

A Search for Dark Photons with the FASER Detector at the LHC

by

Deion Fellers

A dissertation accepted and approved in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
in Physics

Dissertation Committee:

Laura Jeanty, Chair

Eric Torrence, Advisor

Tim Cohen, Core Member

Andrew Marcus, Institutional Representative

University of Oregon

Spring 2024

© 2024 Deion Fellers

This work is licensed under a Creative Commons
Attribution-NonCommercial-NoDerivs (United States) License.



DISSERTATION ABSTRACT

Deion Fellers

Doctor of Philosophy in Physics

Title: A Search for Dark Photons with the FASER Detector at the LHC

The FASER experiment at the LHC is designed to search for light, weakly-interacting particles produced in proton-proton collisions at the ATLAS interaction point that travel in the far-forward direction. FASER is sensitive to probing previously unconstrained dark photon models, which is a theoretical particle that could provide a portal between the standard model of particle physics and a dark sector that contains a dark matter particle. This dissertation presents the first results from a search for dark photons decaying to an electron-positron pair in FASER, using a dataset corresponding to an integrated luminosity of 27.0 fb^{-1} collected at center-of-mass energy $\sqrt{s} = 13.6 \text{ TeV}$ in 2022 in LHC Run 3. No events are seen in an almost background-free analysis, yielding world-leading constraints on dark photons with couplings $\epsilon \sim 2 \times 10^{-5} - 1 \times 10^{-4}$ and masses $\sim 17 \text{ MeV} - 70 \text{ MeV}$.

This dissertation contains previously published as well as unpublished co-authored materials.

CURRICULUM VITAE

NAME OF AUTHOR: Deion Fellers

GRADUATE AND UNDERGRADUATE SCHOOLS ATTENDED:

University of Oregon, Eugene, OR, USA

University of Colorado Boulder, Boulder, CO, USA

DEGREES AWARDED:

Doctor of Philosophy, Physics, 2024, University of Oregon

Bachelor of Science, Physics, 2017, University of Colorado Boulder

Minor, Mathematics, 2017, University of Colorado Boulder

AREAS OF SPECIAL INTEREST:

Dark Matter

Beyond the Standard Model

PROFESSIONAL EXPERIENCE:

Research Assistant, University of Oregon, 2019 – 2024

Graduate Teaching Assistant, University of Oregon, 2018 – 2019

Scientific Student Intern, Los Alamos National Laboratory, 2012 – 2019

GRANTS, AWARDS AND HONORS:

Chamberlain post-doc fellow at Lawrence Berkeley National Laboratory, 2024

SPOT award at Los Alamos National Laboratory, 2014

PUBLICATIONS:

Abreu, Henso, et al. "Search for Dark Photons with the FASER detector at the LHC." *Physics Letters B* 848 (2024): 138378.

ACKNOWLEDGEMENTS

I am fortunate enough to have many people to thank for having a positive impact on my doctorate journey.

First and foremost, I would like to express my sincere gratitude to my PhD advisor, Eric Torrence, who was always kind, wise, and helpful. I couldn't have asked for a better advisor.

I would also like to thank the physics faculty at the University of Oregon for their genuine kindness and contributions to my knowledge of physics.

This dissertation would not be possible without the contributions from all of those that are part of the FASER collaboration. Not only were the people in FASER hard-working and professional, but they were also very friendly and a real pleasure to work with. I feel incredibly lucky to have worked alongside you all. Thank you FASER!

Thank you to the UCN team at LANL, who took me in as a high school student and mentored me for nearly seven years. I developed my love for physics whilst working with the kind people at UCN and I feel eternally grateful for the opportunities that have come my way because of my time spent there.

Finally, I have immense gratitude for the unwavering love and support from my family and friends. I want to especially thank my loving fiancée, parents, siblings, and grandparents for being my foundation. I love you all.

TABLE OF CONTENTS

Chapter		Page
I	INTRODUCTION	22
II	THEORETICAL OVERVIEW	24
	2.1 The Standard Model	24
	2.1.1 SM Phenomenology	24
	2.1.2 The SM Lagrangian	28
	2.1.3 Success of the SM	31
	2.2 Dark Matter	31
	2.2.1 Evidence of Dark Matter	32
	Galaxy Rotation Curves	32
	Bullet Cluster	33
	Cosmic Microwave Background	34
	2.2.2 Dark Matter Models	36
	2.3 Dark Photon	40
	2.3.1 Basic Theory	40
	2.3.2 Production and Decay	43
	2.3.3 Previously Constrained Parameter Space	45

Chapter		Page
III	THE LHC AND DARK PHOTONS AT FASER	50
	3.1 The Large Hadron Collider	50
	3.2 Dark Photons at FASER	57
IV	FASER DETECTOR	64
	4.1 Tracking Spectrometer	65
	4.1.1 Magnets	66
	4.1.2 SCT Module	67
	4.1.3 Tracking Station	70
	4.1.4 Tracker Metrology and Alignment	73
	4.2 Scintillators	76
	4.3 Calorimeter	80
	4.4 Trigger and Data Acquisition	82
	4.4.1 TDAQ Hardware	82
	4.4.2 Trigger Logic	85
V	DATASET AND SIMULATION SAMPLES	88
VI	EVENT RECONSTRUCTION	93
VII	EVENT SELECTION	96

Chapter	Page
VIII	BACKGROUNDS 100
8.1	Veto Inefficiency 100
8.2	Neutral Hadrons 101
8.3	Large-Angle Muons 103
8.4	Neutrinos 104
8.5	Non-collision events 105
8.6	Summary of the Expected Background 107
IX	SYSTEMATIC UNCERTAINTIES ON THE SIGNAL YIELD 108
9.1	Signal Generator 108
9.2	MC Modelling 110
9.3	Summary of the Systematic Uncertainties 114
X	RESULTS 116
XI	CONCLUSION 119
APPENDICES	
A	PMT WAVEFORM RECONSTRUCTION AND DIGITIZATION . . 120
A.1	Waveform Reconstruction 120
A.1.1	Quality Checks 127
A.2	MC Waveform Digitization 132

Chapter	Page
B	CALORIMETER ENERGY CALIBRATION 139
B.1	Test-Beam 139
B.2	Energy Calibration 141
B.2.1	TI12 MIP Normalization 143
B.2.2	PMT HV Gain Curves 147
B.2.3	MC MIP Normalization 149
B.2.4	EM Energy Calibration 150
B.2.5	Correction to MC Calibrated Energy 151
B.3	Uncertainty in Comparing Data and MC Energies 153
B.3.1	Uncertainties in TI12 MIP Normalization 154
B.3.2	Uncertainties in MC MIP Normalization 156
B.3.3	Uncertainties in Correcting MC Calibrated Energy 157
B.3.4	Uncertainties from other differences in MC and data 159
B.3.5	Total Uncertainty 160
C	VALIDATING TWO-TRACK RECONSTRUCTION WITH DELTA RAYS 163
D	B-L GAUGE BOSON 167
	REFERENCES CITED 170

LIST OF FIGURES

Figure	Page
2.1. Particles in the Standard Model of particle physics [1].	25
2.2. The observed and expected rotational velocity curves of spiral galaxy Messier 33 are shown for varying distance from the galactic center. The predicted velocity curve takes into account both the stellar and gaseous matter distributions. Adding a cold dark matter halo to the luminous matter in the galaxy, provides a rotational velocity curve (shown as solid grey line) that closely matches the observed velocities. The image is from Ref [2] and uses the data from Ref [3].	33
2.3. Images of the Bullet Cluster 1E 0657-558 are shown for both (LEFT) visible light and (RIGHT) x-rays. The green lines in both images show the gravitational lensing contours. The white bar indicates a distance of 200 kpc at the location of the cluster. [4]	34
2.4. Map of the temperature differences in the Cosmic Microwave Background as measured by the Planck satellite [5].	35
2.5. The mass range of allowed dark matter candidates is shown. The diagram is not to scale, and the mass ranges are only approximate. Figure is from Ref [6].	37
2.6. The observed upper limits of the spin-independent WIMP-nucleon cross section are shown in the WIMP phase-space for several different direct-detection experiments. The blue shaded region depicts the WIMP phase space that cannot be easily probed as there would be a significant background of irreducible neutrino scattering events. Figure is from Ref [7].	39
2.7. The Feynman diagrams of the dark photon production channels relevant to this analysis are shown, with light neutral meson decays depicted on the LEFT and dark bremsstrahlung radiation shown on the RIGHT. . .	45
2.8. The dark photon decay branching ratios are shown for varying dark photon mass. Plot taken from Ref [8].	45
2.9. The dark photon parameter space is depicted with shaded regions showing the dark photon mass ($m_{A'}$) and kinetic mixing constant (ϵ)	

Figure	Page
that have been excluded from past experiments. The colors of the shaded regions refer to experiments that utilized a similar A' production mechanism where experiments shaded in RED used electron beam dumps, CYAN used proton beam dumps, GREEN used e+e- colliders, BLUE used pp collisions, MAGENTA used meson decays, and YELLOW used electrons on a fixed target. The GREY constraint comes from the precise measurement of the electron magnetic dipole moment. Figure is from Ref [9].	46
2.10. The dark photon one-loop correction to the electron-photon vertex is depicted.	47
3.1. The CERN accelerator complex is shown, where both the accelerator systems and detector locations are depicted. FASER is not shown in this diagram, as it is a new and relatively small experiment. Figure is from Ref [10].	51
3.2. A cross section of an LHC dipole is shown. Figure is from Ref [11]. . . .	54
3.3. A diagram showing two different orientations of quadrupole magnets, the resulting magnetic field lines (black arrows), and the direction of the force (red arrows) applied to a positively charged particle passing through that area of the quadrupole. Figure is from Ref [12].	55
3.4. A schematic showing the location of FASER with respect to the LHC beam and IP1. More detailed views of the LHC beam are shown in the BOTTOM LEFT for distances up to 170 m from the IP, and in the BOTTOM RIGHT for the area around FASER at ~ 480 m. The beam collision axis line-of-sight is shown as a dashed red line.	58
3.5. The production rates of π^0 from the LHC IP, obtained via EPOS-LHC [13], are shown as a function of the meson's angle with respect to the beam axis (θ_π) and momentum (p_π). The angular acceptances for FASER and FASER2 (a proposed future detector) are indicated by the vertical gray dashed lines. Figure is from Ref [14]	59
3.6. Top left: The acceptance for dark photons to decay inside the FASER decay volume. Top right: The fraction of dark photons decaying inside the FASER decay volume that have energy greater than 500 GeV. Bottom: The expected number of dark photon events in FASER for 27.0 fb^{-1} of data, assuming a 50% signal efficiency, on top of the requirement that the A' energy is greater than 500 GeV.	62

