.




Figure 2.7: Extended overview of axion parameter space showing the connection between low-mass axion searches and high-mass axion-like particle searches at colliders. The plot spans from 10^{-12} eV to 10^{12} eV, showing how different experimental approaches cover complementary mass ranges. The theoretical QCD axion band is visible at low masses, while collider searches at the LHC probe much higher masses corresponding to axion-like particles that do not solve the strong CP problem. Figure from [55].



Figure 2.8: Overview of current limits on axions in terms of axion mass and axion-photon coupling, focusing on the theoretically motivated QCD axion region. The theoretical QCD axion predictions are shown for both DFSZ and KSVZ models (yellow band labeled “QCD axion”). The plot highlights the experimental constraints from various detection methods in the mass range where axions can solve the strong CP problem. The gray shaded regions indicate parameter space excluded by various experimental constraints. Figure adapted from [55].

ALP searches at the LHC The QCD axion mass– f_a relationship (Equation 2.15) implies that QCD axions with GeV–TeV masses would have large axion–photon couplings, $g_{a\gamma\gamma} \propto f_a^{-1}$, which are already excluded by non-collider constraints. As a result, collider experiments do not directly search for QCD axions in this mass range. Instead, a broader class of pseudoscalar particles known as Axion-Like Particles (ALPs) can arise in many extensions of the SM, featuring axion-like couplings to photons without satisfying the QCD axion mass– f_a relation [16]. These ALPs do not resolve the strong CP problem but can have GeV–TeV masses with photon couplings small enough to evade existing limits yet large enough to be probed at colliders. At the LHC, ALPs may be produced directly and appear as diphoton resonance signatures, allowing collider searches to set limits on the ALP–photon coupling in this complementary region of parameter space.

Experiment	Method	Mass Range [μeV]	Status
CAST	Solar Helioscope	< 1000	Operating
ALPS-II	Light-Shining-Through-Wall	< 0.1	Commissioning
ADMX	Cavity Haloscope	1.8–10	Operating
HAYSTAC	Cavity Haloscope	15–50	Operating
MADMAX	Dielectric Haloscope	40–400	Prototype
BREAD	Conducting Surface	44–52 (GigaBREAD)	First Results
		1000–1,000,000 (design)	Development
ATLAS/CMS	Collider	10^{14} – 10^{18}	Operating

Table 2.4: Overview of major axion and ALP detection experiments, their methods, target mass ranges, and operational status. Experiments targeting masses above $\sim 10^{-2}$ eV probe ALPs rather than QCD axions that solve the strong CP problem.

Figures 2.8 and 2.7 show how different experimental approaches cover complementary regions of axion parameter space. The extended view in Figure 2.7 demonstrates the enormous mass range probed by different techniques, from ultra-light axions searched for by haloscopes to heavy ALPs accessible at colliders.

2.4 Supersymmetry

Another well-motivated theoretical extension to the SM is Supersymmetry (SUSY) [27, 28, 29, 30, 31, 32], which resolves the hierarchy problem of the Higgs mass by introducing a symmetry between bosons and fermions. In SUSY, every SM particle has a superpartner differing by half a unit of spin: fermions acquire scalar superpartners (sfermions), while bosons acquire fermionic superpartners (sbosons). The loop corrections to scalar masses from particles and their superpartners have opposite signs, such that the large quadratic divergences cancel out. Specifically, fermionic loops contribute negative corrections and bosonic loops contribute positive corrections, stabilizing the Higgs mass (and resolving the hierarchy problem [56, 57]) without fine-tuning.

2.4.1 *The Minimal Supersymmetric Standard Model*

The Minimal Supersymmetric Standard Model (MSSM) represents the minimal extension of the SM to incorporate SUSY [58, 59]. As illustrated in Figure 2.9, its particle content includes squarks and sleptons (scalar partners of quarks and leptons), gauginos (fermionic partners of the gauge bosons: gluino, wino, and bino), and an expanded Higgs sector with two scalar doublets and their fermionic partners (Higgsinos). However, the physical particles accessible in experiments include mass eigenstates known as neutralinos and charginos, which arise from mixing in the electroweak gaugino and Higgsino sectors after electroweak symmetry breaking (see Equation 2.19), analogous to the gauge boson mixing described in Section 2.1.3 for the SM.

Since no superpartners have been observed with masses close to their SM counterparts, SUSY must be a broken symmetry at the TeV scale. This scale is theoretically favored because superpartners with TeV-scale masses can still stabilize the Higgs mass while re-




Figure 2.9: Figure 2.9: An artist's interpretation of the particle content of the MSSM. SUSY doubles the number of fundamental particles: each SM fermion gains a scalar superpartner (sfermion), each gauge boson gains a fermionic superpartner (gaugino), and the Higgs sector includes two doublets with their fermionic partners (Higgsinos). This diagram illustrates the fundamental interaction eigenstates; however, the physical mass eigenstates accessible in experiments are mixtures of these states formed after electroweak symmetry breaking. In particular, the neutralino and chargino sectors arise from mixing the neutral and charged gauginos and Higgsinos, so the superpartners of the photon, W , and Z bosons are not produced as pure states. Figure by Barbara Aulicino, originally published in American Scientist [60].

maining potentially accessible at the LHC through their interactions with SM particles. Beyond their masses, the detection signatures of SUSY particles are determined by R-parity conservation [33, 58], which assigns $R = +1$ to all SM particles and $R = -1$ to their superpartners. This implies that superpartners must be produced in pairs at colliders, and the LSP is stable, providing a natural dark matter candidate [61, 62]. All heavier superpartners decay through cascades that eventually produce the LSP, with the Next-to-Lightest Supersymmetric Particle (NLSP) governing the last visible step of these decays.

In many models, the LSP and NLSP reside in the neutralino sector. Neutralinos are the four neutral mass eigenstates formed by mixing the neutral electroweak gauginos (bino $\tilde{B}$ and neutral wino $\tilde{W}^0$) with the neutral Higgsinos ($\tilde{H}_u^0$ and $\tilde{H}_d^0$) after electroweak symmetry breaking. This mixing is described by a unitary transformation that relates the interaction eigenstates to the mass eigenstates $\tilde{\chi}_{1,2,3,4}^0$, conventionally ordered by increasing mass:

$$\begin{pmatrix} \tilde{\chi}_1^0 \\ \tilde{\chi}_2^0 \\ \tilde{\chi}_3^0 \\ \tilde{\chi}_4^0 \end{pmatrix} = N \begin{pmatrix} \tilde{B} \\ \tilde{W}^0 \\ \tilde{H}_d^0 \\ \tilde{H}_u^0 \end{pmatrix}, \quad (2.19)$$

where N is a 4×4 unitary matrix. The composition of the lightest neutralino $\tilde{\chi}_1^0$ depends on the underlying SUSY parameters and determines both its dark matter properties and experimental signatures. Figure 2.9 shows the fundamental SUSY partners for illustration, but experimental searches target the physical mass eigenstates, such as the neutralinos and charginos, which are formed by mixing the electroweak gauginos and Higgsinos after electroweak symmetry breaking.

2.5 Axion + SUSY Theory: The PQMSSM

When we consider the theoretical framework at the intersection of SUSY and axion physics, a distinctive feature emerges: the axion-photon coupling is suppressed [34]. This combined framework is known as the Peccei–Quinn–Minimal Supersymmetric Standard Model (PQMSSM), which extends the MSSM by incorporating the axion solution to the strong CP problem. As discussed in Chapter 2.3.1, the axion-photon coupling arises from chiral anomaly diagrams—triangle loop diagrams involving charged fermions that mediate interactions between axions and photons (Figure 2.3).

Returning to the axion-photon coupling introduced in Equation 2.14, the critical suppression mechanism in SUSY DFSZ axion models occurs through modification of the electromagnetic chiral anomaly [34]. In non- SUSY DFSZ models, only SM fermions (quarks and leptons) contribute to these triangle loops, giving $E/N = \frac{8}{3}$ [14, 15]. However, in SUSY DFSZ models, fermionic superpartners (higgsinos) also contribute to the electromagnetic chiral anomaly. These additional higgsino contributions cause E/N to converge toward 1.92 [34].

When E/N approaches this critical value of 1.92, the term $(E/N - 1.92)$ in Equation 2.14 approaches zero, causing the axion-photon coupling $g_{a\gamma\gamma}$ to become dramatically suppressed. This can result in a coupling that lies a factor of 17 below the non- SUSY DFSZ axion coupling, and in some cases, through fine-tuning, the coupling may vanish entirely [34, 63]. This suppression renders traditional direct axion detection methods insensitive to axions in SUSY scenarios, as all of these approaches rely on axion-photon conversion.

However, the same theoretical framework also predicts the existence of the axino, the fermionic superpartner of the axion [63]. Unlike the axion, whose suppressed photon coupling may limit axion-photon-based detection strategies, the axino can be probed through its interactions with other SUSY particles and does not rely on $g_{a\gamma\gamma}$ for detectability. When

the axino serves as the LSP, heavier SUSY particles can decay to it, resulting in detectable signatures at colliders [63]. This provides a complementary approach to probe axion physics when direct detection methods are ineffective [34, 63].

2.5.1 Axino as the LSP

By combining SUSY with the axion solution to the strong CP problem in the PQMSSM framework, the axion acquires a fermionic superpartner, the axino ($\tilde{a}$). This particle becomes a fifth neutralino in the theory, extending the neutralino sector from four to five mass eigenstates. In this thesis, we consider the axino as the LSP and label it as $\tilde{\chi}_5^0$, with the lightest MSSM neutralino becoming the NLSP $\tilde{\chi}_1^0$.⁴ The mixing is now described by a 5×5 neutralino mass matrix [64]:

$$\begin{pmatrix} \tilde{\chi}_1^0 \\ \tilde{\chi}_2^0 \\ \tilde{\chi}_3^0 \\ \tilde{\chi}_4^0 \\ \tilde{\chi}_5^0 \end{pmatrix} = N_{5 \times 5} \begin{pmatrix} \tilde{B} \\ \tilde{W}^0 \\ \tilde{H}_d^0 \\ \tilde{H}_u^0 \\ \tilde{a} \end{pmatrix} \quad (2.20)$$

We also assume that the NLSP $\tilde{\chi}_1^0$ is Higgsino-like. The scenario with an axino LSP has been extensively studied in theoretical work, with different types of NLSP configurations considered, including stau NLSP [65], Higgsino NLSP [66], and neutralino NLSP [67] scenarios.

4. By convention, the $\tilde{\chi}_1^0$ state is typically reserved for the LSP. However, since we constructed this model from first principles and wanted to preserve the standard ordering of bino, wino, and Higgsino states, we place the axino as the fifth state with the axino $\tilde{\chi}_5^0$ as the actual LSP and $\tilde{\chi}_1^0$ as the NLSP.

However, these theoretical predictions have been largely unexplored experimentally.⁵ The only previous collider search targeted prompt decays in the DFSZ axino model, assuming axino kinematics similar to a bino LSP [68, 66].

2.5.2 Long-Lived NLSP

In this theory, the Higgsino-like NLSPs decay to the axino LSP through processes mediated by the h or Z boson: $\tilde{\chi}_i^0 \rightarrow \tilde{\chi}_5^0 + h$ or $\tilde{\chi}_i^0 \rightarrow \tilde{\chi}_5^0 + Z$. Since these decays are suppressed by the large PQ scale f_a , the neutralinos become long-lived with macroscopic lifetimes [64]:

$$c\tau_{\tilde{\chi}_i^0} \approx 3 \text{ mm} \cdot \left(\frac{10^{-3} \text{ eV}}{m_a} \right)^2 \sqrt{\frac{1 \text{ TeV}}{m_{\tilde{\chi}_i^0}}} \quad (2.21)$$

Theoretical predictions for neutralino lifetimes as a function of axion mass are shown in Figure 2.10 [64]. In the range $m_a = 95\text{--}1400 \mu\text{eV}$ illustrated in the plot, the resulting proper decay lengths $c\tau$ of the $\tilde{\chi}_1^0$ range from approximately 1 mm to 1 m. This corresponds to decays that would predominantly occur within the tracking volume of collider experiments like ATLAS, making the decay vertices of these long-lived particles reconstructible. These long-lived neutralinos thus provide a unique experimental signature for searching for axinos in the collider environment.

5. For the search discussed in this thesis, we collaborated with phenomenologists at the University of Chicago to develop the first Monte Carlo (MC) simulation for this specific PQMSSM axino model, as no existing event generators included the necessary physics. The model was implemented via **SARAH** and **MadGraph**, validated, and integrated into the ATLAS collaboration’s centralized physics model repository. Much of the theoretical discussion of this model is based on ongoing work for an axino phenomenology study currently being written under the working title “Bridging the divide: axion and axino phenomenology at colliders and direct-detection experiments” by Gabe Hoshino, Keisuke Harigaya, David W. Miller, Jan Offermann, Bianca Pol, Benjamin Rosser, Cecilia Toscari and myself [63].




Figure 2.10: Predicted proper decay length $c\tau$ of the neutralino $\tilde{\chi}_1^0$ as a function of its mass and the axion mass m_a in the axion SUSY theory. Each contour corresponds to a fixed axion mass in the range 95–1400 μeV , with neutralino masses up to 1500 GeV. In this parameter space, neutralino lifetimes span from approximately 1 mm to 1 m. Figure from [64].

2.6 Searching for Axinos with Collider Experiments

These long-lived Higgsino-like neutralinos, which decay into axinos, create a distinctive experimental signature for collider searches targeting axino physics. The neutralinos travel measurable distances through the detector before decaying, creating decay vertices displaced from the primary collision point, while the axino LSP escapes without interacting, carrying away momentum that appears as missing transverse energy [64].

As shown in Figure 2.11, the SM contains several LLPs that create displaced vertices in detectors. However, these SM particles typically have masses below a few GeV. A search targeting Displaced Vertices (DVs) with similar lifetimes but occurring at much higher masses would provide clear discrimination between SM backgrounds and the SUSY signal.

A collider search for this combination of high mass displaced vertices and substantial missing energy would provide a unique window into axion physics in the PQMSSM, where traditional direct detection experiments become ineffective due to the suppressed axion-photon coupling [64, 34]. In this framework, limits set on the axino can be directly translated into constraints on the axion parameter space, particularly in the theoretically motivated μeV to meV mass range where direct detection experiments face complete loss of sensitivity in the PQMSSM.

This thesis presents the first dedicated search for the axino, conducted at the LHC using the ATLAS detector.




Figure 2.11: Mass versus mean proper lifetime for SM particles. The plot shows the characteristic lifetimes of known long-lived particles and indicates the detector regions where different categories of particles typically decay. The region targeted by this analysis (displaced vertices within the tracking detector) is shown, highlighting where new long-lived particles from beyond the SM might be detected. Figure from [69], modified from an original by [70].

CHAPTER 3

THE LARGE HADRON COLLIDER AND ATLAS

Particle physics seeks to understand the fundamental constituents of matter and the forces that govern their interactions. One of its most pressing open questions is the nature of dark matter, and in this thesis, we explore the possibility that dark matter may take the form of axinos—hypothetical supersymmetric partners of axions that would appear as missing transverse energy in collider experiments. At the forefront of this exploration is the LHC, the world’s most powerful particle accelerator, operated by the European Organization for Nuclear Research, *Conseil Européen pour la Recherche Nucléaire* (CERN). This chapter introduces the infrastructure that enables high-energy proton-proton collisions and the detectors used to observe them. Chapter 3.1 gives a brief overview of CERN’s history and mission, while Chapter 3.2 outlines the LHC’s design and operating principles. Chapter 3.3 describes the ATLAS detector and its key subdetectors used for reconstructing physics objects. The gFEX trigger system, which is the focus of my hardware contributions to ATLAS, is introduced separately in Chapter 4.

3.1 CERN

Founded in 1954 during a pivotal period in global history, CERN, emerged as a symbol of international scientific collaboration. In the aftermath of World War II, Europe faced the dual challenges of rebuilding its infrastructure and fostering cooperation between nations that had recently been at war. The scientific community recognized the need for a unified effort to advance fundamental research and reclaim Europe’s leadership in physics, which had shifted significantly to the United States during the war years [71]. Against this backdrop, 12 European countries came together to establish CERN, creating an institution dedicated

to understanding the fundamental laws of nature through particle physics. Today, CERN employs approximately 2,500 staff and hosts over 12,200 users from more than 70 countries [72], highlighting its role as a hub for global scientific collaboration. This cooperation marked a significant step toward using science as a bridge for peace and progress in a divided post-war world.

Since its founding, CERN has become a world leader in advancing not only particle physics but also technology, education, and global collaboration. Among its most groundbreaking achievements is the discovery of the Higgs boson in 2012. CERN is also the birthplace of the World Wide Web, invented in 1989 by Tim Berners-Lee [73]. Beyond physics, CERN's technologies have contributed to practical applications, such as hadron therapy for cancer treatment [74] and advancements in medical imaging [75]. Hosting scientists and engineers from across the globe, CERN stands as a beacon of cooperation, uniting diverse cultures to address fundamental scientific questions.

3.2 The Large Hadron Collider (LHC)

The LHC is one of CERN's most prominent and ambitious projects. It is the world's largest and most powerful particle accelerator, designed to explore fundamental questions about the universe. First turned on in 2008, the LHC consists of a 27-km (16.8-mi) ring of superconducting magnets buried 100 m underground near Geneva, Switzerland, spanning the border between Switzerland and France. The LHC accelerates two beams of hadrons, mostly protons, to very high energies at a rate of around 40 million collisions per second. These collisions occur at four interaction points along the ring, corresponding to the four main LHC experiments: ATLAS, CMS, LHCb, and ALICE [36]. A diagram of the CERN accelerator complex and its subsystems is shown in Figure 3.1.




Figure 3.1: The CERN accelerator complex as of 2016 for LHC Run 2. Figure from [76].

The tunnels used by the LHC were originally dug for the Large Electron-Positron Collider (LEP), which ran from 1989 until 2000 and reached a center-of-mass e^+e^- collision energy of 209 GeV [77, 78]. The LHC was subsequently constructed and installed and began operation in 2008.

At the time of writing this thesis, the LHC has completed two multi-year data-taking runs and is currently in its third. Periods when the LHC is not operational are called Long Shutdowns (LS) and are used for machine and experiment upgrades. Table 3.1 lists the LHC runs and their respective center-of-mass energies.

Run	Dates	$\sqrt{s}$ (TeV)
1	2010–2013	7, 8
2	2015–2018	13
3	2022–2026	13.6
4–6	2029–2041	14

Table 3.1: LHC runs with their operational years and proton-proton collision center-of-mass energies. Runs 4–6 represent planned future operations.

3.2.1 Machine Overview

The LHC is designed to accelerate proton beams to 7 TeV per beam, corresponding to a center-of-mass energy of $\sqrt{s} = 14$ TeV. However, technical issues limited early runs to lower energies, as listed in Table 3.1. The data used in the analysis described in Chapter 6 was collected during Run 2 at $\sqrt{s} = 13$ TeV. The following discussion of the accelerator complex reflects the conditions during Run 2 data-taking.

As shown in Figure 3.1, the LHC is the final stage of the accelerator chain. Prior to 2020, protons for the LHC were produced using Linac2, which directly ionized hydrogen gas to create H^+ ions (bare protons). These protons were accelerated to 50 MeV and injected into the Proton Synchrotron Booster (PSB). This stage of beam production was modified with the introduction of Linac4. Instead of bare protons, Linac4 accelerates H^- ions (negative hydrogen ions, consisting of one proton and two electrons) to 160 MeV. At the PSB injection stage, these ions pass through a carbon stripping foil, which removes the electrons, leaving behind protons.

Once in the PSB, the protons are accelerated to 2 GeV before being injected into the Proton Synchrotron (PS), where they are further accelerated to 26 GeV. From the PS, protons are transferred to Super Proton Synchrotron (SPS) and accelerated to 450 GeV. Finally, protons are injected into the LHC, where each beam is ramped to its final energy of 6.5 TeV

per beam (Run 2) or 6.8 TeV per beam (Run 3) using a combination of RF cavities and superconducting magnets.

Accelerator	Starting Energy	Ending Energy
Linac 2	0	50 MeV
Proton Synchrotron Booster (PSB)	50 MeV	1.4 GeV
Proton Synchrotron (PS)	1.4 GeV	25 GeV
Super Proton Synchrotron (SPS)	25 GeV	450 GeV
Large Hadron Collider (LHC)	450 GeV	6.5 TeV

Table 3.2: Sequential stages of CERN’s proton accelerator complex showing the energy range for each accelerator in the chain leading to the LHC. Protons are accelerated through five stages, from initial acceleration in Linac 2 to final injection into the LHC at 450 GeV per beam.

Once the protons reach the LHC, they are injected into its twin-bore structure, where two independent magnetic channels guide the beams in opposite directions within the same circular tunnel, seen in Figure 3.2. This arrangement ensures that each beam can be independently controlled while sharing the same magnetic and accelerating infrastructure. The beams are accelerated by oscillating electric fields within the LHC’s 16 superconducting RF cavities. The RF cavities generate oscillating electric fields at a frequency of 400 MHz, which interact with the protons as they traverse the accelerator. Proper timing of these fields ensures that protons with the correct energy pass through the cavities without any change, while protons arriving slightly early or late are either accelerated or decelerated [79]. Powerful superconducting magnets are used to steer and focus the proton beams. Dipole magnets bend the beams around the 27-km ring, while quadrupole magnets ensure the beams remain tightly focused. These magnets operate at a field strength of 8.3 T and are cooled to 1.9 K (-271.3°C) using liquid helium, enabling precise control of the beams as they approach the speed of light [80].

Due to the use of RF cavities for acceleration, the LHC beams are not composed of individual protons but instead are structured into a series of dense proton “bunches”, each containing

$\mathcal{O}(10^{11})$ protons, separated by empty intervals. With bunches every 25 ns, it would take 3654 bunches to “fill” the LHC ring, but due to technical limitations, about a quarter of these slots must be kept empty. This results in a beam structure where bunches can occupy up to 2808 positions. These bunches are unevenly distributed around the detector according to a specific filling scheme, which balances the machine’s load and incorporates an “abort gap” of 3 μ s. This abort gap serves as a crucial safety feature of the LHC, allowing the dump kicker magnet sufficient time to fully ramp up and safely deflect the beams into the fixed targets at LHC Point 6, which are specifically designed to absorb the beam’s energy [81].

The two proton beams of the LHC are made to cross paths at fixed interaction points around the ring, surrounding the main experiments, shown in Figure 3.2. Magnetic fields are used to bring bunches from the two beams together in an interaction known as “bunch crossing.” Because the bunches are spaced out in 25 ns intervals, a bunch crossing can occur every 25 ns, or at a rate of 40 MHz. In these bunch crossings, multiple protons from the two bunches can interact with each other, meaning the actual number of pp collisions per bunch crossing is variable. The number of interactions per crossing, μ , is also known as pile-up, and can be partially controlled by adjusting the magnetic fields to change the angle at which the beams intersect. As μ increases, the amount of activity recorded by the detectors also increases, meaning that it becomes harder to identify which specific interaction produced a given observed particle [82].

3.2.2 The LHC Experiments

At each interaction point along the LHC ring are the four main LHC experiments: ATLAS at point 1, A Large Ion Collider Experiment (ALICE) at point 2, Compact Muon Solenoid (CMS) at point 5, and Large Hadron Collider beauty experiment (LHCb) at point 8. ATLAS and CMS are the two general-purpose detectors for studying proton-proton collisions. Both




Figure 3.2: Schematic representation of the LHC ring showing the two counter-rotating proton beams (B1 clockwise, B2 counter-clockwise) and the eight interaction regions (IR1-IR8): the four main experiments ATLAS (IR1), ALICE (IR2), CMS (IR5), and LHCb (IR8), along with specialized systems for RF acceleration (IR4), momentum collimation (IR3), betatron collimation (IR7), and beam extraction (IR6). Figure from [83].

experiments have sub-detectors to reconstruct the kinematics of the particles produced in the pp collisions. These experiments operate independently but cross-check and validate each other's results, as they both perform similar SM measurements and Beyond the Standard Model (BSM) searches. More details on ATLAS will be discussed in Chapter 3.3. The other two experiments at the LHC are more specialized

- LHCb is designed to study particles containing b -quarks (beauty quarks) to explore differences between matter and antimatter and investigate rare decays that may reveal

BSM physics. LHCb focuses on precision measurements in forward regions of particle collisions, using a “forward configuration” optimized for tracking and identifying particles over a narrow angular range with respect to the accelerator beam axis [84].

- ALICE is designed to study heavy-ion collisions rather than the proton-proton (pp) collisions primarily studied by ATLAS and CMS. Using lead ions, ALICE recreates the quark-gluon plasma that existed in the early universe, providing valuable insights into QCD through the dynamics of quark-gluon interactions [85].

3.2.3 Luminosity and Pile-Up

The amount proton-proton collisions per second is commonly referred to as the instantaneous luminosity, given as a function of the colliding beams’ properties:


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

where N_b is the number of particles per bunch, n_b the number of bunches per beam, f_{rev} the revolution frequency, γ_r the relativistic gamma factor, ϵ_n the normalized transverse beam emittance, β^* the beta function at the interaction point, and F accounts for the geometric reduction in luminosity caused by the crossing angle of the beams. At the start of an LHC fill, the instantaneous luminosity typically reaches around $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. However, as collisions occur and the number of particles per bunch diminishes, the instantaneous luminosity gradually declines. Once the beam quality drops below an optimal threshold, the operators dump the beams and refill the LHC. This process—refilling, stabilizing the beams, and beginning collisions—typically takes 1–2 hours. Each fill, which lasts several hours, is repeated throughout physics runs to accumulate the data needed for analysis. Over time, this operational cycle allows experiments to collect a substantial dataset for probing the SM

and searching for phenomena beyond it.

The total or integrated luminosity of collected data is given by integrating the instantaneous luminosity over the running time. For example, during Run 2 (2015–2018), the LHC achieved an integrated luminosity of:

$$\mathcal{L} = \int \mathcal{L}_i dt = 139 \text{ fb}^{-1}. \quad (3.2)$$



(a) Cumulative integrated luminosity delivered by the LHC and recorded by ATLAS over time for each operational year. (b) Distribution of recorded integrated luminosity as a function of mean interactions per bunch crossing for each run period.

Figure 3.3: Integrated luminosity of proton–proton collisions collected by the ATLAS detector. (a) shows the cumulative integrated luminosity delivered by the LHC across Run 1, Run 2, and Run 3. (b) shows the distribution of recorded luminosity as a function of the average number of interactions per bunch crossing (μ) for each run period. Figures from [86].

Figure 3.3 summarizes the integrated luminosity collected by the ATLAS experiment as a function of time and pile-up. 2024 marked the most successful year for ATLAS data collection to date, with the highest integrated luminosity recorded, significantly enhancing the experiment’s potential for precise measurements and new discoveries.

The vast dataset collected by ATLAS provides an essential foundation for testing the SM and exploring potential extensions to it. One effective approach is to compare the observed number of events for specific processes to the predictions made by the SM for a given final

state. Central to these predictions is the concept of the *cross-section* (σ), which quantifies the likelihood of a process occurring given its initial state. Measured in units of area, such as cm^2 , the cross-section has a geometric interpretation as an effective ‘target area’ for interactions. Due to the extremely small values of most cross-sections, the *barn* (b) is used as a practical unit, where $1 \text{ b} = 10^{-24} \text{ cm}^2$. Using the cross-section and the integrated luminosity, the expected number of events for a process can be determined as:

$$N = \sigma \mathcal{L}. \tag{3.3}$$

The cross-section for a process can be theoretically calculated and depends on the energy of the interaction, among other factors. To test these theoretical predictions, we rely on the wealth of data collected by the ATLAS experiment, which provides the statistical power needed to investigate both the SM and potential new physics.

3.3 The ATLAS Experiment

The ATLAS¹ experiment is one of the two general-purpose detectors at CERN. Located 100 m underground at LHC Interaction Point 1, colloquially called Point 1 (P1), ATLAS is designed to explore a wide range of physics, from discovering new particles to making precise measurements of SM processes.

The physical size of ATLAS is enormous, standing 25 m tall, 44 m long, and weighing approximately 7,000 tons (about 1,400 elephants). Cylindrical in shape, with a beam pipe running through the center, where proton-proton collisions occur; it is symmetric to the beam axis.

1. Originally a backronym for *A Toroidal LHC ApparatuS*, but now used as a standalone name




Figure 3.4: Cutaway diagram of the ATLAS detector, with major sub-components labeled. Figure from [87].

Like a high-tech camera, ATLAS captures detailed snapshots of particles produced in high-energy collisions. These ‘pictures’ are reconstructed using a series of layered subdetectors surrounding the interaction point, each tailored to measure specific properties of particles, such as their paths, energies, and identities. This intricate design allows us to precisely reconstruct the momenta, energies, and trajectories of particles emerging from pp interactions. Figure 3.4 shows a cutaway diagram of the detector, with the subdetectors labeled.

Working outward from the interaction point (IP), ATLAS is equipped with several key sub-systems. The Inner Detector (ID) tracks the trajectories of charged particles, the calorimeters measure the energy deposited by particles, and the muon system detects muons that pass through the calorimeters. Additionally, the Trigger and Data Acquisition (TDAQ) system monitors interesting collisions and reads out the relevant data from the detector [89]. Each of




Figure 3.5: Illustration of different types of particles passing through a transverse view of the ATLAS detector. Figure from [88].

these subsystems plays a critical role in reconstructing the momenta, energies, and identities of particles from pp collisions, and will be discussed in detail throughout this chapter.

The ATLAS detector is subdivided into two primary regions: the barrel and the endcaps. The barrel region consists of concentric cylindrical layers of detectors arranged around the beam pipe, extending along the z -axis. This section provides excellent coverage for particles produced close to the interaction point. In contrast, the endcaps are disk-shaped detector structures oriented transversely to the beamline and positioned at both ends of the barrel. These endcap detectors help capture particles produced at larger angles relative to the

beam axis. Both the barrel and endcap regions include detectors from each of the major subsystems, ensuring that ATLAS provides full, 4π coverage for all stable particles typically produced in SM processes, except neutrinos, which interact only weakly and escape without detection.


3.4 ATLAS Geometry

The ATLAS detector has a cylindrical shape, so cylindrical coordinates are commonly used to describe its layout, with the nominal interaction point serving as the origin for the coordinate system. As shown in Figure 3.6b, the z -axis is aligned with the beamline, while the transverse $R - \phi$ plane is perpendicular to it. In the $R - \phi$ plane, the x -axis points to the center of the LHC ring, and the y -axis points upward into space. The radius R defines the distance from the beamline. The detector spans 44 m along the z -axis and has a radius of $R = 12.5$ m. The origin of this coordinate system is defined at the interaction point, which divides the detector into two halves: the positive z side (often referred to as “Side A”) and the negative z side (referred to as “Side C”).² In general, ATLAS is designed to be symmetric with respect to both $\pm z$ and the azimuthal angle ϕ .


Inside the ATLAS detector, physical objects are usually described using spherical coordinates, particularly in terms of their transverse momentum p_T , which refers to the momentum in the transverse (x - y) plane. The full momentum of a particle is described as $\vec{p} = (p_T, p_z)$, but the ATLAS detector does not measure p_z .

The initial momentum of an interaction along the z -axis is unknown, which means that the outgoing system of particles can have a significant boost along the z -axis relative to the lab frame. The polar angle θ between two particles produced in a collision is not Lorentz


2. The nomenclature originates from the orientation of the detector: Side A points toward the airport, while Side C points toward Charly’s Pub in Saint-Genis.



(a) ATLAS Cartesian coordinates



(b) ATLAS cylindrical coordinates



(c) ATLAS η coordinates

Figure 3.6: ATLAS coordinate system

The ATLAS coordinate system, figures modified from ATLAS detector figures from [87]

invariant due to relativistic dilation effects, so θ is not typically used to describe particle trajectories. Instead, the quantity pseudorapidity, η , is defined as :

$$\eta \equiv -\ln \tan \frac{\theta}{2}, \quad (3.4)$$

is used to measure how far an object is “forward” relative to the $z = 0$ axis, instead of using θ directly. Figure 3.6c illustrates that particles perpendicular to the z -axis have $\eta = 0$, while those traveling parallel to the beamline have $\eta \rightarrow \pm\infty$.

The ATLAS detector measures a combination of energy and trajectory, allowing objects to be described using four-momentum, which typically includes η , ϕ , total energy, and p_T . To describe distances between objects in this $\eta - \phi$ space, the distance ΔR is commonly used, defined as:

$$\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}. \quad (3.5)$$

3.5 Magnets

Charged particles traversing a magnetic field $\vec{B}$ experience a force $\vec{F} = q \cdot \vec{v} \times \vec{B}$, where q is the electric charge and $\vec{v}$ is the particle velocity. This force causes particle trajectories to bend, with the curvature (characterized by the sagitta) inversely proportional to momentum. ATLAS employs three superconducting magnet subsystems to generate the required fields, as shown in Figure 3.7.

The central solenoid surrounds the ID and produces a 2 T field parallel to the beamline, bending charged particles in the transverse plane to measure particle p_T . The barrel and endcap toroids are located beyond the calorimeters, and each consists of eight toroidal coils of superconducting Nb-Ti wire arranged symmetrically around the beamline. The endcap toroid coils are offset by 22.5° relative to the barrel coils to provide field uniformity. The




Figure 3.7: A schematic of the ATLAS magnet system. The barrel toroids are shown in blue, the endcap toroids in red, and the central solenoid in green. Figure from [90].

barrel and endcap toroids generate central fields of 0.5 T and 1.0 T, respectively, with field lines circling in the azimuthal direction. This configuration bends muon trajectories in the $R - z$ plane, enabling the Muon Spectrometer (MS) to make precise momentum measurements. All three magnet systems operate at 4.5 K to maintain their superconducting properties. [35]

The following sections provide details on the subdetectors of ATLAS. Here, we discuss detectors designed to track or measure specific objects, such as electrons, muons, jets, and missing transverse energy. The details about what these objects are and how we define them are discussed thoroughly in Chapter 5. Table 3.3 gives a summary of the different components of the subdetectors and their physical location within the detector.

3.6 Inner Detector

The first detector subsystem that a particle leaving the IP encounters is the Inner Detector (ID). The primary function of the ID is to reconstruct the trajectories of charged particles

Sub-Detector	Sub-System	Approx. radial distance from the beam pipe (m)	$ \eta $ Range
Inner Detector (Sec. 3.6)	Pixel	0.05 - 0.12	0 - 2.5
	SCT	0.30 - 0.51	0 - 2.5
	TRT	0.55 - 1.08	0 - 2.5
Calorimeter (Sec. 3.7)	EMB	1.15 - 2.25	0 - 2.5
	EMEC	0.33 - 2.25	2.5 - 3.2
	Tile	2.28 - 4.25	0 - 1.6
	HEC	0.33 - 2.03	1.5 - 3.2
	FCAL	0.07 - 0.45	3.1 - 4.9
Muon Spectrometer (Sec. 3.8)	RPC	4.5 - 11	0 - 1.05
	MDT	4.5 - 11	0 - 2.7
	TGC	1.5 - 12	1.05 - 2.4
	CSC	0.88 - 2.8	2.0 - 2.7

Table 3.3: Summary of the ATLAS sub-detectors, with their radial and η coverage

into tracks. The ID works by detecting the ionization of electrons as charged particles interact with the different layers of the ID. Each of these interactions, or “hits” are used to reconstruct the particle’s path by matching them across multiple layers of the detector into tracks. The tracks can then be traced back to the IP, associating the particle to a specific proton-proton vertex from the Bunch Crossing (BC).

As the particles travel through a magnetic field, the Lorentz force, $\vec{F} = q \cdot \vec{v} \times \vec{B}$, causes them to follow a curved path. The curvature of this trajectory is then used to determine both the particle’s momentum and charge. The ID is surrounded by a 2 T magnetic field, which bends the particles’ paths to enable this precise momentum reconstruction.




Figure 3.8: Cutaway of the ATLAS ID

Cutaway diagram of the ATLAS inner detector. Figure from [35].

The ATLAS ID is shown in Figure 3.8 and is divided into three subsystems: the Pixel Detector, the Semiconductor Tracker (SCT), and the Transition Radiation Tracker (TRT). The innermost subdetectors, Pixel and SCT, use layers of silicon micro-strips in various configurations to detect charged particles. When particles pass through layers of silicon, ionization causes the formation of electron-hole pairs in the semiconductor, generating an electrical signal that is digitized and captured as a hit. Both these subdetectors provide discrete spatial measurements at very high resolution. The TRT uses the ionization produced by charged particles traversing a gas to measure the trajectories of tracks.

The pixel detector is the closest part of the ID to the IP, extending from radius $R = 33 \text{ mm} - 149.6 \text{ mm}$ from the beamline, and covers $|\eta| < 2.5$. The barrel consists of four

concentric layers surrounding the beampipe, covering a length of $z \pm 400.5$ mm. These layers include three from the original detector design, and an additional Insertable B-Layer (IBL) that was added later to improve tracking and vertexing performance by reducing the distance between the first layer of the detector and IP [91]. Each endcap consists of three disks of silicon pixels oriented transversely to the beamline. The pixels are $250\text{ }\mu\text{m}$ thick, with an area of $50 \times 400\text{ }\mu\text{m}$ in the original three layers and endcaps, compared to the smaller $50 \times 250\text{ }\mu\text{m}$ pixels used in the IBL. Across all the barrel and endcap layers, there are approximately 92 million of these pixels [35].

The next concentric layer of the ID is the SCT, which consists of silicon micro-strips [92]. Similar to the Pixel Detector, the SCT uses silicon sensors, but they are segmented into strips rather than pixels. This design results in lower position resolution but allows coverage of a larger surface area, as shown in Figure 3.9. The SCT comprises four cylindrical barrel layers spanning radii from 299 mm to 514 mm, along with nine endcap disks covering $|z|$ values from 854 mm to 2721 mm, totaling over 4,000 modules. Each SCT module consists of two layers of silicon strips glued back-to-back. A single layer of strips can only measure one coordinate, but the inclusion of a second layer, rotated slightly (40 mrad) relative to the first, enables the measurement of a second coordinate for each hit. The dotted lines in Figure 3.9 illustrate example tracks for different values of η , demonstrating that all charged particles with $|\eta| < 2.5$ will nominally produce at least eight hits across a minimum of four SCT layers.




Figure 3.9: Cross section of the ATLAS ID

Layout of one quarter of the Inner Detector, in cylindrical coordinates, showing the locations of the Pixel Detector, the SCT, and the TRT layers. Figure from [93].

The TRT is the outermost component of the ID and uses the ionization produced by charged particles traveling through gas to measure track trajectories [94]. The TRT consists of over 350,000 drift tubes, each 4 mm in diameter, distributed across the barrel and endcap regions. In the barrel region, the TRT is composed of 144 cm long straw tubes that run parallel to the beamline, spanning detector radii from 563 mm to 1,066 mm. In the endcaps, the 37 cm long tubes are radially oriented in wheel-like structures. The geometry of the TRT can be seen in Figure 3.9. These drift tubes are filled with a gas mixture of xenon, carbon dioxide, and oxygen, which can be ionized when traversed by charged particles. A tungsten wire,

acting as an anode, is operated at -1530 V. When a charged particle passes through a tube, it ionizes the gas, and the resulting free electrons drift toward the anode under the influence of the electric field, where they are collected by readout electronics.

The TRT derives its name from the detection of transition radiation, which is produced at the polypropylene fibers and foils located between the straw tubes. When a charged relativistic particle crosses the boundary between two materials with different dielectric constants, it emits transition radiation. The amount of radiation depends on the particle's Lorentz factor, given by $\gamma = E/mc^2$. This means that particles with the same energy but different masses will leave different amounts of transition radiation in the TRT, allowing for particle identification, which is used to differentiate between particles with the same charge, such as electrons and pions.




Figure 3.10: Cross section diagram of the barrel region of the ID. Figure from [95].

A diagram showing a radial cross section of the entire ATLAS ID is shown in Figure 3.10.

3.7 Calorimeters

The purpose of the ATLAS calorimeters is to absorb and measure the energy of electrons, photons, and hadrons emitted from the collision point. Located outside the ID and the central solenoid magnet, the calorimeter system consists of two main components: the (mostly) electromagnetic Liquid Argon (LAr) calorimeter system [96] and the hadronic Tile calorime-

ter [97], shown in Figure 3.11.

Both calorimeters operate as sampling calorimeters, containing alternating layers of absorbing and active media. This design represents a careful compromise: the dense absorber layers excel at initiating particle showers but cannot measure them, while the active layers can measure deposited energy but are inefficient at initiating showers. By interleaving these materials, the calorimeters can both initiate showers in the absorber and measure them in the active material. One trade-off of this sampling design, compared to homogeneous calorimeters that use only active material, is that some shower energy deposited in the absorber layers cannot be directly measured. Consequently, the total shower energy must be estimated from the energy measured in the active layers.

The calorimeter system is arranged in layers of increasing interaction specificity. The Electromagnetic Calorimeter (ECAL) sits closest to the collision point, optimized to measure the energy of electrons and photons through electromagnetic interactions. Hadronic particles typically pass through the ECAL with minimal energy loss and are measured by the surrounding Hadronic Calorimeter (HCAL). Hadronic showers are more complex than electromagnetic ones, involving strong force interactions that produce cascades of secondary particles, including pions. These showers are characterized by the nuclear interaction length (λ_I), which is substantially larger than the radiation length (X_0) governing electromagnetic showers, requiring hadronic calorimeters to be much deeper. Both calorimeter systems incorporate fast electronics to record the energy deposited in each layer, enabling rapid signal processing crucial for trigger decisions, as described in Chapter 3.9.

Component	$ \eta $ Coverage	Technology	Depth	Granularity
Electromagnetic Calorimeter				
EMB (Barrel)	$ \eta < 1.475$	Lead/LAr	$> 22X_0$	Down to 0.025×0.025
EMEC (End-cap)	$1.375 < \eta < 3.2$	Lead/LAr	$> 24X_0$	Varies by layer
Hadronic Calorimeter				
Tile (Barrel)	$ \eta < 1.0$	Steel/Scint.	$7.4\lambda_I$	0.1×0.1 (layers 1,2)
Tile (Extended)	$0.8 < \eta < 1.7$	Steel/Scint.	$7.4\lambda_I$	0.1×0.2 (layer 3)
HEC (End-cap)	$1.5 < \eta < 3.2$	Copper/LAr	$\sim 10\lambda_I$	Varies
FCal	$3.1 < \eta < 4.9$	Cu & W/LAr	$\sim 10\lambda_I$	Coarser (~ 0.2)

Table 3.4: Summary of the ATLAS calorimeter system.




Figure 3.11: A cutaway diagram of the ATLAS calorimeter system. Figure from [98]

3.7.1 Electromagnetic Calorimeter

The ECAL uses LAr as the active medium, cooled to -184°C , lead as an absorbing medium, and Kapton electrodes to measure the size of electromagnetic showers. ECAL are fundamental devices in particle physics that measure the energy of particles through their electromagnetic interactions with matter.

When electrons interact with material, they produce electromagnetic showers primarily through *bremsstrahlung*, a process where photons are emitted as charged particles accelerate through matter. As electrons traverse the dense detector material, they scatter off atomic nuclei, resulting in this radiation emission. For photons, the primary interaction mechanism is pair production, where the photon interacts with an atomic nucleus to produce an electron-positron pair. This process occurs when the photon energy exceeds the pair production threshold of $E_{\gamma} > 2m_e$. These interaction processes create particle cascades with progressively decreasing energy, forming electromagnetic showers.

The ECAL consists of a central barrel region called the Electromagnetic Barrel (EMB), covering $|\eta| < 1.475$, and two Electromagnetic Endcaps (EMECs) from $1.375 < |\eta| < 3.2$. Both sections employ an innovative accordion geometry, where the absorber and active material layers are folded into an accordion shape. This design serves multiple crucial purposes: it ensures complete symmetry in ϕ across the detector volume and eliminates dead zones or “cracks” that could compromise measurement accuracy.