Figure	Page
3.7. FASER's expected sensitivity, with luminosities ranging from 1 to 3000 fb^{-1} , is shown as yellow-shaded areas in the dark photon's parameter space. The gray-shaded regions are excluded by past experiments, and the projected future sensitivities of other experiments are shown as colored contours.	63
4.1. A depiction of the FASER detector, where each color represents a different sub-detector system. The coordinate system used for FASER is also shown in the bottom left, where the particles coming from the ATLAS IP are travelling in the $+\hat{z}$ direction through FASER.	65
4.2. A cross section of a cylindrical dipole magnet in FASER. The arrows indicate the direction of the magnetic field in each of the permanent magnetic blocks. Note that the final orientation of the inner magnetic field is in the $-\hat{x}$ direction, which would correspond with the vertical direction in this figure.	67
4.3. The FEMM [15] simulated 2D magnetic field distribution is shown in a quadrant of the dipole magnet. Note that for this figure, the \hat{x} direction is vertical and the \hat{y} direction is horizontal.	68
4.4. The magnetic field components, measured with the Hall probe method, are shown against z while at the center of one of the 1 m long dipole magnets.	69
4.5. A photograph of a SCT barrel module inside an aluminum test-box. . .	69
4.6. A photograph of a tracker plane, where 8 SCT modules are installed in an aluminum frame.	71
4.7. A tracker station is depicted as an expanded CAD view on the LEFT, and as a photograph of the final assembly on the RIGHT.	71
4.8. The number of hits on track, per mm^2 , is shown across the transverse area of the most upstream tracking station. Both the vertical gaps between the sensors of a module and the staggered horizontal gaps between the modules can be seen. Less hits on track are seen for positions with $r \geq 100$ mm because that region is where the magnetic material lies, which scatters tracks and complicates the track reconstruction.	72
4.9. A picture of the metrology scans being performed on a tracker plane. . .	74
4.10. Comparisons of the reconstructed track residuals are shown for tracks reconstructed from data in an unaligned geometry (red), data in an	

Figure	Page
aligned geometry (black), and from MC simulation with an ideal geometry (orange). In the TOP, the track residual distributions are shown when combining the residuals from all tracker modules. The distributions are normalized to the same number of entries and fit with a Gaussian. The BOTTOM plot shows the means of the track residual distributions for each tracker module. For both plots, the data used has an integrated luminosity of 596 pb^{-1} and only the tracks with momentum greater than 300 GeV and the number of clusters greater than 14 are used.	75
4.11. The VetoNu (TOP-LEFT), Veto (TOP-RIGHT), Timing/Trigger (BOTTOM-LEFT), and Pre-Shower (BOTTOM-RIGHT) scintillator stations are shown.	77
4.12. A picture of the FASER calorimeter is shown, where four LHCb outer ECal modules are placed 2×2 , and the PMTs are coming out of the back of the calorimeter.	81
4.13. The calorimeter module design is depicted.	81
4.14. The calorimeter PMT assembly is shown.	81
4.15. A diagram depicting the FASER TDAQ architecture. Each component is represented by a colorful box and the connections between the components are shown with arrows. The numbers in parentheses indicate the number of channels. The figure is from Ref [16].	83
4.16. The recorded trigger rates for each of the four FASER TLB triggers are shown during LHC Fill 8143 on August 19, 2022. The triggers are further described in Sec. 4.4.2. The structure and decay in rate over time closely follows the LHC luminosity at IP1.	87
5.1. The cumulative luminosity versus time delivered to (yellow) and recorded by FASER (blue) during stable beams are shown.	89
5.2. The luminosity-normalized detector response is shown for all 64 runs used in this analysis. The total number of reconstructed tracks, as well as the number of positively or negatively curved tracks, are shown in the TOP plot; whereas the BOTTOM plot shows the total number of events with energy in the calorimeter that is greater than that of a 100 GeV electromagnetic shower.	90
5.3. The dark photon samples generated for this analysis are shown as red dots in the phase space of the A' mass and coupling (ϵ). The yellow	

Figure	Page
regions show the expected sensitivity of FASER for varying amounts of delivered luminosity. Regions previously excluded by other experiments are shown in grey, and regions that could be excluded in the future by other experiments are shown as colorful dotted lines.	91
6.1. The reconstructed event display of a muon traversing FASER. The event (run 8336, eventID 1477982) occurred on August 23, 2022 and is in time with a collision event. The measured track momentum is 21.9 GeV. The waveforms are shown for signals in the scintillator and calorimeter modules, and are fit using a Crystal Ball function. All PMT waveforms are consistent with a muon passing through the scintillators and one of the calorimeter modules. The detected hits in the semiconductor tracker modules are shown with blue lines and the reconstructed track is shown with a red line. In the title of the waveform plots, left and right is defined facing the downstream direction.	95
7.1. A sketch presenting a side view of the FASER detector, showing the different detector systems as well as the signature of a dark photon (A') decaying to an electron-positron pair inside the decay volume. The white blobs depict where measurements are taken for the A' signal and the solid red lines represent the reconstructed tracks produced by the e^+e^- pair.	96
7.2. The maximum reconstructed charge from the timing station scintillators for data, and MC A' simulated signal. The data is dominated by single muon events. The minimum required charge of 70 pC is shown as a dashed line.	98
8.1. Example plots of the charge distribution for events with good fiducial tracks for the most downstream VetoNu scintillator (left) and the most upstream Veto scintillator (right). The efficiency is calculated as the ratio of the number of events with charge above the 40 pC threshold (indicated on the figures) to the total number of events.	101
8.2. The calorimeter energy in simulated neutrino events passing all signal selection requirements, besides that on the calorimeter energy. GENIE is used to simulate the neutrino interactions. The figure is scaled to a luminosity of 27.0 fb^{-1}	105
8.3. The calorimeter energy distributions are shown for different track requirements in both the (left) cosmic ray control sample and the (right) beam-1 background control sample.	107

Figure	Page
9.1. The energy spectrum of dark photons in FASER produced with meson production modeled by different generators (EPOS-LHC, QGSJET II-04 and SIBYLL 2.3d). Also shown is production from bremsstrahlung with a factor of two variation in the p_T cut off. The bottom panel shows the ratio between the different estimates, and the parameterisation of the uncertainty as a function of energy. A representative signal model (with $m_{A'}=50$ MeV and $\epsilon=3 \times 10^{-5}$) is shown.	109
9.2. (LEFT) The E/P vs. Invariant Mass distribution is shown for the two lowest momentum tracks in three-track events found in data. Requiring an $E/P > 0.5$ is shown to select photon conversion event candidates, whose invariant mass is zero, apart from most other three-track events that have a non-zero invariant mass. (RIGHT) The ratio of the signal in the downstream vs upstream preshower layer is shown for both photon conversion event candidates (black) and other three track events (red) in data, where a preshower ratio cut of > 2 is used to efficiently select 3-track events in data that have an $E/P > 0.5$	112
9.3. (LEFT) The E/P distributions are shown for three-track photon conversion events in both data (black) and Fluka MC simulation (red) for a track momentum range of 75-175 GeV. (RIGHT) The Gaussian-fitted peak position of the E/p in data and MC simulation is shown for several different momentum ranges: $20 \text{ GeV} < p < 35 \text{ GeV}$, $35 \text{ GeV} < p < 75 \text{ GeV}$, $75 \text{ GeV} < p < 125 \text{ GeV}$, $125 \text{ GeV} < p < 175 \text{ GeV}$. For all ranges measured, the E/P peak in both data and MC agree within the conservative 6% uncertainty (blue dashed line) used for the calorimeter energy calibration.	112
9.4. (Top) The two-track reconstruction efficiency versus track separation for overlaid tracks in both data and MC events are shown. The distribution of the separation in e^+e^- tracks of an A' sample is also shown in red with the axis on the right-hand side. (Bottom) The ratio of the overlay tracking efficiencies between MC and data is depicted.	114
10.1. The calorimeter energy distribution for data and three representative MC simulated signal models are shown for (a) all events with at least one good track, (b) events that have no signal in the veto stations and at least one good track, and (c) events that have no signal in the veto stations and exactly two good fiducial tracks. The distributions and expected events from the MC samples are scaled to 27.0 fb^{-1}	117
10.2. 90% confidence level exclusion contours in the dark photon parameter space are shown. Regions excluded by previous experiments are shown	

Figure	Page
in grey (see main text for details). The red line shows the region of parameter space that yields the correct dark matter relic density, with the assumptions discussed in the text.	118
A.1. A typical example of a PMT raw waveform signal coming from the digitizer.	120
A.2. The distribution of ADC values in a given waveform are shown. We fit the distribution around the most common value to obtain the baseline and baseline rms values.	121
A.3. Both figures show a typical waveform after baseline subtraction, inversion, and selection of the reconstruction window. You can see that the reconstruction window differs depending on which sub-detector is being reconstructed, where the (TOP) plot is for a Trigger scintillator and has a reconstruction window from 780 - 900 ns and the (BOTTOM) plot is for a preshower scintillator and has a reconstruction window from 800 - 920 ns.	122
A.4. Crystal Ball fits in blue are shown atop data points in red for a regular waveform pulse (TOP) and a saturated waveform pulse (BOTTOM). . .	124
A.5. The timing resolution of the first preshower layer is shown against the constant-fraction threshold used to measure the timing, as described in the text. The other scintillator and calorimeter channels have a similar shape, and thus a 40% constant-fraction threshold was chosen to measure the pulse time as it minimizes the timing resolution. The signals used to do this measurement were from muons in TI12 collision data.	125
A.6. The mean reconstructed waveform baseline is plotted per FASER run used in the dark photon analysis. The baseline measurements are shown for each calorimeter and scintillator module, and are all normalized to the first run (8725).	128
A.7. The mean reconstructed waveform baseline rms is plotted per FASER run used in the dark photon analysis. The rms measurements are shown for each calorimeter (TOP) and scintillator (BOTTOM) module, and are all normalized to the first run (8725).	128
A.8. The mean reconstructed waveform time is plotted per FASER run used in the dark photon analysis. The time measurements are shown for each calorimeter and scintillator module, and are all normalized to the first run (8725).	129

Figure	Page
A.9. The ratios of the fitted and raw measurements of the waveform peak (TOP) and integral (BOTTOM) are shown for the preshower and calorimeter modules. The entire 2022 data-set was used, where the ratio was calculated for each waveform that had a good fit. Each distribution is normalized such that the sum of it's bins is equal to one.	131
A.10. The digitized MC waveforms of the second preshower layer are shown for varying initial electron energies.	133
A.11. A saturated MC digitized waveform for the Trigger/Timing layer is shown.	134
A.12. Comparisons of the average normalized PMT waveforms are shown between FASER physics data and a digitized MC muon sample for a calorimeter module (LEFT) and a preshower layer (RIGHT).	136
A.13. Comparisons of the reconstructed waveform baseline RMS are shown between FASER physics data and a digitized MC muon sample for a calorimeter module (LEFT) and a preshower layer (RIGHT).	137
A.14. Comparisons of the reconstructed PMT charge distributions are shown between FASER physics data and a digitized MC muon sample for a calorimeter module (TOP) and a preshower layer (BOTTOM).	138
B.1. A diagram of test-beam setup, where the beam-line is coming in from the left.	140
B.2. A photo of test-beam setup in the H2 beam line. During operation, the detector apparatus would be raised such that it is at the same height as the beam-line.	141
B.3. The above plots show the calorimeter channel-0 MIP charge distributions (TOP) and the charge vs position (BOTTOM) for three different track position cuts. The cuts are depicted by the black rectangle in the bottom scatter plots. The distribution above each scatter plot corresponds to the charge distribution using that cut. The bright yellow spot in the bottom plots depict the area where the MIP goes through the PMT and produces a larger signal. The data is from FASER runs 8715 and 8717, where the calorimeter is at high gain and has the light filters installed. The charge distributions are fit with a Landau+Gaussian convolution function, described in the text, over the whole histogram range (Red) and also over -0.5 to +1.0 standard deviations around the MPV (Blue). . . .	144