3.7.1.1 Electromagnetic Barrel Calorimeter

The EMB forms the central component of the ECAL, covering the pseudorapidity range $|\eta| < 1.475$. The EMB is structured as two half-barrels, with a small gap of approximately 4 mm at $z = 0$ separating them. Each half-barrel weighs about 57 tons and has a length of

3.2 m with inner and outer diameters of 2.8 m and 4 m respectively. The EMB is composed of 1024 accordion-shaped lead absorbers interleaved with electrodes, with the electrodes positioned in the center of the LAr gaps by honeycomb spacers. The size of the drift gap on each side of the electrode is 2.1 mm, which corresponds to a total drift time of approximately 450 ns at an operating voltage of 2000 V.

The accordion geometry of the EMB was chosen to provide complete ϕ symmetry without any azimuthal cracks. As shown in Figure 3.12, the EMB is segmented into three primary sampling layers in depth:

- The first layer (front) consists of narrow strips with very fine granularity in η (as fine as $\Delta\eta = 0.0031$ in the central region) but coarser in ϕ ($\Delta\phi = 0.1$). This fine segmentation enables precise η position measurements, helping to discriminate between single photons and pairs of photons from π^0 decays, and complements the ID tracking.
- The second layer (middle) is the thickest section, collecting the majority of the energy from the electromagnetic shower. It has a granularity of $\Delta\eta \times \Delta\phi = 0.025 \times 0.025$.
- The third layer (back) has a coarser granularity of $\Delta\eta \times \Delta\phi = 0.05 \times 0.025$ and is used to collect the tail of the electromagnetic shower and help discriminate between electromagnetic and hadronic showers.

Additionally, for $|\eta| < 1.8$, the EMB is complemented by a presampler detector consisting of a thin active LAr layer (11 mm in the barrel region) placed in front of the calorimeter. The presampler helps correct for energy lost by electrons and photons before they reach the calorimeter.

3.7.1.2 Electromagnetic Endcap Calorimeter (EMEC)

The EMEC extends the electromagnetic coverage beyond the barrel region to $1.375 < |\eta| < 3.2$. Each EMEC consists of two coaxial wheels: an outer wheel covering $1.375 < |\eta| < 2.5$ and an inner wheel covering $2.5 < |\eta| < 3.2$. Like the EMB, the EMEC employs an accordion geometry, but with the accordion waves running radially and the fold angles increasing with radius to maintain a nearly constant LAr gap.

The EMEC structure includes:

- In the region $1.5 < |\eta| < 2.5$, the EMEC is segmented into three longitudinal layers with similar functions as in the EMB. The granularity progressively becomes coarser with increasing η , with the middle section having $\Delta\eta \times \Delta\phi = 0.025 \times 0.025$ at lower η values, degrading to $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ at higher η values.
- In the region $2.5 < |\eta| < 3.2$, the EMEC has a coarser granularity and is segmented into only two longitudinal layers.
- Similar to the barrel region, the endcap regions are equipped with presampler detectors for $1.5 < |\eta| < 1.8$ to account for material effects.

The transition region between the barrel and endcap calorimeters ($1.37 < |\eta| < 1.52$) contains a significant amount of material in front of the first active calorimeter layer. This region, often referred to as the “crack region,” has poorer performance and is typically excluded from analyses requiring precise electromagnetic energy measurements.

Both the EMB and EMEC feature an electronics calibration system that injects calibrated current pulses into the readout chain. Through a multi-stage calibration procedure involving both in-situ and test beam measurements, the ATLAS electromagnetic calorimeter achieves

an energy resolution of $\sigma_E/E \approx 10\%/\sqrt{E} \oplus 0.7\%$ in the central region, where E is in GeV. This excellent energy resolution is crucial for precision measurements involving electrons and photons, such as Higgs boson studies in the $\gamma\gamma$ and four-lepton final states.




Figure 3.12: Diagram of a LAr barrel module, showing the different layers, thickness, and physical resolution. The diagram shows how calorimeter cells are grouped into larger granularity Trigger Towers (with granularity $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$), which serve as the basic input units for the Level-1 Calorimeter Trigger (L1Calo) trigger system. The gFEX system, discussed in Chapter 4, uses an even coarser granularity of gTowers (0.2×0.2) for global trigger calculations. Figure from [35].

As shown in Figure 3.12, the distance inside an electromagnetic calorimeter can be measured in radiation lengths (X_0). A radiation length is the average distance over which an electron's

energy is reduced by a factor of $\frac{1}{e}$, and is dependent on the atomic radius and number density of the nucleus. For the lead absorber used in the detector, the radiation length $X_0(\text{Pb}) = 0.56 \text{ cm}$. The barrel modules extend to a minimum depth of $22X_0$, while the endcap modules reach $24X_0$, in order to contain the majority of electromagnetic showers within the detector volume. Figure 3.12 shows an example diagram of part of the ECAL barrel region, where each of the three distinct layers can be seen.

3.7.2 *Hadronic Calorimeter System*

Complementing the ECAL described above, the ATLAS HCAL system captures and measures energy from hadrons that typically pass through the ECAL with minimal energy loss. The hadronic system employs two distinct technologies optimized for different detector regions: the Tile Calorimeter, which uses steel and scintillator tiles in the central region, and extensions of the LAr system with different absorber materials in the endcap and forward regions. [50]

3.7.2.1 **Tile Calorimeter**

The central component of the HCAL system is the Tile calorimeter, which sits outside the EMB and EMEC calorimeters. It consists of a central barrel covering $|\eta| < 1.0$ and two extended barrels covering $0.8 < |\eta| < 1.7$. Spanning from an inner radius of 2.28 m to an outer radius of 4.25 m, it occupies the space between the ECAL and the muon spectrometer.

Each barrel comprises 64 wedge-shaped modules arranged azimuthally around the beam axis, as shown in Figure 3.11. These modules implement a sampling technique with alternating layers of steel plates (4-5 mm thick) and plastic scintillator tiles (approximately 3 mm thick). The steel serves dual purposes: absorbing and fragmenting incoming hadrons while also

functioning as the return yoke for the solenoid magnetic field.

The Tile calorimeter provides three measurement layers with varying granularity. The inner two layers feature a granularity of $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$, while the outer layer has a coarser granularity of $\Delta\eta \times \Delta\phi = 0.1 \times 0.2$. The total depth ranges from 8 to 12 interaction lengths, depending on the pseudorapidity, ensuring that hadronic showers are largely contained within the detector volume.

3.7.2.2 LAr Hadronic Endcap Calorimeter

In the endcap regions where radiation levels exceed what scintillator-based technology can withstand, ATLAS employs the LAr Hadronic End-Cap Calorimeter (HEC) covering $1.5 < |\eta| < 3.2$. The HEC utilizes the same LAr technology as the electromagnetic calorimeter but substitutes copper for lead as the absorber material, optimizing the calorimeter for hadronic shower measurements.

The HEC consists of two independent wheels on each endcap, with each wheel formed from 32 individual wedge-shaped modules. Each HEC has approximately 40 LAr gaps divided between the two wheels, providing a total depth of roughly 10 interaction lengths. Unlike the accordion geometry of the ECAL, the HEC uses a more conventional parallel-plate design suitable for the endcap region.

When charged particles traverse the LAr, they create electron-ion pairs. The applied electric field causes these electrons to drift toward the copper electrodes, inducing electrical signals proportional to the deposited energy. The entire LAr system operates at approximately -184°C to maintain the argon in its liquid state, sharing the cryogenic system with the electromagnetic calorimeter. [96]

3.7.2.3 Forward Calorimeter

The Forward Calorimeter (FCal) extends calorimeter coverage to the high-pseudorapidity region $3.1 < |\eta| < 4.9$, closest to the beam pipe where particle flux reaches extreme levels. Positioned approximately 4.7 m from the interaction point, the FCal employs specialized geometry with very narrow LAr gaps around electrode rods centered in tubes oriented parallel to the beam axis, enabling faster signal collection in this high-rate environment.

The FCal consists of three sequential modules. The first module, FCal1, functions primarily as an electromagnetic calorimeter using copper as the absorber. The second and third modules, FCal2 and FCal3, are optimized for hadronic measurements and use sintered tungsten as the absorber material [35]. Tungsten's high density and short interaction length provide excellent stopping power in the limited space available in the forward region. Together, these three modules provide approximately 10 interaction lengths to contain forward-going hadrons. This design ensures that hadrons across nearly the full solid angle are measured, crucial for accurately reconstructing jets, measuring MET, and identifying hadronic tau decays throughout the detector volume.

3.8 Muon Spectrometer

Most charged particles fully deposit their energy in either the ECAL or HCAL, but there are exceptions. Most notably, muons interact electromagnetically but are about 200 times more massive than electrons ($m_\mu = 106$ MeV, $M_e = 0.510$ MeV), making them less affected by bremsstrahlung. This means they lose only about 3 GeV of energy as they pass through the calorimeters, earning them the title of minimum-ionizing particles. The Muon Spectrometer (MS) [99], the outermost component of ATLAS, is therefore intended to measure the momenta of muons in the region $|\eta| < 2.7$.




Figure 3.13: The ATLAS muon system. Figure from [35].

Like the ID, the MS is immersed in a magnetic field to cause the muons to curve. A barrel toroid and two endcap toroids generate fields of approximately 0.5 T and 1 T in each region, respectively [35]. The system consists of several layers of drift-tube and wire chambers to detect the passage of muons; tracks are reconstructed by combining hit information from the different layers.

The muon system consists of four types of modules, which can be divided into two subsystems: two types of “precision-tracking” chambers and two types of “trigger” chambers. All four of these chambers use similar technology to the TRT: charged particles cause gas to ionize, and the ionized electrons are then collected and measured to record a hit.

The precision tracking system consists of Monitored Drift Tube (MDT) [100] and Cathode-Strip-Chamber (CSC) [101]. Each MDT contains a tungsten-rhenium wire and is filled with a mixture of argon and carbon dioxide gas. A potential difference of 3,080 V causes electrons to drift when the gas ionizes. These MDTs are organized into three cylindrical barrel layers and three endcap wheels, as shown in Figure 3.13. Collectively, these layers provide tracking coverage out to $|\eta| < 2.7$. However, in the innermost endcap layer, the $2.0 < |\eta| < 2.7$ region sees high rates of activity beyond the capabilities of the MDTs; therefore, CSCs are used in this area instead. CSCs are multi-wire proportional chambers: a single CSC module is divided into cathode strips, with multiple wires running across multiple strips. The cathode strips are read out instead of the anode wires, providing higher resolution than is available with the MDTs.

The other two types of modules, Resistive Plate Chamber (RPC) [102] and Thin Gap Chamber (TGC) [103] form the “trigger” system. RPCs can be found in the barrel, surrounding the MDT layers, and cover the $|\eta| < 1.05$ region, while TGCs are located around the second wheel in each endcap and cover the $1.05 < |\eta| < 2.4$ region. RPCs use a plastic laminate electrode instead of a wire, while TGCs (like CSCs) are multi-wire proportional chambers. As the name suggests, these trigger modules provide muon information to the ATLAS trigger system, indicating whether events should be read out or not. In addition to triggering, data from these modules is also used to complement the information from the precision tracking system when measuring muons.

The ATLAS muon system performance benchmark is a resolution within 10% for a muon with $p_T = 1$ TeV. All together, the muon spectrometer covers an η range of $|\eta| < 2.4$ for the trigger chambers and $|\eta| < 2.7$ for the tracking chambers.

3.9 Trigger and Data Acquisition System

The ATLAS Trigger and Data Acquisition (TDAQ) system serves as the critical decision-making infrastructure that determines which collision events to record for physics analysis. This system addresses a fundamental challenge, the LHC produces proton-proton collisions at approximately 40 MHz, with each event generating about 1 MB of raw data. Without filtering, this would result in roughly 40-50 terabytes of data per second, which would be impossible to process and store with current technology. Moreover, most collisions don't contain interesting physics processes.




Figure 3.14: Cartoon of the trigger rate along the TDAQ path in ATLAS. Image Credit: Maurizio Pierini.

To handle this continuous data flow without missing events, the trigger system must operate at fixed latency. Since new BCs occur every 25 ns, any delay in trigger decisions would cause the system to miss subsequent BCs while still processing earlier ones, creating periods of “dead time” where no data can be recorded. The TDAQ system must stay synchronized with the LHC clock, where each 40 MHz clock cycle corresponds to one potential BC slot. This provides a crucial timing reference that allows all detector and trigger components to correlate their measurements to the same BC, ensuring that particle reconstructions from

different subsystems can be appropriately associated with the collision that produced them. With detector data flowing through processor modules with predetermined timing relative to this clock, the system can process each BC within a fixed time window. When trigger conditions are met, a Level-1 Accept (L1A) signal is issued, instructing all detector modules to read out data for the accepted BC.

The ATLAS TDAQ system employs two stages of event filtering. The first, the Level-1 Trigger (L1) trigger, uses custom hardware to rapidly process calorimeter and MS data, while excluding computationally expensive ID tracking information. Operating within a $2.5\ \mu\text{s}$ latency budget, the L1 trigger system reduces the event rate from 40 MHz to approximately 100 kHz. The L1 trigger comprises three main components: the L1Calo, the Level-1 Muon Trigger (L1Muon), and the Level-1 Central Trigger Processor (L1CTP).

The L1 trigger operates as follows: The L1Calo and L1Muon systems process their respective subdetector inputs and send Region of Interests (ROIs) or Trigger Objects (TOBs) to the L1Topological Trigger and Muon-to-Central Trigger Processor Interface (MuCTPI). These components forward the information to the L1CTP, which makes the final trigger decision and generates a L1A signal at a maximum rate of 100 kHz. The L1A signal prompts all detector components to read out data for the accepted event through the readout system.

The second stage, the High-Level Trigger (HLT), is a software-based system running on commercial computing farms with access to complete detector information, including full granularity and ID tracking data. The HLT focuses on ROIs identified by the L1 trigger and employs sophisticated reconstruction algorithms approaching offline analysis quality. It further reduces the event rate to approximately 1-2 kHz for permanent storage at the CERN Tier-0 computing facility.

3.9.1 Run 2 Trigger System

The ATLAS trigger system underwent significant changes between Run 2 and Run 3 to handle increased luminosity and pile-up conditions. We first discuss the Run 2 implementation before describing the Run 3 upgrades.




Figure 3.15: Overview of the Run 2 Trigger and Data Acquisition System for the ATLAS detector, showing the Level-1 Trigger, the HLT, and the data readout path. Figure from [104].

Figure 3.15 shows the overall Run 2 TDAQ system architecture. The Run 2 L1Calo system consisted of three primary components. The Pre-Processor Module (PPM) digitized and calibrated signals from the calorimeters through a pedestal subtraction scheme. The new

Multi-Chip Module (nMCM) on the PPM processed and summed data into Trigger Towers with a granularity of 0.1×0.1 in $\Delta\eta \times \Delta\phi$. Trigger Towers represent the fundamental granularity unit of the L1Calo system, where multiple calorimeter cells are summed together to create manageable data rates for the trigger algorithms, as illustrated in Figure 3.12. Each Trigger Tower contains the combined EM and hadronic energy deposits from its constituent cells, providing the building blocks for all L1Calo trigger decisions.

The Cluster Processor (CP) used these Trigger Towers to identify electron, photon, and tau candidates through a 4×4 sliding window algorithm that searched for energy clusters meeting specific isolation and shower shape criteria, shown in Figure 3.17a. Meanwhile, the Jet Energy Processor (JEP) identified jets and calculated MET by summing 2×2 Trigger Towers to create jet elements (0.2×0.2 in $\Delta\eta \times \Delta\phi$) and forming jet clusters of various sizes.

A significant enhancement introduced during Run 2 was the Level-1 Topological Trigger (L1Topo), which applied kinematic and angular requirements to trigger objects from both L1Calo and the L1Muon system, thereby improving background rejection and signal acceptance.

The Run 2 HLT ran on a Central Processing Unit (CPU) farm of approximately 28,000 processing units. It employed a sequence of reconstruction algorithms organized in trigger chains, with each chain designed to select specific physics signatures. Events passing the HLT were recorded at a rate of approximately 1 kHz, with an average event processing time of around 4 seconds.

3.9.2 Run 3 Trigger System Upgrades

For Run 3 (2022-present), the LHC has increased its center-of-mass energy from 13 TeV to 13.6 TeV (Table 3.1) and approximately doubled the number of proton bunches, resulting in an expected average of around 60 interactions per BC compared to 25 in Run 2 (Figure 3.3b). To handle these more challenging conditions, the ATLAS TDAQ system underwent significant upgrades, resulting in the system shown in Figure 3.16.




Figure 3.16: Overview of the Run 3 Trigger and Data Acquisition System for the ATLAS detector, showing the major upgrades including the new Feature Extractor modules. Figure from [105].

The L1Calo trigger received one of the most comprehensive upgrades for Run 3. The L1Calo system is designed to identify high-energy electrons, photons, taus, jets, and MET (E_T^{miss}) at level 1 using calorimeter data. Upgrades to the calorimeter systems were primarily motivated by the desire to provide higher granularity information to L1Calo. The movement of the

digitization process earlier in the trigger path also reduces the expected effects of noise and radiation damage, particularly for the High Luminosity–LHC (HL-LHC).

The Run 2 LAr- L1Calo interface was based on the concept of Trigger Towers, which summed energy deposition in an area of $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ over several stages of analog electronics located on-detector. For Run 3, these Trigger Towers were replaced by a new granularity scheme based on supercells, which provide layer-by-layer information with finer segmentation from the LAr EMB and EMEC, as shown in Figure 3.17. The analog signals have been replaced by digital signals transmitted over optical fibers.

The gFEX system, described in detail in Chapter 4, takes a different approach by using an even coarser granularity of 0.2×0.2 (gTowers), enabling global processing of the entire calorimeter on a single processing board.




Figure 3.17: Schematic of the granularity scheme used for LAr-L1Calo interface. Changes from Run 2, using Trigger Towers to using Supercells in Run 3. Figure from [106].

To provide high granularity and high precision digital data to L1Calo, new LAr Trigger Digitizer Boards (LTDBs) were installed along with the LAr front-end electronics. These LTDBs digitize finer granularity information up to $\Delta\eta \times \Delta\phi = 0.025 \times 0.1$ in the first two layers of the ECAL, and also recreate the legacy trigger towers during the transition period.

The digitized signals are processed by the LAr Digital Processing System (LDPS), which converts them to calibrated energies for L1Calo. The LDPS consists of the LAr Trigger Object Mezzanine (LATOME) and the LAr ATCA carrier blade (LArC), with the LATOME firmware handling the bulk of the summing and energy conversion. These LATOMEs serve as the direct interface to the *Feature Extractors* (FEXs) of L1Calo.

For the Tile calorimeter, the TileCal signal needed to be digitized for the new digital L1Calo system, which was incorporated into the Tile Rear Extension Module (TREM). This upgrade to the LAr and Tile Calorimeters enables more sophisticated trigger algorithms and improved particle identification in the high-pile-up environment of Run 3. Modern Field Programmable Gate Array (FPGA) technology with substantially greater processing capabilities now powers the trigger, and the CP and JEP modules have been entirely replaced by new FEX modules.

The Run 3 L1Calo system incorporates three complementary FEX modules. The Electromagnetic Feature Extractor (eFEX) consists of 24 modules that process data from approximately 34,000 supercells to identify isolated electromagnetic clusters characteristic of electrons, photons, and tau leptons. The eFEX replaces the CP and takes advantage of the higher granularity to improve trigger efficiency in high pile-up conditions. Each module contains four Xilinx Virtex Ultrascale+ FPGAs operating within the $2.2 \mu\text{s}$ latency budget.

The Jet Feature Extractor (jFEX) comprises 6 modules that identify small to medium radius jets. The jFEX replaces the JEP and uses a sliding window of size 0.9×0.9 to identify local energy maxima, allowing for improved identification of larger radius jets compared to Run 2.

The gFEX is an entirely new system implemented on a single module equipped with three processing FPGAs plus one System-on-Chip (SoC) FPGA. Unlike the region-based approaches of eFEX and jFEX, the gFEX processes data from the entire calorimeter simultaneously, providing a global view of each collision event. This capability is crucial for identifying large-radius jets ($R > 1.0$) characteristic of boosted objects, such as W/Z bosons

or top quarks, where decay products become collimated and appear as a single, large jet rather than separate objects. The gFEX operates at a reduced granularity of 0.2×0.2 , reading data from all ATLAS calorimeters on a single board.

The data transmission infrastructure has been completely overhauled, transitioning from primarily electrical to optical transmission throughout the trigger chain. New components include the TREX that digitizes the Tile calorimeter data, the Fiber Optic Exchange (FOX) that routes data from the calorimeters to the appropriate FEX modules, the Front End Link Exchange (FELIX) that receives optical output from the FEX systems, and the Software Readout Driver (SWROD) that formats data for forwarding to both the HLT and the CERN Tier-0 data storage facility.

The HLT has also been enhanced for Run 3, with an expanded farm of approximately 40,000 processing cores. The HLT benefits directly from the improved granularity and algorithmic capabilities of the FEX systems, which provide better-defined ROIs with more precise energy and position information than was available in Run 2. This allows the HLT to focus its processing more effectively, resulting in improved trigger efficiency for physics processes of interest.

The final output rate of the HLT in Run 3 has increased to approximately 1.5-2 kHz, up from the previous rate of 1-1.2 kHz, accommodating the broader physics program and higher luminosity while remaining within the constraints of the data acquisition and storage systems.

During the commissioning phase, the legacy system and the upgrade system ran in parallel to ensure a smooth transition. As the new FEX modules were validated, they gradually took over the trigger functions, with the eFEX first handling electron, photon, and tau identification, followed by the jFEX and gFEX systems. As of the end of running in 2024, all legacy items were removed from the trigger menu, making the Phase-I upgraded system the current trigger system for ATLAS.

CHAPTER 4

GFEX - THE GLOBAL FEATURE EXTRACTOR

The Global Feature Extractor (gFEX) is a new component of the ATLAS L1Calo system introduced in Run 3 to improve trigger performance by providing a global view of the calorimeter. It operates on a single Advanced Telecommunications Computing Architecture (ATCA) board, covering the full $\eta - \phi$ range and computes quantities like MET and total energy, and identifies objects like small-radius and large-radius jets. Jets are collimated sprays of particles resulting from the hadronization of quarks and gluons, which appear in the detector as clustered energy deposits and associated tracks (detailed jet reconstruction is described in Chapter 5.2.1). MET represents the momentum imbalance in the transverse plane, indicating the presence of particles that escape detection, such as neutrinos or potential new physics particles (a comprehensive discussion of MET reconstruction is provided in Section 5.3.1).

One of the primary motivations for improving the L1 MET trigger is to detect new physics scenarios involving invisible particles, such as neutrinos, axinos, or other dark matter candidates, that escape the detector and manifest as missing energy. In particular, MET is the primary experimental signature of axinos, the SUSY partners of axions, which leave no direct signal in the detector. Improved MET algorithms with pile-up suppression help maintain high trigger efficiency despite luminosity.

gFEX can also trigger on both large and small-radius jets, providing insight into jet substructure. This helps identify collimated decay products, such as boosted hadronic states from Higgs bosons or top quarks, which were harder to resolve in previous triggers that focused on smaller regions. Figure 4.1 shows how boosted jets become large-radius jets in the detector, while Figure 4.2 demonstrates the significant advantage of a gFEX large-radius jet trigger over the legacy trigger. In the simulated $Z \rightarrow t\bar{t}$ event shown, the legacy trigger

system (red circle) captures only a small fraction of the energy deposits from the boosted top quark decay, missing the W boson decay products, additional b -jet activity, and top radiation. In contrast, a large-radius jet trigger would encompass the entire decay signature within its larger radius, dramatically improving trigger efficiency for these topologies.




Figure 4.1: Example of large-radius jets as boosted objects. Figure from [107].

A key advantage of gFEX is its ability to analyze the entire calorimeter simultaneously, enabling more advanced algorithms to distinguish signal from background. This is increasingly important as LHC luminosity continues to rise in Run 3 and beyond.

4.1 Hardware Design

The gFEX is implemented on a single ATCA board, unlike other L1Calo systems (e.g., eFEX, jFEX) that require multiple boards to cover different detector regions. Consolidating processing on a single board eliminates inter-board communication latency and simplifies the system architecture. Figure 4.3 shows the board and its main components.

The board includes three Xilinx Virtex Ultrascale+ VU9P Processor FPGAs (pFPGAs),




Figure 4.2: Illustration of how large-radius jets improve trigger efficiency for boosted objects. In this simulated $Z' \rightarrow t\bar{t}$ event ($m_{Z'} = 1.75$ TeV), the top quarks decay to W bosons, b -jets, and additional radiation. The legacy small- R online trigger (red circle) captures only a fraction of the energy deposits, while a gFEX large- R jet trigger (black circle) would encompass the full decay signature, significantly improving trigger efficiency for boosted topologies. Figure courtesy of David W. Miller.




Figure 4.3: The gFEX board with main components - the processor pFPGAs, and Zynq Ultrascale+ system on chip - labeled.

labeled A, B, and C. pFPGAs A and B process central calorimeter data ($|\eta| < 2.5$), while C handles the forward region ($2.5 < |\eta| < 4.9$). These devices were selected for their high transceiver bandwidth, large on-chip memory, and extensive digital signal processing (DSP) resources—necessary for implementing complex trigger algorithms within the L1 latency constraint of about $2.5 \mu\text{s}$ from collision to decision [108].

Figure 4.4 shows the division of η - ϕ space among the three pFPGAs. This structure supports parallel processing of separate regions while enabling cross-boundary inter-FPGA communication when required, such as for large-radius jet reconstruction.

In addition to the pFPGAs, the board features a Xilinx Zynq Ultrascale+ Multi-Processor System-on-Chip (MPSoC), which integrates quad-core Arm Cortex-A53 processors, a Cortex-R5F, Mali GPU, and programmable logic. The Zynq runs a custom Linux Operating System (OS) used for system control, monitoring, and slow control tasks that are not latency-sensitive.

The board receives calorimeter data via optical fibers bundled by the FOX and connected through 35 miniPOD transceivers—26 for input and 11 for output. Each pFPGA receives




Figure 4.4: Division of the η - ϕ space across the three pFPGAs. The η space is divided into different bins denoted on the top x-axis as FPGA-A1, FPGA-A2, FPGA-B1, FPGA-B2, FPGA-C N, and FPGA-C P, where the different FPGAs A, B, and C reference different physical FPGAs on the board and are each broken into two halves. Each grid cell represents a gTower with granularity of 0.2×0.2 in $\Delta\eta \times \Delta\phi$. The jet engines within each pFPGA cover 6 bins in η and 1 bin in ϕ , with pFPGA-C primarily providing information sharing to pFPGAs A and B for inter-FPGA communication in large radius jet calculations. Figure from [109].

approximately 100 input channels, each operating at 11.2 Gbps, resulting in a total system input bandwidth exceeding 3 Tbps [105]. The miniPODs convert these optical signals to electrical and route them to the Multi-Gigabit Transceivers (MGTs) on the pFPGAs.

Real-time TOBs are transmitted to the L1Topo processor, while event readout data is sent to the FELIX system upon receipt of a L1A signal. The readout links operate at 9.6 Gbps and support up to 255 bunch crossings per L1A, assuming sufficient spacing between triggers. Figure 4.5 illustrates the input/output data flow between the processing units and external systems.




Figure 4.5: gFEX pFPGA link connections. Figure from [109].

A 40 MHz clock, aligned to the LHC bunch crossing frequency, is distributed to the board by FELIX and cleaned using an SI5345 jitter attenuator. This clock provides the timing reference for all processing stages and ensures synchronization across subsystems. FELIX also delivers control signals such as L1A, Bunch Counter Reset (BCR), and Event Counter Reset (ECR). These control signals have distinct functions critical to synchronization and

event handling:

- Bunch Counter Reset (BCR) — resets the bunch crossing counter from 3563 back to 0 at the start of each LHC orbit.
- Event Counter Reset (ECR) — clears the 24-bit counter that tracks L1A, resetting it to zero.
- Level-1 Accept (L1A) — generated by the L1CTP, this signal instructs all modules to read out data corresponding to the associated Bunch Crossing ID (BCID).

The three pFPGAs communicate directly over high-speed links to exchange energy information for objects that cross η boundaries. Each pFPGA also interfaces with the onboard Zynq SoC, which aggregates MET and the total transverse energy ($\sum E_T$) before forwarding them to L1Topo. This process requires precise synchronization to ensure correct BCID and consistent trigger decision timing.

4.2 Firmware Algorithms

There are two main algorithms implemented in the gFEX firmware: the Large Radius Jet Finder and the “Jets without Jets” (JwoJ) MET algorithm. These algorithms are preceded by the gTower Builder, which, while less algorithmic in nature, plays a crucial role in pre-processing and significantly contributes to latency and overall data handling efficiency.

4.2.1 *gTower Builder*

The gTower Builder firmware operates independently of user control and is implemented on the pFPGAs, each of which is assigned to process a specific η region. This firmware

receives digitized input from the LAr and Tile calorimeters in the form of gCaloTowers, which represent energy sums from multiple Trigger Towers as shown in Figure 3.12.

While the legacy L1Calo system operated on Trigger Towers with granularity $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$, the gFEX system uses a coarser granularity approach. The gTower Builder decodes the relevant bits to extract transverse energy values and saturation flags from these inputs, then constructs gTowers by summing the electromagnetic and hadronic gCaloTowers over a 0.2×0.2 region in $\Delta\eta \times \Delta\phi$ space.

This reduced granularity is a key design choice that enables gFEX to process data from the entire calorimeter simultaneously on a single board, providing the global view necessary for large-radius jet reconstruction and global MET calculations.

Once built, the gTower energies—forming the foundation of all subsequent gFEX processing—are stored in onboard playback memory. These values are then made available to the Jet Finder and MET algorithms, as well as to other subsystems that require access to the raw calorimeter information.

4.2.2 Large Radius Jet Algorithm: Jet Finder

The Jet Finder algorithm is implemented in the pFPGA firmware and is responsible for identifying high- p_T jets by searching for localized energy maxima in the η - ϕ plane. Each pFPGA is subdivided into 64 “jet engines,” with each engine handling one ϕ bin and six η bins. The algorithm begins by computing gBlocks, which are formed from 3×3 sums of gTowers (or 3×2 at the high $|\eta|$ boundaries of the pFPGAs). From these, the four highest transverse momentum (p_T) gBlocks are selected per engine and treated as small-radius jet candidates.

Next, the algorithm constructs large-radius jets, or gJets, by summing 69 gTowers around

each gBlock that exceeds a configurable seed threshold. This threshold, set via the `SEED_THRSH` register, controls whether a gJet is built from a given gBlock. By default, the threshold is set to zero, allowing gJets to be computed for all available gBlocks. To ensure consistency near boundaries between pFPGAs, energy values at the edges are exchanged using high-speed serial links. This inter-FPGA communication is a key function of the jet engines in pFPGA C, which does not construct jets itself but supports the central boundary regions of A and B.

The final output of the algorithm consists of one large-radius gJet and two small-radius gBlocks (with the highest p_T) per half-region of each processing FPGA. These objects are serialized and sent to L1Topo over high-speed optical links, and also made available to the Zynq SoC for readout via FELIX. In addition, the calculated gBlocks are reused as inputs to the JwoJ MET algorithm, avoiding redundant calculations.

4.2.3 MET Algorithm: Jets without Jets

The computation of MET in gFEX is handled by a dedicated firmware module referred to as the MET Finder. The baseline implementation for Run 3 is the JwoJ algorithm, which estimates MET using local calorimeter activity rather than explicit jet reconstruction. By operating on gTowers and gBlocks alone, this method provides broad sensitivity to a wide range of event topologies, including those without clearly resolved jets.

The JwoJ algorithm classifies each gTower based on the transverse energy, E_T , of the surrounding gBlock—a 3×3 set of gTowers centered on the tower in question. If the gBlock energy exceeds a configurable threshold (typically 16 GeV), the gTower is considered part of the hard-scatter process and its energy is added to the MET Hard Term (MHT).¹ Otherwise,

1. This gFEX MHT is distinct from the jet-based Missing Hadronic Transverse energy (MHT) algorithm described in Chapter 5.3.1.

the gTower contributes to the MET Soft Term (MST), which primarily reflects pile-up and underlying event energy. Both hard and soft terms are computed separately in the x and y directions.

To compute vector components, each gTower is assigned a direction based on its azimuthal position, indexed by ϕ . Lookup tables are used to retrieve precomputed values of $\cos \phi$ and $\sin \phi$, which are then used to decompose each gTower's E_T into x and y components. These components are summed independently for MHT and MST. A weighted sum of the two is then formed on each pFPGA using configurable coefficients a and b :

$$E_{T,x(y)} = a \cdot MHT_{x(y)} + b \cdot MST_{x(y)}. \quad (4.1)$$

These coefficients can be tuned per pFPGA but are typically symmetric between the x and y directions. Once calculated, the x and y components are transmitted to the Zynq SoC, which performs a final vector summation to compute the full missing transverse energy magnitude:

$$E_T^{\text{miss}} = \sqrt{E_{T,x}^2 + E_{T,y}^2}. \quad (4.2)$$

This final value is sent to L1Topo for use in L1 trigger decisions and to FELIX for readout. The separation of E_T into MHT and MST enables the application of pile-up mitigation at the firmware level. The soft term, MST, contains the contribution from pile-up and low-energy background, while the hard term, MHT, is expected to correlate with the hard scattering process. By adjusting the weights a and b , the L1 MET measurement can be calibrated to better match the offline reconstruction, especially in high-luminosity conditions where pile-up is significant.

In addition to the baseline JwoJ algorithm, two further MET algorithms have been defined.

The first is the noise-cut MET algorithm, which applies a threshold of 4σ to the gTower E_T , where σ is computed from the RMS of tower energies in minimum bias data. Only towers above this threshold contribute to the final MET vector sum, using a modified version of the existing JwoJ firmware blocks. The second variant, referred to as the ρ +RMS MET, applies pile-up subtraction using the event-specific ρ value and then applies a tighter 3σ noise cut. These versions are still undergoing firmware development and are not running in ATLAS.

All MET algorithms operate downstream of the Jet Finder, reusing the gTower and gBlock energy sums already calculated in the jet path. Since MET is not on the latency-critical path, this reuse strategy introduces no delay while minimizing redundant calculations. The final MET value thus benefits from both efficient firmware design and the contextual information provided by the jet engine data flow.

4.2.4 *gTower Calibration*

The calibration of gTowers is necessary due to differences in detector response across pseudorapidity regions, primarily resulting from variations in dead material in the calorimeter. Without calibration, the energy response can vary significantly, leading to non-uniform trigger thresholds across the detector.

The process for developing our calibration values has followed a data-driven approach using $Z(\rightarrow \mu\mu) + 1$ jet events. Events with a Z boson recoiling against a jet are selected, with requirements that the Z boson p_T exceed 50 GeV and $\Delta\phi(Z, \text{jet}) < 1$. This selection ensures a well-balanced event where the jet energy can be calibrated against the precisely measured Z boson momentum.

For each selected event, the ratio between the offline jet p_T and the sum of gFEX gTower energies within $\Delta R = 0.4$ of the jet is calculated using the formula:

$$CF = \frac{p_{T,\text{leading jet}}}{\sum_{\Delta R < 0.4} \text{gTower } E_T} \quad (4.3)$$

The mean of this ratio is calculated for each η bin (with bin widths of 0.2, matching the gTower granularity). A relative calibration factor is then computed for each η bin by normalizing to the average calibration factor across the central pFPGAs a and b:

$$\text{Relative CF} = \frac{CF(\eta)}{CF(\text{FPGA a \& b})} \quad (4.4)$$

These relative calibration factors are then what get implemented in the gFEX firmware. The actual value programmed into the firmware is $(1024 \times \text{Relative CF})$, where the factor of 1024 is used to convert the floating-point relative calibration factor into an integer value suitable for the fixed-point arithmetic used in the FPGA.

Thirty-six η -dependent calibration factors have been implemented in the gFEX firmware since June 2024, with additional optimization of thresholds ongoing. Following a fix to the LAr calorimeter, these calibration factors have been recalculated and updated to maintain performance.

4.2.4.1 Effects of gTower calibration

The implementation of gTower calibration has had significant positive effects on both trigger rates and efficiency uniformity. Before calibration, the 50% efficiency point for the gLJ140 trigger varied by up to 40 GeV across different η regions. After calibration, these variations were reduced to less than 10 GeV, creating a much more uniform trigger response throughout the detector, as shown in Figure 4.7.




Figure 4.6: Calibration factors as a function of η . Plot by Despoina Sampsonidou.




Figure 4.7: The offline large-R jet 50% jet trigger efficiency point of trigger item gLJ140 as a function of η , shown both with and without the gFEX tower-level calibration applied. Figure courtesy of Despoina Sampsonidou [110]

For MET triggers, calibration sharpened the turn-on curve and reduced rates substantially without compromising efficiency. Comparing the rates of the L1_gXEWOJ100 trigger before and after calibration shows a dramatic improvement: without calibration, the rate was approximately 10 kHz at $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ luminosity, while with calibration, it decreased to approximately 6 kHz at the same luminosity [110]. This 40% rate reduction allows for more efficient bandwidth utilization while maintaining or improving physics performance.

The calibrated trigger achieves high efficiency for events with genuine MET, as verified using $Z \rightarrow \mu\mu$ events where the dimuon system p_T serves as a proxy for the calorimeter MET. The sharper turn-on curve also improves the quality of selected events by reducing the number of spurious triggers caused by calorimeter non-uniformities.

4.2.5 *Pile-Up correction*

To mitigate the impact of multiple proton-proton interactions per bunch crossing (pile-up), the gFEX system applies an event-by-event correction using the Pile-Up Correction (PUC) technique implemented directly in firmware. Pile-up contributes additional energy that can obscure high- p_T signatures and artificially inflate trigger rates. The correction relies on computing the average transverse energy density ρ from all gTowers, where ρ (rho) represents the pile-up density and is also referred to as the “energy density” or “event energy density”:

$$\rho = \frac{\sum_{i \in \text{gTowers}} E_T^i}{\sum_{i \in \text{gTowers}} A_i} \quad (4.5)$$

Each gTower’s corrected energy is then calculated as $E_T^{i,\rho} = E_T^i - A_i \rho$, where A_i is the area of the i th gTower. These corrected values are used in both jet finding and MET calculations, and the computed ρ is included in the gFEX output.

This real-time firmware implementation ensures that pile-up suppression is applied before trigger decisions are made, leading to improved object resolution and more stable trigger rates. Without correction, trigger rates scale non-linearly with instantaneous luminosity (approximately as $L^{1.3}$), but with pile-up subtraction, they scale nearly linearly ($L^{1.05}$). This linear behavior is essential for predictable and efficient trigger operation throughout LHC fills.

Combined with η -dependent calibration, the pile-up correction strategy has been critical for maintaining high trigger performance in the high-luminosity environment of Run 3, allowing the gFEX system to meet its physics objectives while respecting the stringent bandwidth constraints of the L1 system.

The combination of η -dependent calibration and pile-up suppression, both implemented at the firmware level, has been crucial for maintaining excellent trigger performance throughout Run 3 data-taking, particularly in high-luminosity conditions with increased pile-up. These techniques allow the gFEX system to deliver its physics goals while respecting the strict bandwidth limitations of the L1 trigger system.

4.3 gFEX Trigger Objects summary

For each event, gFEX outputs a collection of TOBs that are sent to L1Topo for use in the L1 decision logic. From each half-region of pFPGAs A and B, one large-radius jet (gJet) and two small-radius jets (gBlocks) are selected and transmitted. The global MET value, comprising both x and y components, and the scalar sum of transverse energy ($\sum E_T$) are also included. Additionally, the event-level ρ value computed for pile-up subtraction is sent. A summary of all gFEX algorithm objects is provided in Table 4.1.

These objects enable a variety of trigger strategies targeting high- p_T jets, large MET, or combinations of both. Such flexibility is essential for capturing physics signatures involving boosted objects, invisible particles, or exotic topologies beyond the SM.

4.4 Operating system and low-level monitoring

The Zynq SoC on the gFEX board runs a custom Linux-based OS responsible for system configuration, control, and monitoring. The monitoring infrastructure follows a hierarchical design with multiple integrated layers, each serving specific purposes within the overall monitoring strategy:

1. Sensor-level monitoring: Direct hardware monitoring via the Inter-Integrated Circuit

Object	Description	Size ($\Delta\eta \times \Delta\phi$)
gTower	Basic unit from calorimeter input; sum of EM and hadronic components	0.2×0.2
<i>Jet Algorithm Objects</i>		
gBlock	Sum of 3×3 gTowers: Small-radius jet	0.6×0.6 ($R \approx 0.3$)
gJet	Sum of 69 gTowers around a seed gBlock: Large-radius jet	1.6×1.6 ($R \approx 0.8$)
<i>MET Algorithm Objects</i>		
MHT	Hard term: gTowers with surrounding gBlock $E_T \geq 16$ GeV	
MST	Soft term: gTowers with surrounding gBlock $E_T < 16$ GeV	

Table 4.1: Summary table of gFEX algorithm objects

- (I²C) bus interfaces with temperature, voltage, and current sensors across the board.
2. ATCA integration: The Intelligent Platform Management Controller (IPMC) collects board health information and interfaces with the ATCA Shelf Manager, providing integration with the rack-level monitoring infrastructure.
 3. Control system integration: An Open Platform Communications Unified Architecture (OPC-UA) server hosted on the SoC provides standardized access to board parameters for the ATLAS Detector Control System (DCS).
 4. Operational monitoring: Register values and status information are published to the ATLAS Information Service (IS), making them available to time-series databases and visualization tools like Grafana dashboards.

This multi-layered approach, illustrated in Figure 4.8, ensures comprehensive supervision from the component level through to experiment-wide monitoring systems. The following subsections detail the implementation of the low-level hardware monitoring systems and then cover the data quality (DQ) monitoring used to assess system performance during operation.




Figure 4.8: Overview of the gFEX monitoring systems, including several different graphical interfaces and low-level onboard hardware monitoring. Figure courtesy of Cecilia Tosci.

4.4.1 Custom operating system

The gFEX custom OS is built using the Xilinx-supported versions of Yocto and OpenEmbedded. This setup provides the benefits of an open-source framework along with support for hardware-specific integration. In particular, it allows hardware descriptions and configuration files to be pulled directly into the OS build process, which simplifies the overall workflow and reduces the need for manual setup. The workflow of this build is demonstrated in Figure 4.9.

The OS is built using Bitbake, the main build tool used by Yocto. Bitbake handles tasks like fetching source code, applying patches, configuring packages, and compiling everything into a root filesystem. Each part of the system is described using metadata files, which are grouped into “layers” to organize functionality. Xilinx maintains several of these layers—such as `meta-xilinx` and `meta-xilinx-tools`—which support hardware like the Zynq SoC on the gFEX board. These layers are versioned and released alongside major Yocto versions, making it easier to track compatibility. Table 4.2 shows some of the recent releases used in the project.

The Xilinx Shell Archive (XSA) file is a hardware description package exported from Vivado




Figure 4.9: Workflow for setting up and building the gFEX OS. Figure from [111].

that contains all the information about a design’s components, configuration, and address map, used for generating software platforms and device trees. One major advantage of using this build system is that the XSA file produced by the firmware build can be automatically included in the OS image. This makes it much easier to stay in sync with firmware changes, since the device tree and other hardware configuration files no longer need to be written by hand. In earlier versions of the system, this process was manual and more error-prone.

A custom Yocto layer called `meta-llcalo` was created specifically for use with the gFEX board. It was initially developed by Giordon Stark and later maintained by Emily Smith before I joined the project. This layer contains all the files needed to describe the gFEX hardware, including the device tree and any custom software required to run within the ATLAS trigger system.