- B.4. The results are shown for comparing the original MPV of a toy model with the MPV from two different fits. The toy model is a Landau distribution with a $MPV = 1.45$ pC and width = 0.27 pC, that is then smeared by a Gaussian of varying width. The simple Landau fit shows the bias in the MPV due to the Gaussian noise, whereas the Landau+Gaussian convolution fit described in the text allows us to measure the MPV of the original Landau distribution regardless of the noise. 145
- B.5. The above four plots show the MIP distributions used for the calibration of each TI12 calorimeter module. The data is from FASER runs 8715 and 8717, where the calorimeter is at high gain and has the light filters installed. The distributions are fit with a Landau+Gaussian convolution function over the whole histogram range (Red) and also over -0.5 to +1.0 standard deviations around the MPV (Blue). The Gaussian width used in the fit function was measured 146
- B.6. PMT HV gain curves, described in the text, are shown for the four PMT's used in the TI12 Calorimeter. The HV gain curves were measured with a 75 GeV electron beam (Green) and an LED signal (Red) during the 2021 test-beam period. The data points show the mean PMT charge measurement vs PMT HV and the curves show a 5-degree polynomial fit of the data. 148
- B.7. The initial test-beam energy (E_{TB}) divided by the average MIP-normalized signal size \bar{N}_{MIP}^{TB} in calorimeter module 0 is shown across electron beam energies ranging from 5 to 300 GeV. The dashed blue line shows the 0.33 GeV value used in the simple estimate of E_{EM} 151
- B.8. The normalized calibrated energies for the 2021 test-beam data and MC are shown. Each color of the dots represents the calibrated energy at one of the six calorimeter modules in the test-beam setup, whereas the squares with the line represent the corrected MC calibrated energy. . . . 152
- B.9. The mean calorimeter PMT charge from LED pulses are shown normalized to the first measurement and spanning a two month period. The trend of each PMT response seems to vary in unison with the other calorimeter modules, depicting the LED intensity changing with temperature. The spread of the PMT responses over time depicts the PMT drift, which is $< 1\%$ over this time period. 156
- B.10. This figure depicts the extrapolation done on the test-beam data and MC calibrated energies to estimate the difference at 500 GeV. The average data points were obtained from averaging the calibrated energy of each of the six test-beam calorimeter modules. The MC calibrated energies here

Figure	Page
include the 8.8% correction factor used to make the MC and data agree at an EM energy of 100 GeV. Both data and MC were linearly fit using the range from 50 to 150 GeV, then this fit was used to extrapolate the calibrated energies to 500 GeV, resulting in an approximated difference of 2.46% at 500 GeV.	159
B.11. The above plots depict the difference in test-beam MC with regards to lateral shower development and loss out the edges of a calorimeter module that is centered at X=0 and Y=20 mm. The top plots show the normalized calorimeter responses in both data (Blue) and MC (Green) for 30 GeV electrons traversing the edge of a calorimeter module in X (Left) and Y (Right). The edge in X is adjacent to another calorimeter module that is not being read out, whereas the edge in Y is accompanied by just air. The bottom plots show the % difference in the data and MC vs position. The dashed red line shows the position approximately 20 mm from the edge. The positions are obtained from linearly extrapolating the track segment in the test-beam tracking station to the face of the calorimeter. For the scan across X (left plots) we restrict the Y position to be within -10 and 0 mm, and for the scan across Y (right plots), we restrict the X position to be between 20 and 30 mm.	161
C.1. The reconstructed momentum of the delta-ray candidate track is shown for both data and the Fluka MC sample described in Chapter V. To normalize the two distributions, both were scaled by the inverse of their total number of entries.	163
C.2. The reconstructed track multiplicity in the delta-ray control sample described in the text is shown. The distributions for both data and MC are shown on the TOP. The ratio of MC to data is shown on the BOTTOM plot, where the MC was scaled to match the total number of candidate delta ray events in data.	165
C.3. The separation between the two reconstructed tracks at the first tracking stations, in the delta-ray sample in data and MC simulation. The MC was scaled to match the total number of candidate delta ray events in data.	166
D.1. 90% confidence level exclusion contours in the $B - L$ gauge boson parameter space are shown. Regions excluded by previous experiments are shown in grey.	169

LIST OF TABLES

Table	Page
7.1. Summary of the scintillator PMT charge thresholds used in the A' event selection. Because the Trigger/Timing layer is split into a top and bottom plane, the two-MIP minimum cut on the Trigger/Timing PMT charge accommodates for the case if two particles go into the same plane or separate planes.	98
7.2. Cutflow for a representative signal model with $\epsilon = 3 \times 10^{-5}$ $m_{A'} = 25.1$ MeV, showing the fraction of MC simulated signal events passing the different event selection criteria. The efficiency numbers in the rows marked with a * have been scaled down by 7% following a data driven estimate of the two-track tracking efficiency, as described in Sec. 9.2. The statistical uncertainty is smaller than the last digit shown.	99
8.1. The measured efficiency of the five veto scintillator planes (the number in parenthesis is the statistical uncertainty).	101
8.2. Summary of the different background estimates. Estimates that are negligible are represented with a “-”.	107
9.1. Summary of the systematic uncertainties on the signal yield. For each of the sources of uncertainty, the source and size of the uncertainty is presented. The effect on the signal yield across the full signal parameter space probed is also shown. The numbers in parenthesis indicate the effect on the signals within the parameter space for which this analysis is sensitive.	115
A.1. This table gives the values and descriptions of the waveform status word.	126
A.2. All reconstructed waveform variables are shown and described.	126
A.3. The ratio of waveforms that had a failed fit or an overflow are shown for the 2022 data set.	130
A.4. The digitization scale factor for the Calorimeter are shown for the four different run configurations.	134
A.5. The MC digitization parameters for each detector type are shown.	135

Table	Page
B.1. This table summarizes all uncertainties and their respective values when comparing the calibrated energies of data and MC at 500 GeV.	162

CHAPTER I

INTRODUCTION

This dissertation incorporates and expands upon the article titled “Search for Dark Photons with the FASER Detector at the LHC”, which was published in volume 848 of Physics Letters B in January of 2024 [17]. I am a co-author and lead editor for the article; which was made possible by the entire FASER collaboration, consisting of 93 members from 26 institutions and 10 countries.

The existence of dark matter is strong evidence for new particles beyond the standard model (BSM) of particle physics. Yet, with decades of experimental effort trying to detect dark matter, all that has been learned is what dark matter is not. Expanding the scope of the search, experiments have been increasingly trying to detect other BSM particles that may not account for dark matter, but could still interact with it. These particles are referred to as “portal particles” and could be part of an entire complex dark sector of BSM particles. Detecting portal particles may give us new insights into the particle nature of dark matter, which is of great interest in both particle physics and astrophysics [18, 19].

The dark photon is a theoretical BSM particle that arises from a simple extension to the standard model and could be a portal particle. Certain models of the dark photon and dark matter particle are cosmologically favored due to their ability to explain the observed dark matter abundance through the common freeze-out method. Apart from its relative simplicity and ability to explain cosmological observations, the dark photon is a prominent BSM particle because it often provides visible signatures that can be experimentally probed.

FASER is an experiment that is designed to search for light and weakly-interacting particles, including certain dark photon models, that are produced in the far-forward direction from proton-proton collisions at the Large Hadron Collider (LHC). The detector is positioned 480 m downstream of the proton-proton collisions and aligned with the beam collision axis, where particles must traverse strong LHC magnets and approximately 100 m of rock and concrete before reaching FASER. The novel location of the experiment provides both a low background environment and a high-intensity flux of long-lived particles that allow FASER, a relatively simple and small experiment, to probe previously unconstrained models of the dark photon.

This dissertation will present FASER's search for dark photons using a dataset corresponding to an integrated luminosity of 27.0 fb^{-1} that was collected at a center-of-mass energy of $\sqrt{s} = 13.6 \text{ TeV}$ in 2022 during LHC Run 3. I will first briefly introduce the theory of the standard model, dark matter, and dark photons in Chapter II. I will then give an overview of the LHC and discuss the potential of searching for dark photons with FASER in Chapter III. Chapter IV gives a description of the FASER detector. Next, the data and simulation samples used in this analysis are presented in Chapter V. The event reconstruction and signal selection are then discussed in Chapters VI and VII, respectively. Chapter VIII presents the backgrounds that could mimic the dark photon signal in FASER. The uncertainties in the dark photon signal yield are then estimated in Chapter IX. Finally, the results of the analysis are presented in Chapter X, and the conclusions in Chapter XI.

CHAPTER II

THEORETICAL OVERVIEW

The theoretical motivation behind searching for a dark photon is presented in this chapter. To start, a brief overview of the Standard Model of particle physics is provided in Sec. 2.1. Dark matter, which can not be described by the Standard Model but is supported by numerous cosmological observations, is then discussed in Sec. 2.2. Finally, an overview of the basic theory of the dark photon, which can provide a portal between the Standard Model and a dark sector that contains a dark matter particle, is introduced in Sec. 2.3.

2.1 THE STANDARD MODEL

Building upon special relativity and quantum mechanics, the Standard Model (SM) of particle physics is a quantum field theory describing all known fundamental particles and forces in our universe.

In this section, I will give a brief overview of the phenomenology of the SM, its mathematical description through the Lagrangian, and its success in agreeing with what is observed in nature.

2.1.1 SM Phenomenology

A diagram depicting the particles in the SM can be seen in Fig. 2.1. Not shown in the diagram is that all electrically charged particles have a corresponding anti-particle that is identical except for having the opposite electric charge. For example, the electron has a charge of -1 and its anti-particle is the positron, which has a charge of $+1$ whilst having the same exact mass and spin of the electron.




FIGURE 2.1. Particles in the Standard Model of particle physics [1].

All particles are either a fermion or a boson, which have half-integer and integer spin respectively. Fermions obey Fermi-Dirac statistics and the Pauli Exclusion Principle, which states that no identical fermions can occupy the same quantum state. Bosons, on the other hand, are exempt from this principle and can coexist with other identical bosons in the same quantum state. In the SM, leptons and quarks, generally referred to as the matter particles, are fermions; whereas the force-carrying particles and the Higgs are all bosons.

Leptons have three generations of electron-like particles and their associated neutrinos, including the electron, muon, and tau flavors. Although the muon and tau are just heavier versions of the electron, this also means that they are unstable particles, unlike the electron, and quickly decay to lighter particles via the weak force.

On the other hand, all neutrinos are stable and do not decay, but rather oscillate between one another as they travel through space, because every neutrino is actually a superposition of all three neutrino flavors. Being both stable and nearly massless means that neutrinos are abundant throughout the universe, yet they only interact with other SM particles via the weak force, which occurs extremely rarely.

Quarks are the building blocks of protons, neutrons, and many other composite particles, called hadrons. There are six quark flavors: the up, down, charm, strange, top, and bottom quarks. The up-like quarks (up, charm, and top) all have an electric charge of $+2/3$ and the down-like quarks (down, strange, and bottom) all have an electric charge of $-1/3$. In addition to electric charge, each quark also has any one of three color charges: red, green, or blue. Anti-quarks, which have opposite electric charge, also have opposite color charge: anti-red, anti-green, or anti-blue. When forming hadrons, color confinement demands that the quarks must combine to be color neutral; which means that a color and anti-color pair of quarks canceled each other's color, or there is a triplet of quarks and all three colors are present. The pairs of quark and anti-quark are referred to as mesons and the quark triplets are referred to as baryons. An example of a meson is a pion and is formed from anti-quark and quark pairs of the up and down quarks, where $\pi^0 = u\bar{u}$ or $d\bar{d}$, $\pi^- = d\bar{u}$, and $\pi^+ = u\bar{d}$. The most common baryons are the proton and neutron, which are formed from triplets of the up and down quarks that combine as uud to form a proton and udd to form a neutron, where the electric charges of the quarks add up to make the proton have an electric charge of $+1$ and the neutron to be electrically neutral. Thus, along with the electron, the up and down quark form all of the atoms in the entire universe. The other quarks can also combine to form numerous other hadrons, but these are more rare due to their larger masses and shorter lifetimes. In addition

to the valence quarks in hadrons, which govern the quantum numbers and thus the characteristics of the hadron, there is also a sea of virtual quark and anti-quark pairs that are constantly being produced from, and annihilating into, the gluons that bind the valence quarks together.

There are four known fundamental forces that govern the universe. The electromagnetic force, weak nuclear force, and strong nuclear force are all described by the vector bosons in the SM, who have spin = 1. The fourth fundamental force is that due to gravity, which is not included in the SM, but rather best explained by Einstein's theory of general relativity.

The electromagnetic force is mediated by the photon, a massless and stable boson that interacts with all electrically charged particles. The photon is ubiquitous in our universe and responsible for all things that have to do with electricity and magnetism. Virtual photons can be thought of as constantly shuttling electromagnetic forces between charged particles, and thus govern all atomic/molecular/chemical forces and reactions in nature. In addition, photons that are coherent quantum states of the electromagnetic field are the quanta of light.

Gluons are the mediator of the strong nuclear force, which holds together the quarks in hadrons and also binds the protons and neutrons in atomic nuclei. Similar to how the virtual photon can be thought of as shuttling electromagnetic forces between electrically charged particles, the gluon implements the strong force between color-charged quarks. There are actually eight different gluons, each attributed to different combinations of color and anti-color charges of the quarks. The strong force is unique in that it has a linear effective potential that leads to the gluon energy increasing along with the separation distance between color-charged quarks. Thus, at separation distances of $>\sim 1$ fm, it becomes energetically favorable to pull quarks

out of the sea of virtual quarks and bind with them rather than to continue separating the individual valence quarks. This phenomenon is called color confinement and is why isolated individual quarks have never been observed.

The weak nuclear force is responsible for the radioactive decay of atoms and is mediated by the massive W^\pm and Z bosons. The W^+ and W^- bosons are electrically charged and allow both quarks and leptons to change flavors. For example, beta decay is the process where a neutron decays into a proton, and is mediated by the W^+ boson which allows an up quark to become a down quark.

The Higgs boson [20] is unique in that its field gives the other fundamental particles their masses. This is accomplished through the scalar Higgs field having a non-zero vacuum expectation value, which spontaneously breaks the electroweak symmetry and results in the massive weak gauge bosons (W^\pm and Z) and the massless electromagnetic photon [21]. Through Yukawa interactions, the fundamental fermions also get their mass from the Higgs field. Although the Higgs boson gives the fundamental particles their mass, the vast majority of mass in atomic matter actually comes from the binding energy of the strong force on the quarks in both the proton and neutron.

2.1.2 The SM Lagrangian

Mathematically, the SM is described via quantum field theory, where each particle is an excitation of a quantum field that permeates all of space. The local gauge symmetries $SU(3)_C \times SU(2)_L \times U(1)_Y$ define the SM; where the $SU(3)_C$ symmetry corresponds to the strong nuclear force, while the $SU(2)_L \times U(1)_Y$ symmetries correspond to the electroweak force which is spontaneously broken by the Higgs mechanism in order to generate the seemingly separate weak nuclear and

electromagnetic forces. The probability amplitudes for how the quantum fields of these forces and the SM fermions interact and move through spacetime is calculated by integrating the action S over all possible paths D and field configurations $\Theta(x)$ [21]:

$$\int D \Theta(x) e^{iS[\Theta(x)]/\hbar} \quad (2.1)$$

where \hbar is Planck's constant divided by 2π , x is the space-time four-vector, and the action is just the space-time integral of the Lagrangian density: $S = \int \mathcal{L} d^4x$. Thus the Lagrangian density, defined as the kinetic energy of a system minus the potential energy per unit volume V : $\mathcal{L} = (KE - PE)/V$, is all that is needed to completely define the SM.

The SM Lagrangian is very complex, but for the purpose of this analysis, only a rough and simple overview will be given.

When using natural units and assuming a $(+,-,-,-)$ metric, the SM Lagrangian [21] can be compactly written as:

$$\begin{aligned} \mathcal{L} = & -\frac{1}{4}F^{\mu\nu}F_{\mu\nu} \\ & +\bar{\psi}(i\gamma^\mu D_\mu)\psi + h.c. \\ & +\psi_i y_{ij}\psi_j\Phi + h.c. \\ & +|D_\mu\Phi|^2 - V(\Phi) \end{aligned} \quad (2.2)$$

The first line of Eq. (2.2) corresponds to the kinetic energy of the force carriers, and the $F^{\mu\nu}F_{\mu\nu}$ term can be expressed in terms of the individual strong force and electroweak field strength tensors as:

$$F^{\mu\nu}F_{\mu\nu} = B^{\mu\nu}B_{\mu\nu} + W^{i\mu\nu}W_{\mu\nu}^i + G^{a\mu\nu}G_{\mu\nu}^a \quad (2.3)$$

where $B_{\mu\nu}$, $W_{\mu\nu}^i$, and $G_{\mu\nu}^a$ are the field strength tensors respectively corresponding to the $U(1)_Y$, $SU(2)_L$, and $SU(3)_C$ gauge invariance. The $W_{\mu\nu}^i$ field strength tensor has three gauge fields, each accounted for with $i = \{1, 2, 3\}$. Similarly, $G_{\mu\nu}^a$ represents the strong force and has eight gluons, as discussed in Sec. 2.1.1, with each accounted for by $a = \{1, 2, \dots, 8\}$.

To further understand a field strength tensor, which is important for the theory of the dark photon, as discussed in Sec. 2.3, let us observe the electromagnetic field strength tensor $F_{\mu\nu}^{EM}$ in terms of the electric (E) and magnetic (B) field vector components:

$$F_{\mu\nu}^{EM} = \begin{pmatrix} 0 & -E_x & -E_y & -E_z \\ E_x & 0 & -B_z & B_y \\ E_y & B_z & 0 & -B_x \\ E_z & -B_y & B_x & 0 \end{pmatrix} \quad (2.4)$$

As you can see, the field strength tensor describes the components of the fields in spacetime, and $F^{EM\mu\nu}F_{\mu\nu}^{EM} = 2(B^2 - E^2)$ gives a Lorentz invariant term that is proportional to the total energies of the electric and magnetic fields.

As for the second line of the SM Lagrangian in Eq. (2.2), “ $\bar{\psi}(i\gamma^\mu D_\mu)\psi + h.c.$ ” describes both the kinetic energy of the fermions (ψ) and how they interact with the force carrying particles. The covariant derivative D_μ is defined as:

$$D_\mu = \partial_\mu - \frac{i}{2}g'YB_\mu - \frac{i}{2}g\tau_L^iW_\mu^i - \frac{i}{2}g_s\lambda^aG_\mu^a \quad (2.5)$$

with (g', g, g_s) the coupling strength constants, (Y, τ_L^i, λ^a) the generators, and $(B_\mu, W_\mu^i, G_\mu^a)$ the gauge boson fields of the $U(1)_Y$, $SU(2)_L$, and $SU(3)_C$ symmetries respectively.

Finally, the last two lines of Eq. (2.2) describe how the Higgs boson (Φ) contributes to the SM. The “ $\psi_i y_{ij} \psi_j \Phi + h.c.$ ” term governs how the fermions couple to the Higgs field through the Yukawa coupling matrices y_{ij} , and thus also determines the masses of the fermions. Whereas the “ $|D_\mu \Phi|^2$ ” term shows how the Higgs couples to the force-carrying particles through the covariant derivative, described in Eq. (2.5). Lastly, the Higgs potential “ $V(\Phi)$ ” encompasses the self interactions of the Higgs and the vacuum expectation value.

2.1.3 Success of the SM

Upon inputting the experimentally measured values of the 19 free parameters [22] in the SM (including the fermion and Higgs masses, gauge coupling constants, Higgs vacuum expectation value, etc...), the SM can then be used to predict the behavior of particles as observed in nature, and has been experimentally proven time and time again. Arguably one of the more impressive tests of the SM is the measurement of the electron’s magnetic dipole moment, which agrees with the SM prediction to within 1 part per 100 billion [23]. The SM has also predicted several particles before their discovery. The most recent of which, the Higgs boson, was discovered in 2012 [24, 25] after nearly 50 years from being proposed by Peter Higgs in 1964 [20].

As incredibly successful as the SM is, it is still believed to be incomplete as it fails to explain numerous experimental observations, such as dark matter.

2.2 DARK MATTER

There is strong cosmological evidence for gravitational sources that cannot be described by the SM nor Einstein’s general theory of relativity. This unknown source of gravity is called dark matter. Making up 26.4% of the energy density and 84.4%

of the total mass in the universe [7], the mystery of dark matter has been of prime interest to scientists for many decades [26].

In this section, I will give an overview of the phenomenology of dark matter. Some of the main cosmological observations that point to the existence of dark matter are described in Sec. 2.2.1. Afterwards, an overview of the most popular attempts at describing dark matter are briefly discussed in Sec. 2.2.2.

2.2.1 Evidence of Dark Matter

There are numerous cosmological observations [7] that point to the existence of cold (non-relativistic) dark matter which does not interact with the SM electromagnetic force. Some of the more convincing observations are that of galaxy rotation curves, the bullet cluster, and the cosmic microwave background, which will all be discussed here.