The layer is organized into several folders based on functionality. The `conf/machine` directory contains the files to define the different gFEX “machines,” which are `gfex-production-p1`

Xilinx Branch	Yocto Code Name	Yocto Version	Release Date	Support Level
N/A	Styhead	5.1	October 2024	7 months (May 2025)
rel-v2024.2	Scarthgap	5.0	April 2024	LTS (April 2028)
N/A	Nanbield	4.3	November 2023	EOL
N/A	Mickledore	4.2	May 2023	EOL
rel-v2023.1	Langdale	4.1	October 2022	EOL
rel-v2022.1	Kirkstone	4.0	May 2022	LTS (April 2026)
rel-v2021.2	Gatesgarth	3.2	October 2020	EOL
rel-v2021.1				
rel-v2020.2	Zeus	3.0	October 2019	EOL
rel-v2020.1				
N/A	Dunfell	3.1	April 2020	EOL
rel-v2018.3	Rocko	2.4	October 2017	EOL
rel-v2018.2				
rel-v2018.1				

Table 4.2: Relevant Yocto versions compared to Xilinx Vivado versions, where EOL means some community support on email lists but no commit updates in the repos. Material compiled from [112, 113].

and `gflex-production-stf`, for boards deployed at Point 1 and the L1Calo Surface Test Facility (STF), respectively. This directory is auto-populated by the TCL script described later, except for the `gflex-production.dtsi` file, which tells the hardware how to connect and configure the Ethernet ports and SD card slot so they work correctly when the device starts up. The `recipes-bsp` folder includes the board support XSA file. The `recipes-core` folder defines the full gFEX image and the infrastructure needed to run startup scripts. Additional folders include gFEX-specific applications (`recipes-gflex`), kernel modifications (`recipes-kernel`), and Python modules used by the OS (`recipes-python`). The `recipes-support` folder contains tools like Chrony, which keeps the board’s clock in sync with an external NTP server.

A major update to the build system was completed during the migration from Xilinx release rel-v2020.2 to rel-v2024.2, which corresponds to the Yocto “Scarthgap” release (version 5.0). This change was required to support Ubuntu 22 on the build machines and introduced a new System Device Tree (SDT)–based workflow, replacing the previous xsct method. In the new flow, a TCL script generates SDT files from the hardware design, and then the gen-machineconf tool creates the necessary machine configuration files. According to Xilinx, this approach is intended to simplify getting started with new platforms and align with the updated Yocto project infrastructure [114].

The upgrade involved close collaboration with the Xilinx open-source community, especially with Mark Hatle, who helps maintain the `meta-xilinx` GitHub repository [114]. One of the goals was to ensure that the gFEX system can continue to generate up-to-date OS images even as the build tools evolve. Moving forward, the plan is to adopt smaller, more frequent upgrades to avoid the disruption caused by large version jumps.

4.4.2 I2C Bus Monitoring

The gFEX board uses the I²C protocol, a synchronous multi-controller, multi-target serial interface, to monitor low-level hardware components such as voltage, current, and temperature sensors. These sensors are connected to the Zynq+ SoC, supplying essential data for system protection, supervision, and integration with higher-level infrastructure such as the IPMC. Two independent I²C buses are implemented on the board. The main bus connects sensors on power modules, temperature monitors, and miniPODs, enabling real-time tracking of voltages, current draw, thermal conditions, and signal integrity via Loss of Signal (LOS) indicators. A secondary bus interfaces with the SI5345 clock chip, which supports both configuration and monitoring of the clock generation system, critical for timing synchronization.

Software on the custom OS handles I²C communication using the Python peripheral module to interact with sensors and `/dev/mem` for direct access to memory-mapped FPGA registers. This setup enables comprehensive monitoring of both the external hardware environment and internal firmware states. A dedicated monitoring script continuously polls all available sensors, aggregates the data, and determines the maximum temperature observed across all sources, including miniPODs, power modules, discrete temperature sensors, and the Zynq itself. This value is written to a designated general-purpose General-Purpose Input/Output (GPIO) within the Zynq firmware, where it can be read directly by the IPMC. This design avoids redundant polling and ensures critical thermal information is available to supervisory systems with minimal overhead.

A key feature of this monitoring chain is a built-in safety mechanism that disables the on-board DCDC converter when any temperature exceeds predefined thresholds. This automatic response reduces power consumption and protects sensitive components such as the FPGAs. Beyond local protection, I²C data is also published to the IS, where it feeds into higher-level interfaces such as Grafana dashboards. This layered approach enables shift crews and system experts to track the board’s status in real-time, combining immediate board-level safeguards with long-term performance monitoring across the ATLAS control infrastructure.

4.4.3 IPMC and ATCA Monitoring

The IPMC builds upon the local monitoring infrastructure described above by integrating gFEX into the broader ATCA shelf environment. Rather than reading sensor data directly, the IPMC receives processed temperature measurements from the Zynq OS, which aggregates data from all I²C-connected sensors and exposes the results via dedicated general-purpose GPIOs. This architecture enables a clean separation between low-level data collection and system-wide monitoring, ensuring both local responsiveness and global oversight.

The IPMC firmware reads these GPIO values to obtain maximum temperatures for each sensor class—miniPODs, power modules, onboard temperature sensors, and the Zynq chip itself. Thresholds for each category are defined at multiple severity levels, as shown in Table 4.3, and can trigger escalating responses, including automatic board shutdowns if upper non-recoverable limits are exceeded.

Sensor	Upper Non-Critical	Upper Critical	Upper Non-Recoverable
IPMC Internal	45°C	55°C	70°C
Zynq Temperature	55°C	70°C	85°C
Temperature Sensors	55°C	65°C	70°C
Power Module Sensor	90°C	95°C	100°C
MiniPOD Temperature	65°C	75°C	80°C

Table 4.3: gFEX IPMC temperature sensor threshold values.

A built-in watchdog mechanism ensures robustness: if the Zynq fails to update the GPIO values within a defined timeout window, the firmware substitutes a safe default (0°C), guaranteeing that stale or missing data is flagged immediately. Alerts from the IPMC are propagated through the Shelf Manager to the ATLAS Detector Control System (DCS), which issues notifications to relevant L1Calo experts via email and SMS.

Monitoring data from both the I²C and IPMC systems are integrated into the ATLAS monitoring infrastructure. The Zynq hosts an Open Platform Communications Unified Architecture (OPC-UA) server that interfaces with the DCS, while key register values are published to the IS, where they are visualized through Grafana dashboards in the ATLAS Control Room (ACR). This unified framework ensures both real-time protection and long-term observability across the experiment’s operations.

4.4.4 DCS Integration

The DCS is a central infrastructure within ATLAS that provides unified monitoring and control of all detector subsystems. It monitors operational parameters such as temperatures, voltages, and gas flows across the entire detector, manages alarm handling, archives historical data, and provides operators in the ACR with a consistent interface to oversee the experiment's operation. The DCS features a graphical interface and a client-server architecture that ensures real-time updates to guarantee the safety of hardware components. It also incorporates an alarm system that notifies users of active alerts and implements fail-safe mechanisms that can automatically deactivate modules when necessary. The DCS is critical for safe detector operation, ensuring that all components remain within their specified operating parameters during data-taking.




Figure 4.10: Overview of the DCS system for L1Calo. Figure courtesy of Paul Thompson.

The graphical user interface of the DCS is commonly referred to as the Finite State Machine

(FSM), which defines operational states and transitions for each hardware component. Each component or subsystem is represented by one or more graphical panels within the FSM, which provide color-coded indicators and real-time system data. The architecture of the DCS is built upon the SIMATIC WinCC Open Architecture (WinCCOA) software and the OPC-UA protocol. Panels and data points within the FSM are connected to OPC-UA clients, which gather low-level sensor data from remote or local OPC-UA servers. An overview of this system is shown in Figure 4.10.

gFEX is fully integrated into this ATLAS-wide infrastructure, providing standardized monitoring and control consistent with other detector components. This integration uses OPC-UA servers running on both the gFEX Zynq SoC and dedicated DCS machines. Specifically, gFEX runs a primary OPC-UA server on the SoC to expose board-level sensor data, and an additional OPC-UA server runs on a nearby PC to gather data from the ATCA Shelf Manager. The OPC-UA server on the SoC polls low-level sensors at regular intervals of approximately 30 seconds by reading registers via the gFEX I²C bus, providing a comprehensive view of board health. Figure 4.11 includes classes for the different types of physical sensors and the various types of data being monitored.

The gFEX OPC-UA server is implemented in Python, making it easier to integrate with the Zynq SoC OS. The server uses the same I²C monitoring code that serves the IPMC and onboard monitoring systems to supply low-level sensor data to the OPC-UA server datapoints. The gFEX DCS project builds upon existing libraries created by the ATLAS central DCS group, such as fwATCA, which standardizes ATCA-based OPC-UA servers that utilize IPMC data points delivered by the shelf manager.

The gFEX DCS system includes separate FSM instances for both the P1 and STF boards. I have taken over maintenance of the gFEX DCS system from Emily Smith and, working with Paul Thompson, updated the STF DCS implementation to simultaneously monitor the




Figure 4.11: gFEX OPC-UA server class diagram. Figure courtesy of Emily Smith.

health and safety of two boards in the STF. With this change, we can monitor datapoints for two gFEX boards in the STF simultaneously while keeping the standard board alerts active. This provides an additional layer of monitoring information complementary to the built-in alert systems. This capability is particularly valuable for development and testing, as it enables parallel work on different board configurations without affecting the production system.

The graphical panels provided in the gFEX FSM include:

- **Main gFEX panel:** Overview of values from the ATCA and SoC servers, plus schematic layout and panel navigation buttons, shown in Figure 4.13
- **Temperature panel:** Displays all board temperatures with color-coded status indicators
- **Power module panel:** Shows voltage and current measurements for all power rails




Figure 4.12: gFEX FSM minimap panel at Point 1




Figure 4.13: Main gFEX panel in the STF showing two boards present and labeled.

- **MiniPOD panel:** Combines graphical LOS status with tabular electrical and optical data. Optical alerts are currently disabled by default; however, the values are archived for long-term data on the optical power of the miniPODs, as shown in Figure 4.12.
- **Expert panels:** Available with authentication for specialized configuration tasks

The DCS system archives all data point values to a central database, enabling trend analysis, correlation studies, post-incident investigation, and long-term performance evaluation. Adaptive sampling rates balance storage requirements against data granularity, with critical parameters sampled more frequently.

In parallel with the DCS integration, the gFEX online software also publishes the contents of many status and control registers to the ATLAS IS. These registers are polled at regular intervals and their values updated in IS, where they can be queried by higher-level monitoring interfaces, such as Grafana dashboards, used by the shift crews in the ACR.

This multi-layered monitoring approach—from low-level I²C sensors through the IPMC to the central DCS and IS—provides comprehensive protection for the valuable hardware while ensuring operators have the information needed to maintain system performance. All of the

ATLAS TDAQ DCS implementations are kept in the same repository. Each subsystem has an individual folder to hold its specific implementation of the WinCCOA client and FSM. Individual folders are holding WinCCOA scripts as well as regular bash scripts to facilitate building the various components of the DCS project.

The complete build process for the DCS components includes the ATCA core server, the fwATCA client, the SoC server-based client, and the FSM, creating a fully integrated monitoring solution for the gFEX system.

4.5 Performance Monitoring

In addition to hardware monitoring, offline performance validation ensures that gFEX fulfills its physics goals. Monitoring includes jet trigger efficiency, trigger rates, and validation of algorithm output.

4.5.1 Trigger Efficiency Curves

Trigger efficiency is a fundamental concept in particle physics experiments, particularly in ATLAS where data collision rates far exceed recording capabilities. It quantifies the effectiveness of a trigger system in selecting physics events of interest while rejecting background.

For jet triggers specifically, efficiency measures the probability that a jet with a given transverse momentum (p_T) in offline reconstruction causes the trigger to fire. The offline Particle Flow jets (PFlow jets) used as reference are defined in Chapter 5.2.1.2. While perfect efficiency would mean every jet above a certain threshold triggers event acceptance, real-world factors introduce complications. Resolution effects, calibration differences between online and offline reconstructions, and electronic noise result in a characteristic “turn-on” behavior

instead of a sharp threshold cutoff.

These performance characteristics are visualized through efficiency curves, plotting the trigger acceptance probability against the offline jet p_T . Rather than an ideal step function, real triggers display a gradual transition or “turn-on curve.” The steepness of this transition region reveals the trigger’s resolution, while the efficiency plateau at high p_T indicates any systematic inefficiencies.

Accurate efficiency measurements are crucial for physics analyses, enabling experimenters to determine whether triggers operate at intended thresholds and achieve full efficiency for high-energy objects. The shape of these curves provides immediate feedback on trigger performance, identifying issues that might compromise data quality or limit physics reach.

To calculate trigger efficiency, we use the mathematical definition:

$$\epsilon(p_T) = \frac{N_{\text{Offline}}^{\text{Trigger}}(p_T)}{N_{\text{Offline}}^0(p_T)} \quad (4.6)$$

where $N_{\text{Offline}}^{\text{Trigger}}(p_T)$ is the number of offline jets with transverse momentum p_T that pass the trigger selection, and $N_{\text{Offline}}^0(p_T)$ is the total number of offline jets with the same p_T in the reference sample.

4.5.2 Jet Trigger Efficiency Monitoring Algorithm

To ensure the optimal performance of the ATLAS trigger system during data-taking, continuous monitoring of trigger efficiencies is essential. As part of this thesis work, I developed a monitoring algorithm designed to evaluate the performance of jet triggers in the L1Calo system. This framework, created entirely from scratch, has been fully integrated into the

ATLAS Tier0 processing DQ Monitoring infrastructure, making it accessible through the experiment’s web-based monitoring display system.

The framework systematically analyzes the performance of all jet triggers, including those from both the gFEX and jFEX, by generating trigger efficiency curves for each ATLAS data-taking run. These runs typically occur overnight, with preliminary results available through the “express-express” stream by mid-afternoon the following day, allowing for prompt identification of any potential issues. This rapid turnaround is crucial for maintaining high-quality data collection throughout the experiment’s operation.

The algorithm generates trigger efficiency turn-on curves by populating ROOT TEfficiency objects for each jet trigger. The efficiency calculation is essentially the division of two histograms, demonstrated in Figure 4.14.

- Denominator - histogram of leading offline jet p_T for events that:
 - Pass at least one muon trigger from the reference selection
 - Contain a leading offline jet satisfying basic quality criteria ($p_T > 10$ GeV, $|\eta| < 2.5$)
- Numerator - histogram of leading offline jet p_T for events that:
 - Events in the denominator
 - Pass the specific jet trigger being evaluated

The algorithm only includes the leading offline jet in the reference sample if it meets the specified quality criteria. Since these are single jet triggers, the algorithm only checks whether the trigger fired, without performing any spatial matching between offline jets and trigger objects. This direct approach avoids the computational overhead of accessing and matching individual TOBs, which would slow down the monitoring process.




Figure 4.14: Numerator and denominator histograms that make up the TEfficiency curve for the gLJ100 trigger turn on curve for run 486315.

To produce unbiased efficiency measurements, a reference selection independent of the triggers being evaluated is essential. Without such a reference, we would inadvertently bias our sample toward events that ATLAS is designed to select (high- p_T jets or specific particle topologies favored by existing triggers), preventing an accurate assessment of actual trigger performance across the full kinematic spectrum. There are two main approaches commonly used for reference selections of jet triggers:

- Bootstrap method:** This approach uses low-threshold jet triggers to evaluate the efficiency of higher-threshold ones. While conceptually straightforward, it suffers from significant practical limitations in the ATLAS environment. Low- p_T jet triggers are heavily prescaled, meaning only one out of every N events passing the trigger is actually recorded. This prescaling results in insufficient statistics to properly populate turn-on curves, particularly in the critical low p_T turn-on regions. Additionally, using calorimeter-based triggers to evaluate other calorimeter-based triggers introduces potential self-correlation effects.




Figure 4.15: Resulting TEfficiency curve for the gLJ100 trigger turn on curve for run 486315 from Figure 4.14 components.

- **Muon selection:** This approach, which I implemented in the jet trigger efficiency monitoring algorithm, uses muon-triggered events as an independent reference sample. Muon triggers provide statistically independent event samples, as muons are minimum-ionizing particles that deposit only minimal energy ($\sim 2\text{-}3$ GeV) in the calorimeters, regardless of their momentum. This independence from calorimeter activity makes them ideal for unbiased jet trigger efficiency measurements.

From a physics perspective, this approach leverages several well-understood SM processes that naturally produce both muons and jets in the final state. Z +jets events, where $Z \rightarrow \mu\mu$, provide a clean sample with well-balanced transverse momentum, making them excellent for probing jet triggers across a wide kinematic range. Similarly, $t\bar{t}$ production, with muons from

top quark decays, provides higher jet multiplicity events with energetic b-jets - precisely the kind of complex event topologies relevant for many BSM searches relying on jet triggers.

The muon reference selection in this algorithm includes single muon triggers with various p_T thresholds (e.g., L1_MU14FCH, L1_MU18VFCH), di-muon triggers (L1_2MU8F), multi-muon triggers (L1_3MU3VF, L1_4MU3V), and combined triggers requiring muons with other objects (L1_MU8F_TAU20IM). The inclusion of L1_RDO_FILLED provides a random trigger on filled bunches, offering an additional unbiased reference point. The FCH, VF, and V suffixes denote different geometrical and quality requirements for the muon candidates.

This algorithm populates trigger efficiency curves as a function of offline jet p_T for both small and large radius jet triggers, for gFEX and jFEX. Different jet containers are used for the populating of TEfficiency curves, according to the trigger being investigated:

- **Small-radius jets:** AntiKt4EMPFJet
 - Jets reconstructed using the anti- k_t algorithm [115] with radius parameter $R = 0.4$
 - Uses energy flow objects as inputs with particle flow algorithms [116] to optimize energy measurements
- **Large-radius jets:** AntiKt10UFOCSSKSoftDropBeta100Zcut10Jets
 - Jets with radius parameter $R = 1.0$
 - Uses Unified Flow Object (UFO) [116] as input constituents
 - Applies Constituent Subtraction and Soft Killer (CSSK) pile-up suppression algorithms [117]
 - Soft-drop grooming procedure (parameters $\beta = 1.0$ and $z_{\text{cut}} = 0.10$)

4.5.2.1 Quality Assessment and Web Display Integration

The TEfficiency objects are integrated into the ATLAS DQ web display through modification of the “Han Config File.” A sophisticated quality evaluation system uses a stoplight color scheme (green, yellow, red) to indicate the quality of each trigger’s performance. Each efficiency curve is fitted using a “fermi_fit” function, seen in Equation 4.7, enabling extraction of two critical parameters: maximum efficiency, which is expected to approach 100% for properly functioning triggers, and turn-on position, which is the p_T value at which the trigger reaches 50% efficiency.

$$f(x) = \frac{[0]}{1 + \exp\left(\frac{[1]-x}{[2]}\right)} \quad (4.7)$$

Each efficiency curve is fitted using the “fermi_fit” function (Equation 4.7) [118], which yields two key parameters: the maximum efficiency and the turn-on point, defined as the p_T value at which the efficiency reaches 50%. Both of these fitted values are used to assign a “stoplight” flag – green, yellow, or red – to each trigger on the DQ web display.

The most severe flag from either parameter is applied to each histogram. The classification is as follows:

- Green indicates that everything looks good.
- Yellow is a caution signal—there may be an issue, but it could also stem from something less concerning, like low statistics.
- Red suggests a likely problem that warrants investigation.

The maximum efficiency is evaluated using the simple thresholds summarized in Table 4.4. The turn-on point is compared against trigger-specific expectations defined in Table 4.5.

These expectations were determined through detailed inspection and reflect the known behavior of each individual trigger.

	Green range [%]	Yellow range [%]	Red range [%]
Maximum efficiency	≥ 99	95–99	< 95

Table 4.4: Classification thresholds used to evaluate the fitted maximum efficiency of each trigger.

To provide a high-level summary, the web display also applies the most severe flag color to the dropdown menu label of each algorithm. This means that if all histograms are green, a user can visually confirm healthy performance without even needing to expand the folder of plots. An example of the web display interface is shown in Figure 4.16, which demonstrates the color-coded quality assessment system in action.




Figure 4.16: DQ Tier 0 Web Display of Jet Trigger Efficiency Monitoring Algorithm results for run 486315. Green boxes highlight turn-on curves with no detected issues, and the red box highlights a problem with the turn-on curve for gJ400 discussed in Chapter 4.5.2.2. Web display screen grabbed from the Tier 0 Web Display (available here).

Trigger Name	Green range [GeV]	Yellow range [GeV]	Red range [GeV]
L1_gJ20pETA25	< 55	[55, 75]	> 75
L1_gJ50pETA25	< 110	[110, 180]	> 180
L1_gJ100pETA25	< 195	[195, 300]	> 300
L1_gLJ400pETA25	< 625	[625, 900]	> 900
L1_gJ80pETA25	< 125	[125, 200]	> 200
L1_gLJ100pETA25	< 150	[150, 240]	> 240
L1_gLJ140pETA25	< 210	[210, 315]	> 315
L1_gLJ160pETA25	< 240	[240, 360]	> 360
L1_j30	< 50	[50, 85]	> 85
L1_j40	< 55	[55, 90]	> 90
L1_j50	< 60	[60, 100]	> 100
L1_j60	< 70	[70, 120]	> 120
L1_j80	< 90	[90, 155]	> 155
L1_j90	< 110	[110, 185]	> 185
L1_j125	< 145	[145, 230]	> 230
L1_j140	< 160	[160, 255]	> 255
L1_j160	< 180	[180, 260]	> 260
L1_j180	< 235	[235, 320]	> 320
L1_SC111-CjJ40	< 195	[195, 300]	> 300

Table 4.5: DQ flag ranges for all triggers monitored using the jet trigger efficiency monitoring algorithm, used to quickly determine if a trigger is not turning on as expected. Values are hard-coded in the algorithm and were determined by eye. All thresholds are in GeV.

4.5.2.2 Issue Identification and Resolution

This monitoring algorithm has been a key in identifying real issues within the gFEX firmware implementation. The algorithm has directly contributed to improved system reliability by spotting critical issues that might have otherwise gone undetected during standard commissioning procedures.

A particularly significant discovery was the identification of a bug in the gJ400 trigger, which never reached 100% efficiency despite sufficient jet energy. An example event, where the offline jets are very energetic, but the gFEX TOBs don't have enough energy to fire the trigger, is shown in Figure 4.17. Through a detailed investigation enabled by this monitoring




Figure 4.17: Example of an event where the leading offline jet has greater than 800 GeV, but the event is not passing the gJ400 trigger, due to no gFEX TOB in the event having at least 400 GeV. The leftmost plot shows all gTowers in the event, where the center of the two relevant TOBs is apparent, in the $\eta - \phi$ space. The middle plot shows the gFEX TOBs and their p_T , and the rightmost plot shows the offline small radius jets. Events like this, where there are high p_T jets in the event, but the gJ400 trigger is not firing, prevent gJ400 from reaching 100% efficiency. This is due to an issue with saturated towers in the firmware, which is being implemented.

algorithm, the root cause was traced to the handling of saturated towers in the firmware. For saturated towers, when Analog-to-Digital Converter (ADC)c counts = 2047, the nominal calibration procedure involved adding an offset of 48 and then subtracting it, resulting in a TOB energy of 399.8 GeV, just below the trigger threshold. For highly localized TOBs with towers below saturation, the calibration produced even lower energies (< 400 GeV).

This issue has been addressed in firmware updates that ensure saturated towers always report the maximum ADC value (2047 counts = 409.4 GeV). Additional modifications to the gTower builder provide extra logic for handling maxed towers, along with updated noise cuts for 2025 running.

Through systematic monitoring and quality assessment, the jet trigger efficiency framework continues to play a vital role in ensuring the reliability and performance of the gFEX system and the L1Calo jet triggers as a whole, contributing significantly to the overall data quality of the ATLAS experiment. The ability to visualize trigger performance for each run and

rapidly identify anomalous behavior has made this new algorithm an essential component in the ATLAS DQ infrastructure.

4.5.3 Trigger Rate Monitoring

Trigger rates are a critical aspect of trigger performance monitoring in the ATLAS detector. With the L1 trigger system responsible for reducing the 40 MHz bunch crossing rate to approximately 100 kHz before passing events to the HLT, the rate at which each trigger item accepts events must be carefully monitored and controlled.

For the gFEX system, rate monitoring serves several essential purposes:

- **Bandwidth management:** The finite L1 bandwidth must be shared among all L1Calo submodules. Each trigger competes for a portion of the 100 kHz maximum rate that the HLT can process in Run 3.
- **Performance assessment:** Rate behavior as a function of instantaneous luminosity provides insight into trigger stability. Linear scaling indicates proper pile-up handling, while non-linear behavior may signal algorithmic issues.
- **Comparison with legacy systems:** When evaluating the physics performance of the new gFEX triggers against legacy triggers, rate considerations are essential. Comparing efficiencies at similar rates ensures a fair assessment of detection capability per unit of bandwidth.

The gFEX triggers were designed to maintain comparable or lower rates than their legacy counterparts while providing improved efficiency for physics objects of interest.

Rate monitoring is conducted in real-time in the ACR, where trigger rates are displayed




Figure 4.18: Plot of jFEX, gFEX and legacy met trigger rates for two runs, demonstrating that the gFEX rate (L1_gXEJWOJ110) lowered following an Fcal noise cut firmware update (elog), seen from run 474770 onwards. The jFEX and legacy trigger rates are consistent.

on dedicated monitoring screens through the Grafana interface. This allows for immediate intervention if any trigger begins consuming excessive bandwidth. The monitoring infrastructure also archives rate data, enabling offline analysis of rate behavior across runs and luminosity conditions. This historical data is invaluable for trigger menu optimization and for understanding the long-term stability of the system.

To streamline rate monitoring, I developed a dedicated workflow in Python to generate trigger rate plots. This system extracts per-trigger rates from the **beauty** monitoring framework, allowing for efficient plotting and comparison of trigger behavior across multiple runs [119]. The flexibility of this approach enables rapid checks on whether trigger rates behave as expected, particularly when assessing updated firmware or new calibrations. An example of

this is shown in Figure 4.18, which demonstrates a comparison of gFEX, jFEX, and legacy MET trigger rates for two different runs, one before (closed circles) and one after (hollow circles) a noise cut firmware update in gFEX. This Figure shows that the jFEX and legacy trigger rates are consistent across the two runs, but the gFEX rate drops significantly after the firmware update.


4.6 Performance in Run 3

The gFEX system has been operational since the start of Run 3 in 2022, demonstrating reliable performance and meeting its design objectives. Initially running in parallel with legacy systems for validation, the gFEX triggers have now completely replaced their legacy counterparts in the ATLAS trigger menu, marking a successful transition to the Phase-I upgrade architecture.


Both the jet and MET triggers provided by gFEX have been enabled for physics data-taking, offering improved efficiency and resolution compared to the previous generation of L1 systems. The rate monitoring framework I developed for this system, along with several other custom monitoring frameworks, have proven invaluable in tracking performance, optimizing configurations, and ensuring stable operation within bandwidth limitations.

The small-radius (L1_gJ) and large-radius (L1_gLJ) jet triggers have demonstrated superior performance compared to the legacy L1_J triggers, particularly for events with boosted topologies. The L1_gLJ140 trigger serves as the lowest unprescaled large-radius jet item, and operates at approximately 2.3 kHz at $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ luminosity. This provides efficient coverage of boosted objects while maintaining acceptable trigger rates.

The JwoJ MET algorithm has similarly shown excellent performance. After tower-level calibration, the L1_gXEW0J100 trigger operates at approximately 6 kHz, reduced from 10 kHz



(a) Turn on curves by Lukas Pfaffenbichler



(b) Rates

Figure 4.19: Trigger performance turn on curves and corresponding trigger rates produced for the large- R gFEX jet trigger and legacy system. Efficiencies are measured using $Z \rightarrow \mu\mu$ events as a function of the leading offline large- R jet p_T , and rates are shown as a function of instantaneous luminosity. Jets are reconstructed with $R = 1.0$ and $|\eta| < 2.5$, using UFO with grooming and pile-up suppression applied. Figures from [110].

in its uncalibrated form, without compromising signal efficiency, shown in Figure 4.20. This enhancement improves sensitivity to processes with genuine MET, such as dark matter or SUSY searches, by allowing lower thresholds while maintaining acceptable trigger rates.

4.6.1 Outlook

With the decommissioning of legacy jet and MET triggers now complete, the Phase-I trigger system, comprised of gFEX, jFEX, and eFEX, is fully responsible for collecting ATLAS data for the rest of Run 3. This marks a major milestone for the L1Calo upgrade, as all physics data-taking now relies exclusively on the performance of the new triggers. gFEX plays a critical role in delivering high-efficiency triggers for large-radius jets and missing transverse energy, with particular strength in identifying boosted signatures and events with invisible particles. Its advanced algorithms and global view of the calorimeter have significantly improved both resolution and rate scalability compared to legacy systems.




Figure 4.20: MET efficiency for the JWoJ algorithm as a function of the $\mu\mu$ system p_T using $Z \rightarrow \mu\mu$ events. Efficiencies are presented for the trigger with and without a tower-level energy calibration. Figure courtest of Despoina Sampsonidou [110].

Trigger Name	Type	Threshold	Typical Rate	50% Turn-on
Jet Triggers				
L1_gLJ80	Large-R gFEX Jet	50 GeV	30 kHz	110 GeV
L1_gLJ100	Large-R gFEX Jet	70 GeV	10 kHz	140 GeV
L1_gLJ140	Large-R gFEX Jet	110 GeV	2 kHz	190 GeV
L1_gLJ160	Large-R gFEX Jet	130 GeV	1 kHz	220 GeV
L1_J100	Legacy Jet		4 kHz	170 GeV
MET Triggers				
L1_gXEWOJ100	JwoJ MET	100 GeV	6 kHz	110 GeV
L1_XE50	Legacy MET		7 kHz	110 GeV

Table 4.6: Summary of gFEX and legacy trigger items, including TOB thresholds, rates, and approximate turn-on behavior. All rate values are quoted at an instantaneous luminosity of $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$.

As summarized in Table 4.6, the current gFEX trigger menu provides a comprehensive set of large-radius jet and MET triggers optimized for Run 3 physics analyses. The large-radius jet triggers, ranging from L1_gLJ80 to L1_gLJ160, offer substantially better efficiency than the legacy L1_J100 trigger while maintaining comparable rates. For instance, L1_gLJ140 achieves a 50% turn-on at 190 GeV with only 2 kHz rate, compared to the legacy trigger's 170 GeV turn-on at 4 kHz. Similarly, the L1_gXEW0J100 MET trigger exhibits improved performance, featuring a sharper turn-on curve and better rate control, compared to the legacy L1_XE50 trigger.

The rate monitoring and data quality frameworks I developed continue to serve as key tools for ensuring smooth operation and rapid diagnostics. These systems provide shift crews and experts with immediate feedback on trigger behavior, helping maintain data integrity and enabling ongoing performance optimization. By combining precise rate control with enhanced physics sensitivity, the gFEX system has proven itself to be a key component of the ATLAS Run 3 trigger menu, fulfilling the central goals of the Phase-I upgrade and laying the foundation for future high-luminosity operations.

CHAPTER 5

EVENT RECONSTRUCTION AND OBJECT DEFINITION

Raw data collected by ATLAS undergo multiple stages of processing before they can be interpreted in terms of physics observables. This process, known as reconstruction, transforms detector signals into physics objects that represent the particles produced in collisions, and their observable features and kinematics. These objects — such as tracks, vertices, jets, electrons, photons, and missing transverse energy — form the foundation of physics analyses. Each object type requires specialized reconstruction algorithms to extract the information from the detector subsystems. For example, tracks represent the trajectories of charged particles through tracking detectors, while jets capture the complex hadronic showers resulting from quarks and gluons.

5.1 Tracking

Tracking is the fundamental process of reconstructing the trajectories of charged particles as they traverse the ID [82]. When a charged particle passes through the subdetector elements of the ID, like the Pixel Detector, SCT, or TRT, it creates localized ionization signals referred to as “hits,” as discussed in Chapter 3.6. These hits serve as spatial coordinates marking the particle’s path through the detector volume. Tracking algorithms connect these discrete measurements from the various ID layers to form continuous trajectories, parameterized by variables that describe the particle’s position and momentum. Typically, tracks are characterized by five trajectory parameters defined at their perigee point, or the point of closest approach to the beam spot, including transverse and longitudinal impact parameters (d_0 and z_0), azimuthal and polar angles (ϕ and θ), and the charge-to-momentum ratio (q/p) [120]. The curvature of the track in the ID’s magnetic field provides crucial information about the

particle’s momentum. Standard tracking algorithms are highly efficient for prompt particles originating near the interaction point, but their performance decreases significantly for particles with large impact parameters, necessitating specialized reconstruction techniques for displaced vertices.

Parameter	Definition
d_0 - Transverse impact parameter	Distance in the transverse plane between the track’s perigee and a reference point
z_0 - Longitudinal impact parameter	Longitudinal distance between the track’s perigee and a reference point
ϕ - Azimuthal angle	Angle of track’s momentum in the transverse plane
θ - Polar angle	Angle of track’s momentum in the longitudinal plane
q/p	Charge of the track divided by the magnitude of its momentum

Table 5.1: Definitions of variables used to parameterize tracks. The reference point used to determine the track’s perigee during track reconstruction is the origin of the detector coordinate system.

5.1.1 Standard Tracking

Standard tracking is used to reconstruct particles originating from or near the interaction point [120]. The process begins with clustering charge deposits on adjacent pixels and strip sensors, which represent where particles traversed the detector material. These clusters are transformed into three-dimensional-space-points that precisely locate where particles passed through the active detector components. In the Pixel Detector, each cluster creates a single space-point, while in the SCT, clusters from both sides of a strip layer are combined to form a single space-point with improved spatial resolution.

The track reconstruction follows an “inside-out” approach, where track seeds are formed




Figure 5.1: Illustration of the five global track parameters, defined relative to a reference point, the perigee, or point of closest approach to the beam line. Figure from [121]

from sets of three space-points. A track seed is defined as a combination of space-points that provides an initial estimate of a particle's trajectory. Each space-point in a seed must be located in a different layer of the silicon detectors, specifically from the four layers of the Pixel Detector or the four double layers of the SCT. From these seeds, preliminary trajectories are calculated using the five track parameters illustrated in Figure 5.1. These parameters define the particle's trajectory at its perigee - the point of closest approach to the beam line - and include the transverse impact parameter d_0 , longitudinal impact parameter z_0 , azimuthal angle ϕ , polar angle θ , and the charge-to-momentum ratio q/p .

The track candidates are filtered based on their p_T and impact parameter values, as shown in Table 5.1. For promising seeds, a combinatorial Kalman filter identifies additional compatible space-points along the projected path, which are then added to the track [122].

Since all possible combinations of compatible space-points are used to create track candi-




Figure 5.2: Schematic of LRT and standard tracks in the ATLAS detector. [125]

dates, many candidates may share space-points. An ambiguity-solving algorithm resolves these overlaps by scoring tracks based on quality metrics, including the track χ^2 and the number of missing hits along the trajectory (referred to as “holes”) [123]. After this disambiguation, tracks are limited to at most one shared cluster, and any individual cluster can be shared by no more than two tracks.

In the final step, tracks are extended into the TRT, and compatible hits are incorporated into the track definition. This standard tracking approach achieves reconstruction efficiency exceeding 99% for prompt charged particles [124], while specialized algorithms are required for particles with significant displacement from the interaction point.

5.1.2 Large Radius Tracking

While standard tracking achieves high reconstruction efficiency for prompt charged particles, its performance decreases significantly for long-lived particles that may have large transverse impact parameters relative to the interaction point. To address this limitation, the Large Radius Tracking (LRT) algorithm was developed to reconstruct tracks with larger displacements [126].

LRT is performed after the standard tracking algorithm has completed and operates on silicon detector hits that were not previously associated with any reconstructed track, as illustrated in Figure 5.2. This complementary approach maximizes hit usage and track reconstruction efficiency. The primary difference between LRT and standard tracking is the relaxed geometric constraints applied to track candidates. As shown in Table 5.2, LRT permits significantly larger maximum values of the transverse and longitudinal impact parameters, $|d_0|$ and $|z_0|$, extending the reach to 300 mm and 1500 mm, respectively.

Reconstruction Step	Cut	Standard Tracking	Large-Radius Tracking
Silicon Seeded Tracks	Minimum p_T	500 MeV	900 MeV
	Maximum η	2.7	5.0
	Maximum $ d_0 $	10 mm	300 mm
	Maximum $ z_0 $	250 mm	1500 mm
Clustering	# of silicon hits, not shared	≥ 6	≥ 5
	# of shared hits	≤ 1	≤ 2
TRT-Seeded Tracks	p_T	≥ 1000 MeV	–
	$ d_0 $	≤ 100 mm	–

Table 5.2: Differences between track reconstruction criteria for standard tracking and large-radius tracking [127].

The reconstruction process in LRT begins with the same space-point formation as in standard tracking but uses a sequential Kalman filter instead of a combinatorial one [127]. This choice avoids the exponential growth in track candidates that would otherwise occur due to the larger search space.

After candidate tracks are formed, an ambiguity-solving algorithm is applied using the relaxed criteria shown in Table 5.2. Tracks that pass this step are extended into the TRT where possible, and compatible TRT hits are added. However, unlike in standard tracking, TRT extension is not required for LRT tracks since the TRT only covers $|\eta| < 2.0$.

Combining LRT and standard tracking substantially improves the overall track reconstruc-

tion efficiency. For particles with large displacements, the combined tracking efficiency can be several times higher than standard tracking alone at production radii greater than 50 mm [127].

5.1.3 *Vertexing*

With the track reconstruction techniques described in Chapter 5.1.1 providing trajectories of charged particles in the ID, the next step is to determine where these particles originated. Vertexing algorithms cluster tracks from charged particles that share a common origin point to form vertices. In ATLAS, a Primary Vertex (PV) represents the point where a proton-proton interaction occurred, while Secondary Vertices (SVs) indicate particle decays or interactions with detector material occurring away from the initial collision point.¹

PV reconstruction begins after standard track reconstruction is completed. Tracks meeting quality requirements (including hit multiplicity criteria to reduce contamination from poorly reconstructed tracks) are used to seed vertex candidates. For a track to be considered in PV seeding, it must have at least one hit in either the IBL or B-Layer of the Pixel Detector, since proton-proton interactions occur inside the beampipe.

The vertex finding uses an iterative procedure: after seeding, tracks are added to the vertex, and the position is refit. With each iteration, less compatible tracks are down-weighted to reduce their contribution to the fit. After the vertex position has been determined, tracks deemed incompatible with the vertex are removed and returned to the track collection for use in identifying additional PVs. This procedure continues until all selected tracks have been associated with a vertex or no additional vertices can be constructed from the remaining

1. Multiple primary vertices can exist in a single event due to pile-up interactions. The hard scattering primary vertex refers to the collision of interest, while additional primary vertices originate from simultaneous pp collisions in the same or nearby bunch crossings.

tracks. The hard-scattering PV of an event is identified as the vertex with the highest sum of the squared transverse momenta ($\sum p_T^2$) of its associated tracks, while other vertices are classified as pile-up vertices.

For long-lived particle searches, different algorithms are needed, as this primary vertexing approach is specifically designed to reject vertices with significant displacement from the interaction point [128].

5.1.4 *Secondary Vertexing Algorithms*

While the PV algorithm is designed to identify prompt interaction points near the beamline, additional algorithms are used to reconstruct SVs arising from long-lived particles' decays. These decays can occur at measurable distances from the primary vertex and are indicative of processes like heavy-flavor hadron decays or signatures of new physics. To reconstruct these vertices, two complementary algorithms are employed: the VrtSecInclusiv Vertexing Algorithm (VSI), which targets compact, point-like decays, and the Fuzzy Vertexing Algorithm (FV) algorithm developed by my collaborator Risa Ushida, which is optimized for reconstructing spatially extended decays typical of heavy-flavor topologies, like b-quarks [129].

5.1.4.1 **VrtSecInclusive (VSI) Algorithm**

The VSI algorithm reconstructs point-like vertices from both standard and LRT tracks. What follows is a brief overview of this complex algorithm; the complete details can be found in Ref. [130, 131].


SV reconstruction is computationally expensive, so the algorithm begins by identifying “selected tracks” that meet specific quality criteria to reduce the number of possible combinations for vertex seeding:

- $p_T > 1$ GeV
- Not associated with any primary or pile-up vertices
- $N_{\text{SCT}} \geq 6$ if $N_{\text{Pix}} = 0$
- $N_{\text{TRT}} \geq 7$ if $N_{\text{Pix}} < 2$
- $N_{\text{SCT}} \geq 7$ if $p_T > 25$ GeV
- $N_{\text{TRT}} \geq 20$ if $p_T < 25$ GeV and $|\eta| < 1.7$


Two-track seed vertices are then formed from pairs of selected tracks, if each track satisfies $|d_0| > 2$ mm. Candidate seed vertices then must satisfy $\chi^2/N_{\text{DoF}} < 5$ and radial position less than the SCT/ TRT boundary, $R_{xy} < 563$ mm. Hit pattern checks ensure tracks have hits in tracker layers beyond the vertex position and no hits in layers closer to the interaction point.

The algorithm then combines these two-track seeds into vertices with higher track multiplicity. This is done using an incompatibility graph approach. Think of each track as a node in a graph, where an edge between two tracks means they are incompatible (they don't form a good two-track vertex together). The algorithm identifies sets of tracks that are all compatible with each other (meaning no edges connect them in the graph) and combines them into a single n -track vertex; this is demonstrated in Figure 5.3. At this stage, a single track can be part of multiple vertex candidates.

To resolve these ambiguities, the algorithm assesses each track's contribution to the vertex quality. If a track is shared between multiple vertices, it is removed from the vertex where it contributes most to the χ^2 (indicating a poorer fit). Vertices with similar positions are



(a) Compatibility graph



(b) Incompatibility graph

Figure 5.3: Illustration of the incompatibility graph method used in the VSI algorithm to form n -track vertices. In this example with five tracks (a, b, c, d, e), three tracks (a, b, d) can form a common vertex. (a) The compatibility graph shows edges connecting compatible tracks. (b) The equivalent incompatibility graph shows edges between incompatible tracks. When nodes c and e (and their associated edges) are removed, nodes a, b, and d become isolated, indicating they are mutually compatible and can form a 3-track vertex. Figures from [131]

merged if their spatial separation significance is $S < 10$, where:

$$S = \sqrt{(\vec{v}_1 - \vec{v}_2)^T (\text{Cov}(\vec{v}_1) + \text{Cov}(\vec{v}_2))^{-1} (\vec{v}_1 - \vec{v}_2)} \quad (5.1)$$

where $\vec{v}_1$ and $\vec{v}_2$ are the vertex positions and Cov denotes their covariance matrices. Additionally, vertices within 1 mm of each other are merged, with a series of refits applied. This prevents what would physically be a single decay from being reconstructed as multiple separate vertices.

After seed merging, additional “associated tracks” that didn’t meet the initial strict selection criteria can be attached to existing vertices if they satisfy these looser requirements:

- $p_T > 1$ GeV

- $\chi^2/N_{\text{DoF}} < 20$
- Not associated with another DV
- Has a hit in the next adjacent outer tracker layer with respect to the DV
- Impact parameter significance to the vertex $\text{sig}(\Delta d_0) < 5$
- Longitudinal impact parameter significance $\text{sig}(\Delta z_0) < 5$

The final vertex position is refit with all tracks, and its four-momentum is calculated by summing the four-vectors of all constituent tracks, with each track assigned the charged pion mass. This comprehensive approach enables effective reconstruction of displaced vertices while maintaining strong discrimination against backgrounds.

5.1.4.2 Fuzzy Vertexing Algorithm

The VSI algorithm is highly inefficient at reconstructing decays of LLPs containing heavy-flavor quarks. Because the mesons formed by heavy-flavor quarks can have non-negligible lifetimes of their own, the tracks resulting from the original LLP decay will not necessarily appear to originate from a common point-like vertex. Instead, those tracks will appear to originate from a small volume surrounding the decay. The FV algorithm is designed to reconstruct decay topologies that are spatially extended, like those involving b -hadron decays where tracks do not converge to a single point [129]. Instead of fitting a single common vertex, FV builds a “fuzzy” cluster of related tracks that are statistically likely to originate from the same decay chain.

The algorithm, described in Ref. [132] proceeds through several stages:

1. **Track selection and seed formation:** Tracks with $p_T > 1$ GeV, at least two silicon

hits, and no association to a primary or pile-up vertex are selected. All pairs are evaluated using Boosted Decision Trees (BDTs), trained to classify track pairs based on kinematic and topological features.

2. **Seed identification:** Pairs that score above BDT thresholds ($s_{\text{middle}} > 0.15$ or $s_{\text{long}} > 0.15$) are selected as seeds. These BDTs are trained on simulated samples with varying lifetimes to optimize generalization across different decay distances.
3. **Seed merging:** Seed vertices are grouped into candidate FVs if they are spatially close and kinematically consistent. The merging tolerance increases with radial distance from the beamline (e.g., 1 mm, 3 mm, or 5 mm separation).
4. **Group merging:** Final FVs are formed by merging overlapping seed groups iteratively until convergence.