Galaxy Rotation Curves:

The rotational velocities of galaxies have been measured to be much faster than expected when just taking into account the gravitational forces from visible matter. For example, in a nearby galaxy called Messier 33, the rotational velocities about the galactic center of both starlight and hydrogen gas were measured [3], and shown in Fig. 2.2. When compared to the expected rotational velocity of the galaxy, which takes into account both the stellar and gaseous matter distributions, the measured rotational velocities are significantly larger, especially far from the center of the galaxy. Instead of the rotational velocities decreasing at distances outside the stellar ring of the galaxy, as expected, the observed rotational velocities continue to increase. The discrepancies between the observed and expected rotational velocity curves can

be remedied through the addition of a cold dark matter halo that dominates the luminous matter of the galaxy and whose density decreases radially as $R^{-1.3}$ [3].




FIGURE 2.2. The observed and expected rotational velocity curves of spiral galaxy Messier 33 are shown for varying distance from the galactic center. The predicted velocity curve takes into account both the stellar and gaseous matter distributions. Adding a cold dark matter halo to the luminous matter in the galaxy, provides a rotational velocity curve (shown as solid grey line) that closely matches the observed velocities. The image is from Ref [2] and uses the data from Ref [3].

Bullet Cluster:

The Bullet Cluster is a galaxy cluster merger consisting of two sub-clusters that have undergone a high-speed collision and whose gravitational centers do not coincide with those expected from the visible baryonic matter [4]. Within each galaxy sub-cluster, there are hundreds-to-thousands of galaxies and also an immense amount of hot ionized gas, known as plasma. Upon collision of the two sub-clusters, it is seen that the stellar components of the clusters pass right through one another, whereas the plasma clouds collide and trail behind the stellar components of the clusters.

This is depicted in Fig. 2.3, where the left plot shows the stellar components as observed through visible light, and the right plot shows the locations of the plasma clouds, which emit x-ray radiation. With the mass of the plasma clouds estimated to be greater than the total stellar mass in the galaxy clusters, one would expect the strongest gravitational effects to be centered on the plasma clouds. Yet, upon mapping the contours of the gravitational lensing, the gravitational forces are shown to not coincide with the plasma, but with the matter that did not interact in the collision of the clusters. With the luminous matter being too small to account for this effect, non-interacting dark matter, which constitutes the majority of the mass in the cluster, is needed to explain the observations of the Bullet Cluster.




FIGURE 2.3. Images of the Bullet Cluster 1E 0657-558 are shown for both (LEFT) visible light and (RIGHT) x-rays. The green lines in both images show the gravitational lensing contours. The white bar indicates a distance of 200 kpc at the location of the cluster. [4]

Cosmic Microwave Background:

The Cosmic Microwave Background (CMB) provides a snapshot of the early universe, which shows that dark matter had a large impact on the structure of the universe. Just a fraction of a second after the Big Bang, an isotropic plasma forms where protons, electrons, and photons are sloshing around under extreme

temperatures and pressures that do not allow the particles to bind together. At the same time, the gravitational wells from dark matter increase the local density of the plasma. Upon falling into the gravity wells, the plasma experiences a restoring force from the photon pressure, causing the plasma fluid to oscillate with sound waves at all length scales. As time goes on, the universe is expanding due to Dark Energy and the plasma slowly cools. After $\sim 400,000$ years, the plasma cools down to a temperature of ~ 3000 K, allowing the protons to bind with electrons and create electrically neutral hydrogen and helium atoms. At this instant of recombination, the universe became transparent to photons, which have since been able to travel through space largely undisturbed for billions of years. The photon frequencies emitted from the CMB follow a black-body spectrum, where the intensity of the light is directly related to the temperature of the plasma upon recombination. Thus, detecting these CMB photons from all directions provides us with a map of the plasma temperature distributions in the early universe, as seen in Fig. 2.4 through the Planck satellite [5].




FIGURE 2.4. Map of the temperature differences in the Cosmic Microwave Background as measured by the Planck satellite [5].

Measuring the size and relative frequencies of the temperature anisotropies in the CMB tells us the scales of structure formation in the early universe, which is directly dependent upon the energy densities of baryonic matter, dark matter, and dark energy. Thus, the CMB can tell us a lot about the universe; but with respect to dark matter, the total mass of the universe has to be composed of $\sim 85\%$ dark matter in order to explain the structure formations observed in the CMB [27] [7].

2.2.2 Dark Matter Models

Many theoretical models have been proposed to try and explain dark matter. Today, the leading paradigm of dark matter is that of BSM particles, where there is a plethora of single particle and dark sector theories that can explain all of the cosmological dark matter observations discussed in Sec. 2.2.1.

The only requirements a BSM particle model needs in order to be a candidate for dark matter is that it:

- has a lifetime that is longer than the age of the universe,
- has a non-zero mass such that it interacts gravitationally,
- is non-relativistic (cold) in order to congregate and form large structures,
- and does not significantly interact electromagnetically so as to not absorb or emit light.

Dark matter may be composed of a single particle or of more than one kind of particle, and the dark matter particles may interact only through gravity or also through additional forces. Dark matter therefore motivates a rich variety of ideas for BSM particles.

The allowed dark matter models span a vast parameter space, where similar models can be largely grouped by their dark matter mass range, as shown in Fig. 2.5. For an in-depth review of these different categories of candidate dark matter models, see Ref [6].

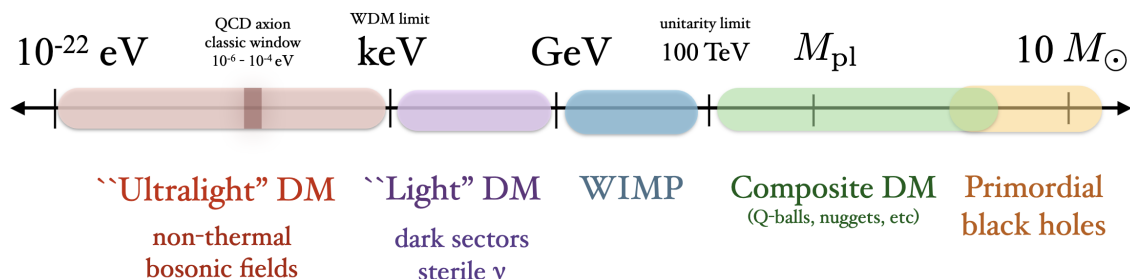


FIGURE 2.5. The mass range of allowed dark matter candidates is shown. The diagram is not to scale, and the mass ranges are only approximate. Figure is from Ref [6].

With so many possible dark matter models, one has to make assumptions of the properties of dark matter in order to justify looking for one particular model over another. One such common assumption is that the abundance of dark matter was set upon freezing out of the thermal bath of the early universe, which is the same mechanism that determined the abundance of many SM particles. Constraints can be set on the mass range and self-annihilation cross sections of dark matter particles needed for a freeze-out mechanism to produce both the large structure size and abundance of dark matter seen today.

Historically, the category of BSM particles that received the most attention as a candidate for dark matter was that of the Weakly Interacting Massive Particles (WIMPs). The popularity of the WIMP is attributed to the coincidence that the freeze-out constrained mass range and self-annihilation cross section of dark matter are at the electroweak-scale [28], which is the same scale that many other theoretically

motivated BSM particles exist. For example, the theory of Supersymmetry (SUSY) postulates that for every SM particle, there is a symmetric BSM partner particle that only differs by having opposite fermi/boson spin-statistics. A big motivation for SUSY particles is that they provide a solution to the hierarchy problem, where instead of needing fine-tuning, the quadratically divergent corrections to the Higgs Boson mass, that come from quantum loop diagrams of virtual particles, can be cancelled out by the quantum loop corrections from the SUSY partner particles [29]. The lightest SUSY particle, the neutralino, turns out to be at the electroweak mass scale, electrically neutral, stable, and weakly interacting; making it a prime WIMP candidate for dark matter. The fact that well motivated theoretical particles, like the neutralino, also fit the characteristics of the WIMP was dubbed the “WIMP Miracle” and caused a flurry of experimental efforts. From indirect detection in particle colliders or astrophysical systems, to direct detection via scattering with atomic nuclei in a detector, scientists have been searching for WIMP signatures for decades. Yet, no signatures have been found [7]. For example, in direct-detection experiments, a large swath of possible WIMP masses and WIMP-nucleon interaction cross sections have been ruled out, as shown in Fig. 2.6. Where, even with more sensitive future detectors, only 1 to 2 orders of magnitude in the WIMP-nucleon cross section can be probed before the detectors are sensitive to an irreducible background of neutrino scattering events that will mask the WIMP signal. Although WIMPs are not entirely ruled out, they have seen decades of thorough experimental efforts produce null results, and thus more attention is now being focused elsewhere, such as on dark sector models.

Dark sector models with light dark matter, whose mass is in the keV to GeV range, can also satisfy the freeze-out constrained mass range and self-annihilation




FIGURE 2.6. The observed upper limits of the spin-independent WIMP-nucleon cross section are shown in the WIMP phase-space for several different direct-detection experiments. The blue shaded region depicts the WIMP phase space that cannot be easily probed as there would be a significant background of irreducible neutrino scattering events. Figure is from Ref [7].

cross section of dark matter when there is an additional BSM particle that mediates the interactions between dark matter and the SM [6]. Instead of assuming that there is a single BSM particle that is responsible for dark matter, dark sector models expand the scope of the search by allowing for any number of dark particles and force carriers to exist. Thus, even if the dark matter particle is truly invisible and does not directly interact with the SM in any way, other dark sector particles could. If these detectable dark sector particles exist, they would provide a portal between the SM and dark sector, allowing for indirect insight into the particle nature of dark matter. One of the leading candidates for a dark sector portal particle is that of the dark photon, which will be described in the following section.

2.3 DARK PHOTON

Just as the SM photon is ubiquitous in our universe, one might suspect that a dark sector has its own photon; where the dark photon is the communicator of a dark charge instead of the SM electric charge.

The theory behind the dark photon is discussed in Sec. 2.3.1, followed by a depiction of the relevant dark photon production and decay channels for this analysis in Sec. 2.3.2, and finally a brief overview of the constrained parameter space of the dark photon is discussed in Sec. 2.3.3.

2.3.1 Basic Theory

The dark photon extension to the SM is accomplished by adding a $U(1)_D$ gauge field to the SM $SU(3)_C \times SU(2)_L \times U(1)_Y$ fields discussed in Sec. 2.1.2. The resulting renormalisable kinetic mixing interaction between the SM and the $U(1)_D$ field takes the form $F^{\mu\nu} F_{\mu\nu}^D$, where $F_{\mu\nu}$ and $F_{\mu\nu}^D$ are the field strength tensors of the SM $U(1)_Y$ field and the dark sector $U(1)_D$ field, respectively. As a result of this interaction, the dark gauge boson mixes with the SM gauge boson, leading to a new particle, the dark photon A' [30].

To further portray the theory behind the dark photon, let us begin with the full kinetic contribution to the Lagrangian from the two $U(1)_Y$ and $U(1)_D$ gauge bosons:

$$\mathcal{L}_{kin} = -\frac{1}{4}F_Y^{\mu\nu}F_{Y\mu\nu} - \frac{1}{4}F_D^{\mu\nu}F_{D\mu\nu} - \frac{\epsilon}{2}F_Y^{\mu\nu}F_{D\mu\nu} \quad (2.6)$$

Where ϵ is defined as the kinetic mixing coupling constant. The field strength tensors can be expanded in terms of the gauge boson field A_μ as $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$, where ∂_μ denotes the partial derivative with respect to the spacetime coordinate x^μ .