Unlike VSI, which fits a point-like vertex, FV constructs a cluster of tracks that define a vertex region. Figure 5.4 compares the two approaches: VSI attempts to trace back all tracks to a single point, while FV represents the vertex as a more general space that captures the more complex decay topology of the intermediate heavy-flavor hadrons.




Figure 5.4: Illustration of how an LLP ($\tilde{\chi}_2^0$) decaying via b -quarks would be reconstructed by the VSI and FV algorithms. Blue circles represent the reconstructed vertices. Figure courtesy of Jan Offerman

Both VSI and FV significantly expand the range of decay topologies that can be reconstructed, with each algorithm optimized for a different class of signatures. The complementary nature of these algorithms enhances the ability to capture displaced decays occurring at varying distances.

5.2 Object Definitions

Beyond tracks and vertices, several higher-level physics objects are reconstructed from detector signals to represent the particles produced in collisions. These reconstructed objects serve as the basis for interpreting proton-proton interactions in terms of the underlying physical processes. The reconstruction algorithms transform low-level detector responses into physics objects with well-defined kinematic properties.

The ATLAS detector is designed to identify a wide range of particles by exploiting their characteristic interactions with various subdetectors. Charged particles produce tracks in the ID; electrons and photons deposit energy primarily in the ECAL; hadrons create more extended showers in the HCAL; and muons traverse all detector layers and are measured in the muon spectrometer. By combining signals from these subdetectors, a complete picture of the event can be built.

This section outlines the reconstruction and identification procedures for jets, electrons, photons, and muons, including the algorithms used and the key selection criteria that define high-quality, well-measured objects. Procedures for resolving ambiguities between objects and suppressing overlaps are also discussed.

5.2.1 Jets

In high-energy particle collisions, quarks and gluons produced in the hard scattering process cannot be observed directly due to color confinement. Instead, they undergo hadronization, producing collimated sprays of hadrons known as jets [82]. A jet can be thought of as the visible “splash” in the detector resulting from an invisible quark or gluon. In the ATLAS detector, jets are reconstructed from clustered energy deposits in the calorimeters and, in some cases, associated tracks in the ID.

5.2.1.1 Jet Clustering Algorithm

At the core of jet reconstruction is the anti- k_t algorithm, which defines how individual particles or energy deposits are combined to form jets [133]. The algorithm works by calculating a distance measure between pairs of objects:

$$d_{ij} = \min(p_{T,i}^{-2}, p_{T,j}^{-2}) \frac{\Delta_{ij}^2}{R^2} \quad (5.2)$$

and the distance between an object i and the beam as

$$d_{iB} = p_{T,i}^{-2}, \quad (5.3)$$

where $\Delta_{ij}^2 = (y_i - y_j)^2 + (\phi_i - \phi_j)^2$, with $p_{T,i}$, y_i , and ϕ_i representing the transverse momentum, rapidity and azimuth of object i . The parameter R controls the size of the resulting jets (typically 0.4 for small radius jets and 1.0 for large radius jets).

For each object i , the algorithm sorts the distances to all other objects and to the beam. If the smallest distance is to another object, the two objects are clustered together, while if the smallest distance is to the beam, the object i is classified as a jet and removed from

the list of objects. After recalculating all distances, this procedure repeats until no objects remain. This approach produces regular, cone-like jets centered on the hardest (highest p_T) particles, making it less susceptible to soft radiation and pile-up effect [133].

5.2.1.2 Jet Types

Different jet collections are reconstructed depending on the input objects used by the anti- k_t algorithm [133]. The analysis presented in Chapter 6 makes use of both electromagnetic topological jet (EMTopo jet) jets and track jets.

EMTopo Jets EMTopo jets are reconstructed from topologically clustered calorimeter energy deposits, known as topological clusters (topo-clusters) [134]. A topo-cluster is a three-dimensional grouping of neighboring calorimeter cells that represent localized energy deposits. The clustering algorithm, often called the “4-2-0” scheme, begins with seed cells having $|E/\sigma| > 4$, adds neighbors with $|E/\sigma| > 2$, and then includes surrounding cells with any $E/\sigma > 0$. This efficiently reconstructs particle showers while suppressing calorimeter noise.

Jets are then formed by applying the anti- k_t algorithm (typically with $R = 0.4$) to the resulting topo-clusters. Due to fundamental differences in interaction mechanisms, electromagnetic and hadronic showers exhibit distinct energy deposition characteristics in the calorimeters. Hadronic showers lose energy to invisible processes, resulting in lower detector response than EM showers of equal energy. Since the particle type initiating each shower cannot be determined with certainty on a cluster-by-cluster basis, all energy deposits are initially calibrated at the EM energy scale. This assumption introduces a systematic bias in jet energy measurements, particularly for jets with significant hadronic content. Therefore, a comprehensive set of calibrations is applied to correct the energy measurement and reduce

the associated uncertainty [134].

Track Jets Track jets are reconstructed by applying the anti- k_t algorithm with $R = 0.4$ to charged-particle tracks in the ID [135]. They provide excellent pile-up suppression and vertex association, and are especially useful in regions with limited calorimeter coverage.

Only standard (non- LRT) tracks that pass strict quality criteria are used: $p_T > 1$ GeV, $|\eta| < 2.5$, $|d_0| < 2$ mm, at least one pixel hit, six silicon strip hits, and at least seven total silicon detector hits. Additional requirements on shared modules and holes help suppress fakes. Tracks are associated to a PV if they satisfy $|z_0/\sigma(z_0)| < 3$, or $|z_0| < 0.5$ mm if not. If associated with multiple vertices, the one with the highest $\sum p_T^2$ is chosen. Jets are retained only if $p_T > 10$ GeV, $|\eta| < 2.5$, and they contain at least two constituent tracks.

Compared to calorimeter jets, track jets offer better angular and momentum resolution for low- p_T charged particles but are limited to $|\eta| < 2.5$ and cannot reconstruct neutral particles. The analysis presented in Chapter 6 uses these track jets for background estimation due to their pile-up rejection properties.

Particle Flow Jets PFlow jets combine tracking and calorimeter information to improve resolution, particularly at low p_T . Charged-particle tracks are matched to calorimeter clusters, and their expected energy is subtracted from those clusters. The track momenta are then used instead, and the remaining neutral energy deposits are retained. The resulting set of Particle Flow objects is clustered using the anti- k_t algorithm. This hybrid method leverages the precision of tracking and the wide acceptance of calorimetry ($|\eta| < 4.9$), improving jet resolution in high-pile-up environments [116]. PFlow jets are used for the gFEX validation studies discussed in Chapter 4.5.

5.2.1.3 *b*-tagging

In most cases, jet reconstruction algorithms do not provide information about the type of particle that initiated the jet. However, jets originating from *b*-quarks can be identified due to the relatively long lifetime of *B*-hadrons, which travel several millimeters before decaying. These decays produce displaced vertices and tracks with large impact parameters.

Dedicated algorithms, referred to as *b*-taggers, combine tracking information to determine whether a jet is likely to have come from a *b*-quark [136]. These algorithms are used for the EMTopo jet employed in the background estimate analysis discussed in Chapter 6.

5.2.1.4 Jet Energy Scale and Resolution Uncertainties

Jets reconstructed in the detector must be accurately calibrated to relate their measured energy to the true particle-level jet energy. This calibration is encapsulated in the Jet Energy Scale (JES), while the precision of the measurement is quantified by the Jet Energy Resolution (JER) [134]. The JES defines the average response of reconstructed jets relative to the true energy, whereas the JER describes the spread of this response due to detector effects.

Both the JES and JER are subject to several sources of uncertainty that must be carefully evaluated and propagated to physics analyses. These uncertainties directly impact measurements involving jets and MET.

The JES uncertainty is determined from a combination of simulation-based methods and in-situ measurements. Significant contributions include uncertainties from in-situ calibration techniques, variations in flavor composition and jet response, pile-up effects, η intercalibration, and high- p_T extrapolation. The uncertainty varies with jet p_T and η , generally being

smaller for central jets with moderate transverse momentum and larger for forward jets or jets with very low or very high p_T [134].

5.2.2 *Electrons and Photons*

Electrons and photons are reconstructed from topo-clusters in the ECAL, as described in Chapter 5.2.1.2. These clusters are matched to tracks in the ID to distinguish between the two: electrons, being charged, typically leave tracks, while photons, being neutral, do not unless they undergo conversion to e^+e^- pairs in detector material. Clusters matched to tracks are identified as electrons, while unmatched clusters or those matched only to conversion vertices are classified as photons. For electrons, tracks are refit to account for bremsstrahlung energy loss; for photons, conversion vertices are reconstructed to correctly identify the original photon [137].

The energy of these objects is measured using calibrated windows in the second ECAL layer: a 3×7 cell area in the barrel region and a 5×5 cell area in the end-cap region. Energy scales are calibrated using $Z \rightarrow e^+e^-$ and $J/\psi \rightarrow e^+e^-$ events [137].

5.2.3 *Muons*

Muons are identified by their ability to traverse the whole detector with minimal energy loss, leaving tracks in both the ID and the MS. Due to their relatively long lifetime and lack of strong interactions, they typically escape the detector without undergoing hadronic interactions or depositing significant energy in the calorimeters [82].

Muon reconstruction begins by identifying track segments in individual stations of the MS, using information from the MDT, CSC, RPC, and TGC layers. Segments from different

layers are then combined to form track candidates pointing toward the interaction point. The trajectory through the magnetic field is fit using a global χ^2 fitter, which accounts for multiple scattering and energy loss [138]. If a track candidate in the MS can be matched to a corresponding track in the ID, a combined fit is performed using hits from both detectors to reconstruct a combined muon (CB).

If a matching ID track is available but does not pass the requirements for a full combined fit, the muon is classified as a *segment-tagged muon*, based on the match between the ID track and one or more MS segments. An alternative reconstruction strategy, referred to as inside-out combined muon (IO), begins with a high-quality ID track that is extrapolated outward to the MS and matched to at least three aligned hits. This approach is useful in regions with limited MS granularity or coverage.

In areas where tracking coverage is poor, additional muon types may be reconstructed. calorimeter-tagged muons are formed when an ID track is associated with a calorimeter energy deposit consistent with a minimum ionizing particle, typically within $|\eta| < 0.1$ and $15 < p_T < 100$ GeV. Extrapolated muons consist of standalone MS tracks that cannot be matched to any ID track but can be extrapolated back to the interaction region, providing low-efficiency coverage in otherwise inaccessible regions [138].

Muon identification is calibrated using muon-triggered data, typically from $Z \rightarrow \mu\mu$ events. Several working points (Loose, Medium, Tight, High- p_T) are defined based on track quality, detector matching, and isolation variables to suppress backgrounds from hadronic decays and misreconstructed tracks [138]. In the analysis described in Chapter 6, muons are primarily used for overlap removal, missing energy reconstruction, and validation of trigger efficiency.

5.3 Overlap Removal

The reconstruction of physics objects in ATLAS is performed independently by different algorithms, each optimized for specific signatures, such as tracks, energy deposits, or muon spectrometer hits. As a result, a single physical signature in the detector can sometimes be reconstructed as multiple distinct objects. For instance, a track might be simultaneously associated with both a muon and an electron candidate, or an energy cluster in the calorimeter could be classified as both a jet and a photon.

Overlap removal is the procedure used to resolve these ambiguities. It defines a hierarchy of object identification and removes duplicate reconstructions to ensure that each detector signature is assigned to only one physics object. This process serves several critical purposes: it prevents double-counting of energy deposits and tracks, improves the accuracy of derived quantities like MET, clarifies the true identity of overlapping detector signatures, and ultimately leads to cleaner event topologies for downstream analyses.

The procedure typically follows a predefined sequence in which objects are compared pairwise based on specific criteria, often using a spatial separation threshold like ΔR . When an overlap is identified, a decision is made to retain one object and discard the other based on factors like identification confidence, isolation, and physics priorities. The specific overlap removal strategy used in the analysis of this thesis is described in Chapter 6.

5.3.1 *Missing Transverse Energy*

The objects discussed previously all interact with the ATLAS detector in some way, which enables their reconstruction. However, some particles—such as neutrinos or hypothetical weakly interacting particles—pass through the detector without leaving a direct signal. Miss-

ing Transverse Energy (MET), E_T^{miss} , is a key observable that helps identify these invisible particles by measuring energy imbalance in the transverse plane.

Since the initial state of a proton-proton collision has approximately zero transverse momentum, any imbalance in the final-state visible momentum indicates the presence of undetected particles. The missing transverse momentum is defined as the negative sum of all reconstructed objects' transverse momenta:

$$\vec{E}_T^{\text{miss}} = - \left(\sum \vec{E}_T^{\text{jets}} + \sum \vec{E}_T^{\text{electrons}} + \sum \vec{E}_T^{\text{muons}} + \sum \vec{E}_T^{\text{photons}} + \sum \vec{E}_T^{\text{soft}} \right), \quad (5.4)$$

where $\sum E_T^{\text{soft}}$ accounts for low-momentum tracks and calorimeter deposits not associated with any high-level object. The scalar magnitude of this energy imbalance defines the missing transverse energy [82]:

$$\left| \vec{E}_T^{\text{miss}} \right| = \sqrt{(E_x^{\text{miss}})^2 + (E_y^{\text{miss}})^2} \quad (5.5)$$

Triggers based on MET cannot use the full E_T^{miss} calculation because it requires objects to be calibrated and undergo overlap removal. Four main algorithms were implemented in the HLT during Run 2 data collection, each offering different approaches to MET trigger-level calculation with varying trade-offs between accuracy and speed [139]:

- **cell:** This algorithm computes E_x^{miss} and E_y^{miss} by directly summing over all calorimeter cells that exceed noise thresholds. These thresholds are carefully configured based on the average pile-up conditions of the data-taking period and the electronic noise observed in individual cells. This approach is particularly sensitive to pile-up effects.

- **tc-lcw:** Topo-cluster algorithm first reconstructs three-dimensional clusters of topologically connected calorimeter cells with significant energy deposits. These topo-clusters undergo local cell weighting (lcw) calibration, which applies different weights to EM and hadronic energy deposits before calculating E_T^{miss} . This improves the energy resolution compared to the cell-based approach.
- **mht:** Missing Hadronic Transverse energy (MHT) algorithm calculates MET as the negative vector sum of all anti- k_t EMTopo jet with $p_T > 7$ GeV. These jets undergo pile-up subtraction and JES calibration before entering the MET calculation. This approach provides good pile-up resilience but misses energy from objects not clustered into jets.²
- **pufit:** pile-up fit (pufit) algorithm combines topo-clusters into patches corresponding to a jet of radius $R = 0.4$, then performs a statistical fit to estimate and subtract the contribution from pile-up. The algorithm makes the physics-motivated assumption that the vector sum of pile-up contributions to E_T^{miss} should be approximately zero. This method is best for high-pile-up conditions.

The performance of these algorithms varies with pile-up conditions, with pufit demonstrating the best rejection of fake E_T^{miss} at high luminosity. Notably, muons are not included in any trigger-level E_T^{miss} calculations, causing them to appear as invisible particles. This feature is exploited in evaluating E_T^{miss} trigger efficiencies using events selected with high-efficiency muon triggers.

2. This jet-based MHT algorithm is distinct from the gFEX MHT, which represents the hard term in the Jets without Jets algorithm and includes gTowers whose surrounding gBlock exceeds a seed threshold, as described in Chapter 4.2.3.

CHAPTER 6

DISPLACED VERTICES AND E_T^{miss} ANALYSIS

Many extensions of the SM predict the existence of LLPs that decay at measurable distances from the primary proton-proton collision point. Such decays can give rise to distinctive signatures within the ATLAS detector, including DVs formed by charged particle tracks that do not extrapolate back to the interaction point. In scenarios where LLPs decay into stable, weakly interacting particles, such as the axino, the SUSY partner of the axion, the resulting events are also characterized by significant Missing Transverse Energy.

The Displaced Vertices and E_T^{miss} (DV + MET) search targets events containing both a high-mass, high-track-multiplicity DVs and large MET, optimizing sensitivity to models where LLP decays are accompanied by non-interacting particles that escape the detector. SM backgrounds for such signatures are extremely low, as no known process yields DVs with these characteristics alongside large MET. However, non-collision backgrounds, detector effects, and rare misreconstruction processes can produce vertices that mimic the signal, necessitating a fully data-driven approach to background estimation.

This analysis is performed using the full Run 2 dataset collected by the ATLAS detector between 2016 and 2018, corresponding to an integrated luminosity of approximately 137 fb^{-1} at a center-of-mass energy of 13 TeV. To maximize reconstruction efficiency for displaced decays, several specialized techniques are employed, including LRT and the SV algorithms VSI and FV.

The overall search strategy utilizes a data-driven approach to background estimation, relying on photon-triggered control regions to model the correlation between track density and DV production. Signal Regions (SRs) are defined by stringent requirements on vertex mass and track multiplicity, with additional selections to reject vertices arising from material

interactions and detector artifacts. This approach allows for robust background modeling with minimal reliance on simulation.

While broadly sensitive to any new physics model predicting massive LLPs decaying within the ID volume in association with large MET, the analysis is interpreted in several BSM scenarios. The results presented in this thesis focus specifically on the DFSZ Axino interpretation, as motivated in Chapter 2.5. No significant excess is observed above the predicted background. The analysis establishes stringent constraints on long-lived Higgsinos decaying to axinos via Higgs- and Z-mediated channels. The analysis establishes stringent constraints on long-lived Higgsinos decaying to axinos via Higgs- and Z-mediated channels. The search excludes axino production for axion masses ranging from approximately $1\ \mu\text{eV}$ to $2.5\ \text{meV}$, where the analysis is most sensitive around $100\text{-}200\ \mu\text{eV}$ with cross-section limits reaching $0.1\text{-}0.2\ \text{fb}$. This represents the first experimental bounds on this mass range for DFSZ Axino models.

Overall, this chapter introduces several theoretical models that motivate displaced vertex signatures with missing transverse energy. While all of these signal models are considered in the analysis, the discussion throughout this chapter is framed primarily around the DFSZ Axino, which serves as the focus of my thesis. The sections that follow describe the MC simulation strategy, define the relevant SRs, and outline the analysis framework used to determine final results in the form of cross-section limits. Although the analysis includes multiple signal models, the presentation emphasizes those most relevant to the DFSZ Axino scenario, with others discussed where appropriate.

6.1 Motivation and Previous Searches

Building on the theoretical frameworks introduced in Chapter 2, this analysis targets SUSY models predicting LLPs that produce both DVs and significant MET. This signature arises when LLPs decay into a combination of visible and invisible particles – such as neutralinos or axinos – offering a powerful, background-suppressed way to search for new physics.

DVs are particularly attractive experimental signatures because SM particles with measurable lifetimes (such as K^0 , B -mesons, and Λ baryons) typically produce low-mass, low-multiplicity vertices that are easily distinguished from the high-mass, high-multiplicity DV topologies targeted in this analysis. This creates a naturally background-suppressed environment for new physics searches. The first ATLAS DV + MET search was performed using 32.8 fb^{-1} of 13 TeV proton-proton collision data collected in 2016 [140], targeting only long-lived gluino decays in a split- SUSY scenario with the standard VSI vertexing algorithm.

The first ATLAS DV + MET search used 32.8 fb^{-1} of 13 TeV proton-proton data collected in 2016 [140], targeting only long-lived gluinos reconstructed with the standard VSI vertexing algorithm. This analysis extends the previous DV + MET search in three key ways. First, we utilize the complete Run 2 dataset of 137 fb^{-1} , more than quadrupling the available statistics. Second, we employ a newly developed fuzzy vertexing algorithm, as discussed in Chapter 5.1.4.2, specifically designed to improve reconstruction efficiency for displaced decays involving heavy-flavor hadrons, which are challenging for the standard VSI algorithm. And third, we broaden the physics scope by introducing new signal model interpretations, including the first dedicated search optimized specifically for the DFSZ Axino signature.¹

1. A prior all-hadronic electroweakino search [68] included an axino interpretation among multiple signal models considered without dedicated MC simulation. This analysis represents the first search specifically designed and optimized for axino detection with a custom DFSZ Axino model implementation and full MC generation.

The Bino-Wino Coannihilation model is motivated by dark matter relic density constraints. In this scenario, long-lived neutralinos arise from pair-produced charginos and neutralinos as the NLSPs. These neutralinos decay via off-shell Higgs bosons to b -quark pairs, producing DVs with tracks that appear spread out due to the additional displacement from decaying b -hadrons. This scattered track pattern necessitates the fuzzy vertexing algorithm for efficient reconstruction.

The Higgs portal model connects the SM to a hidden sector via exotic Higgs decays, producing events with significant missing energy. In this case, the Higgs can decay into pairs of long-lived scalars, each decaying to $b\bar{b}$. These LLP decays often produce DVs with tracks that appear spread out due to the additional displacement from decaying b -hadrons, necessitating the use of fuzzy vertexing for efficient reconstruction. This process typically leads to multiple DVs per event, with the E_T^{miss} arising from associated production or undetected particles.

Finally, the DFSZ Axino model extends the MSSM with the axino as the LSP, as detailed in Chapter 2.5. Long-lived neutralinos decay to visible particles plus the escaping axino, producing both displaced vertices and MET. R-parity conservation (Chapter 2.4.1) ensures axino stability, and this analysis presents the first dedicated search for displaced axino production at ATLAS.

6.2 Analysis Strategy

This analysis employs a “cut-and-count” approach across several strategically defined SRs, each with event selection criteria optimized explicitly for the signal models under investigation. The analysis methodology prioritizes data-driven background estimation techniques to minimize reliance on simulation for background modeling, which can be challenging for these

DV signatures. The full analysis employs three distinct SRs: the standard displaced vertex (SDV) SR, targeting models such as the Gluino R -hadron and DFSZ Axino that typically produce one well-separated displaced vertex per event; and the fuzzy displaced vertex (FDV) SR, which is subdivided into two bins - the single fuzzy vertex (1 FDV) bin and the two fuzzy vertex (2 FDV) bin. The 1 FDV bin is optimized for the Bino-Wino Coannihilation model, where decays to b -quarks and subsequent secondary decays often smear the DV structure. In contrast, the 2 FDV bin is designed for the Higgs portal model featuring pairs of displaced decays producing b -quark final states.

It is important to note that the SDV SR and FDV SR are not explicitly made orthogonal to each other, as they use different vertexing algorithms and are analyzed independently. The 1 FDV and 2 FDV bins are treated as two bins of the same SR in the statistical analysis, with results statistically combined for the full FDV SR interpretation.

For background estimation, the analysis employs a data-driven approach that assumes backgrounds originate from three primary sources: Hadronic Interaction (HI), where particles interact with detector material; Accidental Crossing (AX), where unrelated tracks are incorrectly attached to vertices; and Merged Vertex (MV), where multiple vertices are erroneously combined into a single higher-mass, higher-multiplicity vertex.

Multiple complementary Control Regions (CRs) and Validation Regions (VRs) are established to support this approach, providing essential cross-checks of the background estimation methods. The background estimation is performed independently for each SR, with methods optimized for the specific characteristics of each region.

6.3 Event and Object Selection

This section describes the selection criteria applied at both the event and object levels in this analysis. Events are selected using MET triggers and must satisfy quality and kinematic requirements to suppress backgrounds from non-collision sources. Physics objects are reconstructed using the standard ATLAS algorithms described in Chapter 5, and are required to satisfy additional analysis-specific criteria optimized for the topology of this analysis.

6.3.1 Data Collection and Trigger Strategy

This analysis targets displaced vertex signatures in events with large missing transverse energy (MET), using data from Run 2 of the LHC and employing MET triggers as the primary handle for event collection.

6.3.1.1 Run 2 Data Sample

This analysis uses pp collision data recorded by ATLAS during the 2016–2018 data-taking periods, corresponding to an integrated luminosity of 137 fb^{-1} at $\sqrt{s} = 13 \text{ TeV}$.²

To ensure high-quality data, all events must pass the Good Runs List (GRL) requirement, which certifies that all relevant ATLAS detector subsystems were operational during data-taking. Events must also include a well-reconstructed PV with at least three associated tracks and longitudinal position $|z_{\text{PV}}| < 200 \text{ mm}$.

Standard data-quality cleaning procedures are applied to reject events affected by calorimeter

2. The 2015 dataset, totaling 3.2 fb^{-1} , is not included because the `DRAW_RPVLL` filter (explained in detail in Chapter 6.3.1.2) did not include MET triggers during that data-taking period. This results in a 2.3% reduction in the total integrated luminosity compared to the full Run 2 dataset.

noise bursts, readout corruption, or single-event upsets in the SCT. Events with high- p_T jets ($p_T > 50$ GeV) pointing to non-functional calorimeter regions are removed if $\Delta\phi(\text{MET}, \text{jet}) < 0.3$.

Additional background rejection based on non-collision activity is described in Chapter 6.4.

6.3.1.2 E_T^{miss} Trigger Selection

Events were collected using a set of E_T^{miss} triggers whose threshold increased throughout the data-taking period to accommodate increasing instantaneous luminosity. As discussed in Chapter 4.5, triggers with very low thresholds are typically prescaled — a technique where only one out of every N events passing the trigger is actually recorded — to keep data acquisition rates manageable. The lowest un-prescaled E_T^{miss} triggers included in the Derivation of the RPVLL format for Analysis Workflow (DRAW_RPVLL) filter for the 2016-2018 data-taking periods, listed in Table 6.1, represent the lowest E_T^{miss} threshold at which every single event meeting the criteria is saved, rather than just a fraction of them. These un-prescaled triggers effectively draw a line between where we collect only some events (lower thresholds) and where we collect all events, giving us complete statistics about events with this amount of online MET, while keeping our overall trigger rates within the allowed bandwidth. The efficiency of these triggers reaches 100% for events reconstructed with offline E_T^{miss} above 180 GeV.

The DRAW_RPVLL filter plays a crucial role in selecting events for specialized reconstruction. After an event passes a trigger, detector data are sent to the CERN Tier-0 computing center for reconstruction. While most events are reconstructed with standard algorithms, those passing the DRAW_RPVLL filter are reprocessed using the more computationally intensive techniques optimized for LLP searches, like LRT and both standard (VSI) and fuzzy displaced vertexing. To pass the filter, events must satisfy one of the MET triggers listed in Table 6.1

Year	Period	L (pb ⁻¹)	E_T^{miss} Trigger
2016	A-D3	6226.2	HLT_xe90_mht_L1XE50
	D4-	27176.0	HLT_xe110_mht_L1XE50
2017	B	5427.3	HLT_xe90_pufit_L1XE50 HLT_xe110_pufit_L1XE55
	C	2385.1	HLT_xe100_pufit_L1XE55 HLT_xe100_pufit_L1XE50 HLT_xe110_pufit_L1XE55
			HLT_xe110_pufit_L1XE55
			HLT_xe110_pufit_L1XE55
	D1-D5	5150.1	HLT_xe110_pufit_L1XE55 HLT_xe110_pufit_L1XE50
	D6-K	31668.1	HLT_xe110_pufit_L1XE50
			HLT_xe110_pufit_L1XE50
2018	B	3783.5	HLT_xe110_pufit_xe70_L1XE50 HLT_xe120_pufit_L1XE50
	C-J	17479.1	HLT_xe110_pufit_xe65_L1XE50 HLT_xe110_pufit_xe70_L1XE50 HLT_xe120_pufit_L1XE50
			HLT_xe110_pufit_xe65_L1XE50
			HLT_xe110_pufit_xe70_L1XE50
	K-	37529.0	HLT_xe110_pufit_xe65_L1XE50
			HLT_xe110_pufit_xe70_L1XE50

Table 6.1: Lowest un-prescaled MET triggers included in the `DRAW_RPVLL` filter, described in Chapter 6.3.1.2, by data-taking period. The integrated luminosity L in pb⁻¹ is shown for each period. In the trigger names, the last component denotes the L1 trigger used in the trigger chain. For example, a trigger name ending in `L1XE55` requires that the event pass the L1 trigger requiring 55 GeV of MET. The first component of the trigger name denotes the HLT trigger used in the trigger chain. These components are of the form `HLT_xe{X}_{Y}`, where $\{X\}$ is the minimum MET required by algorithm $\{Y\}$. Triggers containing an additional `xe{Z}` term are required to additionally have at least $\{Z\}$ GeV of MET as determined by the cell algorithm.

and have offline MET above 180 GeV as computed by the `LocHadTopo` (LHT) algorithm (`MET_LocHadTopo`), which uses locally calibrated hadronic topo-clusters to reconstruct the transverse momentum imbalance.

Events meeting these criteria undergo enhanced reconstruction and are stored in the `DRAW_RPVLL` format. For the final analysis, additional requirements are applied: standard offline MET must exceed 150 GeV, and `MET_LocHadTopo` must again be greater than 180 GeV. The latter assumes all calorimeter energy clusters originate from hadrons, thereby suppressing contri-

butions from identified leptons and photons.

Events that satisfy these E_T^{miss} -based trigger and offline selection criteria constitute the event pool called the MET-triggered dataset (MTR). This MTR serves as the foundation from which both our SRs and VRs are defined, with additional selections applied. The MTR requirements ensure that all events considered for the analysis have significant MET, which is a key signature for the models this analysis targets.

The analysis also utilizes supporting data samples selected with photon triggers that require a transverse momentum $p_T(\gamma) > 140$ GeV to validate the background estimation methodology. These events make up the Photon-Triggered Control Region (PTR) and populate orthogonal CRs with negligible signal contamination but similar kinematic properties to the SRs.

6.4 Non-Collision Background Veto

Non-Collision Backgrounds (NCBs) pose a significant challenge for this analysis, as they can produce apparent MET and fake DVs. To mitigate these backgrounds, a specialized NCB veto is implemented based on the properties of the leading jet in the event after overlap removal. This veto utilizes several jet observables to discriminate between collision-like and non-collision-like signatures:

- f_{max} : the maximum fraction of the jet energy deposited in any single layer of the calorimeter, which is sensitive to beam-halo and noise events that often show highly localized energy deposits
- f_{EM} : the fraction of the jet energy deposited in the electromagnetic calorimeter, common for noise bursts
- f_{ch} : the ratio of the scalar sum of p_T of tracks associated to the jet via ghost association

to the jet p_T

- $f_{\text{ch}}/f_{\text{max}}$: the ratio of charged track momentum to total jet momentum in the most energetic layer, which provides additional discrimination against jets arising from non-collision sources
- SuperLooseBadLLP: a loose jet cleaning working point designed to eliminate fake jets produced by calorimeter noise

Two versions of the NCB veto are applied depending on the SR, as summarized in Table 6.2. The standard NCB veto is used in the SDV SR, while the stricter version is applied in the FDV SRs to provide enhanced protection against non-collision backgrounds, which is crucial for maintaining control of backgrounds in the low-background search environment. Events with no jets automatically pass both veto variations. The combination of these vetoes significantly reduces the NCB contamination, as evidenced by the substantial suppression of the peaks at $\phi = 0$ and $\phi = \pm\pi$ in the MET ϕ distribution, while having minimal impact on potential signals.

Criterion	Standard NCB Veto	Strict NCB Veto
$f_{\text{max}} > 0.8$ (leading jet)	Yes (OR)	Yes (AND)
$f_{\text{EM}} > 0.96$ (leading jet)	Yes (OR)	Yes (AND)
$f_{\text{ch}}/f_{\text{max}} < 0.1$ (leading jet)	No	Yes (AND)
SuperLooseBadLLP fail (any jet)	Yes (OR)	Yes (OR)

Table 6.2: Standard and strict NCB veto definitions.

6.5 Displaced Vertex Algorithms and Track Selections

This analysis targets events containing both DVs and significant MET. DVs are the key experimental signature of LLPs that decay at macroscopic distances within the ID.

Standard tracking algorithms, discussed in Chapter 5.1.1 are optimized for prompt particles from the primary IP, making them inefficient for reconstructing the large impact parameter tracks characteristic of LLP decays. To address this, two complementary vertex reconstruction algorithms are employed in this analysis, supported by dedicated track cleaning and baseline selection procedures to ensure the quality of the vertex.

6.5.1 *Vertex Reconstruction Algorithms*

This analysis employs both the SDV and FDV algorithms to capture a broad range of LLP topologies. While the primary focus is on the SDV algorithm, the FDV algorithm provides complementary sensitivity, particularly for decays involving heavy-flavor hadrons. Detailed descriptions of both algorithms are provided in Sections 5.1.4.1 and 5.1.4.2.

- The SDV algorithm excels at reconstructing well-defined vertices with tracks originating from a common point, making it particularly effective for the Gluino R -hadron model and for the DFSZ Axino model, where decay products emerge from a single point.
- The FDV algorithm is optimized for LLP decays involving heavy-flavor hadrons (particularly $b\bar{b}$ b -hadrons), where secondary decays can create a more diffuse vertex structure. This algorithm employs a BDT to evaluate tracks for their likelihood of originating from an LLP and performs better for the Bino-Wino Coannihilation and Higgs portal models.

6.5.2 Track Cleaning Procedures

A crucial component of DV reconstruction is the application of track cleaning procedures, which remove poorly reconstructed tracks that could form fake vertices. After DV reconstruction, these cleaning selections are applied to all tracks associated with each candidate vertex, as summarized in Table 6.3.

Criteria	SDV Track Cleaning	FDV Track Cleaning
p_T threshold (min)	2 GeV (all), 3–4 GeV (attached)	> 1.1 GeV
$ d_0 /\sigma(d_0)$ threshold	10–15	> 5
Upstream hit veto	No hits with $r < r_{DV}$	No hits with $r < r_{DV}$
Angle α cut	> 0.02 (all), > 0.2 (specific)	> 0.15 (varies with r)
ϕ_{PV-DV} cuts	< 3 if $d_0 < 1$ mm	< 1.75 or < 1 (varies with r)

Table 6.3: Track cleaning selections for SDV and FDV algorithms.

The SDV algorithm applies stricter p_T and impact parameter significance criteria than the FDV algorithm, reflecting their different sensitivities to track properties. Both approaches implement an upstream hit veto and place requirements on the angle between tracks and the vector from the primary vertex to the displaced vertex.

6.5.3 Baseline Vertex Selections

After track cleaning, baseline quality criteria are applied to all candidate vertices, as detailed in Table 6.4.

These baseline requirements ensure that candidate vertices are well-reconstructed and located within the active tracking volume. The material veto, implemented as a cylindrical histogram of material-rich regions, is particularly important for suppressing backgrounds from HIs with detector components.

Requirement	SDV Baseline	FDV Baseline
Fiducial volume	$R_{xy} < 300 \text{ mm}, z < 300 \text{ mm}$	$R_{xy} < 300 \text{ mm}, z < 300 \text{ mm}$
Transverse DV-PV separation	$R_{\text{DV-PV}}^{xy} > 4 \text{ mm}$	$R_{\text{DV-PV}}^{xy} > 4 \text{ mm}$
Vertex fit quality	$\chi^2/N_{\text{DoF}} < 5$	–
Photon veto	$\Delta R(\text{DV}, \gamma) > 0.1$	$\Delta R(\text{DV}, \gamma) > 0.1$
Material veto	Strict veto	Strict veto
Max p_T fraction	$< 0.95 \times \text{total } p_T$	$< 0.9 \times \text{total } p_T$

Table 6.4: Baseline displaced vertex selection requirements.

For SDV candidates, the vertex fit quality requirement ensures that the fitted tracks are consistent with a common spatial origin. The photon veto removes vertices that may arise from photon conversions in detector material. The maximum p_T fraction requirement helps eliminate vertices dominated by a single high- p_T track, which often indicates a mismeasurement.

After applying these baseline selections and track cleaning procedures, vertex properties, including mass and track multiplicity, are recalculated. These refined vertex candidates form the foundation for the subsequent event selection and signal region definitions that will be detailed in later sections.

6.5.4 Standard Physics Objects

Baseline physics objects are reconstructed using the algorithms described in Chapter 5, with additional selection criteria as summarized in Tables 6.5–6.8. These objects are primarily used in the E_T^{miss} calculation and the overlap removal procedure.

Electrons are reconstructed from energy deposits in the ECAL matched to tracks in the ID. Baseline electrons are required to have $p_T > 10 \text{ GeV}$, $|\eta| < 2.47$, and satisfy the `LooseAndBLayerLLH` identification criteria. This likelihood-based discriminant combines

calorimeter shower shapes and tracking information to provide high efficiency ($\sim 95\%$) while requiring a hit in the innermost pixel layer to reject photon conversions. No explicit veto is applied to the calorimeter crack region ($1.37 < |\eta| < 1.52$).

Muons are reconstructed from tracks in both the ID and MS. Baseline muons must have $p_T > 10$ GeV, $|\eta| < 2.7$, and satisfy the Medium identification working point. This means that tracks are required to have minimum hit multiplicities in both detector systems, achieving $\sim 97\%$ efficiency while maintaining good background rejection and minimizing systematic uncertainties. No track-to-vertex association requirement is imposed.

Photons are identified as electromagnetic calorimeter energy deposits without matching tracks. Baseline photons are required to have $p_T > 25$ GeV, $|\eta| < 2.37$, and satisfy the Tight identification working point. This strict selection applies stringent requirements on electromagnetic shower shapes across multiple calorimeter layers, distinguishing prompt photons from neutral meson decays and hadronic showers. Additionally, looser photons with $p_T > 60$ GeV are used for displaced vertex veto purposes.

Jets are reconstructed using the anti- k_t algorithm with radius parameter $R = 0.4$, as described in Chapter 5.2.1. Baseline jets must have $p_T > 20$ GeV and $|\eta| < 4.5$. To suppress pile-up contributions, jets with $|\eta| < 2.5$ must satisfy the Medium Jet Vertex Tagger (JVT) working point (value > 0.64), while forward jets ($2.5 < |\eta| < 4.5$) must satisfy the Tight forward JVT working point (value > 0.4). Events containing jets that fail the **SuperLooseBadLLP** quality criteria are rejected to remove non-collision backgrounds from beam-induced effects and detector noise that could mimic displaced signals.

Requirement	Value
Transverse momentum p_T	> 10 GeV
Pseudorapidity $ \eta $	< 2.47
Identification	LooseAndBLayerLLH
Crack veto ($1.37 < \eta < 1.52$)	Not applied
Usage	Overlap removal, E_T^{miss} calculation

Table 6.5: Electron selection requirements.

Requirement	Value
Transverse momentum p_T	> 10 GeV
Pseudorapidity $ \eta $	< 2.7
Identification	Medium
Track-to-vertex association	Not required
Usage	Overlap removal, E_T^{miss} calculation

Table 6.6: Muon selection requirements.

Requirement	Value
Transverse momentum p_T	> 25 GeV
Pseudorapidity $ \eta $	< 2.37
Identification	Tight
Additional photons for DV vetoes	$p_T > 60$ GeV, no ID requirement
Usage	Overlap removal, E_T^{miss} calculation

Table 6.7: Photon selection requirements.

Requirement	Value
Transverse momentum p_T	> 20 GeV
Pseudorapidity $ \eta $	< 4.5
JVT (for $ \eta < 2.5$)	> 0.64 (Medium WP)
fJVT (for $2.5 < \eta < 4.5$)	> 0.4 (Tight WP)
SuperLooseBadLLP cleaning	Applied

Table 6.8: Jet selection requirements.

6.5.4.1 Overlap Removal

To resolve ambiguities when reconstructed objects share detector signatures, overlap removal is applied following the general procedure described in Chapter 5.3. The specific sequence used in this analysis is:

1. If an electron and a photon share the same cluster, the photon is removed.
2. If an electron and a muon share the same track, the electron is removed.
3. If a jet is within $\Delta R < 0.2$ of an electron, the jet is removed; otherwise, if an electron is within $\Delta R < 0.4$ of a jet, the electron is removed.
4. If a jet is within $\Delta R < 0.2$ of a muon and has fewer than three associated tracks, the jet is removed; otherwise, if a muon is within $\Delta R < 0.4$ of a jet, the muon is removed.
5. If a jet is within $\Delta R < 0.4$ of a photon, the jet is removed.

6.6 Material Map

A significant background for DV searches arises from HIs between SM particles and detector material. When charged hadrons scatter off dense layers in the ID, they can produce low-mass, collimated vertices that mimic LLP decays. These hadronic interaction vertices frequently occur in high-density regions such as the pixel support structures or the SCT service layers.

To suppress this background, a *material veto* is applied using a cylindrical histogram map of the ID material distribution, as shown in Figure 6.1. This map is derived from the spatial distribution of low-mass vertices in data, with bins classified as either “inside” or “outside” material based on the density of reconstructed low-mass vertices.




Figure 6.1: Map showing a slice at $z = 0$ of the physical material comprising the ATLAS detector's beam pipe and tracking detector, used to define the material veto. Figure from [141]

For Monte Carlo simulations, the same map is applied with alignment corrections to account for slight x - y shifts in pixel layers due to their attachment to the beam pipe. Two versions of the veto are used. The *loose veto* excludes displaced vertices whose reconstructed positions fall within bins tagged as material-rich. The more stringent *strict veto* additionally excludes vertices in bins adjacent to material regions, as well as vertices whose position uncertainty ellipsoids intersect any material-containing bins. These definitions help account for boundary effects and uncertainties in reconstruction.

The material veto plays a dual role in the analysis. In signal regions, it suppresses background by removing vertices in known material-rich zones. When inverted—selecting only vertices inside the material—it defines validation regions enriched in hadronic interactions, providing essential control samples for background estimation.

6.7 Signal Models and Monte Carlo Simulation

While this analysis is experimentally defined by the signature of DVs and large MET, it is theoretically motivated by several BSM models that predict LLPs decaying within the ATLAS detector. These models span a diverse range of production and decay mechanisms, and each maps to distinct regions of phase space, allowing the analysis to remain broadly sensitive to various classes of new physics. Four signal models are considered: the Gluino R -hadron, the Bino-Wino Coannihilation model, the Higgs portal model, and the DFSZ Axino model, which forms the primary motivation for this search.

6.7.1 Gluino R -hadron

The Gluino R -hadron model arises in split-SUSY scenarios [142, 143], where the supersymmetric partners of the SM particles have a hierarchical mass spectrum. In this framework,

the gluino ($\tilde{g}$) is the lightest colored superpartner and can become long-lived due to suppressed decay through off-shell squarks, which are assumed to be very heavy. The gluinos hadronize shortly after production into bound states known as R -hadrons, as illustrated in Figure 6.2. These R -hadrons then decay to two SM quarks and the lightest neutralino ($\tilde{\chi}_1^0$), assumed to be stable.




Figure 6.2: Feynman diagram for the Gluino R -hadron model. Figure from [141].

The Gluino R -hadron signal grid explores the parameter space summarized in Table 6.9, with gluino masses between 400 and 3000 GeV, either fixed $\tilde{\chi}_1^0$ mass at 100 GeV or with mass splittings $\Delta m_{\tilde{g}, \tilde{\chi}_1^0}$ ranging from 10 to 100 GeV, and gluino lifetimes between 0.01 and 30 ns.

Parameter	Range
Gluino mass $m_{\tilde{g}}$	400–3000 GeV
Neutralino mass $m_{\tilde{\chi}_1^0}$	100 GeV (fixed) or $\Delta m = 10$ –100 GeV
Gluino lifetime $\tau_{\tilde{g}}$	0.01–30 ns

Table 6.9: Parameter space explored for the Gluino R -hadron signal model.

6.7.2 Wino-Bino co-annihilation

In this signal model, we consider the co-annihilation of Wino- and Bino-like neutralinos, as introduced in Chapter 2.4.1. In SUSY, “wino” refers to the superpartner of the W boson, while “bino” is the superpartner of the B boson, which mix to form the physical neutralino mass eigenstates we observe. Specifically, we examine neutralino-chargino production: $pp \rightarrow \tilde{\chi}_1^\pm \tilde{\chi}_2^0$, as shown in Figure 6.3. This neutralino mixing is detailed in Chapter 2.4.1.

The chargino $\tilde{\chi}_1^\pm$ (a wino-like particle) promptly decays to the Bino-like $\tilde{\chi}_1^0$ through off-shell $W^\pm$ emission. The neutralino $\tilde{\chi}_2^0$ (also wino-like) is long-lived and also decays to $\tilde{\chi}_1^0$ via off-shell Higgs or Z -boson emission. We target the case where the Higgs boson decays to a bottom quark-antiquark pair [144, 145]. The resulting final state contains displaced b -jets and significant MET from the undetected $\tilde{\chi}_1^0$.




Figure 6.3: Feynman diagram for the Wino-Bino co-annihilation model. Figure from [141].

The signal grid explores the parameter space detailed in Table 6.10, scanning $\tilde{\chi}_2^0$ masses from 200 to 1100 GeV, with mass splittings of 20–50 GeV relative to $\tilde{\chi}_1^0$, and neutralino lifetimes of 0.01–3 ns.

Parameter	Range
$\tilde{\chi}_2^0$ mass	200–1100 GeV
Mass splitting $\Delta m_{\tilde{\chi}_2^0, \tilde{\chi}_1^0}$	20–50 GeV
$\tilde{\chi}_2^0$ lifetime	0.01–3 ns

Table 6.10: Parameter space explored for the Wino-Bino co-annihilation signal model.

6.7.3 Higgs portal

The Higgs portal scenario extends the SM by allowing the Higgs boson to decay into a pair of scalar LLPs (S), as illustrated in Figure 6.4 [146]. These scalars decay into $b\bar{b}$ pairs at a macroscopic distance from the primary interaction point, leading to displaced vertices. The associated MET arises from vector bosons or top quarks produced in association with the Higgs boson and decaying leptonically [147].




Figure 6.4: Feynman diagram for the Higgs portal model. Figure from [141].

Parameter	Range
Scalar mass (m_S)	5–55 GeV
Scalar lifetime (τ)	0.0033–3.3 ns

Table 6.11: Parameter space explored for the Higgs portal signal model.

The signal grid explores the parameter space summarized in Table 6.11, including scalar

masses between 5 and 55 GeV and scalar proper lifetimes ranging from 0.0033 to 3.3 ns.

Parameter	Range
$\tilde{\chi}_1^0$ mass	200–1000 GeV
$\tilde{a}$ mass	0–500 GeV
Axion decay constant f_a	$5 \times 10^9, 5 \times 10^{10}$ GeV
$\tilde{\chi}_1^0$ lifetime	0.00058–6 ns

Table 6.12: Parameter space explored for the DFSZ Axino signal model.

6.7.4 DFSZ Axino

The DFSZ Axino signal provides the primary physics motivation for this thesis, building on the theoretical framework detailed in Chapter 2.5. This model incorporates the DFSZ axion into a SUSY framework [148, 149, 150, 151], where supersymmetry provides the axion with a fermionic superpartner, the axino ($\tilde{a}$), which serves as the LSP. The MSSM is extended to include the axino as a fifth neutralino mass eigenstate ($\tilde{\chi}_5^0$), realized through perturbative mixing of up- and down-type Higgsinos [64]. In this setup, the lightest Higgsino-like neutralino ($\tilde{\chi}_1^0$) becomes long-lived and decays to the axino plus either a Higgs or Z boson, both of which decay hadronically, as illustrated in Figure 6.5.³

The signal grid explores the parameter space summarized in Table 6.13, spanning $\tilde{\chi}_1^0$ masses from 200 to 1000 GeV, axino masses from 0 to 500 GeV, and two values of the axion decay constant f_a , resulting in neutralino lifetimes ranging from 0.00058 to 6 ns.

3. Although it is conventional for $\tilde{\chi}_1^0$ to denote the lightest neutralino in the model, we use $\tilde{a}$ for the new, mostly-axino neutralino to avoid confusion with the standard MSSM neutralinos.




Figure 6.5: Feynman diagram for the DFSZ Axino model. Figure from [141].

Parameter	Range
$\tilde{\chi}_1^0$ mass	200–1000 GeV
$\tilde{a}$ mass	0–500 GeV
Axion decay constant f_a	$5 \times 10^9, 5 \times 10^{10}$ GeV
$\tilde{\chi}_1^0$ lifetime	0.00058–6 ns

Table 6.13: Parameter space explored for the DFSZ Axino signal model.

6.7.5 Monte Carlo Generation

MC simulations are used to model both the expected signal processes and relevant Standard Model backgrounds. The signal processes described in Sections 6.7.1- 6.7.4 are generated using either MADGRAPH (Gluino R -hadron , Bino-Wino Coannihilation, DFSZ Axino) or POWHEG (Higgs portal), with the decays of the primary particles provided by MADSPIN (Bino-Wino Coannihilation, DFSZ Axino), or by PYTHIA (Gluino R -hadron , Higgs portal). Finally, PYTHIA is used to simulate the parton showering and hadronization for all of the signal processes. In each case, the LLP decays are characterized by kinematic and lifetime distributions dictated by the underlying BSM models described in the relevant sections. Background processes, including $Z \rightarrow \nu\nu + \text{jets}$ and $\gamma + \text{jets}$, are also simulated in order to study the MET trigger efficiency turn-on curve and background estimation. All generated events for both signal and background processes are passed through a detailed GEANT4 simulation, which models the interactions of particles with the ATLAS detector, including the propagation of R -hadrons and realistic collision conditions with pile-up representative of each run period. The simulated events then undergo standard ATLAS digitization, reconstruction, and an additional tracking step, with two secondary vertexing algorithms applied to identify both SDVs and FDVs (see Chapter 5).

For all signal and background models, the simulation includes realistic collision conditions with pile-up profiles matching the 2016-2018 data-taking periods. The detailed physics descriptions of each model are provided in the previous sections, while Table 6.14 summarizes the technical details of Monte Carlo generation.

For all signal and background models, the simulation includes realistic collision conditions with pile-up profiles matching the 2016-2018 data-taking periods. The detailed physics descriptions of each model are provided in the previous sections, while Table 6.14 summarizes

Model	Generator	Decay Simulation	Cross-section order	Mass / Lifetime Ranges
Gluino R -hadron Model	MADGRAPH 2.6.2 (A14 tune, NNPDF2.3lo)	PYTHIA 8.24	NNLO+NNLL	$m_{\tilde{g}}$: 400–3000 GeV Δm : 10–2900 GeV τ : 0.01–30 ns
Bino-Wino Coannihilation Model	MADGRAPH 2.9.5 (A14 tune, NNPDF23LO)	MADSPIN + PYTHIA 8.212	NLO+NLL	$m_{\tilde{\chi}_2^0}$: 100–1100 GeV Δm : 20–50 GeV τ : 0.01–3 ns
Higgs portal Model	POWHEG 2 (A14 tune, NNPDF23LO)	PYTHIA 8.307/8.212	NLO	m_S : 5–55 GeV τ : 0.0033–3.3 ns
DFSZ Axino Model	MADGRAPH 2.9.16 (A14 tune, NNPDF23LO)	MADSPIN + PYTHIA 8.307	NLO	$m_{\tilde{a}}$: 0–500 GeV $m_{\tilde{\chi}_1^0}$: 200–1000 GeV m_a : 56–5690 μeV τ : 0.00058–6.0 ns
$Z \rightarrow \nu\nu + \text{jets}$	SHERPA 2.2.11	SHERPA 2.2.11	NLO	–
$Z \rightarrow \mu\mu + \text{jets}$	SHERPA 2.2.11	SHERPA 2.2.11	NLO	–
$\gamma + \text{jets}$	SHERPA 2.2.2	SHERPA 2.2.2	NLO	–

Table 6.14: Summary of MC generation details for signal and background samples.

the technical details of Monte Carlo generation.

6.7.5.1 Gluino R -hadron MC Production

The matrix element calculation for gluino pair-production was performed with MadGraph 2.6.2 at leading-order precision [152]. PYTHIA 8.24 was then used to simulate the decays of the R -hadrons, the parton showering, and the underlying event [153]. The A14 set of tuned underlying event and shower parameters [154] was used with the NNPDF2.3LO [155] parton distribution function set to generate the R -hadron samples. The gluino pair-production cross sections and their uncertainties have been calculated at approximate Next-to-Next-to-Leading Order (NNLO) with Next-to-Next-to-Leading Logarithmic (NNLL) accuracy in

accordance with the PDF4LHC15 recommendations [156, 157]. The gluinos in the simulated samples have mass $m_{\tilde{g}}$ between 400–3000 GeV, mean proper lifetime τ between 0.01–30 ns, and neutralino mass $m_{\tilde{\chi}_1^0}$ ranging from 100 GeV up to $m_{\tilde{g}} - 10$ GeV.

6.7.5.2 Bino-Wino Coannihilation MC Production

The matrix element for the pair-production of wino-like electroweakinos $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ is calculated using MadGraph 2.9.5 at the Next-to-Leading Order (NLO) with Next-to-Leading Logarithmic (NLL) order accuracy. The decay of $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^0$ are simulated using MADSPIN [158]. The events are then interfaced with PYTHIA 8.212 to model parton showering, hadronization, and the underlying events with A14 tune and NNPDF23LO [155] Parton Distribution Function (PDF) sets. The branching ratios $\tilde{\chi}_1^\pm \rightarrow \tilde{\chi}_1^0 W^*$ and $\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 h^*$ are set to 100%. For off-shell Higgs boson with mass between 20–50 GeV, the branching ratio $h^* \rightarrow b\bar{b}$ is theoretically predicted to be approximately 90% [159], and is set to 100% in this analysis for simplicity. The simulated samples have $m_{\tilde{\chi}_2^0}$ between 100–1100 GeV, τ between 0.01–3 ns, and $\Delta m_{\tilde{\chi}_2^0, \tilde{\chi}_1^0}$ between 20–50 GeV.

6.7.5.3 Higgs portal MC Production

MC samples were generated for Higgs boson production associated with vector bosons (VH), top-quark pairs ($t\bar{t}H$), and via gluon-gluon fusion (ggH) to support the Higgs portal search. W and Z bosons from VH production and from top quark decays can produce neutrinos in the event, contributing to MET. Due to this, the $t\bar{t}H$ mode dominates sensitivity at most lifetimes, with the higher cross-section ggH mode only dominating at the higher lifetimes when there is a higher chance of LLPs escaping detection. VH and ggH samples were produced using POWHEG 2 (revision number 3133) and PYTHIA 8.212 at next-to-leading

order and $t\bar{t}H$ samples at leading order using POWHEG 2 (revision number 3133). For the Higgs portal process, PYTHIA 8.307 was used to decay the Higgs bosons from these samples into a pair of scalar LLPs which subsequently decay to $b\bar{b}$. Samples were generated with S masses between 5–55 GeV and mean proper lifetime $c\tau$ between 0.003–3 ns.

6.7.5.4 DFSZ Axino MC Production

Pair-production of neutralinos $\tilde{\chi}_1^0 \tilde{\chi}_2^0$ is generated with MadGraph 2.9.16, with prompt $\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 Z^*(\nu\bar{\nu})$ decay where $\Delta m = 2.5$ GeV between the neutralinos. We assume the neutralinos $\tilde{\chi}_1^0$ and $\tilde{\chi}_2^0$ to be purely Higgsino states, and neutralinos $\tilde{\chi}_3^0$ and $\tilde{\chi}_4^0$ to be purely Wino and Bino states.

While this analysis is not the first to feature an axino interpretation, it breaks new ground by developing a fully custom MC simulation of the DFSZ Axino signal. Working closely with phenomenologists at the University of Chicago, including Gabriel Hoshino and Keisuke Harigaya, we constructed the DFSZ Axino model Lagrangian and implemented it using the SARAH framework. Model implementation was led by Jan Offermann, who derived the axino interactions and generated the corresponding UFO model files.

The $\tilde{\chi}_1^0$ is the LLP in this model, and its mean lifetime is determined by the axion decay parameter f_a – which is related to the axion and neutralino masses, but not the axino mass (which is a free parameter). The primary neutralinos in the simulated samples have masses between 200–1000 GeV, whereas the axino mass ranges from massless up to 500 GeV. For the QCD axion model, the axion mass, m_a , is determined from the axion decay parameter f_a , for which two values are adopted in this analysis: $f_a = 5 \times 10^9$ GeV, 5×10^{10} GeV. The corresponding axion masses are thus $m_a = 113 \mu\text{eV}$, $1138 \mu\text{eV}$.

We consider the cases of subsequent $\tilde{\chi}_1^0 \rightarrow \tilde{a}h$ and $\tilde{\chi}_1^0 \rightarrow \tilde{a}Z$ decays, which are performed

with MadSpin. For efficiency during generation, the $h \rightarrow b\bar{b}$ and $Z \rightarrow q\bar{q}$ branching ratios are set to 100%, and later rescaled by their respective SM ratios (approximately 58% and 70% respectively). Parton showering, hadronization, and underlying event generation are subsequently handled by PYTHIA 8.307, using the A14 tune and NNPDF23LO parton distribution function set.

This model underwent several rounds of validation to ensure consistency with theoretical expectations. Kinematic distributions such as p_T , η , and MET were checked, and the distribution of DVs was confirmed to match the expected lifetimes. A handful of these kinematic distributions can be seen in Appendix C. Once validated, the model was integrated with the GEANT4 detector simulation and successfully added to the central ATLAS model repository. This “theory to code” approach, from first principles through to detector-level simulation, provides unprecedented control over the modeling of exotic signatures and represents the first implementation of an axino model specifically designed to study displaced signatures in a collider experiment.

6.8 Signal Regions

This analysis defines two primary SRs based on the two different vertexing algorithms: the SDV SR, using the VSI vertexing algorithm, and the FDV SR, using the FDV algorithm. The FDV SR is further divided into two orthogonal bins: the single fuzzy vertex bin (*1 FDV*) targeting models with exactly one LLP decay (Bino-Wino Coannihilation), and the two fuzzy vertices bin (*2 FDV*) designed for models involving pair production of LLPs decaying to two or more distinct FDVs (Higgs portal). These bins are analyzed independently, and their results are statistically combined for the final FDV SR result.

6.8.1 Signal Region Selection Criteria

After passing baseline vertex selection criteria, candidates DVs must meet additional selections to satisfy the SRs, as summarized in Table 6.15. For each SR, specific requirements are applied to the following variables:

- Scalar p_T sum: Maximum fraction of total vertex momentum carried by the highest- p_T track ($p_T^{\max}/\Sigma|p_T| < 0.9$).
- Track multiplicity ($N_{\text{Tracks}}^{\text{FDV}}$): Minimum number of tracks associated with the vertex ($N_{\text{tracks}} \geq 5$).
- Invariant mass (m_{FDV}): Minimum invariant mass of the displaced vertex ($m_{\text{DV}} > 10$ GeV).
- Track pseudorapidity spread: Maximum difference in pseudorapidity between any single track and the sum of the other tracks ($\max(\Delta\eta(p_{\text{trk}}, p_{\text{tot}} - p_{\text{trk}})) < 3.5$ for 1 FDV only).
- Number of two-track seeds ($N_{\text{Seeds}}^{\text{FDV}}$): For 1 FDV, requires at least $N_{\text{tracks}}C_2 - 1$ seed vertices⁴. For 2+ FDV, requires at least $N_{\text{tracks}}^{\text{FDV}}$ seed vertices.
- Jet separation ($\Delta R_{\text{FDV, jet}}$): Maximum angular separation between the FDV and closest jet (for 2+ FDV only, < 0.4).
- Pseudorapidity separation ($|\eta_{\text{PV-FDV}}|$): Maximum pseudorapidity separation between the primary vertex and the FDV (for 2+ FDV only, < 2.5).
- Number of vertices ($N_{\text{DV}_{\text{sat}}}$): Minimum number of vertices satisfying the criteria.

4. Where C_2 represents the number of ways to choose 2 tracks from N_{tracks} tracks: $C_2 = \frac{N_{\text{tracks}}!}{(N_{\text{tracks}}-2)!2!}$.

For the 1 FDV selection, it is additionally required that there are no other FDVs satisfying the 2+ FDV selections, ensuring orthogonality between the SRs.

Full SR Vertex Selection	SDV SR	FDV SR	
		1 FDV	2+ FDV
Vertices	SDV	FDV	FDV
Scalar p_T sum	$p_T^{\max}/\Sigma p_T < 0.9$	$p_T^{\max}/\Sigma p_T < 0.9$	$p_T^{\max}/\Sigma p_T < 0.9$
N_{tracks}	≥ 5	≥ 5	≥ 4
m_{DV} [GeV]	> 10	> 10	> 1.5
$\max(\Delta\eta(p_{\text{trk}}, p_{\text{tot}} - p_{\text{trk}}))$	–	< 3.5	–
N_{seeds}	≥ 2	$\geq N_{\text{tracks}}C_2 - 1$	$\geq N_{\text{tracks}}$
$\Delta R_{\text{FDV, jet}}$	–	–	< 0.4
$ \eta_{\text{PV-FDV}} $	–	–	< 2.5
$N_{\text{DV}_{\text{sat}}}$	≥ 1	$\geq 1^5$	≥ 2

Table 6.15: Full vertex selection criteria for SDV SR and FDV SR (including the 1 FDV and 2+ FDV subregions).

6.8.2 Signal Efficiencies

Once we have defined our SRs with selection criteria carefully optimized to maximize sensitivity to specific signal models,⁶ the natural next question is: how well do these targeted selections actually work? To answer this, we take our MC signal samples and run them through exactly the same gauntlet of reconstruction algorithms and selection cuts that we apply to real data in each respective SR. This gives us the signal efficiency - simply the frac-

6. The SDV SR for the Gluino R -hadron and DFSZ Axino models, the 1 FDV SR for the Bino-Wino Coannihilation model, and the 2 FDV SR for the Higgs portal model. Note that the DFSZ Axino model was developed after the SR definitions were finalized, and while it was initially expected that the $h \rightarrow b\bar{b}$ channel would perform better in the FDV SR, both axino channels ultimately achieved higher sensitivity in the SDV SR.

tion of events that make it through all requirements and end up in the SR where that model is expected to perform best. We then multiply this efficiency by the theoretical production cross-section and our integrated luminosity of 137 fb^{-1} to get the expected number of signal events we should see if nature cooperates.

For LLP searches, the efficiency calculation is particularly crucial due to its strong dependence on particle lifetime. Particles with lifetimes that are too short decay before producing displaced signatures, while those with lifetimes that are too long escape the detector fiducial volume before decaying. The particle masses also significantly influence both the production cross-section and the kinematics of decay products, affecting reconstruction efficiency. These predicted yields are essential for interpretation: regions of parameter space with higher expected yields allow for more stringent exclusion limits if no significant excess is observed above background expectations in the corresponding SR.

Model	Parameters	Lifetime (ns)	Efficiency (%)
Gluino R -hadron SDV SR	$m_{\tilde{g}} = 1600 \text{ GeV},$ $m_{\tilde{\chi}_1^0} = 100 \text{ GeV}$	0.01	6.4
		0.1	41
		3.0	6.9
DFSZ Axino (hh) SDV SR	$m_{\tilde{a}} = 400 \text{ GeV},$ $m_{\tilde{\chi}_1^0} = 800 \text{ GeV}$	0.0045	5.6
		0.097	25
		1.5	3.1
DFSZ Axino (ZZ) SDV SR	$m_{\tilde{a}} = 400 \text{ GeV},$ $m_{\tilde{\chi}_1^0} = 800 \text{ GeV}$	0.269	31
		0.752	17
		1.471	9.9
Bino-Wino Coannihilation 1 FDV SR	$m_{\tilde{\chi}_1^0} = 650 \text{ GeV},$ $m_{\tilde{\chi}_2^0} = 700 \text{ GeV}$	0.03	0.37
		0.3	0.20
		3.0	0.037
Higgs portal 2 FDV SR	$m_S = 55 \text{ GeV}$	0.0033	2.2
		0.033	5.7
		3.3	0.018

Table 6.16: Examples of SR selection efficiencies for representative grid points from each signal model. These show how efficiency varies with lifetime for typical parameter choices in each model’s optimized signal region.

Table 6.16 illustrates how signal efficiency typically peaks at intermediate lifetimes (0.03-0.1 ns) and decreases for both shorter and longer lifetimes. This behavior is expected: very short-lived particles decay before reaching sensitive detector regions, while very long-lived particles often decay outside the detector volume.

The DFSZ Axino is the primary focus of this thesis. We consider two decay channels: the hh channel where $\tilde{\chi}_1^0 \rightarrow \tilde{a}h$ with $h \rightarrow b\bar{b}$, and the ZZ channel where $\tilde{\chi}_1^0 \rightarrow \tilde{a}Z$ with $Z \rightarrow q\bar{q}$.

Expected signal yields for both axino channels are shown in Figure 6.6 across a subset of our signal grid. These results are for $f_a = 5 \times 10^9$ GeV, or a corresponding axion mass of $m_a = 1138$ μeV . The yields are already scaled by the appropriate SM branching ratios for each decay channel.

Figure 6.6a shows yields for the hh channel. The yields peak around 0.1 ns lifetime. This is because shorter lifetimes result in DVs that fail the 4 mm transverse DV- PV separation requirement (Table 6.4), while longer lifetimes cause decays outside the detector’s fiducial volume.

The ZZ channel yields in Figure 6.6b are consistently higher than the hh channel across most of the parameter space. This difference reflects the better reconstruction efficiency of the SDV algorithm for light-quark jets from Z decays, which produce clean, collimated vertex structures compared to b -jets that create complex vertex topologies due to secondary B -hadron decays spanning several millimeters.

The yields are highest for lower neutralino and axino masses, with sensitivity decreasing as either mass parameter increases due to reduced production cross-sections and phase space suppression effects.

An important finding is that both axino channels perform better in the SDV SR than the FDV SR. This indicates that the FDV algorithm, optimized for the Bino-Wino Coannihila-

tion with off-shell Higgs bosons and their characteristic kinematics, is less effective for the on-shell Higgs decays and different kinematic distributions in the axino model.




Figure 6.6: Expected signal yields for the DFSZ Axino in the SDV SR for (a) the hh channel and (b) the ZZ channel. Results are for $f_a = 5 \times 10^9$ GeV ($m_a = 1138$ μ eV) and normalized to 137 fb $^{-1}$. Figures from [160].

6.8.3 Validation Regions

To validate the background estimation methods, several VRs are defined for each SR. These VRs are kinematically similar to the SRs while maintaining orthogonality, allowing us to test the background estimation techniques in phase spaces adjacent to the SRs.

As shown in Table 6.17, the VRs are constructed by modifying key selection criteria such as track multiplicity ($N_{\text{Tracks}}^{\text{DV}}$), DV invariant mass (m_{DV}), seed vertex count ($N_{\text{Seeds}}^{\text{DV}}$), and material flag requirements. For the SDV SR, VRs are defined with exactly four tracks or by selecting vertices with lower mass ranges (5–10 GeV). For the FDV SR, various combinations of track multiplicity, mass ranges, and seed requirements create VRs for both the single FDV and multiple FDV analysis channels.

For the SDV SR, a set of VRs is defined with various combinations of mass and track

VR Selection	SDV VRs	FDV VRs	
		1FDV	2FDV
Vertices	SDV	FDV	
$N_{\text{Tracks}}^{\text{DV}}$	$= 4, > 4$	≥ 5	≥ 4
$m_{\text{DV}} [\text{GeV}]$	5-10	2-5, 5-10, ≥ 10	2-5, 5-10, ≥ 10
$N_{\text{Seeds}}^{\text{DV}}$	–	Nominal, Low seed: 2-4, 3	Nominal, 2- $N_{\text{Tracks}}^{\text{FDV}}$
Material Flag	Outside	Outside, Inside	Outside, Inside
N_{DVsat}	≥ 1	$\geq 1^{\dagger}$	≥ 2

[†] It is required not to have any additional FDVs satisfying the multiple FDV selections.

Table 6.17: Vertex selection criteria for the various VRs for each of the SRs.

multiplicity. These regions include VRs with low mass (2–5 GeV) and exactly 4 tracks, intermediate mass (5–10 GeV) with 4 or more tracks, and higher mass regions (10–15 GeV, 15–20 GeV, and >20 GeV) with exactly 4 tracks. This approach ensures comprehensive validation across different regions of phase space.

For the 1 FDV SR, VRs are categorized in three ways: those inside material (where the material map veto is inverted), those with a modified $N_{\text{Seeds}}^{\text{DV}}$ requirement, and a nominal region with standard requirements. The inside material regions provide a background-dominated sample to validate the background estimation methods, as they contain negligible signal contamination due to the dominance of HIs with detector material in these regions.

For the 2 FDV SR, VRs include those with $N_{\text{Tracks}} = 3$ instead of ≥ 4 , as well as regions with modified $N_{\text{Seeds}}^{\text{DV}}$ requirements. For some regions in this configuration, a looser material veto is applied, and the m_{DV} requirement is relaxed to $m_{\text{DV}} > 0$ due to potential signal contamination.

The background estimation procedure is applied independently to each VR, and the comparison between predicted and observed numbers of events serves as a validation of the methodology. Where discrepancies are observed between prediction and observation, appropriate non-closure systematic uncertainties are assessed to account for potential mismodeling.

For the background estimates, non-closure uncertainties of 22% and 20% are applied to the 1 FDV and 2 FDV SRs, respectively.

6.9 Photon Triggered Control Region

Another key element of the background estimation strategy for this analysis is the PTR, which serves as the primary CR. This region consists of events that fail one or more of the MET-based trigger or offline MET criteria but pass a high-threshold single-photon trigger, `HLT_g140_loose`, requiring an isolated photon with $p_T^\gamma > 140$ GeV.

Since LLP signals are not expected to be produced in association with high- p_T isolated photons, the PTR is orthogonal to the SRs while retaining similar detector backgrounds. This makes the PTR an ideal region for studying background processes without signal contamination. The PTR provides a statistically rich sample for modeling non-collision and instrumental backgrounds using jet and track observables, allowing for the derivation of probabilities essential to the data-driven background estimation techniques employed in this analysis. For instance, the Event-DV probability (EDP), used for the inclusive background estimate, is calculated as the ratio of DVs in events with specific track density properties to the total number of events with those properties in the PTR.

The statistical power of this CR enables precise measurement of the probability of forming signal-like DVs as a function of event properties, which can then be applied to the MTR to predict background contributions in the SRs. To account for differences in the pile-up distribution between the PTR and the MTR, events in the PTR are reweighted to match the pile-up profile of the MTR when calculating these probabilities, with the difference between weighted and unweighted estimates taken as a systematic uncertainty.

6.10 Background Estimation

The analysis employs two complementary data-driven methods to estimate backgrounds in the signal regions, providing cross-validation of the predicted background estimates. The primary approach is the inclusive method, which is used for both SDV and FDV SRs. This method models the probability of forming high-mass DVs as a function of jet and track densities in the event. It begins by deriving an EDP from the PTR, representing the likelihood of finding a signal-like DV in an event with specific characteristics. This probability is then applied to events in the MTR that pass the common event-level selections. An extrapolation factor is used to account for differences in vertex requirements between the CR and the SR, allowing for an accurate prediction of the signal region background yield.

The second approach, the component-based method, takes a more granular view by separately estimating contributions from the distinct expected sources of background. This estimation includes HI, where particles interact with detector material to produce collimated tracks and low-mass vertices; AX, where unrelated tracks from the primary vertex or pile-up are erroneously attached to a DV, artificially increasing its mass and track multiplicity; and MV, where separate physical vertices are mistakenly combined during reconstruction into a single vertex with higher mass and track count. Each component is modeled with dedicated techniques: HIs are parameterized based on the vertex mass distribution, AX are estimated using kinematic properties of candidate background tracks, and MV are studied by artificially overlapping events in data. The consistency between these two fundamentally different estimation strategies provides significant confidence in the overall background prediction, as they approach the problem from distinct starting points, yet should converge on similar results if properly implemented.

6.10.1 SDV Background Estimate

Searches for SDVs with high invariant mass and multiple tracks face minimal background due to the rarity of SM processes producing such signatures. Instead, background SDVs primarily arise from:

- **Hadronic Interactions**, where particles interact with detector material, producing collimated tracks and low-mass vertices.
- **Accidental Crossings**, where unrelated tracks from the primary vertex or pile-up attach to a SDV, increasing its apparent mass and track count.
- **Merged Vertices**, where separate SDVs are mistakenly combined into a single vertex with higher mass and multiplicity.

All three sources of background are estimated inclusively with a data-driven technique that exploits the relationship between DVs and the track density of the event to predict the number of background vertices. The estimated number of background events in the SR is computed using an EDP: $P(\text{vertex} | N_{\text{Tracks}}^{\text{Event}}, N_{\text{b-tag}})$, where $N_{\text{Tracks}}^{\text{Event}}$ is the total number of tracks associated with any Collision Vertex (CV) and $N_{\text{b-tag}}$ is the number of b -tagged-jet-count per event. Defined in Equation 6.1, the EDP represents the probability of producing a SR-like vertex satisfying the baseline vertex selections with the additional requirement: $m_{\text{DV}} > 5 \text{ GeV}$ and $N_{\text{Tracks}}^{\text{DV}} \geq 4$, which is measured in the PTR. This selection on m_{DV} and $N_{\text{Tracks}}^{\text{DV}}$ is intentionally less stringent than the SDV SR requirement, as the stricter criteria result in too few events to generate a non-trivial estimate. The extrapolation factor f , from Equation 6.2, corrects for this by extrapolating to the SR selection values. The EDP is evaluated across all events passing the common event-level selections described in – for the SR region, the background estimate is then computed as in Equation 6.3, where the sum is

taken over the SR.

$$P(\text{vertex} | N_{\text{Tracks}}^{\text{Event}}, N_{\text{b-tag}}) = \frac{N_{\text{Vertex}}^{\text{CR}}(N_{\text{Tracks}}^{\text{Event}}, N_{\text{b-tag}})}{N_{\text{Events}}^{\text{CR}}(N_{\text{Tracks}}^{\text{Event}}, N_{\text{b-tag}})} \quad (6.1)$$

$$f = \frac{N_{\text{vertex}}^{\text{CR}}(m_{\text{DV}} > 10 \text{ GeV}, N_{\text{Tracks}}^{\text{DV}} \geq 5)}{N_{\text{vertex}}^{\text{CR}}(m_{\text{DV}} > 5 \text{ GeV}, N_{\text{Tracks}}^{\text{DV}} \geq 4)} \quad (6.2)$$

$$N_{\text{Bkg}} = f \cdot \sum_{i=0}^{N_{\text{Events}}} P(\text{vertex} | N_{\text{Tracks}}^{\text{Event}}, N_{\text{b-tag}})_i \quad (6.3)$$

This method predicts an expected background of $0.56_{-0.39}^{+0.40}$ events in the SDV SR.

6.10.1.1 SDV Background Estimate Validation

The predictions of the background estimation method are tested in three statistically independent sideband VRs. Events in the VRs are required to satisfy all the 1 SDV SRs, except for the relaxed conditions of $5 \text{ GeV} < m_{\text{SDV}} \leq 10 \text{ GeV}$ or $N_{\text{Tracks}}^{\text{SDV}} \geq 4$. The full background estimate procedure, including calculating an EDP and an extrapolation factor f , is performed for each VR. The details of the VRs and a comparison of the predicted and observed number of events in each VR are shown in Figure 6.7. The predictions of the background estimate are found to agree with the number of observed events in the data for all three of the VRs within uncertainty.

The total expected background obtained in each VR and the SDV SR agree within statistical uncertainty with the expected background from the inclusive method.

In addition to validating the background estimate in sideband regions, a separate background estimate is performed where the contribution from each source – HIs, MVs and AXs – is es-




Figure 6.7: Estimated and observed background count in the sideband VRs of the SDV SR. The number of observed SDVs in each VR is shown in black, and the number of predicted SDVs in each VR are shown in red. The bottom portion of the plot shows the ratio of the observed to the estimated number of SDVs. Figure from [141].

estimated independently. Each contribution is estimated using data-driven techniques that construct mass templates describing the shape of m_{SDV} for background SDVs from the particular background source. These methods are described and validated in previous ATLAS searches for SDVs [161, 162].

6.10.1.2 Hadronic Interaction Background Estimate

HIIs occur when particles from the proton-proton collision scatter off nuclei in the detector material itself, discussed in Chapter 6.6. When a high-energy particle hits a structure, such as a support or cable, in the ID, it can create multiple secondary charged particles that are reconstructed as what appears to be a DV. These fake vertices tend to have low masses and highly collimated tracks because they originate from a single energetic particle hitting stationary material.

Even though our SR requires SDVs to be reconstructed outside of detector material, we can still have residual backgrounds from HIIs. These arise in several ways: the material map doesn't capture every piece of detector hardware (such as cables, cooling lines, and small

supports), vertex reconstruction uncertainties can place a material-originated vertex outside the mapped material regions, and secondary particles from material interactions can travel before producing the reconstructed tracks.

While we use MC simulations of HI processes to understand and validate our methods, the actual background estimate relies entirely on data. This is important because the detector material isn't perfectly modeled in simulation; there are cables, cooling lines, and other components that don't make it into the simulated MC geometry.

To estimate how many HI backgrounds remain in our SR, we fit the mass distribution of SDVs with the function:

$$N_{\text{SDV}}(m; b, C, l) = \left(\frac{1}{C(m-b)} + \exp \frac{m-b}{l} \right)^{-1} \quad (6.4)$$

where m represents the SDV mass, and b , C , and l are parameters that we determine by fitting to the photon-triggered CR in data. This function captures the characteristic shape we expect from HIs - a sharp rise at low masses followed by an exponential fall off at higher masses.

We perform these fits in regions where we know HIs dominate (like areas flagged as inside material), then extrapolate to predict how many such backgrounds survive in our signal region with $m_{\text{SDV}} > 10$ GeV and $N_{\text{tracks}} \geq 5$. Based on this estimate, we conclude that the material veto performs exceptionally well. We find that HI contributes less than 10^{-4} events to our SR, making it essentially negligible compared to other background sources like AX and MV.

We validate this approach by comparing our fits to independent data samples and checking closure in regions where we can apply additional selections that preferentially remove HIs. The good agreement we observe gives us confidence that this background is well under

control.




Figure 6.8: Fits performed to estimate the HI contribution to the background. Note that fits are performed in both the inside and outside material regions, and for the outside material region, the SR is blinded. Figure from [160].


6.10.1.3 Accidental Crossing Background Estimate

AX is a potential background source for this analysis because it is possible that unrelated tracks from different physical processes (such as pile-up vertices or PV tracks) can cross the trajectory of genuine SDVs, becoming artificially associated with them and creating vertices with inflated mass and track multiplicity that may satisfy the SR.


The AX background estimation employs a data-driven technique that directly samples both the rate and kinematics of accidental crossings observed in data. The methodology is performed separately in four radial regions of the ID to account for the varying background properties as a function of displacement.

The procedure begins by identifying candidate 2- and 3-track vertices consistent with $K_S^0 \rightarrow \pi^+ \pi^-$ decays in data. For three-track vertices where one track pair has an invariant mass within 50 MeV of the K_S^0 mass (498 MeV), the remaining track is flagged as an accidentally

1. Find K_s^0 candidates with an accidentally-crossing track.



2. Create a mass template of DVs with accidentally-crossing tracks.



3. Weight template by accidental crossing probability, estimated using K_s^0 candidates.

$$\frac{\# \left(m_{K_s^0} \{ \text{blue circle with red line} \} \right)}{\left(\# \left(m_{K_s^0} \{ \text{blue circle with red line} \} \right) + \# \left(m_{K_s^0} \{ \text{blue circle with red line} \} \right) \right)}$$

(= AX probability)

Figure 6.9: An overview of the AX background estimation method. We use K_0^S candidates both to collect accidentally crossing tracks and to estimate the rate at which accidental crossing occurs. Figure from [160]

crossing track. These 2- and 3 track candidates are stored in a database along with their z and R_{xy} positions of the crossed K_S^0 , as well as properties such as $\Delta\eta$ and $\Delta\phi$ with respect to the vector pointing from the primary vertex to the vertex, to preserve the relevant kinematics of the crossing.

The accidental crossing rate is then measured by comparing the yields of 2-track and 3-track K_S^0 candidates in data. The number of K_S^0 decays is determined by fitting the invariant mass distributions with signal and background components. For 2-track candidates, a Crystal Ball function models the K_S^0 peak, while for 3-choose-2-track candidates, a Gaussian function is used. The crossing rate is calculated as:

$$f_C = \frac{N_{K_S^0}^{3\text{-track}}}{N_{K_S^0}^{2\text{-track}} + N_{K_S^0}^{3\text{-track}}} \quad (6.5)$$

where $N_{K_S^0}^{2\text{-track}}$ and $N_{K_S^0}^{3\text{-track}}$ represent the fitted yields of 2- and 3-track K_S^0 candidates, respectively.

To create templates for the AX background, tracks from the database are artificially attached to baseline SDVs found in data that pass all selection criteria except the SR mass and track multiplicity requirements. For each baseline vertex, a random track from the appropriate radial region is selected and attached, and the vertex mass is recalculated. This process generates $(n+1)$ -track mass templates from n -track vertices. The resulting mass distributions are then normalized by the measured crossing factors to reflect the actual rate of accidental crossings expected in the data. The background estimate for each validation or signal region is obtained by integrating these scaled mass templates over the relevant mass and track multiplicity ranges.

Systematic uncertainties on the AX estimate arise from several sources: statistical uncertainties on the K_S^0 fits used to determine crossing rates, uncertainties in the choice of track




Figure 6.10: An example mass template produced as part of the AX background estimate, showing the integration of the different regions to produce the estimated numbers of background events. Figure from [160].

selection criteria for building the database, and potential biases in the artificial attachment procedure. These uncertainties are propagated through the full estimation chain and combined in quadrature to yield the final systematic uncertainty on the background prediction. The AX background estimate contributes approximately 0.76 ± 0.15 events to the expected background in the SDV SR, representing the dominant component of the combined background estimate for this analysis.

6.10.1.4 Merged Vertices Background Estimation

MVs represent another potential background source in this analysis. This background arises when the VSI algorithm merges two unrelated SDVs, creating a single vertex with higher track multiplicity and mass that can pass signal region selections. The VSI algorithm includes a merging step designed to correctly reconstruct LLPs that produce multiple charged decay products. When a massive LLP decays, it may create several tracks that could be reconstructed as separate low-mass vertices if not properly combined. To address this, the algorithm merges nearby vertices when their distance significance satisfies $S < 10$, where:

$$S = \sqrt{(\vec{v}_1 - \vec{v}_2)(\text{Cov}(\vec{v}_1) - \text{Cov}(\vec{v}_2))^{-1}(\vec{v}_1 - \vec{v}_2)^T} \quad (6.6)$$

where $\vec{v}_i$ represents the position vector of the i -th DV and $\text{cov}(\vec{v}_i)$ is its covariance matrix from the vertex fitting algorithm.

However, this merging procedure can also inadvertently combine two genuinely separate background vertices that happen to be reconstructed near each other. These unrelated vertices typically originate from SM processes such as heavy-flavor hadron decays. When merged, the resulting vertex contains more tracks and has a higher invariant mass than either original vertex alone, potentially satisfying the high-mass, high-multiplicity selection criteria

of our SR despite not originating from BSM physics. This unwanted merging of unrelated background vertices constitutes the MV background contribution that must be estimated.

The MV background estimate is entirely data-driven, using real vertex pairs in the data to measure the actual merging rate and construct realistic mass templates for merged vertices. The estimate follows these steps:

1. **Merging rate calculation:** We calculate the merging rate by comparing distance significance distributions between vertex pairs from the same event versus pairs from different events (mixed events). Figure 6.11 shows a sample dataset with several DVs, while Figure 6.12 illustrates how we construct pairs of DVs from both same-event and mixed-event combinations. The VSI algorithm merges vertices when their distance significance $S < 10$. For vertex pairs with $S > 10$, where merging is forbidden, the same-event and mixed-event distributions should be identical, as demonstrated in the high- S region of Figure 6.13a. However, for $S < 10$, we observe a deficit in the same-event distribution, Figure 6.13b, indicating that merging has already occurred during reconstruction. This deficit quantifies our merging rate.
2. **Mass template construction:** We create mass templates by artificially merging pairs of DVs from different events that satisfy $S < 10$. The invariant mass of these artificially merged vertices is calculated from the four-momenta of all constituent tracks, assuming each track corresponds to a charged pion. Separate templates are maintained for each combination of input vertex track multiplicities (e.g., 2+2 tracks, 2+3 tracks, etc.).
3. **Template normalization:** We normalize the mass templates using two correction factors: a scale factor to account for combinatorial effects when artificially overlaying events, and the measured merging rate from step 1. The normalization ensures that the total number of artificially merged vertices equals the deficit observed in the same-event

distance significance distribution:

$$N_{\text{norm}} = \frac{\int_0^{10} H_{S,\text{mixed}} dS - \int_0^{10} H_{S,\text{same}} dS}{\int H_{\text{mass}} dm} \quad (6.7)$$

where H_{mass} is the unnormalized mass template, $H_{S,\text{mixed}}$ and $H_{S,\text{same}}$ are the distance significance distributions for mixed and same events, respectively.

4. **Track cleaning application:** We apply the complete track cleaning procedure to the artificially merged vertices. This step is essential as it replicates the effect of our analysis selection criteria on genuine MVs in data. After cleaning, we recalculate both the invariant mass and track multiplicity of the vertices and re-bin them into the appropriate mass templates based on their post-cleaning properties.
5. **Background estimate:** We estimate the MV background contribution by integrating the normalized, post-cleaning mass templates over our SR requirements ($m_{\text{SDV}} > 10$ GeV and $N_{\text{SDV}}^{\text{tracks}} \geq 5$). Statistical uncertainties are propagated from the Poisson uncertainties on the number of vertex pairs in both the same-event and mixed-event $S < 10$ regions to derive the final systematic uncertainty on the background estimate.






Figure 6.11: Sample data set including 3 events with one or two DVs each. Figure from [160]

Using this method, we estimate a MV contribution of 0.183 ± 0.029 events in the SDV




(a) Example of the same-event-SDV-pairings from the sample dataset (Figure 6.11), that the same-event-pairing S^2 is calculated from.



(b) Example of the mixed-event-DV-pairings from the sample dataset (Figure 6.11), that the mixed-event-pairing S^2 is calculated from.

Figure 6.12: The first step of calculating the merging rate in data is to compare the distance significance (S^2) distributions between same-event-pairs and mixed-event-pairs of DVs. This demonstrates the population of these S^2 distributions that fill the histogram of Figure 6.13. Figures from [160].



(a) Number of DV pairs with 2+3 tracks vs. S^2

(b) $S < 10$ region only

Figure 6.13: The distance significance squared (S^2) between pairs of DVs taken from the same event and mixed events from events outside material. The mixed-event-pairing distribution is normalized to match the same-event-pairing distribution in the $S^2 > 100$ region, where accidental merging is not possible with the VSI algorithm. The deficit in the same-event distribution for $S^2 < 100$ (b) provides a direct measurement of the merging rate. Figures from [160].

SR. Compared to the other background components, MVs contribute a subdominant but non-negligible amount to the combined background estimate.

6.10.1.5 Summary of SDV SR Background Estimates

We compare the combined background estimate (formed by summing individual contributions from MVs, AXs, and HIs) with the inclusive method. The results show excellent agreement across all approaches, as shown in Table 6.18. The combined estimate yields 0.94 ± 0.15 expected background events in the SDV SR, which agrees within statistical uncertainty with the inclusive estimate of 0.56 ± 0.40 events (Table 6.18).




Figure 6.14: The combined MV, AX, and HI background estimates in the SR and VRs. Figure from [160].

6.10.2 1FDV Background Estimate

The background estimate for 1 FDV SR is done the same as the inclusive SDV SR background estimate, described in Chapter 6.10.1, using the FDVs. The expected background is predicted to be 0.81 ± 0.49 events in the 1 FDV SR. The dominant uncertainty comes from the statistical uncertainty of the number of FDVs in the CR. The uncertainty also includes the effect of the difference in the distributions of the number of pile-up collisions between the PTR and MTR and the “non-closure” uncertainty assigned based on the validation results described in Chapter 6.10.2.1.

Estimation Method	Predicted Background
Inclusive	0.56 ± 0.40
Combined	0.94 ± 0.15
SDV SR Combined Components	
Merged Vertex	0.18 ± 0.03
Accidental Crossings	0.76 ± 0.15
Hadronic Interactions	$(1.8 \pm 5.9) \times 10^{-5}$

Table 6.18: Background estimates for the SDV SR comparing the combined estimate and component contributions.