In order to have two independent fields that do not mix, one can rotate the A_μ^Y and A_μ^D gauge fields with the following general redefinition [31]:

$$\begin{pmatrix} A_\mu^D \\ A_\mu^Y \end{pmatrix} = \begin{pmatrix} \frac{1}{\sqrt{1-\epsilon^2}} & 0 \\ \frac{-\epsilon}{\sqrt{1-\epsilon^2}} & 1 \end{pmatrix} \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} A'_\mu \\ A_\mu \end{pmatrix} \quad (2.7)$$

such that Eq. (2.6) is now diagonalized in terms of the A'_μ and A_μ boson fields, which represent the dark photon and SM photon (after electroweak symmetry breaking) respectively. The orthogonal rotation defined by θ is determined upon requiring that the SM photon is massless and the dark photon is massive [31]; in which case $\theta = 0$, and Eq. (2.7) becomes:

$$\begin{pmatrix} A_\mu^D \\ A_\mu^Y \end{pmatrix} = \begin{pmatrix} \frac{1}{\sqrt{1-\epsilon^2}} & 0 \\ \frac{-\epsilon}{\sqrt{1-\epsilon^2}} & 1 \end{pmatrix} \begin{pmatrix} A'_\mu \\ A_\mu \end{pmatrix} \quad (2.8)$$

The dark photon's mass can be generated either through spontaneous symmetry breaking with a dark Higgs or through the Stueckelberg mechanism [32], where the Stueckelberg mechanism can be described as an extreme limit of the Higgs mechanism with the Higgs mass going to infinity and decoupling from the action while the vacuum expectation value remains finite. Although the dark photon can be either massless or massive, only the massive model is discussed here, as it is the focus of this dissertation.

Now, assuming that the original A_μ^Y and A_μ^D gauge fields only couple to the SM particles or to the dark sector particles, respectively, the interactions can be written as the following:

$$\mathcal{L}_{int} = eJ^\mu A_\mu^Y + e'J'^\mu A_\mu^D \quad (2.9)$$

Where the four-current $J^\mu = \sum_f \bar{f} \gamma^\mu f$ defines the interactions with all SM fermions (f) of charge e, and similarly J'^μ defines the interactions with all dark fermions of dark charge e'. Substituting the gauge field redefinitions (Eq. (2.8)) into Eq. (2.9), the interaction Lagrangian becomes:

$$\mathcal{L}_{int} = \left(\frac{-e\epsilon}{\sqrt{1-\epsilon^2}} J^\mu + \frac{e'}{\sqrt{1-\epsilon^2}} J'^\mu \right) A'_\mu + e J^\mu A_\mu \quad (2.10)$$

It is now easy to see that the SM photon A_μ couples only to the SM current J^μ , whereas the dark photon A'_μ interacts with both the SM and dark currents. The SM-to-A' portal term: $\frac{-e\epsilon}{\sqrt{1-\epsilon^2}} J^\mu A'_\mu$ can be further simplified in the limit that the kinetic mixing coupling is small ($\epsilon \ll 1$), becoming $-e\epsilon J^\mu A'_\mu$.

Finally, including the kinetic energy of the massive dark photon, with mass $m_{A'}$, the interactions between the SM and the dark photon are defined through the Lagrangian terms:

$$\mathcal{L} \supset \frac{1}{2} m_{A'} A'^2 - \epsilon e \sum_f q_f A'^\mu \bar{f} \gamma_\mu f, \quad (2.11)$$

where the sum is over all SM fermions f with SM electric charge q_f . Thus, the dark photon can interact with the SM both gravitationally and through the kinetic mixing with the SM photon to SM fermion and anti-fermion pairs. Although, the kinetic mixing must be weak, with coupling $\epsilon \ll 1$, or we would have already detected the dark photon's affects on the SM by now.

Being a weakly interacting massive particle, the dark photon could be a candidate for dark matter itself, but direct detection experiments have already ruled out the thermally generated dark photon parameter space that could account for the abundance of dark matter [33].

On the other hand, the dark photon could still provide a portal to the dark sector and may couple to the dark matter particle χ . In this analysis, it is assumed that $2m_e < m_{A'} < 2m_\chi$ such that the dark photon decays visibly to SM particles. Thermal freeze-out is then determined by the processes $\chi\chi \leftrightarrow A' \leftrightarrow f\bar{f}$. For light masses $m_{A'} \sim \text{MeV} - \text{GeV}$ and loop-induced or otherwise suppressed couplings $\epsilon \sim 10^{-6} - 10^{-3}$, the dark matter particle's thermal relic density is in the right range to be a significant fraction of cosmological dark matter [34, 35, 36]. These values of $m_{A'}$ and ϵ are therefore cosmologically favoured and provide a well-defined thermal relic target in the dark photon parameter space for experimental searches.

This study presents the results of a search for dark photons in this cosmologically favoured parameter space using the FASER detector at the LHC, both of which will be discussed in Chapter III.

2.3.2 Production and Decay

Kinetic mixing between the SM photon and dark photon potentially allows for numerous dark photon production and decay channels, where $m_{A'}$ governs which channels are kinematically-allowed. For production, a dark photon can come from any off-shell SM photon whose energy is greater than $m_{A'}$. Similarly, a dark photon can decay via an off-shell SM photon to any fermion anti-fermion pairs whose total mass is less than $m_{A'}$. Although there can be many channels for the dark photon to interact with the SM, they rarely (if ever) occur as both its production and decay are proportional to ϵ^2 , where $\epsilon \ll 1$.

Only the relevant dark photon production and decay channels for this analysis are discussed further.

At the LHC, with only light and highly boosted dark photons produced in the far forward direction of proton-proton collisions reaching FASER, this analysis is only sensitive to dark photons with masses $\sim 10 - 100$ MeV, as described in Sec. 3.2.

For FASER, the dominant sources of dark photon production are SM neutral meson decay and dark bremsstrahlung:

- Neutral pion decay $\pi^0 \rightarrow A'\gamma$: This mode is accessible for $m_{A'} < m_{\pi^0} \simeq 135$ MeV. The branching fraction is $B(\pi^0 \rightarrow A'\gamma) = 2\epsilon^2(1 - m_{A'}^2/m_{\pi^0}^2)^3 B(\pi^0 \rightarrow \gamma\gamma)$ where $B(\pi^0 \rightarrow \gamma\gamma) \simeq 0.99$ [37].
- Eta meson decay $\eta \rightarrow A'\gamma$: This mode is open for $m_{A'} < m_\eta \simeq 548$ MeV. The branching fraction is $B(\eta \rightarrow A'\gamma) = 2\epsilon^2(1 - m_{A'}^2/m_\eta^2)^3 B(\eta \rightarrow \gamma\gamma)$ where $B(\eta \rightarrow \gamma\gamma) \simeq 0.39$ [37].
- Dark bremsstrahlung $pp \rightarrow ppA'$: In this process, a dark photon is emitted via initial or final state radiation from colliding protons in a coherent way. This mode is open for dark photon masses up to $\mathcal{O}(2 \text{ GeV})$ [38].

These processes are depicted in Fig. 2.7. Neutral pion decay is typically the leading signal contribution, but η decay can be comparable for $m_{A'} \sim 100$ MeV, and dark bremsstrahlung can be comparable near the boundary of FASER's sensitivity [38]. Other production mechanisms include the decays of heavier mesons (such as η' or ω) and direct Drell-Yan production $q\bar{q} \rightarrow A'$, but these are subdominant and are neglected.

As for the SM decay of the dark photon, there are possibly numerous channels available, as seen in Fig. 2.8. Yet, because FASER is only sensitive to $m_{A'} \sim 10 - 100$ MeV, which is $2m_e < m_{A'} < 2m_\mu \simeq 211$ MeV, then dark photons decay to electron-positron pairs with $B(A' \rightarrow e^+e^-) = 100\%$.




FIGURE 2.7. The Feynman diagrams of the dark photon production channels relevant to this analysis are shown, with light neutral meson decays depicted on the LEFT and dark bremsstrahlung radiation shown on the RIGHT.




FIGURE 2.8. The dark photon decay branching ratios are shown for varying dark photon mass. Plot taken from Ref [8].

2.3.3 Previously Constrained Parameter Space

Numerous experiments have already searched for, and did not find, a massive dark photon, resulting in regions of the dark photon parameter space ($m_{A'}$, ϵ) being ruled out. The constrained dark photon parameter space is depicted in Fig. 2.9, which also shows which experiments utilized similar A' production mechanisms. Even more

generally, the previous experiments can be largely grouped into three categories via their similarities in constraining the dark photon parameter space:

1. the electron magnetic dipole moment measurement,
2. experiments searching for prompt or invisible dark photon decays,
3. and experiments searching for visible decays of long-lived dark photons.

The remainder of this section will only broadly discuss these three categories of experiments that constrain the dark photon parameter space; for references to all of the experiments depicted in Fig. 2.9, and for a brief summary of the dark photon production and decay channels that each experiment used, see Appendix C of Ref [9].




FIGURE 2.9. The dark photon parameter space is depicted with shaded regions showing the dark photon mass ($m_{A'}$) and kinetic mixing constant (ϵ) that have been excluded from past experiments. The colors of the shaded regions refer to experiments that utilized a similar A' production mechanism where experiments shaded in RED used electron beam dumps, CYAN used proton beam dumps, GREEN used $e+e-$ colliders, BLUE used pp collisions, MAGENTA used meson decays, and YELLOW used electrons on a fixed target. The GREY constraint comes from the precise measurement of the electron magnetic dipole moment. Figure is from Ref [9].

One way to probe the parameter space of the dark photon is via the measurement of the electron magnetic dipole moment. The existence of the dark photon would affect the electron magnetic dipole moment through a one-loop correction to the electron-photon vertex, as shown in Fig. 2.10, whose affect scales with $\frac{e^2}{m_{A'}^2}$ [39]. Yet, because the measurement of the electron magnetic dipole moment agrees so well with the SM [23], constraints on the dark photon parameter space can be made [39] that result in the exclusion of the grey region in Fig. 2.9.




FIGURE 2.10. The dark photon one-loop correction to the electron-photon vertex is depicted.

Another way to probe the parameter space of the dark photon is to perform bump hunts in the invariant mass spectra (or missing mass spectra) from prompt visible (or invisible) dark photon decays, such that the signal can be distinguished apart from numerous SM backgrounds that are also present at the production point. Prompt decay searches have mostly been performed with experiments that are sensitive to detecting particles coming directly from e^+e^- collisions, proton-proton collisions, meson decays, and electron-nucleon interactions; which constrain the green, blue, magenta, and yellow shaded regions of the dark photon parameter space in Fig. 2.9, respectively. For constructing the invariant mass spectra from prompt visible dark photon decays, the reconstruction of the decay products from $A' \rightarrow e^+e^-$, $\mu^+\mu^-$, and $\pi^+\pi^-$ have been used. Whereas, to construct the missing mass spectra, a process

such as $e^+e^- \rightarrow A'\gamma$ and $A' \rightarrow \textit{invisible}$ is used, where the initial state is known and the visible final state particles can all be detected. The sensitivity of these prompt decay searches to the mass of the dark photon is largely constrained via the signal reconstruction efficiency, background intensity, and/or kinematic suppression of A' production; which all strongly depend upon $m_{A'}$, but not ϵ . Thus, the constrained dark photon parameter space from a prompt decay search is typically rectangular in shape and reaches towards decreasing ϵ , where the rate of dark photon production and decay both scale with ϵ^2 .