6.10.2.1 1FDV Background Estimate Validation

Events in the VR are required to meet the same criteria as the 1 FDV SR, except that $N_{\text{Tracks}}^{\text{FDV}}$ requirements are relaxed from $N_{\text{Tracks}}^{\text{FDV}} > 5$ to $N_{\text{Tracks}}^{\text{FDV}} > 4$, and m_{FDV} requirements are changed from $m_{\text{FDV}} > 10$ GeV to $2 \text{ GeV} < m_{\text{FDV}} < 5 \text{ GeV}$, which defines the nominal VR. Another set of sideband VRs are developed by loosening the conditions on $N_{\text{Seeds}}^{\text{FDV}}$. To validate the effects of background HIs (Chapter 6.6), a third VR is defined by requiring vertices to be located inside the material. The selection criteria for these VRs are summarized in Figure 6.15. For each VR, the full background estimation procedure described in Chapter 6.10.1 is applied, and the predicted and observed numbers of FDVs are compared. In several regions, a discrepancy greater than 1σ of uncertainty is observed. As a result, a maximum relative error of 22% is assigned as an additional non-closure uncertainty.

6.10.3 2FDV Background Estimate

The estimated number of background events in the 2 FDV SR follows a similar inclusive approach to the SDV SR and 1 FDV SR, leveraging correlations between FDVs and jet properties rather than track density. This estimate is obtained using a probability distribu-




Figure 6.15: Estimated and observed background in the sideband VRs of the FDV VRs if the FDV SR. The number of observed FDVs in each VR is shown in black, and the number of predicted FDVs in each is shown in red. The bottom portion of the plot shows the ratio of the observed to the estimated number of FDVs. Figure from [141].

tion over the jets in the E_T^{miss} -triggered region, defined by the probability of an individual event containing a FDV near a jet. This probability is parameterized based on the jet's p_T and b -tag value. The individual jet probabilities, referred to as the jet-FDV probability (JFP), are calculated in the photon-triggered CR using the formula in Equation 6.8.

$$\text{JFP}_{(n)}(\text{FDV} \mid \text{jet}(x, y)) = \frac{N_{\text{FDV}}^{\text{CR}}(\text{n matched jets} \mid b_{\text{tag}} = x, p_T = y)}{N_{\text{Events}}^{\text{CR}}(b_{\text{tag}} = x, p_T = y)} \quad (6.8)$$

In the CR, the JFP is determined using jets spatially matched to FDVs within $\Delta R \leq 0.4$. If multiple FDVs are matched to the same jet, it contributes to the estimation of $\text{JFP}_{(\geq 2)}$, whereas if a jet only matches a single FDV it contributes to $\text{JFP}_{(1)}$.

The number of background events with 2 or more FDVs is then calculated as

$$N_{\text{background}}^{\geq 2\text{FDV}} = N_{\text{events}} - N_{\text{background}}^{1\text{FDV}} - N_{\text{background}}^{0\text{FDV}}, \quad (6.9)$$

with the number of background events with n background FDVs, $N_{\text{background}}^{n\text{FDV}}$, given by

$$N_{\text{background}}^{0\text{FDV}} = \sum_{\text{events}} \prod_{\text{jets}} \left(1 - \text{JFP}_{(1)} - \text{JFP}_{(\geq 2)} \right), \quad (6.10)$$

$$N_{\text{background}}^{1\text{FDV}} = \sum_{\text{events}} \left(\prod_{\text{jets}} \left(1 - \text{JFP}_{(1)} - \text{JFP}_{(\geq 2)} \right) \times \left(\sum_{\text{jets}} \frac{\text{JFP}_{(1)}}{1 - \text{JFP}_{(1)} - \text{JFP}_{(\geq 2)}} \right) \right). \quad (6.11)$$

An expected background of 1.29 ± 0.64 events is predicted in the 2 FDV SR.

6.10.3.1 2FDV Background Estimate Validation

Events in the 2 FDV VR are required satisfy the 2 FDV SR conditions, except for the relaxed requirement of $N_{\text{Tracks}}^{\text{SDV}} = 3$ or 4, $m_{\text{FDV}} > 1.5$ GeV and $N_{\text{Seeds}}^{\text{FDV}} \geq N_{\text{Tracks}}^{\text{SDV}}$ or $[2, N_{\text{Tracks}}^{\text{SDV}})$, creating sideband VRs.

In addition to the $N_{\text{Tracks}}^{\text{SDV}}$, m_{FDV} , and $N_{\text{Seeds}}^{\text{FDV}}$ - sideband VRs, the estimate is validated using both the strict inside material and strict outside material flags defined in Chapter 6.6. Exceptionally, the outside material, 4+ tracks, and low seed validation regions were loosened to increase statistics by loosening the material veto and being inclusive of the entire mass range ($m_{\text{FDV}} > 0$ GeV). The selection criteria for these VR are shown in Figure 6.16. The $N_{\text{background}}^{\geq 2\text{FDV}}$ is calculated for each VR using the background estimation method outlined in Chapter 6.10.3. Figure 6.16 compares each VR's predicted and observed number of events




Figure 6.16: Estimated and observed backgrounds in the sideband validation regions for the 2FDV signal region. Observed FDVs are shown in black, predicted FDVs in red, and the ratio of observed to predicted FDVs is displayed in the lower panel. Figure from [141]

with 2 or more FDVs. The discrepancy between observed and predicted 1σ is seen in several regions, resulting in a maximum relative error of 20% assigned as a non-closure uncertainty for the 2 FDV SR.

6.10.4 Summary of Background Estimates

Background Estimate Summary	
SR	Background Estimate
SDV	0.56 ± 0.40
1 FV	0.81 ± 0.49
2 FV	1.29 ± 0.59

Table 6.19: A summary of the background estimates for all SRs.

All background estimates presented in this analysis are data-driven, using inclusive background estimation techniques that avoid reliance on MC simulations. For the SDV SR, we

additionally validated our inclusive estimate with an independent combined estimate that separately models HIs, MVs, and AXs; finding good agreement between the two approaches. All estimates were further validated in dedicated VRs, where the predicted event counts generally agreed well with observations. Table 6.19 summarizes the final background estimates for all SRs.

6.11 Unblinding and Observed Events

6.11.1 *Blinded Analysis Strategy*

This analysis was performed in a blinded fashion until all analysis procedures, background estimation methods, and systematic uncertainties were finalized. Blinding is a critical component of experimental particle physics searches as it prevents unconscious bias from influencing analysis decisions based on the observed data. The SR selections, background estimation techniques, and systematic uncertainty evaluations were all developed and validated using CRs and VRs, without examining the final SR yields.

6.11.2 *Unblinding Results*

Upon unblinding the data, we found 1 event in the SDV SR, 3 events in the 1 FDV SR, and 2 events in the 2 FDV SR. A summary of the observed and expected event counts is shown in Table 6.20. The observed yields are consistent with the predicted SM background within uncertainties, indicating no significant evidence for BSM physics. Detailed properties of all observed events are provided in Tables 6.21 and 6.22.

Table 6.21 shows the event-level properties, including missing transverse momentum, run numbers, and data-taking years for each observed event.

Signal Region	Observed Events	Predicted Background
SDV SR	1	$0.56^{+0.40}_{-0.39}$
1 FDV SR	3	0.81 ± 0.49
2 FDV SR	2	1.29 ± 0.59

Table 6.20: Observed event yields compared to predicted background in each SR.

Signal Region	E_T^{miss} [GeV]	Run Number	Year
SDV SR	211	338897	2017
1 FDV SR	169.7	329869	2017
	196.8	329964	2017
	199.9	360161	2018
2 FDV SR	162.2	309759	2016
	278.7	363400	2018

Table 6.21: Event-level properties for observed events in each SR.

The vertex-level properties for each DV are detailed in Table 6.22, including radial position, invariant mass, track multiplicity, and jet associations.

6.11.3 Event Characterization

Given the small number of observed events, we were able to investigate the event properties of each event and assess their cause. The observed events are summarized in Tables 6.21 and 6.22.

The single event in the SDV SR appears consistent with an accidental crossing background. This event has a DV reconstructed inside the beam pipe near a b -tagged jet, with one track having significantly different η and ϕ compared to the other four tracks. When the invariant mass is recalculated excluding this suspicious track, $m_{DV} = 2.09$ GeV, well below the SR requirement of 10 GeV. This suggests a SM b -hadron decay contaminated by an unrelated track accidentally crossing the vertex region, matching the dominant predicted background

Signal Region	R_{xy} [mm]	DV Mass [GeV]	DV N_{tracks}	Jet Association
SDV SR	5.88	15.1	5	b -jet
1 FDV SR	15.5	14.1	5	b -jet
	5.63	14.5	5	None
	4.50	12.0	5	None
2 FDV SR	8.25	2.64	4	b -jet (shared)
	6.05	5.73	6	
	169.78	6.48	4	Light jet (shared)
	243.75	7.87	4	

Table 6.22: Properties of each DV for observed events in each SR. Jet association is determined by $\Delta R < 0.4$ matching criterion with the nearest jet. For the 2 FDV SR events, “shared” indicates both FDVs in the same event are matched to the same jet.

source. Figure 6.17 shows the event display for this event, illustrating both the detailed detector geometry and the broader event topology.

The three events in the 1 FV SR have varied properties. The first event is associated with a nearby b -jet and was recorded in 2017, while the second and third events show no clear jet association. All three events have similar DV masses around 12-15 GeV and identical track multiplicities of 5, suggesting they may represent a mix of b -hadron decays and other background processes.

The two events in the 2 FV SR represent distinct background scenarios. The first event contains two DVs close to each other (both inside the beam pipe) with low masses (2.64 and 5.73 GeV) that are both matched to the same b -jet, consistent with a regular b -hadron decay. The second event has two DVs far apart (at radii of 170 and 244 mm) with higher masses (6.48 and 7.87 GeV) matched to the same light jet, suggesting a cascaded double hadronic interaction with detector material.

All observed events appear consistent with SM background processes, with no clear signal-like properties that would suggest BSM physics.




Figure 6.17: Event display of the single event observed in the SDV SR. Primary vertices are shown in blue, secondary b -tagging vertices in yellow, and VSI vertices in red. The signal-like SDV is located near the beampipe, close to jets (shown as yellow cone structures). 6.17a shows a detailed view of the beampipe region featuring the Pixel, SCT, and TRT detector layers, while 6.17b provides a broader view of the inner detector volume showing the spatial relationship between different vertex types and the beam pipe. Figures from [160].

6.11.4 Statistical Interpretation

With no significant excess observed over the predicted SM background, we proceed to set upper limits on the production cross-sections of our BSM signal models. To establish these limits, we must consider both the expected signal yields for our benchmark models and the systematic uncertainties that affect these predictions.

The expected signal yields, which depend on the efficiency of our selection criteria for different models, are detailed in Chapter 6.8.2. The various sources of systematic uncertainty that affect both our background estimates and signal predictions are comprehensively described in Chapter 6.12. These uncertainties include experimental effects such as tracking inefficiencies and trigger performance, as well as theoretical uncertainties on the signal cross-sections and modeling.

The statistical framework used to combine our observed data with the predicted backgrounds

and signal yields, along with the resulting 95% confidence level exclusion limits on our benchmark models, will be presented in Chapter 6.13. These limits will constrain the allowed parameter space for our BSM physics.

6.12 Uncertainties

The reliability of any search for BSM physics depends critically on a thorough understanding of both statistical and systematic uncertainties. In this analysis, we face a unique challenge with background estimation: the DV signatures we search for occur so rarely in SM processes that conventional MC simulations lack sufficient statistics to model them accurately. This rarity necessitates the development of specialized, data-driven background estimation techniques, each with its own unique uncertainty considerations. Simultaneously, we must account for various theoretical and experimental uncertainties affecting our signal models. This section outlines the workflow used to evaluate and propagate these uncertainties throughout our analysis.

6.12.1 *Signal Yield Statistical Uncertainties*

The statistical uncertainty on our signal yield predictions comes from the limited size of the simulated signal samples. This uncertainty was estimated using the standard approach of considering the binomial uncertainty on the signal selection efficiency:

$$\sigma_{\text{stat}} = \sqrt{\frac{\varepsilon(1 - \varepsilon)}{N}} \quad (6.12)$$

where ε is the selection efficiency and N is the total number of simulated events. For all signal models, these statistical uncertainties are small (typically below 1%) due to the large

size of our simulated samples and therefore have a negligible impact on the final results.

6.12.2 *Background Estimation Statistical Uncertainties*

The statistical uncertainties on our background estimates arise from the limited statistics in the control regions used to calculate the EDP (Equation 6.1). The EDP is parameterized by event-level variables such as the total number of tracks and number of b -tagged jets, creating multi-dimensional histograms where many bins have low or zero counts. While event histograms with sufficient counts can use Gaussian approximations with $\sigma_i = \sqrt{N_i}$, the DV histograms have such low statistics that Poisson treatment is required for proper handling of the discrete counting statistics.

To accurately estimate these statistical uncertainties, we implemented a toy MC method with 10,000 pseudo-experiments. This technique generates many simulated datasets by randomly varying the observed data within its statistical uncertainties. For each experiment:

1. DV counts in each bin of the EDP numerator are varied by drawing from Poisson distributions: $N'_i \sim \text{Poisson}(N_i)$
2. Event counts in low-statistics bins are similarly varied using Poisson statistics
3. For extended statistics methods, the f -factor (Equation 6.2) is varied by drawing the total number of target DVs from $\text{Poisson}(N_{\text{target}})$
4. The background estimate is recalculated using all varied inputs

The standard deviation of the background estimates across all pseudo-experiments provides the statistical uncertainty. This approach correctly handles empty bins, which can fluctuate to non-zero values, and accounts for the asymmetric uncertainties inherent in low-count Poisson statistics.

6.12.3 Background Estimate Systematic Uncertainties

Beyond statistical fluctuations, we considered several sources of systematic uncertainties for background estimation:

6.12.3.1 Pile-Up Uncertainty

Events in the MTR systematically have higher pile-up than those in the PTR, as shown in Figure 6.18. To account for this difference, we reweighted events in the PTR to match the pile-up distribution in the MTR before computing our background estimates. The difference between estimates with and without this reweighting was assigned as a systematic uncertainty, which ranges from 2% to 5% depending on the SR.




Figure 6.18: Comparison of the number of pile-up vertices in the MTR and PTR. Figure from [160].

6.12.3.2 Non-linearity Uncertainty

For the SDV SR, we assessed whether our EDP parameterization fully captured all relevant dependencies by studying the correlation between the EDP and variables not used in the parameterization (such as $\sum p_T^2$, number of track jets, and number of EMTopo jets). By fitting the fraction of events with signal-like SDVs as a function of these variables, we identified that only the number of track jets showed a significant residual dependence, contributing an asymmetric uncertainty of $^{+16\%}_{-0\%}$ to the background estimate.

6.12.3.3 Non-closure Uncertainty

For the FDV signal regions, we evaluated the performance of our background estimation method in VRs, comparing predicted and observed counts. The maximum relative discrepancy observed was assigned as a systematic uncertainty, resulting in a 22% non-closure uncertainty for the 1 FDV SR and 20% for the 2 FDV SR.

6.12.4 *Signal Yield Systematic Uncertainties*

The systematic uncertainties on our signal yield predictions fall into two main categories: experimental uncertainties related to detector performance and reconstruction, and theoretical uncertainties related to signal modeling.

6.12.4.1 Track Reconstruction Uncertainties

The largest source of systematic uncertainty on the predicted signal yields arises from track reconstruction efficiency, particularly for the large impact parameter tracks that are essential for DV reconstruction. The LRT algorithm, discussed in Chapter 5.1.2, which is used to

recover these displaced tracks, has lower efficiency than standard tracking, and this efficiency degradation is not perfectly modeled in simulation.

To quantify this uncertainty, we use $K_S^0 \rightarrow \pi^+\pi^-$ candidate decays as a control sample to compare tracking efficiency between data and MC simulation [124, 127]. The method involves:

1. Identifying two-track K_S^0 candidates in both data and simulation using the same selection criteria as for the accidental crossing background estimate
2. Normalizing MC yields to match data in the inner detector region ($R_{xy} < 38$ mm) where standard tracking efficiency is well-understood
3. Measuring the data/ MC ratio as a function of radial position R_{xy}
4. Interpreting any deficit in the data/ MC ratio as tracking inefficiency not captured in simulation

The data/ MC comparison reveals increasing discrepancies at larger radii, as shown in Figure 6.19. Any excess in the MC simulation K_S^0 yield is associated with a tracking inefficiency that exists in data but is not modeled correctly in simulation. The resulting R_{xy} -dependent tracking inefficiency is derived from this comparison and combined with a baseline uncertainty of 1.8%, determined by the ATLAS tracking combined performance group.

The uncertainty is propagated to signal yields through a “track-killing” study where tracks are randomly removed from simulated DVs according to the measured inefficiency rates shown in Figure 6.20. After track removal, DV properties are recalculated and selection criteria reapplied. The standard deviation of the resulting yields across multiple iterations defines the systematic uncertainty.




Figure 6.19: Comparison of K_S^0 candidate yields between data (black points) and normalized MC simulation (red points) as a function of the DV radial position R_{xy} . The MC is normalized to match the data in the innermost bin, where the standard tracking efficiency is well understood. The bottom panel shows the data/ MC ratio, with deviations from unity indicating tracking inefficiencies not modeled in simulation. Figure from [160].




Figure 6.20: Estimated tracking inefficiency as a function of DV radial position R_{xy} , derived from the data/MC comparison of K_S^0 candidates. The inefficiency increases with radius, ranging from approximately 2% near the beam pipe to over 20% at the outer edge of the tracking volume. A baseline uncertainty of 1.8% has been added in quadrature. Figure from [160].




Figure 6.21: Illustration of the track-killing study methodology. Signal-like DVs from MC simulation (left) have tracks randomly removed according to the measured tracking inefficiency rates (right). The DV mass and track multiplicity are then recalculated to determine the impact on signal selection efficiency. Figure from [160].

This procedure yields tracking uncertainties ranging from 2-18% depending on the signal model, with longer-lived particles experiencing larger uncertainties due to their increased reliance on LRT. For the benchmark models, tracking systematic uncertainties reach 13.3% for the Gluino R -hadron model and 9.3% for the DFSZ Axino model, making this the dominant experimental systematic uncertainty. The track-killing methodology illustrated in Figure 6.21 ensures that the full impact of tracking inefficiencies on displaced vertex reconstruction is properly accounted for in the signal yield predictions.

6.12.4.2 Jet and MET Uncertainties

The JES, JER, and MET reconstruction all affect our signal yields through the event-level selection requirements. These uncertainties were evaluated following standard ATLAS procedures by varying the corresponding nuisance parameters up and down by one standard deviation and recalculating the signal yields [163, 164]. Four components were considered in this evaluation: the JES, the JER, the E_T^{miss} soft term scale, and the E_T^{miss} soft term resolution. The resulting uncertainties were combined in quadrature, assuming they are uncorrelated. These uncertainties are generally small (0%–2.5%) compared to the tracking uncertainties, with the most significant values observed for the Higgs portal model where the MET often comes from leptonic decays of vector bosons, making it more sensitive to the soft term uncertainties.

6.12.4.3 MET Trigger Uncertainties

To assess whether MET trigger efficiency is accurately modeled in MC simulation, we evaluated potential differences between data and MC for the lowest unprescaled MET triggers across the 2016-2018 data-taking periods [139]. Since the analysis uses an offline MET requirement of $E_T^{\text{miss}} > 150$ GeV, inaccurate trigger modeling in MC could bias the predicted

signal yields. This approach used muon-triggered events as a reference sample, measuring MET trigger efficiencies as a function of MET with muons treated as invisible particles (METNoMu). By recalculating the offline MET without including muon contributions, we create events that appear to have genuine missing transverse momentum from the trigger’s perspective, since the trigger does not correct for muons and they appear as MET.⁷ This method is effective because muons are not corrected for in online MET trigger calculations, muon triggers have high efficiency and well-understood performance, and events with muons appear indistinguishable from events with genuine MET at the trigger level.

The MET trigger performance study applies the following selections:

- $\text{MET}_{\text{LocHadTopo}} > 180 \text{ GeV}$
- $\text{METNoMu} > 50 \text{ GeV}$
- At least one muon in the event, determined by passing the reference trigger criteria

Trigger efficiency curves were produced for each data-taking period using the primary triggers shown in Table 6.23. For 2016 data, we analyzed two different triggers in separate periods⁸ since none remained unrescaled for the entire year. For 2017 and 2018, we used single triggers that remained unrescaled throughout each respective year, enabling full-year efficiency studies with improved statistical precision.

All triggers showed the expected plateau at 100% efficiency by approximately 180 GeV. Following the methodology described in Chapter 4.5.2, trigger efficiency turn-on curves were

7. Muons typically deposit only 2-4 GeV as minimum ionizing particles when traversing the calorimeter system. While offline MET calculations include corrections for muon energy losses, online trigger-level MET calculations at both L1 and HLT use only calorimeter information and do not correct for muons. The METNoMu variable therefore closely reproduces the MET value seen from the trigger’s perspective.

8. Data-taking periods refer to subdivisions of each year based on detector operating conditions, trigger configurations, and other experimental factors. The integrated luminosity for each period is shown in Table 6.1.


Year	Period	Lowest Unprescaled MET Trigger
2016	A-D3	HLT_xe90_mht_L1XE50
	D4-	HLT_xe110_mht_L1XE50
2017	All year	HLT_xe110_pufit_L1XE50
2018	All year	HLT_xe110_pufit_xe70_L1XE50

Table 6.23: Lowest unprescaled MET triggers used for the MET turn-on curve study. For 2016, prescales were added and removed from different triggers throughout the year, requiring separate period-dependent analyses. For 2017 and 2018, single triggers remained unprescaled throughout the data-taking years.


produced by comparing data and simulated $Z \rightarrow \mu\mu$ events. The data-to- MC efficiency ratios remained within 2% of unity for all bins with $\text{METNoMu} > 150$ GeV across all years and triggers, as demonstrated in Figure 6.22.

We also checked whether limiting our trigger study to only the lowest unprescaled triggers gave us the complete picture of trigger performance. The analysis has two MET requirements: first, any MET trigger from the complete `DRAW_RPVLL` list, discussed in Chapter 6.3.1.2, must fire to include the event in our data sample, and second, events must have offline $E_T^{\text{miss}} > 150$ GeV to enter the signal region. We only studied the subset of lowest unprescaled triggers, so we needed to verify this didn’t miss SR candidate events. We tested this by comparing signal efficiencies calculated using the complete trigger list with those calculated using only the lowest unprescaled triggers from Table 6.23. The resulting signal efficiencies were identical across all signal models and grid points. This makes sense because while events could enter our SR event selection by firing heavily prescaled MET triggers with lower thresholds, those events would need offline $E_T^{\text{miss}} > 150$ GeV to reach the signal region, which means they would also probably fire the lowest unprescaled triggers we studied.


This validation confirms that our trigger efficiency studies, using only the lowest unprescaled triggers, provide a comprehensive assessment of trigger performance for events that reach our SRs.




(a) 2016 periods A-D3



(b) 2016 periods D4 onwards



(c) 2017 all year



(d) 2018 all year

Figure 6.22: MET trigger efficiency turn-on curves as a function of METNoMu for data (blue) and simulated $Z \rightarrow \mu\mu$ events (red) across all data-taking years. The lowest unprescaled triggers used are: HLT_xe90_mht_L1XE50 (2016 A-D3), HLT_xe110_mht_L1XE50 (2016 D4 onwards, 2017), and HLT_xe110_pufit_xe70_L1XE50 (2018). The numbers below each histogram bin indicate the number of events in the passed histogram for data (blue) and MC (red), providing a sense of the statistics available for the efficiency measurement. The bottom panel of each plot shows the ratio of data to MC efficiency, demonstrating agreement within 2% in the region above 150 GeV across all triggers and periods. Figures from [160]




Figure 6.23: Percent difference in signal efficiency for the R-hadron model when using the complete DRAW_RPVLL trigger list versus using only the lowest unprescaled triggers from Table 6.23. The signal efficiency difference is at the level of $\sim 10^{-14}\%$ across all grid points, confirming that limiting the trigger efficiency study to only the lowest unprescaled triggers captures the complete trigger performance for SR candidate events. Figure from [160]

To validate the trigger efficiency measurement in the specific context of our analysis’s event topology, we performed an additional study requiring events to contain at least one DV passing the baseline selection criteria, including fiducial volume requirements, primary vertex separation, vertex fit quality, and material map veto. While the statistical precision was reduced, particularly for 2016 data periods, the data-to- MC agreement remained consistent with the all-inclusive studies, confirming that trigger efficiency modeling is accurate even in the presence of DV signatures.




Figure 6.24: MET trigger efficiency turn-on curves for HLT_xe110_pufit_L1XE50 as a function of METNoMu for 2017 data (blue) and simulated $Z \rightarrow \mu\mu$ events (red), with the additional requirement of at least one SDV with at least 2 tracks passing baseline selection criteria. The numbers below each histogram bin indicate the number of events in the passed histogram for data (blue) and MC (red). While the statistics are substantially reduced compared to the inclusive studies due to the rarity of long-lived particle signatures, the bottom panel shows the ratio of data to MC efficiency continues to demonstrate good agreement in the region above 150 GeV. Figure from [160]

Trigger efficiency was also validated using signal MC samples for both the Gluino R -hadron and Bino-Wino Coannihilation models. Using the same muon trigger reference selection method, we measured MET trigger efficiency as a function of offline MET across represen-

tative samples spanning the range of masses and lifetimes considered in the analysis. Since signal events contain genuine missing transverse energy from neutralinos rather than muons, this validation uses the full offline MET rather than METNoMu. Given the analysis requirement of $E_T^{\text{miss}} > 150$ GeV, signal events predominantly populate the trigger plateau region where efficiency approaches 100%, so this study validates the plateau modeling rather than investigating turn-on behavior. The validation showed agreement with data at the better-than-3% level for $E_T^{\text{miss}} > 150$ GeV, confirming that trigger modeling in signal MC accurately represents the observed data performance.




Figure 6.25: MET trigger efficiency as a function of offline MET for 2017 data (blue) and simulated R-hadron signal events (red) using HLT_xe110_pufit_L1XE50. Unlike the $Z \rightarrow \mu\mu$ validation studies, this uses the full offline MET and shows plateau behavior rather than turn-on curves due to the analysis requirement of $E_T^{\text{miss}} > 150$ GeV. The bottom panel shows the ratio of data to MC efficiency, demonstrating agreement within 3% in the high-MET plateau region. Figure from [160].

Based on the consistent agreement between data and MC in the turn-on region across all triggers, selections, and validation studies, no scale factor correction was applied to the MC samples. The systematic uncertainty associated with MET trigger modeling was determined

to be negligible compared to other sources of uncertainty in the analysis, and no additional uncertainty is added.

6.12.4.4 Theoretical Uncertainties

Theoretical uncertainties on signal production cross-sections and acceptance are evaluated following PDF4LHC15 recommendations [165]. These include cross-section uncertainties calculated at NNLO+ NNLL for the Gluino R -hadron model, NLO+ NLL for the Bino-Wino Coannihilation model, and NLO for DFSZ Axino and Higgs portal models; scale variations (factors of two); and PDF/ α_s uncertainties evaluated using on-the-fly event weights. Signal acceptance uncertainties range from 0.3% to 10.8%, depending on the model and parameters, as shown in Table 6.25, while cross-section uncertainties increase with particle mass. These uncertainties are calculated using event weights generated during MC production with PYTHIA [153] and NNPDF parton distribution functions [155].

6.12.4.5 Additional Uncertainties

Several additional sources of uncertainty on the predicted signal yields were considered:

- **pile-up reweighting:** To account for differences between the pile-up profile in simulation and data, signal events were reweighted to match the data distribution. The uncertainty on signal yields from this procedure was found to be negligible.
- **Luminosity:** The uncertainty on the integrated luminosity measurement (0.84%) affects the normalization of predicted signal yields but is small compared to other systematic uncertainties [166].

6.12.5 Summary of Uncertainties

6.12.5.1 Background Uncertainty Summary

Table 6.24 summarizes the uncertainties on our background estimates. In all SRs, the statistical uncertainty dominates, reflecting the challenge of estimating extremely rare background processes.

Region	Estimate	Uncertainty components		
		Statistical	Non-linearity/Non-closure	Pile-Up
SDV SR	$0.56^{+0.40}_{-0.39}$	$\pm 70\%$	$+16\%$ -0%	$\pm 5\%$
1 FDV SR	0.81 ± 0.49	$\pm 57\%$	$\pm 22\%$	$\pm 2\%$
2 FDV SR	1.29 ± 0.64	$\pm 45\%$	$\pm 20\%$	$\pm 4\%$

Table 6.24: Summary of background estimate uncertainties by signal region

6.12.5.2 Signal Uncertainty Summary

Table 6.25 summarizes the ranges of systematic uncertainties for each signal model, grouped by SR. The values represent the 16th to 84th percentile interval across all relevant grid points.

Systematic uncertainty	Gluino R -hadron SDV	DFSZ Axino (hh) SDV	Higgs portal 2 FDV	Bino-Wino Coannihilation 1 FDV
JES + JER	0.00%–0.16%	0.00%–0.14%	0.00%–2.51%	0.00%–0.55%
MET	0.00%–0.13%	0.00%–0.13%	0.00%–1.68%	0.00%–0.30%
Theory	4.95%–10.76%	0.30%–1.45%	1.24%–10.49%	1.14%–5.51%
Tracking	2.10%–13.26%	1.76%–9.31%	6.40%–18.28%	2.25%–9.62%

Table 6.25: Ranges of systematic uncertainties (16th–84th percentile) per signal model

6.12.6 *Impact of Uncertainties on Results*

For signal models, tracking uncertainties and theoretical uncertainties dominate, particularly for models with longer lifetimes. For the background estimate, statistical uncertainties are the dominant factor, reflecting the extreme rarity of background processes that can mimic our signal.

It's worth noting that while these systematic uncertainties are essential for proper modeling, the statistical uncertainty from the limited number of observed events in our signal regions (1, 3, and 2 events in the SDV, 1 FDV, and 2 FDV SRs, respectively) ultimately dominates our sensitivity in this search.

6.13 Statistical Analysis and Limit Setting

In particle physics searches, when no significant excess of events is observed over the expected background, we set upper limits on the production cross section of the hypothetical signal process. More specifically, since we did not observe the signal we were searching for, we can definitively state that the production cross section cannot exceed a certain threshold value—because if it were above this limit, we would have detected a statistically significant number of signal events in our experiment.

This statistical procedure allows us to constrain or exclude regions of the theoretical model parameter space by comparing three key quantities:

1. **Observed cross section limit** ($\sigma_{\text{obs}}^{\text{limit}}$): The maximum cross section our experiment can rule out based on the actual data collected
2. **Expected cross section limit** ($\sigma_{\text{exp}}^{\text{limit}}$): The maximum cross section we would have

been able to rule out if we had observed exactly the predicted background, demonstrating our experimental sensitivity under background-only assumptions

3. **Theory predicted cross section** (σ_{theory}): The production rate that theoretical models predict should occur for specific combinations of model parameters

Exclusion occurs when the theory-predicted cross section exceeds our observed limit, i.e., when $\sigma_{\text{theory}} > \sigma_{\text{obs}}^{\text{limit}}$ —meaning the theoretical prediction is incompatible with our non-observation of the signal. Conversely, parameter regions where $\sigma_{\text{theory}} < \sigma_{\text{obs}}^{\text{limit}}$ remain experimentally allowed, as such low production rates would generate too few events for our experiment to have detected them above background fluctuations. This systematic comparison between experimental limits and theoretical predictions allows us to map out which regions of BSM model parameter space are definitively ruled out by our search sensitivity.

6.13.1 Maximum Likelihood Fitting

The statistical analysis employs a maximum likelihood approach to compare the observed data with predictions under different signal hypotheses. We construct a likelihood function [167] that describes the probability of observing our data given a particular signal strength:

$$\mathcal{L}(\text{data}|\mu, \boldsymbol{\theta}) = \text{Poisson}(n|\mu S + B) \cdot \prod_i f_{\theta,i}(\theta_i), \quad (6.13)$$

where:

- μ is the signal strength parameter, with $\mu = 0$ corresponding to the background-only hypothesis and $\mu = 1$ to the nominal signal hypothesis

- S and B are the expected numbers of signal and background events
- n is the observed number of events
- $\boldsymbol{\theta}$ represents nuisance parameters encoding systematic uncertainties
- $f_{\theta,i}(\theta_i)$ are constraint terms for the systematic uncertainties

The nuisance parameters $\boldsymbol{\theta}$ are auxiliary parameters that account for systematic uncertainties affecting our signal and background predictions, such as the tracking efficiency uncertainties, jet energy scale variations, and background estimation uncertainties detailed in Chapter 6.12. While we are primarily interested in constraining the signal strength μ , we must simultaneously account for these nuisance parameters to avoid biased results. Each nuisance parameter θ_i is constrained by auxiliary measurements through the terms $f_{\theta,i}(\theta_i)$, typically Gaussian distributions centered on the measured values with widths determined by the corresponding systematic uncertainties. During the statistical analysis, these nuisance parameters are profiled, meaning their values are adjusted to maximize the likelihood for each tested signal strength, allowing the data to determine the most compatible values of the systematic uncertainties while properly propagating their impact on our conclusions [167].

The signal strength μ can then be interpreted as a scaling factor on the theoretical production cross section, such that the actual cross section being tested is $\sigma = \mu \cdot \sigma_{\text{theory}}$.

The statistical analysis is implemented using the `pyhf` 0.7.6 [168, 169] and `cabinetry` 0.6.0 [170] software packages following standard ATLAS procedures.

6.13.2 Test Statistic and CLS Method

To quantify the compatibility between a given signal hypothesis and the observed data, we construct a profile likelihood ratio test statistic:

$$q_\mu = -2 \ln \frac{\mathcal{L}(\mu, \hat{\boldsymbol{\theta}}(\mu))}{\mathcal{L}(\hat{\mu}, \hat{\boldsymbol{\theta}})} \quad (6.14)$$

where $\hat{\mu}$ and $\hat{\boldsymbol{\theta}}$ are the values that maximize the likelihood, and $\hat{\boldsymbol{\theta}}(\mu)$ maximizes the likelihood for a fixed value of μ .

Due to the low event counts in our SRs, we cannot rely on asymptotic approximations for the distribution of the test statistic. Instead, we perform 100,000 pseudo-experiments for each value of μ , varying the nuisance parameters within their uncertainties to determine the probability density function $f(q_\mu|\mu)$.

From these distributions, we calculate p-values for the signal-plus-background (p_{s+b}) and background-only (p_b) hypotheses. We then apply the CL_s method [171, 172] to avoid excluding signal models to which the analysis has insufficient sensitivity:

$$\text{CL}_s = \frac{p_{s+b}}{1 - p_b} \quad (6.15)$$

The CL_s approach prevents us from excluding signal models in cases where our analysis lacks sufficient sensitivity. This is crucial when the expected signal is much smaller than the background uncertainty—a situation where a simple p-value test could lead to spurious exclusions. By normalizing p_{s+b} with $(1 - p_b)$, we only exclude signals when there is genuine tension between the theoretical prediction and our observations.

A signal hypothesis at strength μ is excluded at 95% confidence level if $\text{CL}_s < 0.05$.

6.13.3 Upper Limit Setting

The limit-setting procedure determines whether each combination of signal model parameters is compatible with our observed data by systematically testing different signal strength hypotheses and identifying the boundary between excluded and allowed regions in parameter space. For each point in a model’s parameter grid, we apply the statistical framework described above to find the critical signal strength $\mu_{95\%}$ where $\text{CL}_s = 0.05$, corresponding to the maximum signal rate that can be excluded at 95% confidence level. This process transforms our discrete observations into continuous exclusion contours that map out the sensitivity reach of our search across the full range of theoretical predictions.

6.13.3.1 Procedure for Individual Grid Points

For each grid point in a model’s parameter space, we perform the following procedure:

1. **Theoretical prediction:** The model parameters yield a theoretical cross section σ_{theory} and corresponding expected signal yield in our detector, determined using simulated signal events and accounting for detector acceptance and selection efficiency, discussed in Chapter 6.8.2.
2. **Signal strength scan:** We scan over signal strength hypotheses μ , where $\mu = 0$ corresponds to the background-only hypothesis and $\mu = 1$ corresponds to the nominal theoretical prediction.
3. **Critical signal strength:** Using the statistical framework described above, we determine $\mu_{95\%}$ —the maximum signal strength compatible with our data at 95% confidence level, defined by $\text{CL}_s = 0.05$.

4. **Exclusion determination:** The grid point is excluded if $\mu_{95\%} < 1$, meaning even a signal weaker than the theoretical prediction is incompatible with our observations.

This yields an upper limit on the production cross section:

$$\sigma_{95\%CL} = \mu_{95\%CL} \cdot \sigma_{\text{theory}} \quad (6.16)$$

Grid points where $\sigma_{\text{theory}} > \sigma_{95\%CL}$ (equivalently, $\mu_{95\%} < 1$) are excluded at 95% confidence level.

For example, if an DFSZ Axino grid point predicts $\sigma_{\text{theory}} = 0.5$ fb and we find $\mu_{95\%} = 0.7$, then $\sigma_{95\%CL} = 0.35$ fb. Since $0.5 > 0.35$ fb, this parameter point is excluded.

6.13.3.2 Constructing Exclusion Contours

By scanning across the complete model parameter space and applying this procedure to each grid point, we construct exclusion limits as a function of the model parameters. Parameter values where the theoretical cross section exceeds our observed limit ($\sigma_{\text{theory}} > \sigma_{95\%CL}$, equivalent to $\mu_{95\%} < 1$) are excluded at 95% confidence level.

6.13.4 Setting Limits on the DFSZ Axino

Our search for DFSZ Axino signatures in the ATLAS detector did not reveal a significant excess of events above the expected SM background. In the SDV SR, we observed 1 event compared to an expected background of $0.56^{+0.40}_{-0.39}$ events estimated using the data-driven methods detailed in Chapter 6.10. This observation allows us to constrain the DFSZ Axino model and set upper limits on the cross section for producing DV plus MET signatures

through long-lived neutralino decays to axinos.

For the DFSZ Axino model introduced in Chapter 6.7, the SDV SR provides the highest expected signal yields and therefore the strongest sensitivity among our defined SRs.

6.13.5 DFSZ Axino Model Results and Interpretation

For each combination of DFSZ Axino model parameters (neutralino mass $m_{\tilde{\chi}_1^0}$, axion mass m_a , and axino mass $m_{\tilde{a}}$), we apply the limit-setting workflow described above to determine upper limits on the DFSZ Axino cross section.

The exclusion limits shown in Figures 6.26 and 6.27 demonstrate significant constraints on the DFSZ Axino model parameter space, providing the first direct collider limits on this DFSZ Axion- SUSY framework. Our search excludes cross sections typically ranging from 10^{-1} to 10^1 fb across most of the explored parameter combinations for both hh and ZZ decay channels.

The characteristic shape of our limits, where sensitivity is strongest at intermediate axion masses around $m_a \sim 1000 \mu\text{eV}$, reflects the fundamental physics of searches DV signatures. This behavior arises from the relationship between axion mass, neutralino lifetime, and detector acceptance:

- Short lifetimes: Neutralinos decay too promptly, producing vertices too close to the interaction point
- Optimal lifetimes: Neutralinos decay at $\mathcal{O}(10 - 100)$ mm, ideal for reconstruction within the ATLAS ID
- Long lifetimes: Neutralinos travel beyond the ID before decaying




Figure 6.26: Exclusion limits on the DFSZ Axino model at 95% confidence level as a function of axion mass with fixed neutralino mass of 800 GeV, and axino mass of 400 GeV. Both hh and ZZ decay channels are shown. Solid lines represent observed limits, while dotted lines indicate expected limits with uncertainty bands. Regions where the theoretical cross section (gray line) exceeds the observed limit (solid colored lines) are excluded at 95% confidence level. Figure from [141].




Figure 6.27: Exclusion limits on the DFSZ Axino model at 95% confidence level as a function of axino mass with fixed neutralino mass of 600 GeV. Both hh and ZZ decay channels are shown. Solid lines represent observed limits, while dotted lines indicate expected limits with uncertainty bands. Regions where the theoretical cross section (gray line) exceeds the observed limit (solid colored lines) are excluded at 95% confidence level. Figure from [141].

The sharp uptick in the ZZ channel limits (blue curve) between 400 and 500 GeV in Figure 6.27 is due to kinematic thresholds for the neutralino decay $\tilde{\chi}_1^0 \rightarrow \tilde{a} + Z$. With the neutralino mass fixed at 600 GeV and requiring the Z boson to be on-shell ($m_Z = 91$ GeV), the kinematic limit occurs at $m_{\tilde{a}} < 509$ GeV. As the axino mass approaches this threshold, the available energy for the $\tilde{a}$ decreases dramatically, resulting in very low MET and in turn reduced signal acceptance. The apparent sharp discontinuity between 400 and 500 GeV is an artifact of sparse grid spacing; the line connects two available grid points at these masses. The hh channel (orange curve) terminates at 400 GeV due to the heavier Higgs boson ($m_h = 125$ GeV) reaching its kinematic limit earlier.

Overall, this search probes axion decay constants in the range $f_a \sim 10^9 - 10^{12}$ GeV, an intermediate regime between laboratory axion searches and astrophysical constraints.

6.13.6 Gluino R Hadron Model Results

Following the same statistical framework, we derive exclusion limits for the Gluino R -hadron model using the SDV SR, where the high track multiplicity from gluino decays provides optimal sensitivity.

As shown in Figure 6.28, the left panel demonstrates the characteristic lifetime dependence of displaced vertex searches, with optimal sensitivity occurring at intermediate lifetimes around 0.1 ns where the gluino decay length matches our detector acceptance. The exclusion reaches its maximum strength in this optimal lifetime region, with sensitivity decreasing for both very short lifetimes (where vertices are too close to the interaction point) and very long lifetimes (where gluinos escape the detector before decaying).

The right panel shows exclusion limits as a function of gluino and neutralino masses for the optimal lifetime of 0.1 ns. The kinematically forbidden region where $m_{\tilde{\chi}_1^0} > m_{\tilde{g}}$ is clearly




Figure 6.28: 95% confidence level observed and expected upper limits on the cross section of the Gluino R -hadron model. The left plot shows the exclusion as a function of gluino lifetime and mass, for R -hadrons with the lightest neutralino mass $m_{\tilde{\chi}_1^0} = 100$ GeV. The right plot shows the exclusion as a function of gluino and neutralino mass, for gluino lifetime $\tau(\tilde{g}) = 0.1$ ns, with the kinematically forbidden region $m_{\tilde{\chi}_1^0} > m_{\tilde{g}}$ also indicated. The dashed black line and yellow band show the expected limit and uncertainty band from the fit, while the solid and dashed red lines show the calculated observed limit and uncertainty from varying the R -hadron production cross section. The area below the contour is excluded. Figure from [141].

indicated, and the limits become less stringent for compressed mass spectra where the mass splitting between the gluino and neutralino is small.

These results represent substantial improvements over previous ATLAS displaced vertex searches, extending the excluded parameter space across most of the explored gluino mass and lifetime range.

6.13.7 Bino-Wino Coannihilation Model Results

The Bino-Wino Coannihilation model described in Chapter 6.7 performs best in the FV SRs, as expected due to the b -jets in the final state. The 1 FV SR was specifically optimized for this model's signature.

As shown in Figure 6.29, the left panel demonstrates the characteristic lifetime dependence




Figure 6.29: 95% confidence level observed and expected upper limits on the cross section of the Bino-Wino Coannihilation model, with exclusion contours as described in Figure 6.28. The left plot shows the exclusion as a function of $\tilde{\chi}_2^0$ lifetime and mass with a mass splitting $\Delta m = 50$ GeV between $\tilde{\chi}_2^0$ and $\tilde{\chi}_1^0$. The right plot shows the exclusion as a function of mass splitting and $\tilde{\chi}_1^0$ mass with a lifetime $\tau(\tilde{\chi}_2^0) = 0.03$ ns. The weaker constraint on the observed limit compared to the expected limit is due to a small excess in the FV SR. Figure from [141].

for this model, with optimal sensitivity occurring at intermediate lifetimes around 0.1 ns. We exclude $\tilde{\chi}_2^0$ masses up to approximately 800 GeV for lifetimes in this range, with sensitivity decreasing for very short or very long lifetimes due to acceptance effects of LLPs in ATLAS.