The experiments searching for visible decays from long-lived dark photons have thus far mostly searched for the $A' \rightarrow e^+e^-$ signal and utilized beam-dump/fixed-target production of dark photons, which constrain the red and cyan regions in Fig. 2.9. These experiments require the dark photon to decay within a finite volume that is some distance (L) away from the point of production. Assuming that $E_{A'} \gg m_{A'} \gg m_e$, then the decay length (d) of a dark photon is proportional to $\frac{E_{A'}}{\epsilon^2 m_{A'}^2}$ [14]. Thus, the dark photon parameters that these experiments are sensitive to is largely determined by d being similar to L such that a significant portion of the dark photons decay inside the sensitive region of the detector. For example, as you increase both ϵ and $m_{A'}$ such that $d \ll L$, then the vast majority of dark photons decay before they reach the detector; similarly, as you decrease both ϵ and $m_{A'}$ such that $d \gg L$, then the vast majority of dark photons decay after the detector. This sensitivity to the dark photon decay length, combined with both the production and decay of dark photons scaling with ϵ^2 , largely describes the shape of the constrained parameter space from these long-lived dark photon searches.

There are also experiments that utilize proton-proton colliders, instead of beam-dumps/fixed-targets, to generate and search for visible decays from long-lived dark

photons. For example, LHCb found sensitivity to long-lived dark photons using a displaced-vertex signature from $A' \rightarrow \mu^+\mu^-$ and excluded a small region [40] around $\epsilon \sim 10^{-5}$ and $m_{A'} \sim 250$ MeV, as seen in Fig. 2.9. LHCb's sensitivity for long-lived dark photons is expected to greatly expand in the future with improved statistics and the inclusion of a displaced-vertex signature of $A' \rightarrow e^+e^-$ from rare charm meson decays [41]: $D^{*0} \rightarrow D^0 A'$. FASER also utilizes proton-proton collisions at the LHC to search for long-lived dark photons, but does so in a similar way to the beam-dump/fixed-target experiments. The expected dark photon sensitivity of FASER will be discussed in detail in the following chapter.

CHAPTER III

THE LHC AND DARK PHOTONS AT FASER

The LHC is a circular particle accelerator with a circumference of 27 km, that collides protons together head-on at 99.9999991% the speed of light [42]. With center-of-mass energies of 13.6 TeV, the interactions between the two colliding protons are energetic enough to produce even the heaviest fundamental SM particles. The quarks and gluons that make up the proton are what are fundamentally interacting in these collisions; and since the quarks interact with all force-carrying particles, as discussed in Sec. 2.1, then every fundamental SM particle can be produced in the collisions at the LHC. Thus, if BSM particles, like the dark photon, interact with the SM and are not too heavy, they will be produced in the LHC collisions as well. Discovering dark photons at the LHC is then only a question of how rarely are they produced and can they be detected amongst a plethora of SM particle signatures. FASER is a new experiment at the LHC that looks for dark photons travelling in the far-forward direction of LHC pp collisions; where FASER's distant location provides a relatively quiet environment and sensitivity to long-lived dark photons with masses and couplings that were previously unconstrained.

In this chapter, first the LHC is described in more detail in Sec. 3.1, and then the potential for detecting dark photons with the FASER experiment is presented in Sec. 3.2.

3.1 THE LARGE HADRON COLLIDER

The LHC, located 100 m underground at CERN and straddling the border of Switzerland and France, is the world's largest and most powerful particle collider.

Having already gone through two prior multi-year periods of operation, the LHC is currently in its third operational period (Run-3), where the protons are accelerated to center-of-mass energies of $\sqrt{s} = 13.6$ TeV before colliding. Prior to being accelerated to these energies in the LHC, the protons first pass through a series of other systems in the CERN accelerator complex, depicted in Fig. 3.1.

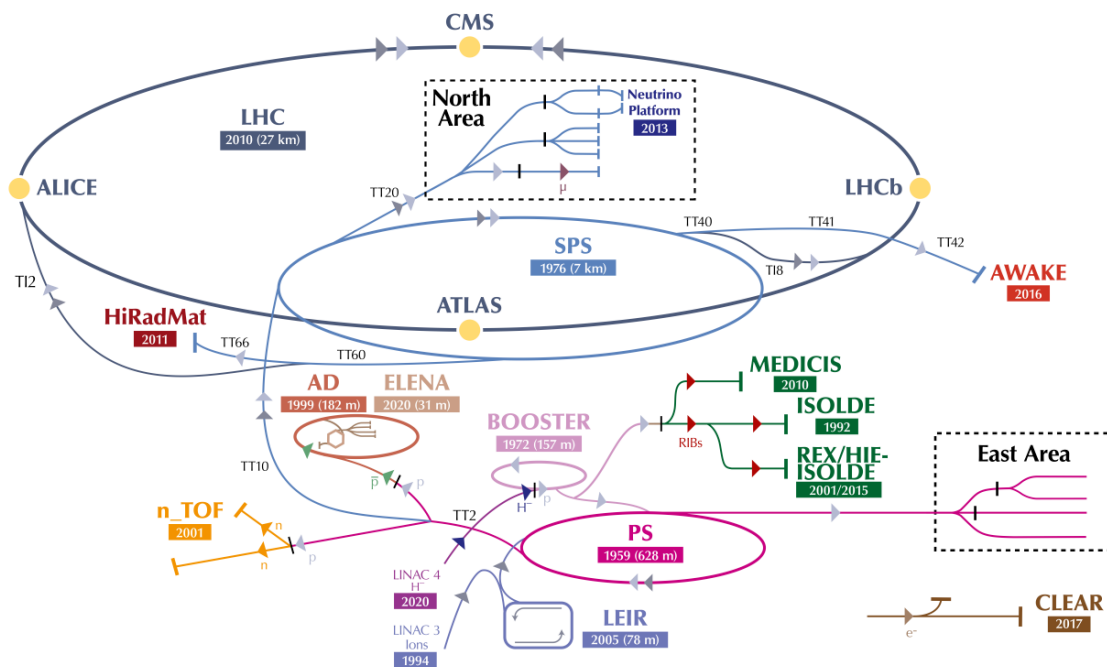


FIGURE 3.1. The CERN accelerator complex is shown, where both the accelerator systems and detector locations are depicted. FASER is not shown in this diagram, as it is a new and relatively small experiment. Figure is from Ref [10].

At the start of the LHC injection chain, H_2 molecules from a bottle of Hydrogen gas pass through a plasma and strong electric field that strips the electrons away and produces H^+ (protons), which are then made to interact with a metal cathode to pick up two electrons and become H^- ions [43]. These ions are then pulsed into the first accelerator stage, a linear accelerator called LINAC-4, where they reach energies of 160 MeV. The energetic H^- ions are then shot through a thin carbon foil, where the electrons are stripped away to leave just protons. From there, the

protons are accelerated to 2 GeV in the Proton Synchrotron Booster (PSB), then to energies ranging from 14 to 26 GeV in the Proton Synchrotron (PS), and finally up to energies of 450 GeV in the Super Proton Synchrotron (SPS) before entering the LHC as two beams travelling in opposite directions around the 27 km ring [44].

At all accelerator stages up to and including the LHC, radio frequency (RF) cavities and superconducting magnets produce electromagnetic fields that accelerate and guide the beams of protons, whose reactions to the fields are governed by the relativistic Lorentz force:

$$\vec{F} = \frac{1}{\sqrt{1 - \frac{v^2}{c^2}}} q(\vec{E} + \vec{v} \times \vec{B}) \quad (3.1)$$

where q and \vec{v} are the charge and velocity of the proton respectively, \vec{E} is the electric field, and \vec{B} is the magnetic field. As you can see, electric fields can apply a force aligned with the protons velocity, whereas magnetic fields can only apply a force perpendicular with the velocity. Thus, strong electric fields generated in the RF cavities are used to speed up the protons, whereas the magnets are used to steer and focus them.

The basic design concept of an RF cavity is that of a metallic chamber that has a strong resonant electric field oscillating inside of it. This oscillating electric field can be thought of as positive and negative charges congregating on opposite sides of the RF cavity, which then swap sides in a sinusoidal fashion. When timed right, the RF cavity will have a negative charge on it's upstream side as a proton approaches it, such that the proton feels an attractive force downstream. Then, as the proton enters the RF cavity, the charges in the RF cavity swap sides such that the proton feels a positive charge pushing it from behind and a negative charge

pulling it from the front. Finally, when the proton leaves the RF cavity, the charges must oscillate again such that there is a positive charge on the downstream side of the cavity, pushing the proton from behind. The result of a carefully timed string of RF cavities is an acceleration of the proton beam while also congregating the protons into bunches that are spaced out in time by the inverse of the RF cavity oscillating frequency. The RF cavities at the LHC operate at 400 MHz such that there are 2.5 ns long buckets where proton bunches could exist. Yet, by convention a maximum of one out of ten buckets are ever filled, resulting in a bunch spacing of 25 ns and a maximum of 3564 bunches per beam, but not all bunches are filled and only 2808 filled bunches circulate the LHC. There are 16 RF cavities total in the LHC, 8 for each beam, which have their oscillating electric fields driven by high-power klystrons [45] and are cooled to superconducting temperatures that result in a maximum voltage of 2 MV per RF cavity. Simplistically, this would mean that 16 MeV of energy is imparted into the protons for each lap they make around the LHC, and it would only take ~ 35 seconds to accelerate the LHC beams from their injection energies of 450 GeV to their desired energies of 6.8 TeV, as the protons complete 11,245 laps around the LHC per second. Yet in reality, and in part due to the protons not always experiencing the maximum voltage of the RF cavities, it takes about 20 minutes to actually ramp the LHC beams to full power [46].

Superconducting dipole electromagnets are used to incrementally bend the beams around the LHC. The 1,232 dipole magnets in the LHC are each 15 meters long and account for the majority of the length of the LHC. Each LHC dipole magnet actually consists of two nearby dipole magnetic fields, one for each beam, which are both housed in a single structure as shown in Fig. 3.2. Superconducting coils are placed on both sides of the beam pipe and produce a uniform magnetic field of 8.3 T.




FIGURE 3.2. A cross section of an LHC dipole is shown. Figure is from Ref [11].

The last main component of the LHC is that of the superconducting quadrupole electromagnets, which focus the protons into tighter bunches. There are 392 quadrupole magnets, each 5–7 meters long, in the LHC. With a design very similar to the LHC dipole magnet, the quadrupole utilizes superconducting coils positioned around each quadrant of the beam-pipe such that the direction of current through the coils can be configured to produce two north poles on opposing corners and two south poles on the other opposing corners. The two quadrupole configurations and their effects on a proton passing through them are shown in Fig. 3.3. As you can see, a proton bunch will get squeezed either horizontally or vertically by a quadrupole magnet. Thus, having two successive quadrupole magnets with different configurations, allows for a proton bunch to get squeezed in both directions. This is important for not only routinely fighting the repulsive forces between the protons in

a bunch, but also for increasing the number of proton-proton collisions by focusing the transverse size of the bunch down by a factor of 10 right before the interaction point.