The right panel shows the exclusion as a function of mass splitting between $\tilde{\chi}_2^0$ and $\tilde{\chi}_1^0$ for a fixed short lifetime of 0.03 ns. The observed limits are slightly weaker than the expected limits due to a small excess observed in the FV SRs, though this excess is not statistically significant. The FV approach demonstrates clear advantages for this model, achieving significantly better sensitivity than would be possible with standard displaced vertex reconstruction alone.

6.13.8 Higgs portal model Results

The Higgs portal model achieves the strongest constraints in the FV SRs, particularly benefiting from the two-vertex topology when both scalar particles decay within the detector.




Figure 6.30: Expected and observed limits on the branching ratio of $h \rightarrow SS \rightarrow 4b$ in the Higgs portal model at 95% confidence level, shown as a function of the lifetime τ of the new heavy scalar S . Limits are evaluated for several different values of the scalar mass m_S , and compared to a 10.7% branching ratio for the Higgs boson decaying into BSM final states [173]. Each color indicates a specific value of m_S . The weaker constraint on the observed limit compared to the expected limit is due to a small excess in FV SR. Figure from [141].

As demonstrated in Figure 6.30, we set limits on the branching ratio $\mathcal{B}(h \rightarrow SS \rightarrow 4b)$ ranging from approximately 2-17% depending on the scalar mass and lifetime, where each point represents the maximum branching ratio that can be excluded for the given parameter combination. The strongest constraints are achieved for intermediate lifetimes around $\tau_S \sim 0.03$ ns, where the scalar decay length is well-matched to our detector acceptance. For scalar masses above 35 GeV and optimal lifetimes, we exclude branching ratios as low as 2-4%, well below the 10.7% threshold for exotic Higgs decays beyond the Standard Model.

The observed limits are consistently weaker than expected due to the small excess observed

in the FV SRs, though this deviation is not statistically significant. These results demonstrate the power of DV signatures for constraining exotic Higgs decay scenarios and provide important constraints on Higgs portal models across a wide range of parameter space.

6.13.9 Final Remarks

Our search for BSM physics with DV and MET signatures has provided strong constraints on several theoretical scenarios. This analysis demonstrates how data-driven background estimation methods can be effective even when traditional Monte Carlo simulations lack sufficient statistics to accurately model rare background processes.

We didn't see any clear signs of new physics. While this may seem disappointing, it actually provides us with meaningful limits on various theoretical models. We've set the first direct collider constraints on DFSZ Axino models, pushed existing limits on Gluino R -hadron scenarios much further, and established competitive constraints on Bino-Wino Coannihilation and Higgs portal models.

We used data-driven background estimates that don't rely on MC simulations, and our observed sensitivity closely matches these estimates. The analysis utilizes most of the Run 2 dataset (137 fb^{-1} collected from 2016-2018) and benefits from advanced reconstruction techniques, including the new FV algorithm designed for final states with heavy-flavor quarks.

Looking forward, the continuation of this search is promising. Run 3 data collection is already well underway, with the new gFEX triggers described in Chapter 4, and the planned HL-LHC will provide even larger datasets. Improved detectors, better triggers designed specifically for displaced signatures, and higher collision rates will allow future searches to probe much deeper into parameter space, hopefully uncovering the new physics that continues to evade detection.

Beyond setting the first direct collider constraints on the DFSZ Axino model, this analysis demonstrates how collider searches can complement broader axion searches by probing regions of parameter space that are otherwise difficult to access. Figure 6.31 shows the axion-photon coupling parameter space as a function of axion mass m_a and decay constant f_a . The two grey shaded regions at the top represent parameter space excluded by direct detection experiments: the lower grey band corresponds to the scenario without SUSY, while the higher band shows how these limits shift when the axion-photon coupling is suppressed, as expected in supersymmetric models like the PQMSSM, discussed in Chapter 2.5.

Direct detection experiments lose sensitivity in this regime of suppressed coupling, leaving a greater range of parameter space unprobed. This is where collider searches offer unique reach by probing the axion parameter phase space through the axino: the blue, yellow, and orange bands show the limits set by this DV + MET analysis for the representative benchmark grid points $(m_{\tilde{a}}, m_{\tilde{H}})$ of (0 GeV, 400 GeV), (300 GeV, 600 GeV), and (500 GeV, 1000 GeV), respectively, where $m_{\tilde{H}}$ denotes the Higgsino mass.⁹ These limits are derived from our ATLAS analysis results, which are reinterpreted in the axion parameter space through the theoretical connections between axino and axion properties in the PQMSSM framework.

This figure highlights that the collider-accessible lifetime regime overlaps with regions left unprobed by direct detection, demonstrating how displaced vertex searches can extend coverage of the axion parameter space even when the axion-photon coupling is suppressed by SUSY.

These axino-derived limits therefore not only test SUSY axion models like the PQMSSM, but also extend experimental coverage of axion decay constants f_a . The complementary nature of collider searches and direct detection experiments becomes evident when considering that

9. Here, the Higgsino mass is equivalent to the mass of the lightest neutralino, $m_{\tilde{\chi}_1^0}$.




Figure 6.31: Axion-photon coupling parameter space as a function of axion mass m_a and decay constant f_a , showing experimental constraints and theoretical predictions. The shaded curves labeled “Photon Coupling (SUSY)” and “Photon Coupling (No SUSY)” correspond to models with and without supersymmetric contributions, respectively. Markers indicate representative DFSZ Axino benchmark points for which displaced vertex signatures are accessible at the LHC via neutralino decays. This preliminary figure illustrates how the collider-accessible lifetime regime maps onto the axion mass scale, allowing reinterpretation of our exclusion limits in the broader axion parameter space. Figure from [63].

different experimental approaches are sensitive to different regions of the parameter space, with collider searches offering unique access to these specific lifetime and mass ranges that may be challenging for other experimental techniques to probe.

CHAPTER 7

CONCLUSION

This thesis presented a search for LLPs decaying to DVs and MET, motivated by supersymmetric models in which the DFSZ axino serves as both a dark matter candidate and the LSP. The analysis was performed using the full Run 2 dataset collected by the ATLAS detector at the LHC, corresponding to 137 fb^{-1} of pp collisions at $\sqrt{s} = 13 \text{ TeV}$. Candidate events were selected using a MET trigger and required to contain a displaced vertex reconstructed using either the VSI or fuzzy vertexing algorithms.

Backgrounds were estimated using data-driven methods with dedicated CRs and VRs targeting specific sources such as HIs, AXs, and MVs. No significant excess was observed above the SM expectation, and upper limits were placed on several benchmark signal models. Notably, this analysis set the first collider limits on the displaced PQMSSM axino model, excluding new regions of parameter space for long-lived neutralinos decaying to QCD axinos.

Beyond the physics analysis, this work also contributed to the commissioning and performance monitoring of the gFEX system. The gFEX performs real-time reconstruction of global event quantities such as MET and large-radius jets and is a core component of the Run 3 L1Calo trigger upgrade. Firmware monitoring, calibration validation, and efficiency tracking tools developed during this work helped ensure stable and reliable performance, with particular emphasis on maintaining high efficiency for MET-based triggers targeting exotic signatures.

This search provides a novel probe of the axion parameter space via the axino channel. By targeting displaced neutralino decays to axinos, the analysis accesses a previously unexplored region of axion sensitivity. The ability to translate collider constraints on axino production into limits on the axion decay constant f_a demonstrates how particle collider experiments

can probe the axion parameter space through supersymmetric channels. This axino-based approach opens a new avenue for exploring axion dark matter, complementing traditional detection techniques with sensitivity to well-motivated supersymmetric scenarios.

As Run 3 progresses and the ATLAS experiment prepares for the HL-LHC, the methods developed in this thesis provide a robust foundation for future searches. Continued improvements in trigger systems, vertex reconstruction, and background modeling will enable even deeper exploration of long-lived particles and dark sector physics at the energy frontier.

Appendices

APPENDIX A

DIRECT DETECTION AXION SEARCHES

This appendix provides additional detail on the primary experimental approaches for direct axion detection, expanding on the brief overview presented in Chapter 2.3.1.3. While the main body of this thesis focuses on axino searches, these direct axion detection techniques provide important context for understanding the broader theoretical framework that motivates axino phenomenology.

Four distinct experimental strategies have emerged, each employing different technologies and innovations to detect axions across complementary mass ranges: solar searches exploit axions produced in stellar environments, laboratory light-shining-through-wall experiments create and detect axions through photon-axion conversion, cavity resonance methods probe galactic dark matter axions, and emerging conducting surface techniques enable broadband frequency coverage.

A.1 Solar Searches

Solar axion searches rely on the Primakoff process, through which axions are produced in the Sun via electromagnetic interactions with charged particles in strong Coulomb fields. This creates a continuous flux of axions reaching Earth, with energies peaking between 1–10 keV [174].

Helioscope experiments detect these solar axions through the inverse Primakoff process, where axions convert back into X-ray photons within strong laboratory magnetic fields. The conversion probability scales with both magnetic field strength and interaction path length. These instruments track the Sun’s apparent motion to maximize exposure during




Figure A.1: Schematic of the CAST experiment, which detects solar axions by enhancing their conversion probability using a powerful perpendicular magnetic field. The resulting X-ray photons are detected at the far end of the experimental setup. Figure from [175].

data collection periods.

The primary helioscope experiments employing this method are CAST and the upcoming BabyIAXO.

A.1.1 CAST Experiment

The CERN Axion Solar Telescope (CAST) operates with a 9 T magnetic field over a 9.26 m interaction length. As the first helioscope to incorporate X-ray focusing optics, CAST established the world's most stringent laboratory constraints on axion-photon coupling: $g_{a\gamma\gamma} < 0.66 \times 10^{-10} \text{ GeV}^{-1}$ for axion masses below 0.02 eV [176, 54].

Since 2015, CAST has expanded beyond solar axion searches to investigate chameleons and other dark sector particles, operating as both helioscope and haloscope with RADES and CAST-CAPP sub-detectors for galactic dark matter axion searches [177, 178].

A.1.2 Next-Generation Helioscopes

BabyIAXO, under construction at Deutsches Elektronen-Synchrotron (DESY) with data collection anticipated between 2025–2028, features a 10-meter toroidal magnet with dual 70 cm diameter bores operating at 2.5 T, targeting sensitivity improvements of approximately $100\times$ over CAST [179]. The Tokyo Axion Helioscope continues operations with its 4 T, 2.3-meter design, achieving tracking precision of ± 0.5 mrad [180].

A.2 Laboratory Photon-Axion Conversion Searches

Light-shining-through-wall (Light Shining through Walls (LSW)) experiments test axion-photon coupling under controlled laboratory conditions. These setups employ two optical cavities separated by an opaque barrier, each within strong perpendicular magnetic fields. Laboratory laser light undergoes partial photon-to-axion conversion in the first cavity, with axions penetrating the barrier while photons are blocked. The second cavity converts axions back into detectable photons. This approach allows axions to be produced and detected entirely within the laboratory [174].

The fundamental challenge lies in the extremely low probability of each conversion step, requiring two consecutive coupling events for signal detection. The primary experiments using this approach are ALPS-II, Optical Search for QED Vacuum Bifringence, Axions and Photon Regeneration (OSQAR), and the completed CERN Resonant Weakly Interacting sub-eV Particle Search (CROWS) experiment.




Figure A.2: Schematic of the light shining through walls experimental method. Laboratory laser light undergoes partial photon-to-axion conversion in a strong magnetic field. Only axions penetrate the optical barrier, where they convert back into photons in the second magnetic field region for detection. Figure from [181].

A.2.1 ALPS Experiments

ALPS-II at DESY represents the most advanced LSW experiment, completing its inaugural science run in mid-2023. Utilizing 24 recycled HERA dipole magnets providing 5.3 T fields across dual 120-meter optical cavities, the experiment achieves cavity storage times of 7 milliseconds and detection sensitivities down to $2 \times 10^{-5} \text{ s}^{-1}$ photon flux. ALPS-II targets axion-photon coupling sensitivity of $2 \times 10^{-11} \text{ GeV}^{-1}$ for masses below 0.1 meV [182].

A.2.2 Other LSW Experiments

OSQAR at CERN employs two LHC spare dipole magnets at 9 T over 28.6 meters total interaction length, achieving coupling limits below $8.0 \times 10^{-8} \text{ GeV}^{-1}$ [183]. CROWS completed microwave-based measurements in 2013, improving sensitivity to hidden sector photons [184].

A.3 Cavity Resonance Searches

Cavity haloscope experiments search for galactic dark matter axions by exploiting resonant conversion within high-quality-factor microwave cavities permeated by strong magnetic fields. When the cavity resonance frequency matches the axion mass, dark matter axions from the galactic halo may convert into detectable microwave photons. This requires systematic frequency scanning across target mass ranges, with the cavity tuned in small, overlapping steps to ensure comprehensive coverage [174].

The leading experiments using this technique include ADMX, Center for Axion and Precision Physics Research (CAPP), HAYSTAC, QUaerere AXion (QUAX), and Oscillating Resonant Group AxioN (ORGAN).

A.3.1 ADMX

ADMX has operated since 1995, targeting the μeV axion mass range through progressively improved cryogenics and detector technologies. The experiment employs moveable copper or dielectric rods for cavity tuning and has evolved from pumped ^4He cooling with 2 K noise amplifiers to current $^3\text{He}/^4\text{He}$ dilution refrigeration targeting 100 mK base temperatures [185, 186, 187, 188, 19].

Recent ADMX searches covered 1.10–1.31 GHz (4.55–5.43 μeV), achieving DFSZ theoretical sensitivity benchmarks required for detecting QCD axions [189].

A.3.2 Other Cavity Experiments

ADMX established the benchmark for cavity-based axion searches, with subsequent experiments building on this foundation to explore different mass ranges and implement technological improvements.

CAPP (South Korea) achieved the first DFSZ sensitivity above 1 GHz using multi-cell cavity architectures with 18 T high-temperature superconducting magnets [190]. HAYSTAC (Yale) pioneered squeezed-state receivers, achieving scan rate improvements through quantum noise reduction in the 20–25 μeV range [191]. QUAX (Italy) explores the 36–37 μeV range using high-Q sapphire cavities [192]. ORGAN (Australia) targets high-mass regions around 110 μeV using innovative pipe-organ resonator designs [193].

A.4 Conducting Surface Experiments

Conducting surface experiments enable broadband axion detection by exploiting axion-to-photon conversion at interfaces between materials with different electromagnetic properties. Unlike cavity methods requiring sequential frequency scanning, these experiments can simultaneously probe wide mass ranges. When axions encounter conducting or dielectric boundaries in magnetic fields, they convert to electromagnetic radiation with efficiencies enhanced through optimized material arrangements.

The primary experiments using this approach are MADMAX and BREAD.




Figure A.3: Schematic of the MADMAX experimental method. Multiple dielectric disks positioned within a strong magnetic field induce axion-photon conversion at each interface. Signal amplification through the disk array enables broadband detection across the μeV mass range. Figure from [194].

A.4.1 MADMAX

MADMAX targets the 40–400 μeV mass range using arrays of dielectric disks within strong magnetic fields. Incident axions convert to photons at each disk interface, with signal amplification through constructive interference across the disk array. The resulting THz radiation is collected by broadband corrugated horn antennas and detected using heterodyne mixing techniques [194].

In 2024, MADMAX completed its first axion search using the Morpurgo magnet at CERN. The Closed Booster 200 prototype employed three sapphire disks, demonstrating boost factors of 10^3 – 10^4 and establishing coupling limits down to $|g_{a\gamma}| \sim 2 \times 10^{-11} \text{ GeV}^{-1}$ in the 76–79 μeV range [195].

MADMAX and BREAD represent the two major conducting surface experiments, each employing different approaches to broadband axion detection. MADMAX exploits axion-photon conversion at material interfaces, while BREAD utilizes novel dish antenna designs. More details on BREAD are provided in Appendix B.

Direct detection experiments have made significant progress in probing the axion parameter space, exploring complementary regions in mass and coupling through a variety of techniques. The continued development of novel methods and improved sensitivities highlights the strong theoretical motivation to search for axions across a wide range of energies.

Appendix B introduces the BREAD experiment, which contributes to this broader effort with a new approach to broadband axion detection. Its development reflects the growing interest in axion dark matter and reinforces the value of pursuing axion-related signatures, both in dedicated detection experiments and in collider-based searches for related particles like the axino.

APPENDIX B

THE BREAD EXPERIMENT

This appendix provides additional detail on BREAD, expanding on the brief overview presented in Chapter 2.3.1.3. This appendix describes both the overall BREAD experimental program and my specific contributions to detector characterization and spectroscopic instrumentation development.

BREAD employs a novel broadband dish antenna approach that differs fundamentally from traditional cavity haloscopes [24]. Rather than scanning slowly through frequencies one at a time, BREAD searches across multiple decades of axion mass simultaneously, providing comprehensive coverage of theoretically motivated regions where QCD axion dark matter might exist.

B.1 Experimental Design

The BREAD detector consists of a cylindrical metallic barrel housed within a solenoid magnet, with a Hershey’s Kiss-shaped parabolic reflector positioned coaxially inside the cylinder [24]. The parabolic surface focuses photons from the large cylindrical conversion surface onto a detector positioned at the focal point. The cylindrical barrel serves as the axion-to-photon conversion surface, requiring the strong magnetic field described in Section 2.3.1.

The BREAD detector consists of a cylindrical metallic barrel housed within a solenoid magnet, with a Hershey’s Kiss-shaped parabolic reflector positioned coaxially inside the cylinder [24]. The parabolic surface focuses photons from the large cylindrical conversion surface onto a detector positioned at the focal point. For axion searches, the cylindrical barrel serves as the axion-to-photon conversion surface in the presence of strong magnetic fields—the col-




Figure B.1: 3D rendering of the GigaBREAD experimental setup at Argonne National Laboratory, showing the cylindrical reflector barrel positioned inside the 3.9 T solenoid magnet. The parabolic reflector focuses converted photons onto the coaxial horn antenna at the focal point. The human figure provides scale reference for this room-temperature proof-of-concept implementation. Rendering by David W. Miller.

laboration utilizes repurposed MRI magnets which can provide fields up to 10 Tesla as described in Section 2.3.1.

This cylindrical geometry provides optimal compatibility with solenoid magnets, which naturally produce fields parallel to the cylinder walls, ensuring proper alignment between the magnetic field and the metal surface where axion conversion occurs [24]. Previous dish antenna experiments [23, 196] faced significant challenges with this field-surface alignment problem, making BREAD’s cylindrical design a key practical advantage.

The BREAD experimental design addresses specific limitations found in other experimental strategies. While CAST [54] requires solar tracking, BREAD operates continuously for galactic dark matter searches. Where traditional haloscopes like ADMX [197] must scan slowly through narrow frequency ranges, BREAD searches broad frequency bands simultaneously.

The BREAD experiment also provides versatility for dark photon searches, which can operate without magnetic fields [25] since dark photons convert to regular photons through kinetic mixing alone. Dark photons are another well-motivated bosonic DM candidate below the eV scale, with the key experimental advantage that their conversion to photons occurs independent of mass, overcoming the narrowband limitations of resonant cavity approaches. This allows the same experimental apparatus to search for multiple dark matter candidates across broad frequency ranges simultaneously.

B.2 The BREAD Multi-Stage Approach

The BREAD experimental program employs a multi-stage approach, with different implementations targeting specific axion mass ranges using the most appropriate available detector technology [24, 26, 198]. Each stage addresses distinct frequency regimes where different sensor technologies excel, building toward comprehensive coverage of theoretically motivated




Figure B.2: Projected sensitivity of BREAD in the coupling vs mass plane for different photosensing technologies and NEP assumptions. These assume no backgrounds beyond the intrinsic detector noise. The upper axis corresponds to detected photon frequency. Shaded blue regions denote existing astrophysical and haloscope exclusions. The shaded green region shows the viable QCD axion models and the dark green line denotes the KSVZ scenario. Figure from [24].

QCD axion parameter space. This approach allows the collaboration to make immediate scientific contributions while developing advanced technologies for the most challenging detection regimes [24].

Stage	Mass Range	Frequency	Sensors	Status
GigaBREAD	10–100 μeV	2–25 GHz	RF receivers	Results: 44–52 μeV [26]
TeraBREAD	1–100 meV	0.3–30 THz	SNSPDs, TES	In development
InfraBREAD	0.1–1 eV	25–250 THz	SNSPDs, TES	In development

Table B.1: Summary of the staged BREAD experiment implementations.

- **GigaBREAD:** Microwave frequencies targeting roughly 10–100 μeV (2–25 GHz), with first results published for 44–52 μeV (10.7–12.5 GHz) using room-temperature RF receivers. This stage demonstrates the dish antenna approach and has achieved the most stringent laboratory constraints on axion-photon coupling in this initial mass range (Figures B.3a and B.3b), with plans for upgrades using quantum-limited amplifiers and higher magnetic fields [26, 24].
- **TeraBREAD:** Terahertz frequencies targeting 1–100 meV (0.3–30 THz) using quantum sensors such as SNSPDs, Transition-Edge Sensor (TES) bolometers, and other specialized detectors at sub-Kelvin temperatures. This stage addresses the challenging *terahertz gap*—a frequency region where traditional electronic devices become inefficient due to parasitic effects while optical devices cannot operate effectively due to insufficient photon energies—where extending existing detector capabilities to lower energies remains a key technological challenge [198, 24].
- **InfraBREAD:** Infrared frequencies targeting 0.1–1 eV using established quantum detectors like SNSPDs and TES bolometers at millikelvin temperatures. This stage leverages mature infrared detector technology to target theoretically favored QCD

axion mass ranges [24].

This multi-stage approach enables BREAD to pursue parallel development paths, making scientific progress with existing technologies while advancing detector capabilities for the most technically demanding frequency ranges [24]. A summary of these staged implementations is provided in Table B.1.

B.3 GigaBREAD First Results

GigaBREAD demonstrated the first ALP search using a dish antenna, utilizing a 3.9 T MRI magnet at Argonne National Laboratory [26]. Figure B.1 shows the experimental configuration, with the cylindrical reflector positioned inside the solenoid magnet bore. The experiment implemented FPGA-based data acquisition with real-time Fast Fourier Transform (FFT) processing and frequency hopping to mitigate radio interference.

The results, shown in Figure B.3, established the most stringent laboratory constraints on axion-photon coupling in the 44 to 52 μeV mass range, excluding coupling strengths around $4 \times 10^{-11} \text{ GeV}^{-1}$ [26].

B.4 Detecting Single Photons

The BREAD experiments face the experimental challenge of detecting single photons from axion-photon coupling, at rates estimated around one photon per day for QCD axions with standard couplings [24]. The primary challenge is distinguishing this genuine axion signal from thermal blackbody radiation and other background sources, requiring sensors with exceptional noise equivalent power and sophisticated background rejection techniques over months of operation.




Figure B.3: ALP exclusion limits from GigaBREAD. a Detailed view of the exclusion limits in the 44–54 μeV mass range, where the green region shows the parameter space excluded by this first dish antenna axion search. The fluctuating limit reflects the frequency-dependent sensitivity of the measurement. b Results shown in the broader context of other haloscope experiments and astrophysical constraints, demonstrating how the dish antenna approach fills a critical gap in the experimental landscape. Figures from [26].

The broad frequency range targeted by BREAD implementations requires multiple detector technologies, each optimized for specific energy ranges and operating conditions. Table B.2 summarizes the candidate photosensor technologies under consideration, ranging from room-temperature RF receivers used in GigaBREAD to ultra-sensitive quantum detectors operating at millikelvin temperatures for the most demanding applications. The choice of detector technology depends critically on the target photon energy, required sensitivity, and achievable operating temperature. The most sensitive technologies, like TES bolometers, Kinetic Inductance Detector (KID) arrays, and SNSPDs, require millikelvin operation and complex cryogenic infrastructure but can approach the fundamental limits needed for QCD axion detection.

Technology	Energy Range [meV]	Operating Temp [K]	NEP/DCR
Pyroelectric	[0.4, 120]	293	$10^{-8} \text{ W/Hz}^{1/2}$
Bolometer	[0.24, 248]	1.6	$5 \times 10^{-14} \text{ W/Hz}^{1/2}$
TES	[0.2, 125]	0.3	$2 \times 10^{-19} \text{ W/Hz}^{1/2}$
KID	[0.2, 125]	0.3	$2 \times 10^{-19} \text{ W/Hz}^{1/2}$
SNSPD	[124, 830]	0.3	10^{-4} Hz (DCR)
QCDet	[2, 125]	0.015	$3 \times 10^{-21} \text{ W/Hz}^{1/2}$

Table B.2: Comparison of candidate photosensor technologies and their typical performance characteristics relevant for BREAD implementations. Table originally from [24].

B.4.1 Terahertz Spectroscopy Development

Future BREAD experiments will need to work at much higher frequencies than GigaBREAD, requiring careful testing of optical components and sensors across broad frequency ranges. To support this development, I designed and built a terahertz spectroscopy system at the University of Chicago [199]. This instrument allows us to measure how well different components transmit and reflect light at the frequencies where future axion searches will operate.




Figure B.4: Schematic diagram of the terahertz spectroscopy system that I built for testing BREAD optical components and sensors. Figure from [199].

The system uses an interferometer design that can precisely measure how materials behave at terahertz frequencies—the gap between microwave and infrared light that is particularly challenging to work with. The instrument can resolve fine spectral details across more than a factor of ten in frequency [199].

This capability supports several critical needs for BREAD: testing how well the reflector surfaces focus light, characterizing filters that block unwanted thermal radiation, measuring how sensitive different photon detectors are across various frequencies, and providing calibration standards for data analysis. These measurements become essential as BREAD moves to higher frequencies where individual photons carry more energy and are harder to detect.




Figure B.5: Testing the accuracy of our spectroscopy system by measuring known materials. The measured data (cyan) matches published reference values (orange) very well, confirming that our instrument works correctly. The absorption peaks correspond to specific molecular vibrations in polyethylene plastic. Figure from [199].

Figure B.5 shows how we validated the system’s accuracy by measuring well-known reference materials. Our measurements match published literature values extremely well, with absorp-

tion features appearing at the expected frequencies and with the correct peak widths. This confirms that our instrument is properly calibrated and can make reliable measurements.




Figure B.6: Demonstration of the system’s broad frequency coverage, showing measurements across two orders of magnitude in frequency. The agreement with theoretical predictions (dashed line) confirms accurate calibration across this wide range. Figure from [199].

The broad frequency coverage shown in Figure B.6 demonstrates that this system can test components across the full range needed for future BREAD experiments. This work establishes the measurement capabilities that will be essential for developing the more advanced detector technologies needed to search for axions at higher frequencies.

B.4.2 Superconducting Nanowire Single Photon Detectors

SNSPDs offer promising capabilities for near-infrared frequencies [200, 201]. These devices consist of ultra-thin superconducting nanowires patterned in meandering geometries across the detector active area. When operated at mK temperatures, the nanowires carry bias current close to the critical current threshold.




Figure B.7: Scanning electron microscope image of an SNSPD showing the meandering nanowire pattern. The device consists of ultra thin superconducting wires arranged to maximize the active detection area while maintaining the necessary electrical connectivity. The inset shows the detailed wire structure with typical dimensions. Figure from [202].

SNSPD photon detection operates through the hotspot model [203]: photon absorption breaks Cooper pairs and creates a localized normal region. This hotspot forces current redistribution, creating a measurable voltage pulse before superconducting recovery within nanoseconds. SNSPDs can achieve ultra-low dark count rates, Dark Count Rate (DCR) $< 10^{-4}$ Hz, high detection efficiency, $> 90\%$ for photons above threshold, and scalability to millimeter dimensions [200, 201], making them promising candidates for BREAD applications.

SNSPDs were originally designed for hundreds of terahertz (near infrared) frequencies, while BREAD requires functionality in the single to double digit terahertz ranges that we aim to target. This necessitated characterization work to understand the utility of SNSPDs at the specific frequency ranges where we plan to search for axions.




Figure B.8: Detection mechanism of an SNSPD showing the hotspot formation process. (i) The device is biased below the critical current. (ii) A photon is absorbed, creating a hotspot. (iii) The hotspot expands as current density increases. (iv) A resistive barrier forms across the nanowire. (v) Current is blocked, creating a detectable voltage pulse. (vi) Superconductivity recovers and the device resets. Figure from [200].


B.4.3 SNSPD Characterization Studies

In preparation for their use in BREAD, I conducted characterization studies of SNSPDs, focusing on extending their frequency bandwidth to lower frequencies than their original design targets. SNSPDs require mK temperatures from Adiabatic Demagnetization Refrigerator (ADR) systems [204] for several reasons: to maintain superconductivity below their critical temperature, suppress thermal dark counts from Cooper pair breaking, and minimize blackbody radiation backgrounds that would otherwise overwhelm the detector. These refrigerators use paramagnetic salts that cool when demagnetized, reaching temperatures below 20 mK for the long operation times that axion searches require.

The goal of my work was to make three critical measurements: current-voltage characterization to determine optimal operating points, dark count rate characterization as a function of bias current to establish fundamental noise floors, and detection efficiency measurements



(a) ADR system closed



(b) ADR system open

Figure B.9: The ADR fridge system at Fermilab used for SNSPD characterization. The closed system showing the external vacuum chamber. a The open fridge revealing the internal temperature stages.




Figure B.10: SNSPD characterization workflow showing the two-step measurement process. (a) Preliminary current-voltage characteristic from my characterization studies, showing the hysteretic switching behavior between superconducting and normal states. The switching current determines the maximum usable bias current for photon detection. (b) Example dark count rate measurement as a function of bias current that we aimed to achieve (reproduced from Ref. [202]). This type of measurement establishes the fundamental noise floor and optimal operating conditions for the device.

across relevant photon energies. The experimental setup, shown conceptually in Figure B.11, requires careful bias control and low-noise amplification with readout chains including bias-T coupling and both cryogenic and room temperature amplification stages.

The first step involves current-voltage characterization to identify the switching current and determine the operating regime. The switching current is the minimum current required to drive the superconducting nanowire from its zero-resistance superconducting state to its finite-resistance normal state. We find this current by slowly ramping the bias current until the device spontaneously switches, even without photon absorption. This measurement is critical because it establishes the maximum current we can apply while maintaining superconductivity. We typically operate at 90-95% of this switching current to ensure that photon absorption will trigger a detectable transition to the normal state while avoiding false triggers from thermal fluctuations.




Figure 1 conceptual experimental apparatus

Figure B.11: Experimental setup for SNSPD characterization showing the device at mK temperatures with low-noise readout electronics. Schematic by Tony X. Zhou.

Figure B.10a shows a preliminary IV curve that I measured, revealing the hysteretic switching behavior between superconducting and normal states. The switching current establishes the critical current threshold for setting appropriate bias points in subsequent dark count rate measurements. The ultimate goal of this characterization work was to achieve the type of dark count rate measurements shown in Figure B.10b [202]. These measurements establish the fundamental noise floor and optimal operating conditions for the sensors.

This work established characterization capabilities at Fermilab that were useful for BREAD to evaluate whether SNSPDs could be viable for our quantum detector technology needs in future high-frequency implementations targeting QCD axion dark matter. Interestingly, these characterization techniques have since found applications beyond dark matter searches. Recent work has demonstrated that superconducting microwire devices can efficiently detect high-energy particles from accelerator beams [205], showing that the quantum sensors we characterized for dark matter searches have found new applications in high-energy physics experiments.




Figure B.12: Photograph of an SNSPD device glued to the cryogenic mount that can be wired to resistors and readout electronics inside the refrigerator.

B.5 Future Plans

GigaBREAD’s successful demonstration of the first axion search using a dish antenna [26] validates the BREAD approach and paves the way for future implementations that could target the theoretically favored mass ranges for QCD axion dark matter at higher frequencies. The next step involves scaling up from GigaBREAD’s 3.9 T magnet to a 10 T magnetic field for significantly improved sensitivity.

Future BREAD implementations will target progressively more sensitive parameter space through improved magnetic fields, larger reflector areas, and advanced sensor technologies [24]. The collaboration continues developing the infrastructure needed for larger experiments while reducing systematic uncertainties demonstrated in the GigaBREAD analysis [26].

BREAD’s unique combination of broadband sensitivity and geometric focusing provides a new experimental avenue for axion detection that complements existing approaches and extends sensitivity into previously unexplored mass ranges. The successful demonstration through GigaBREAD establishes the foundation for future implementations targeting theoretically motivated axion parameter space, representing a promising new direction in the search for dark matter.

APPENDIX C

OBSERVED EVENT KINEMATICS IN THE DV + MET ANALYSIS

This appendix provides detailed kinematic distributions comparing the observed events in each signal region with representative signal models. While the main analysis focuses on statistical interpretation and limit setting, these kinematic comparisons offer valuable insight into the consistency between observed events and expected signal characteristics across different model parameters.

C.1 Event Properties and Signal Comparisons

In five of the six observed events reported in Table 6.22, the reconstructed DVs are located quite close to or inside the beam pipe. Further, in four of the six events the DVs are found very close to a b -tagged jet, suggesting possible consistency with SM background processes involving heavy-flavor hadron decays.

The kinematic distributions presented in this section demonstrate several key features across the different signal models. The general shapes of the distributions remain consistent across different lifetimes for each model, indicating robust signal characteristics. For all included variables MET, DV mass, and DV track multiplicity, the greatest number of signal events occurs at intermediate lifetimes, highlighting that the chosen signal grids effectively cover the complete span of lifetime sensitivity. Both very prompt and extremely long-lived decays reach the edges of the signal region acceptance, as expected from the geometric constraints of the ATLAS ID.

C.1.1 Standard Displaced Vertex Signal Region

Figures C.1 and C.2 show the MET, vertex m_{DV} , and $N_{DVtracks}$ distributions for the observed event in the SDV SR compared to both axino signal models. Figure C.1 presents the comparison with $\tilde{a} \rightarrow hh$ decays, while Figure C.2 shows the corresponding distributions for $\tilde{a} \rightarrow ZZ$ decays. Both comparisons include variations in the axino mass m_a , which directly affects the lifetime τ of the neutralino $\tilde{\chi}_1^0$. The kinematic features are broadly similar between the two axino decay channels, reflecting the comparable mass scales of the Higgs boson and Z boson in the final states.

Figure C.3 presents a similar comparison for the gluino R-hadron model, showing how the distributions vary with the gluino lifetime τ .

C.1.2 Single Fuzzy Displaced Vertex Signal Region

A comparison between the three observed events in the 1 FDV SR and the wino-bino coannihilation model is presented in Figure C.4. This model is particularly well-suited to the fuzzy vertexing algorithm due to the presence of b -hadrons in the final state from the off-shell Higgs boson decay.

C.1.3 Two Fuzzy Displaced Vertex Signal Region

Figure C.5 shows the kinematic comparison for the two observed events in the 2 FDV SR and the Higgs portal model. This region targets pair production of long-lived particles that each decay within the ID, producing two distinct fuzzy vertices.




Figure C.1: Comparison of m_{DV} a, N_{tracks} b, and E_T^{miss} c between the observed event in the SDV region and a selection of $\tilde{a} \rightarrow hh$ signals in which the axino mass m_a , and as a result the lifetime τ of the neutralino $\tilde{\chi}_1^0$, is varied. Figure from [160].







Figure C.2: Comparison of m_{DV} a, N_{tracks} b, and E_T^{miss} c between the observed event in the SDV region and a selection of $\tilde{a} \rightarrow ZZ$ signals in which the axion mass m_a , and as a result the lifetime τ of the neutralino $\tilde{\chi}_1^0$, is varied. Figure from [160].



(a)



(b)



(c)

Figure C.3: Comparison of m_{DV} a, N_{tracks} b, and E_T^{miss} c between the observed event in the SDV region and a selection of gluino R-hadron signals in which the lifetime τ of the gluino is varied. Figure from [160].




Figure C.4: Comparison of m_{FDV} a, N_{tracks} b, and E_T^{miss} c between the three observed events in the 1FDV region and a selection of wino-bino coannihilation signals in which the lifetime τ of $\tilde{\chi}_2^0$ is varied. Figure from [160].




Figure C.5: Comparison of m_{FDV} a, N_{tracks} b, and E_T^{miss} c between the two observed events in the 2FDV region and a selection of Higgs portal signals in which the lifetime $c\tau$ of the new scalar S is varied. Figure from [160].

BIBLIOGRAPHY

- [1] G. Aad et al. “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”. In: *Phys. Lett. B* 716 (2012), pp. 1–29. DOI: 10.1016/j.physletb.2012.08.020 (cit. on p. 1).
- [2] S. Chatrchyan et al. “Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC”. In: *Phys. Lett. B* 716 (2012), pp. 30–61. DOI: 10.1016/j.physletb.2012.08.021 (cit. on p. 1).
- [3] N. Aghanim et al. “Planck 2018 results. VI. Cosmological parameters”. In: *Astron. Astrophys.* 641 (2020), A6. DOI: 10.1051/0004-6361/201833910 (cit. on pp. 1, 23, 24).
- [4] Gianfranco Bertone, Dan Hooper, and Joseph Silk. “Particle dark matter: Evidence, candidates and constraints”. In: *Phys. Rept.* 405 (2005), pp. 279–390. DOI: 10.1016/j.physrep.2004.08.031 (cit. on p. 1).
- [5] R.D. Peccei and Helen R. Quinn. “CP Conservation in the Presence of Instantons”. In: *Phys. Rev. Lett.* 38 (1977), pp. 1440–1443. DOI: 10.1103/PhysRevLett.38.1440 (cit. on pp. 2, 25, 29).
- [6] Frank Wilczek. “Problem of Strong P and T Invariance in the Presence of Instantons”. In: *Phys. Rev. Lett.* 40 (1978), pp. 279–282. DOI: 10.1103/PhysRevLett.40.279 (cit. on pp. 2, 29).
- [7] John Preskill, Mark B. Wise, and Frank Wilczek. “Cosmology of the Invisible Axion”. In: *Phys. Lett. B* 120 (1983), pp. 127–132. DOI: 10.1016/0370-2693(83)90637-8 (cit. on pp. 2, 25).
- [8] L. F. Abbott and P. Sikivie. “A Cosmological Bound on the Invisible Axion”. In: *Phys. Lett. B* 120 (1983), pp. 133–136. DOI: 10.1016/0370-2693(83)90638-X (cit. on pp. 2, 25).
- [9] M. Dine and W. Fischler. “The Not So Harmless Axion”. In: *Phys. Lett. B* 120 (1983), pp. 137–141. DOI: 10.1016/0370-2693(83)90639-1 (cit. on pp. 2, 25).
- [10] N. Aghanim et al. “Planck 2018 results. I. Overview and the cosmological legacy of Planck”. In: *Astron. Astrophys.* 641 (2020), A1. DOI: 10.1051/0004-6361/201833880. arXiv: 1807.06205 [astro-ph.CO] (cit. on p. 2).
- [11] M. Tanabashi et al. “Review of Particle Physics”. In: *Phys. Rev. D* 98 (2018), p. 030001. DOI: 10.1103/PhysRevD.98.030001 (cit. on p. 2).
- [12] Jihn E. Kim. “Weak Interaction Singlet and Strong CP Invariance”. In: *Phys. Rev. Lett.* 43 (1979), p. 103. DOI: 10.1103/PhysRevLett.43.103 (cit. on pp. 2, 27, 29).

- [13] Mikhail A. Shifman, Arkady I. Vainshtein, and Valentin I. Zakharov. “Can Confinement Ensure Natural CP Invariance of Strong Interactions?” In: *Nucl. Phys. B* 166 (1980), pp. 493–506. DOI: 10.1016/0550-3213(80)90209-6 (cit. on pp. 2, 25, 27, 29).
- [14] Michael Dine, Willy Fischler, and Mark Srednicki. “A Simple Solution to the Strong CP Problem with a Harmless Axion”. In: *Phys. Lett. B* 104 (1981), pp. 199–202. DOI: 10.1016/0370-2693(81)90590-6 (cit. on pp. 2, 27, 29, 40).
- [15] A.R. Zhitnitsky. “On Possible Suppression of the Axion Hadron Interactions. (In Russian)”. In: *Sov. J. Nucl. Phys.* 31 (1980), p. 260 (cit. on pp. 2, 27, 29, 40).
- [16] Joerg Jaeckel and Andreas Ringwald. “The Low-Energy Frontier of Particle Physics”. In: *Ann. Rev. Nucl. Part. Sci.* 60 (2010), pp. 405–437. DOI: 10.1146/annurev.nucl.012809.104433. arXiv: 1002.0329 (cit. on pp. 2, 36).
- [17] Henry Primakoff. “Photo-Production of Neutral Mesons in Nuclear Electric Fields and the Mean Life of the Neutral Meson”. In: *Phys. Rev.* 81 (1951), pp. 899–899. DOI: 10.1103/PhysRev.81.899 (cit. on pp. 2, 27).
- [18] N. Du et al. “A Search for Invisible Axion Dark Matter with the Axion Dark Matter Experiment”. In: *Phys. Rev. Lett.* 120 (2018), p. 151301. DOI: 10.1103/PhysRevLett.120.151301. arXiv: 1804.05750 (cit. on p. 2).
- [19] N. Du et al. “A Search for Invisible Axion Dark Matter with the Axion Dark Matter Experiment”. In: *Phys. Rev. Lett.* 120.15 (2018), p. 151301. DOI: 10.1103/PhysRevLett.120.151301. arXiv: 1804.05750 [hep-ex] (cit. on pp. 2, 31, 250).
- [20] L. Zhong et al. “Results from phase 1 of the HAYSTAC microwave cavity axion experiment”. In: *Phys. Rev. D* 97 (2018), p. 092001. DOI: 10.1103/PhysRevD.97.092001. arXiv: 1803.03690 (cit. on p. 2).
- [21] L. Zhong et al. “Results from phase 1 of the HAYSTAC microwave cavity axion experiment”. In: *Phys. Rev. D* 97.9 (2018), p. 092001. DOI: 10.1103/PhysRevD.97.092001. arXiv: 1803.03690 [hep-ex] (cit. on pp. 2, 31).
- [22] Jia Liu, Zhen Liu, and Lian-Tao Wang. “Long-lived particles at the LHC: catching them in time”. In: *Phys. Rev. Lett.* 122 (2019), p. 131801. DOI: 10.1103/PhysRevLett.122.131801. arXiv: 1805.05957 [hep-ph] (cit. on p. 3).
- [23] D. Horns et al. “Searching for WISPy Cold Dark Matter with a Dish Antenna”. In: *JCAP* 1304 (2013), p. 016. DOI: 10.1088/1475-7516/2013/04/016. arXiv: 1212.2970 [hep-ph] (cit. on pp. 3, 256).
- [24] Jesse Liu et al. “Broadband Solenoidal Haloscope for Terahertz Axion Detection”. In: *Phys. Rev. Lett.* 128 (2022), p. 131801. DOI: 10.1103/PhysRevLett.128.131801. arXiv: 2111.12103 [hep-ex] (cit. on pp. 3, 32, 33, 254, 256–259, 261, 270).

- [25] Stefan Knirck et al. “First Results from a Broadband Search for Dark Photon Dark Matter in the 44 to 52 μeV Range with a Coaxial Dish Antenna”. In: *Phys. Rev. Lett.* 132 (2024), p. 131004. arXiv: 2310.13891 [hep-ex] (cit. on pp. 3, 256).
- [26] Gabe Hoshino and BREAD Collaboration. “First Axion-Like Particle Results from a Broadband Search for Wave-Like Dark Matter in the 44 to 52 μeV Range with a Coaxial Dish Antenna”. In: *Phys. Rev. Lett.* 134 (2025), p. 171002. DOI: 10.1103/PhysRevLett.134.171002. arXiv: 2501.17119 [hep-ex] (cit. on pp. 3, 256, 258–260, 270).
- [27] Y. A. Golfand and E. P. Likhtman. “Extension of the Algebra of Poincare Group Generators and Violation of P Invariance”. In: *JETP Lett.* 13 (1971), pp. 323–326 (cit. on pp. 3, 37).
- [28] D. V. Volkov and V. P. Akulov. “Is the neutrino a goldstone particle?” In: *Phys. Lett. B* 46 (1973), pp. 109–110. DOI: 10.1016/0370-2693(73)90490-5 (cit. on pp. 3, 37).
- [29] J. Wess and B. Zumino. “Supergauge transformations in four dimensions”. In: *Nucl. Phys. B* 70 (1974), pp. 39–50. DOI: 10.1016/0550-3213(74)90355-1 (cit. on pp. 3, 37).
- [30] J. Wess and B. Zumino. “Supergauge invariant extension of quantum electrodynamics”. In: *Nucl. Phys. B* 78 (1974), pp. 1–13. DOI: 10.1016/0550-3213(74)90112-6 (cit. on pp. 3, 37).
- [31] S. Ferrara and B. Zumino. “Supergauge invariant Yang-Mills theories”. In: *Nucl. Phys. B* 79 (1974), pp. 413–421. DOI: 10.1016/0550-3213(74)90559-8 (cit. on pp. 3, 37).
- [32] A. Salam and J. Strathdee. “Super-symmetry and non-Abelian gauges”. In: *Phys. Lett. B* 51 (1974), pp. 353–355. DOI: 10.1016/0370-2693(74)90226-3 (cit. on pp. 3, 37).
- [33] G. R. Farrar and P. Fayet. “Phenomenology of the production, decay, and detection of new hadronic states associated with supersymmetry”. In: *Phys. Lett. B* 76 (1978), pp. 575–579. DOI: 10.1016/0370-2693(78)90858-4 (cit. on pp. 3, 39).
- [34] Kyu Jung Bae, Howard Baer, and Hasan Serce. “Prospects for axion detection in natural SUSY with mixed axion-higgsino dark matter: back to invisible?” In: *JCAP* 06 (2017), p. 024. DOI: 10.1088/1475-7516/2017/06/024. arXiv: 1705.01134 [hep-ph] (cit. on pp. 4, 40, 41, 44).
- [35] The ATLAS Collaboration. In: 3.08 (2008), S08003. DOI: 10.1088/1748-0221/3/08/S08003 (cit. on pp. 4, 61, 63, 64, 73, 76, 77).
- [36] Lyndon Evans and Philip Bryant. “LHC Machine”. In: *JINST* 3.08 (2008), S08001. DOI: 10.1088/1748-0221/3/08/S08001 (cit. on pp. 4, 47).

- [37] E. Noether. “Invariant Variation Problems”. In: *Transport Theory and Statistical Physics* 1 (1918), p. 186. DOI: 10.1080/00411457108231446 (cit. on p. 14).
- [38] J. H. Christenson et al. “Evidence for the 2π Decay of the K_2^0 Meson”. In: *Phys. Rev. Lett.* 13 (1964), pp. 138–140. DOI: 10.1103/PhysRevLett.13.138 (cit. on p. 20).
- [39] B. Aubert et al. “Observation of CP violation in the B^0 meson system”. In: *Phys. Rev. Lett.* 87 (2001), p. 091801. DOI: 10.1103/PhysRevLett.87.091801. eprint: hep-ex/0107013 (cit. on p. 20).
- [40] K. Abe et al. “Observation of large CP violation in the neutral B meson system”. In: *Phys. Rev. Lett.* 87 (2001), p. 091802. DOI: 10.1103/PhysRevLett.87.091802. eprint: hep-ex/0107061 (cit. on p. 20).
- [41] Makoto Kobayashi and Toshihide Maskawa. “CP Violation in the Renormalizable Theory of Weak Interaction”. In: *Prog. Theor. Phys.* 49 (1973), pp. 652–657. DOI: 10.1143/PTP.49.652 (cit. on p. 20).
- [42] C. Abel et al. “Measurement of the permanent electric dipole moment of the neutron”. In: *Phys. Rev. Lett.* 124.8 (2020), p. 081803. DOI: 10.1103/PhysRevLett.124.081803. arXiv: 2001.11966 [hep-ex] (cit. on p. 21).
- [43] Gerard ’t Hooft. “Naturalness, chiral symmetry, and spontaneous chiral symmetry breaking”. In: *Recent Developments in Gauge Theories*. Ed. by G. ’t Hooft et al. Plenum Press, 1980, pp. 135–157 (cit. on p. 22).
- [44] V. C. Rubin and W. K. Ford Jr. “Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions”. In: *Astrophys. J.* 159 (1970), pp. 379–403. DOI: 10.1086/150317 (cit. on p. 22).
- [45] D. Clowe et al. “A direct empirical proof of the existence of dark matter”. In: *Astrophys. J.* 648 (2006), pp. L109–L113. DOI: 10.1086/508162. arXiv: astro-ph/0608407 [astro-ph] (cit. on p. 22).
- [46] Gerard t Hooft. “Computation of the Quantum Effects Due to a Four-Dimensional Pseudoparticle”. In: *Phys. Rev. D* 14 (1976), pp. 3432–3450. DOI: 10.1103/PhysRevD.14.3432 (cit. on p. 25).
- [47] Particle Data Group and R. L. Workman et al. “Axions and other similar particles”. In: *Prog. Theor. Exp. Phys.* 2022.8 (2022), p. 083C01. DOI: 10.1093/ptep/ptac097 (cit. on p. 25).
- [48] Ian P. Stern. *Axion Dark Matter Searches*. Technical Report, University of Florida. 2016. URL: <https://www.osti.gov/servlets/purl/1273235> (cit. on p. 26).

- [49] P. Sikivie. “Experimental Tests of the Invisible Axion”. In: *Phys. Rev. Lett.* 51 (1983). Ed. by M.A. Srednicki, pp. 1415–1417. DOI: 10.1103/PhysRevLett.51.1415 (cit. on p. 26).
- [50] Particle Data Group. “Review of Particle Physics”. In: *Progress of Theoretical and Experimental Physics* 2022.8 (Aug. 2022), p. 083C01. ISSN: 2050-3911. DOI: 10.1093/ptep/ptac097 (cit. on pp. 27, 74).
- [51] Luca Di Luzio, Federico Mescia, and Enrico Nardi. “Redefining the Axion Window”. In: *Phys. Rev. Lett.* 118.3 (2017), p. 031801. DOI: 10.1103/PhysRevLett.118.031801. arXiv: 1610.07593 [hep-ph] (cit. on p. 29).
- [52] The MADMAX Working Group. “Dielectric Haloscopes: A New Way to Detect Axion Dark Matter”. In: *Phys. Rev. Lett.* 118 (2017), p. 091801. DOI: 10.1103/PhysRevLett.118.091801. arXiv: 1611.05865 (cit. on p. 31).
- [53] R Bähre et al. “Any light particle search II — Technical Design Report”. In: *Journal of Instrumentation* 8.09 (Sept. 2013), T09001–T09001. ISSN: 1748-0221. DOI: 10.1088/1748-0221/8/09/t09001 (cit. on p. 31).
- [54] V. Anastassopoulos et al. “New CAST Limit on the Axion-Photon Interaction”. In: *Nature Phys.* 13 (2017), pp. 584–590. DOI: 10.1038/nphys4109. arXiv: 1705.02290 [hep-ex] (cit. on pp. 32, 247, 256).
- [55] C. Antel et al. “Feebly Interacting Particles: FIPs 2022 workshop report”. In: *Eur. Phys. J. C* 83 (2023), p. 1122. arXiv: 2305.01715 [hep-ph] (cit. on pp. 34, 35).
- [56] N. Sakai. “Naturalness in supersymmetric GUTS”. In: *Z. Phys. C* 11 (1981), p. 153. DOI: 10.1007/BF01573998 (cit. on p. 37).
- [57] S. Dimopoulos, S. Raby, and F. Wilczek. “Supersymmetry and the scale of unification”. In: *Phys. Rev. D* 24 (1981), pp. 1681–1683. DOI: 10.1103/PhysRevD.24.1681 (cit. on p. 37).
- [58] Stephen P. Martin. *A Supersymmetry Primer*. Advanced Series on Directions in High Energy Physics. World Scientific, 1998. URL: http://dx.doi.org/10.1142/9789812839657_0001 (cit. on pp. 37, 39).
- [59] A. Djouadi et al. *The Minimal supersymmetric standard model: Group summary report*. 1998. arXiv: hep-ph/9901246 [hep-ph] (cit. on p. 37).
- [60] Barbara Aulicino. “Going Nowhere Fast”. In: *American Scientist* 105.6 (2017), pp. 336–341. URL: <https://www.americanscientist.org/article/going-nowhere-fast> (cit. on p. 38).
- [61] H. Goldberg. “Constraint on the Photino Mass from Cosmology”. In: *Phys. Rev. Lett.* 50 (1983), pp. 1419–1422. DOI: 10.1103/PhysRevLett.50.1419 (cit. on p. 39).

- [62] J. Ellis et al. “Supersymmetric relics from the big bang”. In: *Nucl. Phys. B* 238 (1984), pp. 453–476. DOI: 10.1016/0550-3213(84)90461-9 (cit. on p. 39).
- [63] Gabe Hoshino et al. “Bridging the divide: axion and axino phenomenology at colliders and direct-detection experiments”. 2025 (cit. on pp. 40–42, 241).
- [64] Gabriel Hoshino. *Searching for the DFSZ Axino in Collider Experiments*. <https://indico.cern.ch/event/1218225/contributions/5384880/>. 2023 (cit. on pp. 41–44, 173).
- [65] Frank Daniel Steffen. *Collider Signatures of Axino and Gravitino Dark Matter*. 2005. arXiv: hep-ph/0507003 [hep-ph]. URL: <https://arxiv.org/abs/hep-ph/0507003> (cit. on p. 41).
- [66] Gabriela Barenboim et al. “Implications of an axino LSP for naturalness”. In: *Physical Review D* 90.3 (Aug. 2014). ISSN: 1550-2368. DOI: 10.1103/physrevd.90.035020 (cit. on pp. 41, 42).
- [67] C. S. Redino and D. Wackerroth. “Exploring the hadronic axion window via delayed neutralino decay to axinos at the LHC”. In: *Physical Review D* 93.7 (Apr. 2016). ISSN: 2470-0029. DOI: 10.1103/physrevd.93.075022 (cit. on p. 41).
- [68] ATLAS Collaboration. “Search for charginos and neutralinos in final states with two boosted hadronically decaying bosons and missing transverse momentum in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”. In: *Phys. Rev. D* 104.11 (2021), p. 112010. DOI: 10.1103/PhysRevD.104.112010. arXiv: 2108.07586 [hep-ex] (cit. on pp. 42, 154).
- [69] Izaak Neutelings. *Particle mass versus lifetime & decay length*. https://tikz.net/sm_particles_masses/ (cit. on p. 45).
- [70] Henriette Elvang. *Mass vs Lifetime Plot of SM Particles* (cit. on p. 45).
- [71] CERN. *CERN: Who We Are - History*. 2024. URL: <https://home.cern/about/who-we-are/history> (cit. on p. 46).
- [72] CERN. *CERN: Our People*. 2024. URL: <https://home.cern/about/who-we-are/our-people> (cit. on p. 47).
- [73] CERN. *The birth of the Web*. URL: <https://home.cern/science/computing/birth-web> (cit. on p. 47).
- [74] Kristiane Bernhard-Novotny. *Fighting cancer with LIGHT*. URL: <https://home.cern/news/news/knowledge-sharing/fighting-cancer-light> (cit. on p. 47).

- [75] Ranveig Strom. *Spectral imaging: from CERN to medical technologies*. URL: <https://home.cern/news/news/knowledge-sharing/spectral-imaging-cern-medical-technologies> (cit. on p. 47).
- [76] Esma Mobs. “The CERN accelerator complex. Complexe des accélérateurs du CERN”. In: (2016). URL: <https://cds.cern.ch/record/2197559> (cit. on p. 48).
- [77] *LEP design report*. Report. Geneva: CERN, 1984. URL: <https://cds.cern.ch/record/102083> (cit. on p. 48).
- [78] In: *Physics Letters B* 565 (July 2003), pp. 61–75. ISSN: 0370-2693. DOI: 10.1016/s0370-2693(03)00614-2 (cit. on p. 48).
- [79] CERN. *Accelerating: Radiofrequency cavities*. URL: <https://home.cern/science/engineering/accelerating-radiofrequency-cavities> (cit. on p. 50).
- [80] CERN. *Pulling Together: Superconducting Electromagnets*. URL: <https://home.cern/science/engineering/pulling-together-superconducting-electromagnets> (cit. on p. 50).
- [81] Oliver Sim Brüning et al. *LHC Design Report*. CERN Yellow Reports: Monographs. Geneva: CERN, 2004. DOI: 10.5170/CERN-2004-003-V-1 (cit. on p. 51).
- [82] The ATLAS Collaboration. *Expected Performance of the ATLAS Experiment - Detector, Trigger and Physics*. 2009. arXiv: 0901.0512 [hep-ex]. URL: <https://arxiv.org/abs/0901.0512> (cit. on pp. 51, 131, 143, 147, 150).
- [83] R Bruce et al. “Reaching record-low β^* at the CERN Large Hadron Collider using a novel scheme of collimator settings and optics”. In: *Nucl. Instrum. Methods Phys. Res., A* 848 (2017), pp. 19–30. DOI: 10.1016/j.nima.2016.12.039 (cit. on p. 52).
- [84] CERN. *LHCb*. 2025. URL: <https://home.cern/science/experiments/lhcb> (cit. on p. 53).
- [85] CERN. *ALICE*. 2025. URL: <https://home.cern/science/experiments/alice> (cit. on p. 53).
- [86] ATLAS Collaboration. *Luminosity Public Results for Run-3 of the LHC*. URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun3> (cit. on p. 54).
- [87] Joao Pequeno. *Computer generated image of the whole ATLAS detector*. URL: <https://cds.cern.ch/record/1095924> (cit. on pp. 56, 59).
- [88] Joao Pequeno and Paul Schaffner. “How ATLAS detects particles: diagram of particle paths in the detector”. 2013. URL: <https://cds.cern.ch/record/1505342> (cit. on p. 57).

- [89] ATLAS Collaboration. “The ATLAS Experiment at the CERN Large Hadron Collider”. In: *Journal of Instrumentation* 3 (2008), S08003. DOI: 10.1088/1748-0221/3/08/S08003. arXiv: 0901.0512 [physics.ins-det] (cit. on p. 56).
- [90] *ATLAS Magnet System Diagram*. URL: <https://www.quantumdiaries.org/wp-content/uploads/2011/05/exp-magnets.png> (cit. on p. 61).
- [91] M Capeans et al. *ATLAS Insertable B-Layer Technical Design Report*. Tech. rep. 2010. URL: <https://cds.cern.ch/record/1291633> (cit. on p. 64).
- [92] The ATLAS Collaboration. “The silicon microstrip sensors of the ATLAS semiconductor tracker”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 578.1 (2007), pp. 98–118. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2007.04.157> (cit. on p. 64).
- [93] The ATLAS Collaboration. “Study of the material of the ATLAS inner detector for Run 2 of the LHC”. In: *Journal of Instrumentation* 12.12 (2017), P12009. DOI: 10.1088/1748-0221/12/12/P12009 (cit. on p. 65).
- [94] The ATLAS Collaboration. “The ATLAS Transition Radiation Tracker (TRT) proportional drift tube: design and performance”. In: *JINST* 3 (2008), P02013. DOI: 10.1088/1748-0221/3/02/P02013 (cit. on p. 65).
- [95] The ATLAS Collaboration. “Alignment of the ATLAS Inner Detector in Run-2”. In: *Eur. Phys. J. C* 80.12 (2020), p. 1194. DOI: 10.1140/epjc/s10052-020-08700-6. arXiv: 2007.07624 (cit. on p. 67).
- [96] *ATLAS liquid-argon calorimeter: Technical Design Report*. Technical design report. ATLAS. Geneva: CERN, 1996. DOI: 10.17181/CERN.FWRW.F00Q (cit. on pp. 67, 75).
- [97] *ATLAS tile calorimeter: Technical Design Report*. Technical design report. ATLAS. Geneva: CERN, 1996. DOI: 10.17181/CERN.JRBJ.7028 (cit. on p. 68).
- [98] Joao Pequeno. “Computer Generated image of the ATLAS calorimeter”. In: (2008). URL: <https://cds.cern.ch/record/1095927> (cit. on p. 69).
- [99] *ATLAS muon spectrometer: Technical Design Report*. Technical design report. ATLAS. Geneva: CERN, 1997. URL: <https://cds.cern.ch/record/331068> (cit. on p. 76).
- [100] F. Bauer et al. “Construction and test of MDT chambers for the ATLAS muon spectrometer”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 461.1 (2001), pp. 17–20. ISSN: 0168-9002. DOI: [https://doi.org/10.1016/S0168-9002\(00\)01156-6](https://doi.org/10.1016/S0168-9002(00)01156-6) (cit. on p. 78).

- [101] T. Argyropoulos et al. “Cathode strip chambers in ATLAS : Installation, commissioning and in situ performance”. In: *2008 IEEE Nuclear Science Symposium Conference Record*. 2008, pp. 2819–2824. DOI: [10.1109/NSSMIC.2008.4774958](https://doi.org/10.1109/NSSMIC.2008.4774958) (cit. on p. 78).
- [102] G. Aielli et al. “The RPC first level muon trigger in the barrel of the ATLAS experiment”. In: *Nuclear Physics B - Proceedings Supplements* 158 (2006), pp. 11–15. ISSN: 0920-5632. DOI: <https://doi.org/10.1016/j.nuclphysbps.2006.07.031> (cit. on p. 78).
- [103] S. Majewski et al. “A thin multiwire chamber operating in the high multiplication mode”. In: *Nuclear Instruments and Methods in Physics Research* 217.1 (1983), pp. 265–271. ISSN: 0167-5087. DOI: [https://doi.org/10.1016/0167-5087\(83\)90146-1](https://doi.org/10.1016/0167-5087(83)90146-1) (cit. on p. 78).
- [104] The ATLAS collaboration. “Operation of the ATLAS trigger system in Run 2”. In: *Journal of Instrumentation* 15.10 (Oct. 2020), P10004–P10004. ISSN: 1748-0221. DOI: [10.1088/1748-0221/15/10/p10004](https://doi.org/10.1088/1748-0221/15/10/p10004) (cit. on p. 81).
- [105] The ATLAS collaboration. “The ATLAS trigger system for LHC Run 3 and trigger performance in 2022”. In: *Journal of Instrumentation* 19.06 (June 2024), P06029. ISSN: 1748-0221. DOI: [10.1088/1748-0221/19/06/p06029](https://doi.org/10.1088/1748-0221/19/06/p06029) (cit. on pp. 83, 92).
- [106] M (CERN) Aleksa et al. *ATLAS Liquid Argon Calorimeter Phase-I Upgrade: Technical Design Report*. Tech. rep. 2013. URL: <https://cds.cern.ch/record/1602230> (cit. on p. 84).
- [107] Luca Cadamuro. *Taggers for boosted HH searches within the ATLAS experiment*. 2024. URL: <https://cds.cern.ch/record/2891237/files/ATL-PHYS-SLIDE-2024-025.pdf> (cit. on p. 88).
- [108] The ATLAS Collaboration. “The ATLAS Level-1 Calorimeter Trigger”. In: *CERN Document Server* (2008). URL: <https://cds.cern.ch/record/1080560/files/daq-pub-2008-001.pdf> (cit. on p. 90).
- [109] Michael Begel et al. *Global Feature Extractor of the Level-1 Calorimeter Trigger: ATLAS TDAQ Phase-I Upgrade gFEX Final Design Report*. Tech. rep. Geneva: CERN, 2016. URL: <https://cds.cern.ch/record/2233958> (cit. on pp. 91, 92).
- [110] ATLAS Collaboration. *Level-1 Calorimeter Trigger Public Results*. 2025. URL: https://twiki.cern.ch/twiki/bin/view/AtlasPublic/L1CaloTriggerPublicResults#ATLAS_Level_1_calorimeter_gFEX_j (cit. on pp. 100, 128, 129).
- [111] Yocto Project. *Building Images - Yocto Project Development Tasks Manual*. <https://docs.yoctoproject.org/dev-manual/building.html>. 2023 (cit. on p. 105).
- [112] Xilinx Inc. *Yocto*. Xilinx Wiki. 2025. URL: <https://xilinx-wiki.atlassian.net/wiki/spaces/A/pages/18841883/Yocto> (visited on 04/08/2025) (cit. on p. 106).

- [113] Yocto Project. *Releases*. Yocto Project Wiki. 2025. URL: <https://wiki.yoctoproject.org/wiki/Releases> (visited on 04/08/2025) (cit. on p. 106).
- [114] Xilinx. *meta-xilinx SDT BSP README*. <https://github.com/Xilinx/meta-xilinx/blob/master/meta-xilinx-standalone-sdt/README.sdt.bsp.md> (cit. on p. 107).
- [115] M. Cacciari, Gavin P. Salam, and Gregory Soyez. “The anti-k(t) jet clustering algorithm”. In: *JHEP* 04 (2008), p. 063. DOI: 10.1088/1126-6708/2008/04/063. arXiv: 0802.1189 [hep-ph] (cit. on p. 120).
- [116] ATLAS Collaboration. “Jet reconstruction and performance using particle flow with the ATLAS Detector”. In: *The European Physical Journal C* 77.7 (). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-017-5031-2 (cit. on pp. 120, 145).
- [117] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. “SoftKiller, a particle-level pileup removal method”. In: *Eur. Phys. J. C* 75 (2015), p. 59. DOI: 10.1140/epjc/s10052-015-3267-2. arXiv: 1407.0408 [hep-ph] (cit. on p. 120).
- [118] ATLAS Collaboration. *RootFitTEff.cxx in atlas/athena GitLab repository*. <https://gitlab.cern.ch/atlas/athena>. 2024. URL: https://gitlab.cern.ch/atlas/athena/-/blob/main/DataQuality/dqm_algorithms/src/RootFitTEff.cxx%5C#L34-36 (visited on 04/10/2025) (cit. on p. 121).
- [119] ATLAS TDAQ Software Group. *beauty: ATLAS Trigger and DAQ Monitoring Framework*. <https://gitlab.cern.ch/atlas-tdaq-software/beauty>. 2025 (cit. on p. 126).
- [120] *Track Reconstruction Performance of the ATLAS Inner Detector at $\sqrt{s} = 13$ TeV*. Tech. rep. Geneva: CERN, 2015. URL: <https://cds.cern.ch/record/2037683> (cit. on pp. 131, 132).
- [121] T G Cornelissen et al. *Updates of the ATLAS Tracking Event Data Model (Release 13)*. Tech. rep. Geneva: CERN, 2007. URL: <https://cds.cern.ch/record/1038095> (cit. on p. 133).
- [122] R. Frühwirth. “Application of Kalman filtering to track and vertex fitting”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 262.2 (1987), pp. 444–450. ISSN: 0168-9002. DOI: [https://doi.org/10.1016/0168-9002\(87\)90887-4](https://doi.org/10.1016/0168-9002(87)90887-4) (cit. on p. 133).
- [123] T G Cornelissen et al. “The global χ^2 track fitter in ATLAS”. In: *Journal of Physics: Conference Series* 119.3 (July 2008), p. 032013. DOI: 10.1088/1742-6596/119/3/032013 (cit. on p. 134).
- [124] ATLAS Collaboration. “Performance of the ATLAS track reconstruction algorithms in dense environments in LHC Run 2”. In: *The European Physical Journal C* 77.10

- (Oct. 2017). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-017-5225-7 (cit. on pp. 134, 214).
- [125] Jackson Burzynski. “Improved Track Reconstruction Performance for Long-lived Particles in ATLAS”. In: *Connecting The Dots 2022*. June 2022. URL: <https://indico.cern.ch/event/1103637/contributions/4821867/attachments/2454227/4206124/CTD2022.pdf> (cit. on p. 134).
 - [126] ATLAS Collaboration. *Search for the dimuon decay of the Higgs boson in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*. ATLAS-CONF-2017-014. 2017. URL: <https://cds.cern.ch/record/2257726> (cit. on p. 134).
 - [127] *Performance of the reconstruction of large impact parameter tracks in the ATLAS inner detector*. Tech. rep. Geneva: CERN, 2017. URL: <https://cds.cern.ch/record/2275635> (cit. on pp. 135, 136, 214).
 - [128] ATLAS Collaboration. “Reconstruction of primary vertices at the ATLAS experiment in Run 1 proton–proton collisions at the LHC”. In: *The European Physical Journal C* 77.5 (Oct. 2017). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-017-4887-5 (cit. on p. 137).
 - [129] Risa Ushioda. “Search for Long-lived Supersymmetric Particles Using Fuzzy Displaced Vertices in pp Collision at $\sqrt{s} = 13$ TeV with the LHC-ATLAS Detector”. Ph.D. thesis. Tokyo Institute of Technology, 2024. URL: <https://repository.cern/records/e19t1-jpq48> (cit. on pp. 137, 140).
 - [130] Hideyuki Oide. *VrtSecInclusive: an inclusive displaced vertex reconstruction package in release 21*. Tech. rep. Geneva: CERN, 2018. URL: <https://cds.cern.ch/record/2303549> (cit. on p. 137).
 - [131] ATLAS Collaboration. *Performance of vertex reconstruction algorithms for detection of new long-lived particle decays within the ATLAS inner detector*. ATL-PHYS-PUB-2019-013. 2019. URL: <https://cds.cern.ch/record/2669425> (cit. on pp. 137, 139).
 - [132] *Performance of vertex reconstruction algorithms for detection of new long-lived particle decays within the ATLAS inner detector*. Tech. rep. Geneva: CERN, 2019. URL: <https://cds.cern.ch/record/2669425> (cit. on p. 140).
 - [133] Matteo Cacciari, Gavin P Salam, and Gregory Soyez. “The anti-ktjet clustering algorithm”. In: *Journal of High Energy Physics* 2008.04 (Apr. 2008), pp. 063–063. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2008/04/063 (cit. on pp. 143, 144).
 - [134] ATLAS Collaboration. “Jet energy scale and resolution measured in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”. In: *The European Physical*

- Journal C* 81.8 (). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-021-09402-3 (cit. on pp. 144–147).
- [135] ATLAS Collaboration. *Flavor Tagging with Track Jets in Boosted Topologies with the ATLAS Detector*. Tech. rep. ATL-PHYS-PUB-2014-013. CERN, 2014. URL: <https://cds.cern.ch/record/1750681> (cit. on p. 145).
 - [136] ATLAS Collaboration. “ATLAS b-jet identification performance and efficiency measurement with $t\bar{t}$ events in pp collisions at $\sqrt{s} = 13$ TeV”. In: *The European Physical Journal C* 79.11 (). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-019-7450-8 (cit. on p. 146).
 - [137] ATLAS Collaboration. “Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton-proton collision data”. In: *Journal of Instrumentation* 14.12 (). ISSN: 1748-0221. DOI: 10.1088/1748-0221/14/12/p12006 (cit. on p. 147).
 - [138] ATLAS Collaboration. “Muon reconstruction and identification efficiency in ATLAS using the full Run 2 pp collision data set at $\sqrt{s} = 13$ TeV”. In: *The European Physical Journal C* 81.7 (). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-021-09233-2 (cit. on p. 148).
 - [139] ATLAS Collaboration. “Performance of the missing transverse momentum triggers for the ATLAS detector during Run-2 data taking”. In: *Journal of High Energy Physics* 2020.8 (). ISSN: 1029-8479. DOI: 10.1007/jhep08(2020)080 (cit. on pp. 150, 217).
 - [140] ATLAS Collaboration. “Search for long-lived, massive particles in events with displaced vertices and missing transverse momentum in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector”. In: *Phys. Rev. D* 97.5 (2018), p. 052012. DOI: 10.1103/PhysRevD.97.052012. arXiv: 1710.04901 [hep-ex] (cit. on p. 154).
 - [141] ATLAS Collaboration. “Search for displaced decays of long-lived, massive particles in events with missing transverse momentum in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector”. In: *To be submitted to JHEP* (2025) (cit. on pp. 168, 170–172, 174, 190, 203, 205, 233, 234, 236–238).
 - [142] G. F. Giudice and A. Romanino. “Split supersymmetry”. In: *Nucl. Phys. B* 699 (2004), pp. 65–89. DOI: 10.1016/j.nuclphysb.2004.08.001. eprint: hep-ph/0406088 (cit. on p. 169).
 - [143] Asimina Arvanitaki et al. “Mini-Split”. In: *arXiv preprint arXiv:1210.0555* (2012). DOI: 10.48550/arXiv.1210.0555 (cit. on p. 169).
 - [144] Howard Baer et al. “Exploring the BWCA (Bino-Wino Co-Annihilation) Scenario for Neutralino Dark Matter”. In: *arXiv preprint hep-ph/0511034* (2005). DOI: 10.48550/arXiv.hep-ph/0511034 (cit. on p. 171).

- [145] Guang Hua Duan et al. “Probing bino-wino coannihilation dark matter under the neutrino floor at the LHC”. In: *arXiv preprint arXiv:1804.05238* (2018). DOI: 10.48550/arXiv.1804.05238 (cit. on p. 171).
- [146] R. E. Shrock and M. Suzuki. “Invisible decays of Higgs bosons”. In: *Physics Letters B* 110.3 (1982), pp. 250–254. DOI: 10.1016/0370-2693(82)91247-3 (cit. on p. 172).
- [147] David Curtin et al. “Exotic decays of the 125 GeV Higgs boson”. In: *Phys. Rev. D* 90 (2014), p. 075004. DOI: 10.1103/PhysRevD.90.075004. eprint: 1312.4992 (cit. on p. 172).
- [148] A. R. Zhitnitsky. “On possible suppression of the axion hadron interactions”. In: *Sov. J. Nucl. Phys.* 31 (1980), p. 260 (cit. on p. 173).
- [149] M. Dine, W. Fischler, and M. Srednicki. “A simple solution to the strong CP problem with a harmless axion”. In: *Phys. Lett. B* 104 (1981), pp. 199–202. DOI: 10.1016/0370-2693(81)90590-6 (cit. on p. 173).
- [150] F. Staub. “SARAH 4: A tool for (not only SUSY) model builders”. In: *Comput. Phys. Commun.* 185 (2014), pp. 1773–1790. DOI: 10.1016/j.cpc.2014.02.018. eprint: 1309.7223 (cit. on p. 173).
- [151] J. Alwall et al. “The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations”. In: *JHEP* 07 (2014), p. 079. DOI: 10.1007/JHEP07(2014)079. eprint: 1405.0301 (cit. on p. 173).
- [152] J. Alwall et al. “The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations”. In: *Journal of High Energy Physics* 2014.7 (July 2014). ISSN: 1029-8479. DOI: 10.1007/jhep07(2014)079 (cit. on p. 176).
- [153] Torbjörn Sjöstrand et al. “An introduction to PYTHIA 8.2”. In: *Computer Physics Communications* 191 (June 2015), pp. 159–177. ISSN: 0010-4655. DOI: 10.1016/j.cpc.2015.01.024 (cit. on pp. 176, 224).
- [154] *ATLAS Pythia 8 tunes to 7 TeV data*. Tech. rep. Geneva: CERN, 2014. URL: <https://cds.cern.ch/record/1966419> (cit. on p. 176).
- [155] Richard D. Ball et al. “Parton distributions with LHC data”. In: *Nuclear Physics B* 867.2 (Feb. 2013), pp. 244–289. ISSN: 0550-3213. DOI: 10.1016/j.nuclphysb.2012.10.003 (cit. on pp. 176, 177, 224).
- [156] Wim Beenakker et al. “NNLL resummation for squark and gluino production at the LHC”. In: *Journal of High Energy Physics* 2014.12 (Dec. 2014). ISSN: 1029-8479. DOI: 10.1007/jhep12(2014)023 (cit. on p. 177).

- [157] Jon Butterworth et al. “PDF4LHC recommendations for LHC Run II”. In: *Journal of Physics G: Nuclear and Particle Physics* 43.2 (Jan. 2016), p. 023001. ISSN: 1361-6471. DOI: 10.1088/0954-3899/43/2/023001 (cit. on p. 177).
- [158] Pierre Artoisenet et al. “Automatic spin-entangled decays of heavy resonances in Monte Carlo simulations”. In: *Journal of High Energy Physics* 2013.3 (Mar. 2013). ISSN: 1029-8479. DOI: 10.1007/jhep03(2013)015 (cit. on p. 177).
- [159] CERN. *CERN Yellow Reports: Monographs, Vol 2 (2017): Handbook of LHC Higgs cross sections: 4. Deciphering the nature of the Higgs sector*. en. 2017. DOI: 10.23731/CYRM-2017-002 (cit. on p. 177).
- [160] Zubair Bhatti et al. *Search for displaced decays of long-lived, massive particles in events with missing transverse energy and displaced vertices in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector*. Tech. rep. ANA-SUSY-2022-15-INT1. ATLAS Collaboration, Feb. 2025. URL: <https://cds.cern.ch/record/2886709> (cit. on pp. 184, 192, 193, 195, 198–201, 209, 212, 215, 216, 220–223, 273–277).
- [161] ATLAS Collaboration. “Search for long-lived, massive particles in events with displaced vertices and missing transverse momentum in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector”. In: *Phys. Rev. D* 97 (2018), p. 052012. DOI: 10.1103/PhysRevD.97.052012. arXiv: 1710.04901 [hep-ex] (cit. on p. 190).
- [162] ATLAS Collaboration. “Search for long-lived, massive particles in events with displaced vertices and multiple jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”. In: (2023). arXiv: 2301.13866 [hep-ex] (cit. on p. 190).
- [163] ATLAS Collaboration. “Performance of missing transverse momentum reconstruction with the ATLAS detector using proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *Eur. Phys. J. C* 78 (2018), p. 903. DOI: 10.1140/epjc/s10052-018-6288-9. arXiv: 1802.08168 [hep-ex] (cit. on p. 217).
- [164] ATLAS Collaboration. “Jet energy scale and resolution measured in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”. In: *The European Physical Journal C* 81.8 (Aug. 2021). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-021-09402-3 (cit. on p. 217).
- [165] Jon Butterworth et al. “PDF4LHC recommendations for LHC Run II”. In: *J. Phys. G* 43 (2016), p. 023001. DOI: 10.1088/0954-3899/43/2/023001. arXiv: 1510.03865 [hep-ph] (cit. on p. 224).
- [166] ATLAS Collaboration. “Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC”. In: *Eur. Phys. J. C* 83 (2023), p. 982. DOI: 10.1140/epjc/s10052-023-11747-w. arXiv: 2212.09379 [hep-ex] (cit. on p. 224).

- [167] G. Cowan et al. “Asymptotic formulae for likelihood-based tests of new physics”. In: *The European Physical Journal C* 71 (2011). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-011-1554-0 (cit. on pp. 227, 228).
- [168] Lukas Heinrich, Matthew Feickert, and Giordon Stark. *pyhf: v0.7.6*. Version 0.7.6. 2023. DOI: 10.5281/zenodo.1169739 (cit. on p. 228).
- [169] Lukas Heinrich et al. “pyhf: pure-Python implementation of HistFactory statistical models”. In: *Journal of Open Source Software* 6.58 (2021), p. 2823. DOI: 10.21105/joss.02823 (cit. on p. 228).
- [170] Alexander Held et al. *scikit-hep/cabinetry: v0.6.0*. Version 0.6.0. 2023. DOI: 10.5281/zenodo.8360833 (cit. on p. 228).
- [171] A. L. Read. “Presentation of search results: the CLs technique”. In: *Journal of Physics G: Nuclear and Particle Physics* 28.10 (2002), pp. 2693–2704. DOI: 10.1088/0954-3899/28/10/313 (cit. on p. 229).
- [172] Glen Cowan et al. “Asymptotic formulae for likelihood-based tests of new physics”. In: *The European Physical Journal C* 71.2 (Feb. 2011). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-011-1554-0 (cit. on p. 229).
- [173] ATLAS Collaboration. “Combination of searches for invisible decays of the Higgs boson using 139 fb⁻¹ of proton-proton collision data at $\sqrt{s} = 13$ TeV collected with the ATLAS experiment”. In: *Physics Letters B* 842 (July 2023), p. 137963. ISSN: 0370-2693. DOI: 10.1016/j.physletb.2023.137963 (cit. on p. 238).
- [174] Peter W. Graham et al. “Experimental Searches for the Axion and Axion-Like Particles”. In: *Annual Review of Nuclear and Particle Science* 65.1 (2015), pp. 485–514. DOI: 10.1146/annurev-nucl-102014-022120 (cit. on pp. 246, 248, 250).
- [175] CAST Collaboration. “New CAST limit on the axion–photon interaction”. In: *Nature Physics* 13 (2017), pp. 584–590. DOI: 10.1038/nphys4109 (cit. on p. 247).
- [176] M. Arik et al. “Search for Solar Axions by the CERN Axion Solar Telescope with ³He Buffer Gas: Closing the Hot Dark Matter Gap”. In: *Phys. Rev. Lett.* 112.9 (2014), p. 091302. DOI: 10.1103/PhysRevLett.112.091302. arXiv: 1307.1985 [hep-ex] (cit. on p. 247).
- [177] A. Álvarez Melcón et al. “First results of the CAST-RADES haloscope search for axions at 34.67 μ eV”. In: *JHEP* 10 (2021), p. 075. DOI: 10.1007/JHEP10(2021)075 (cit. on p. 247).
- [178] C. Adair et al. “Search for dark matter axions with CAST-CAPP”. In: *Nat. Commun.* 13 (2022), p. 6180. DOI: 10.1038/s41467-022-33913-6 (cit. on p. 247).

- [179] E. Armengaud et al. “Axion search with BabyIAXO in view of IAXO”. In: *arXiv preprint arXiv:2012.06634* (2020). URL: <https://arxiv.org/abs/2012.06634> (cit. on p. 248).
- [180] S. Moriyama et al. “Tokyo Axion Helioscope”. In: *arXiv preprint arXiv:1004.1308* (2010). URL: <https://arxiv.org/abs/1004.1308> (cit. on p. 248).
- [181] R. Battxpesti et al. “High magnetic fields for fundamental physics”. In: *Phys. Rep.* (2018). DOI: 10.1016/j.physrep.2018.07.005. arXiv: 1803.07547 [ins-det] (cit. on p. 249).
- [182] ALPS Collaboration. “Towards the first axion search results of the Any Light Particle Search II experiment”. In: *arXiv preprint arXiv:2401.14773* (2024). URL: <https://arxiv.org/abs/2401.14773> (cit. on p. 249).
- [183] OSQAR Collaboration. “Results and Perspectives for Laboratory Search of Weakly Interacting Sub-eV Particles with the OSQAR Experiment”. In: *CERN Document Server* (2013). URL: <http://cds.cern.ch/record/1553025> (cit. on p. 249).
- [184] M. Betz et al. “First results of the CERN Resonant WISP Search (CROWS)”. In: *Phys. Rev. D* 88 (2013), p. 075014. DOI: 10.1103/PhysRevD.88.075014 (cit. on p. 249).
- [185] S. J. Asztalos et al. “Large scale microwave cavity search for dark matter axions”. In: *Phys. Rev. D* 64 (2001), p. 092003. DOI: 10.1103/PhysRevD.64.092003 (cit. on p. 250).
- [186] S. J. Asztalos et al. “Experimental constraints on the axion dark matter halo density”. In: *Astrophys. J. Lett.* 571 (2002), pp. L27–L30. DOI: 10.1086/341130. arXiv: astro-ph/0104200 (cit. on p. 250).
- [187] S.J. Asztalos et al. “A SQUID-based microwave cavity search for dark-matter axions”. In: *Phys. Rev. Lett.* 104 (2010), p. 041301. DOI: 10.1103/PhysRevLett.104.041301. arXiv: 0910.5914 [astro-ph.CO] (cit. on p. 250).
- [188] A. Wagner et al. “A Search for Hidden Sector Photons with ADMX”. In: *Phys. Rev. Lett.* 105 (2010), p. 171801. DOI: 10.1103/PhysRevLett.105.171801. arXiv: 1007.3766 [hep-ex] (cit. on p. 250).
- [189] C. Bartram et al. “Search for Axion Dark Matter from 1.1 to 1.3 GHz with ADMX”. In: *arXiv preprint arXiv:2504.07279* (2024). URL: <https://arxiv.org/abs/2504.07279> (cit. on p. 250).
- [190] S. Ahn et al. “Extensive Search for Axion Dark Matter over 1 GHz with CAPP’S Main Axion Experiment”. In: *Phys. Rev. X* 14 (2024), p. 031023. DOI: 10.1103/PhysRevX.14.031023 (cit. on p. 251).

- [191] M. J. Jewell et al. “Dark matter axion search with HAYSTAC phase II”. In: *arXiv preprint arXiv:2409.08998* (2024). URL: <https://arxiv.org/abs/2409.08998> (cit. on p. 251).
- [192] A. Rettaroli et al. “Search for axion dark matter with the QUAX-LNF tunable haloscope”. In: *Phys. Rev. D* 110 (2024), p. 022008. DOI: 10.1103/PhysRevD.110.022008 (cit. on p. 251).
- [193] A. P. Quiskamp et al. “Near-quantum-limited axion dark matter search with the ORGAN experiment around 26 μeV ”. In: *arXiv preprint arXiv:2407.18586* (2024). URL: <https://arxiv.org/abs/2407.18586> (cit. on p. 251).
- [194] A. Caldwell et al. “Dielectric Haloscopes: A New Way to Detect Axion Dark Matter”. In: *Phys. Rev. Lett.* 118.9 (2017), p. 091801. DOI: 10.1103/PhysRevLett.118.091801. arXiv: 1611.05865 [physics.ins-det] (cit. on p. 252).
- [195] B. A. dos Santos Garcia et al. “First search for axion dark matter with a Madmax prototype”. In: *arXiv preprint arXiv:2409.11777* (2024). URL: <https://arxiv.org/abs/2409.11777> (cit. on p. 252).
- [196] J. Suzuki et al. “Experimental Search for Hidden Photon CDM in the eV mass range with a Dish Antenna”. In: *JCAP* 09 (2015), p. 042. DOI: 10.1088/1475-7516/2015/09/042. arXiv: 1504.00118 [hep-ex] (cit. on p. 256).
- [197] C. Bartram et al. “Search for invisible axion dark matter in the 3.3-4.2 μeV mass range”. In: *Phys. Rev. Lett.* 127 (2021), p. 261803. DOI: 10.1103/PhysRevLett.127.261803. arXiv: 2110.06096 [hep-ex] (cit. on p. 256).
- [198] S. S. Dhillon et al. “The 2017 terahertz science and technology roadmap”. In: *J. Phys. D* 50 (2017), p. 043001. DOI: 10.1088/1361-6463/50/4/043001 (cit. on pp. 256, 258).
- [199] K. Dona et al. “Design and performance of a multi-terahertz Fourier transform spectrometer for axion dark matter experiments”. In: *Journal of Instrumentation* 17.06 (2022), P06014. DOI: 10.1088/1748-0221/17/06/P06014. arXiv: 2104.07157 [physics.ins-det] (cit. on pp. 261–263).
- [200] Chandra M Natarajan, Michael G Tanner, and Robert H Hadfield. “Superconducting nanowire single-photon detectors: physics and applications”. In: *Superconductor Science and Technology* 25.6 (2012), p. 063001. DOI: 10.1088/0953-2048/25/6/063001. arXiv: 1204.5560 [quant-ph] (cit. on pp. 263–265).
- [201] Francesco Marsili et al. “Detecting single infrared photons with 93% system efficiency”. In: *Nature Photonics* 7.3 (2013), pp. 210–214. DOI: 10.1038/nphoton.2013.13. arXiv: 1209.5774 [physics.optics] (cit. on pp. 263, 264).

- [202] Ilya Charaev et al. “Large-area microwire MoSi single-photon detectors at 1550 nm wavelength”. In: *Applied Physics Letters* 116.24 (2020), p. 242603. DOI: 10.1063/5.0005439 (cit. on pp. 264, 267, 268).
- [203] Joel KW Yang et al. “Modeling the electrical and thermal response of superconducting nanowire single-photon detectors”. In: *IEEE Transactions on Applied Superconductivity* 17.2 (2007), pp. 581–585. DOI: 10.1109/TASC.2007.898660 (cit. on p. 264).
- [204] Frank Pobell. *Matter and Methods at Low Temperatures*. 3rd. Berlin, Heidelberg: Springer, 2007. ISBN: 978-3-540-46360-3. DOI: 10.1007/978-3-540-46360-3 (cit. on p. 265).
- [205] Cristián Peña et al. “Superconducting nanowire single-photon detectors for high-energy physics applications”. In: *JINST* 20 (2025), P01008. DOI: 10.1088/1748-0221/20/01/P01008. arXiv: 2410.12194 [physics.ins-det] (cit. on p. 268).