FIGURE 3.3. A diagram showing two different orientations of quadrupole magnets, the resulting magnetic field lines (black arrows), and the direction of the force (red arrows) applied to a positively charged particle passing through that area of the quadrupole. Figure is from Ref [12].

Measuring the integrated luminosity (L), which represents the total number of proton-proton collisions at an interaction point, is crucial for every analysis as it relates the number of events produced (N) with the cross section (σ) of a given interaction, where $N = \sigma L$. The cross section of a given particle interaction is a measure of how likely the interaction is to occur and can be calculated from the quantum field theory. The units of the cross section are that of area and are often expressed in barns (b), where $1 b = 10^{-28} m^2$, and thus L is expressed as inverse-barns such that N is unit-less. The instantaneous luminosity per colliding bunch

(L_b) can be calculated from the beam parameters:

$$L_b = \frac{fn_1n_2}{2\pi\Sigma_x\Sigma_y}, \quad (3.2)$$

where f is the frequency that the bunches collide, n_1 and n_2 are the numbers of protons in the two colliding bunches, and Σ_x and Σ_y are the convolved beam sizes in the horizontal and vertical directions. These beam parameters can be measured by scanning the colliding beams across one another in both the transverse x and y directions, which is performed at low luminosities in a van der Meer (vdM) scan and provide the absolute luminosity scale. Measuring the activity of numerous detectors, each with different and complementary sources of systematic uncertainties, is necessary to accurately extrapolate the absolute luminosity scale from the low luminosity conditions of the vdM scans to the high-luminosities seen in regular physics data-taking. At the LHC interaction point 1 (IP1), where FASER gets its proton-proton collisions from, sub-detectors in ATLAS are used to measure the luminosity and the methodology is discussed in Ref. [47].

There are nine experiments studying the particles emitted in the LHC collisions. The four main experiments (ATLAS, CMS, LHCb, and ALICE) each have their own designated interaction point in the LHC, while the five smaller experiments (TOTEM, LHCf, MoEDAL-MAPP, SND@LHC, and FASER) are positioned nearby the main experiments in order to study particles that the main experiments are not sensitive to.

3.2 DARK PHOTONS AT FASER

FASER, which first saw data in 2022, is one of the newest experiments at the LHC and is designed to search for light, weakly-interacting, long-lived particles, including dark photons [14, 38, 48].

The FASER detector is located in the TI12 tunnel, which served as a transfer line connecting the Super Proton Synchrotron (SPS) to the Large Electron-Positron (LEP) collider but has since been decommissioned, approximately 480 m from the ATLAS interaction point (IP1) along the beam collision axis line-of-sight (LOS), as shown in Fig. 3.4. However, due to the crossing angle in IP1, the LOS is offset vertically by 6.5 cm with respect to the centre of the detector, which is properly accounted for in the simulation. In order for particles coming from the IP to get to FASER, they will have to invisibly pass through:

- the “D1” dipole magnet, which bends the LHC beam and is strong enough to deflect even the most energetic of charged particles coming from the IP,
- the TAN absorber [49], which is ~ 30 Tons of copper, steel, and marble that is designed to absorb the flux of forward neutral particles that could otherwise interact in downstream superconducting magnets and cause them to quench,
- and ~ 100 m of rock and concrete that is between the main LHC beam-line tunnel and the TI12 tunnel.

Thus, only weakly-interacting, long-lived particles travelling in the far-forward direction will make it to FASER.

If dark photons are light and weakly (or feebly) interacting, they are long-lived and can be produced in large numbers in the proton-proton collisions at the

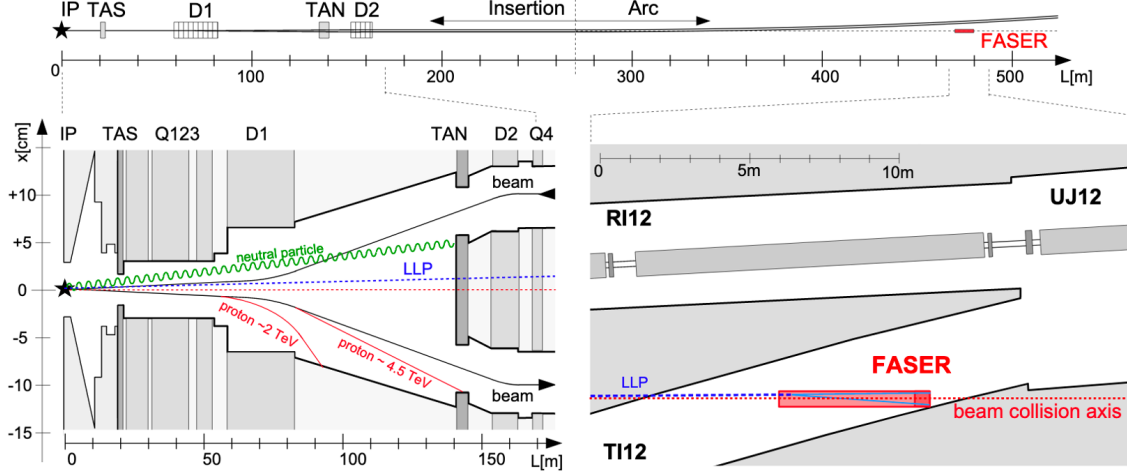


FIGURE 3.4. A schematic showing the location of FASER with respect to the LHC beam and IP1. More detailed views of the LHC beam are shown in the BOTTOM LEFT for distances up to 170 m from the IP, and in the BOTTOM RIGHT for the area around FASER at ~ 480 m. The beam collision axis line-of-sight is shown as a dashed red line.

LHC [50]. As discussed in Sec. 2.3.2, the dominant source of dark photons in FASER is from neutral pion decay ($\pi^0 \rightarrow A'\gamma$). Using the EPOS-LHC [13] Monte Carlo event generator, the energy and angle distributions of π^0 production from the IP can be simulated, as shown in Fig. 3.5. As you can see, there is an immense amount of highly energetic neutral pions produced in the far-forward direction, where the forward $10^{-6}\%$ of solid angle sees $\sim 1\%$ of all pions with $E \geq 10$ GeV. With $300 fb^{-1}$ of data, $\sim 10^{16}$ pions are predicted to be produced in the far-forward direction, towards FASER [38]. The decay products of these far-forward, \sim TeV energy, pions will also be highly collimated, as the angle between the decay products can be approximated by $\theta \approx p_T/E \approx m_{\pi^0}/\text{TeV} \approx 10^{-4}$ radians and thus will only spread out ~ 1 cm after 100 m of travel. This process produces a high-intensity and collimated beam of dark photons in the far-forward direction along the beamline. Thus, with a dark photon kinetic mixing of $\epsilon \sim 10^{-4}$, there could be as many as 10^8 dark photons produced in the direction of FASER [38], assuming $300 fb^{-1}$ of data.




FIGURE 3.5. The production rates of π^0 from the LHC IP, obtained via EPOS-LHC [13], are shown as a function of the meson's angle with respect to the beam axis (θ_π) and momentum (p_π). The angular acceptances for FASER and FASER2 (a proposed future detector) are indicated by the vertical gray dashed lines. Figure is from Ref [14]

Once produced, dark photons then may invisibly travel a macroscopic distance before decaying, leading to a striking signal of high-energy e^+e^- particles far from the IP.

Apart from the rare neutrino event, the only other SM LLP expected to reach FASER is the muon, which must be generated from particle interactions in material after the D1 magnet, or else the muon would have been swept away by the dipole magnet. Measurements of the muon flux at a location similar to FASER, but on the opposite side of the IP, were performed in Run-2 of the LHC with a pilot emulsion

detector, resulting in $\sim 1.9 \times 10^4$ muons per $fb^{-1}cm^2$ [51] and agreed with estimates from detailed FLUKA [52, 53] simulations. With $27 fb^{-1}$ of data in this analysis, and the sensitive area in FASER being about $\sim 300 cm^2$ large, we roughly expect $\sim 1.5 \times 10^8$ muon events traversing FASER. Although a large flux of muons is expected in FASER, they are not expected to mimic the striking dark photon signal. FASER is therefore well suited to search for light and weakly-interacting dark photons in a very low background environment.

FASER's dark photon sensitivity is largely determined by its location, where if the dark photon decay length is too short or too far compared to the distance to FASER, then the ratio of dark photons decaying in the FASER volume will be too small to detect. The probability that a dark photon produced at the IP with momentum p and angle θ , with respect to the beam axis, will decay within the FASER detector is

$$P(p, \theta) = (e^{-(L-\Delta)/d} - e^{-L/d})\Theta(R - L \times \tan\theta), \quad (3.3)$$

where d is the exponential decay length of the dark photon, Θ is the Heaviside step function, and the decay volume in FASER (as described in Sec. 4.1.1) is cylindrical with radius $R = 10$ cm, a length $\Delta = 1.5$ m, and is a distance of $L \approx 480$ m from the IP. For $E_{A'} \gg m_{A'} \gg m_e$, the decay length of a dark photon with lifetime τ travelling at speed $\beta = v/c$ is [14]

$$d = c\beta\tau\gamma \approx (80 \text{ m}) \left[\frac{10^{-5}}{\epsilon} \right]^2 \left[\frac{E_{A'}}{\text{TeV}} \right] \left[\frac{100 \text{ MeV}}{m_{A'}} \right]^2. \quad (3.4)$$

Thus, with the simulated dark photon p and θ distributions at the IP, the geometrical likelihood that a given dark photon model will decay inside FASER can be calculated, as shown in the top left plot of Fig. 3.6.

A rigorous estimate of FASER’s sensitivity in the dark photon parameter space can be calculated by combining the following:

- the simulated distribution and number of forward-going π^0 from the IP,
- the likelihoods of A' production and decay, which both scale as ϵ^2 and are discussed in Sec. 2.3.2,
- the geometrical likelihood of A' decaying in FASER,
- the likelihood of $E_{A'} > 500$ GeV (shown in the top right plot of Fig. 3.6),
 - * a cut applied to separate the dark photon signal apart from the SM LLPs (muons and neutrinos) that also make it to FASER
- and assume a 50% signal efficiency in addition to the 500 GeV cut.

The resulting numbers of expected A' events in FASER, scaled to a luminosity of 27.0 fb^{-1} , are shown across the A' phase space in the bottom plot of Fig. 3.6. As you can see, FASER can expect to be sensitive to detecting dark photon models in a region extending out to $\epsilon \sim 10^{-5}$ and $m_{A'} \sim 100$ MeV. In overlaying this A' sensitivity region of FASER with the regions previously excluded by past experiments, as depicted in Fig. 3.7, it is shown that FASER is expected to be sensitive to dark photons with masses and couplings that were previously unconstrained.

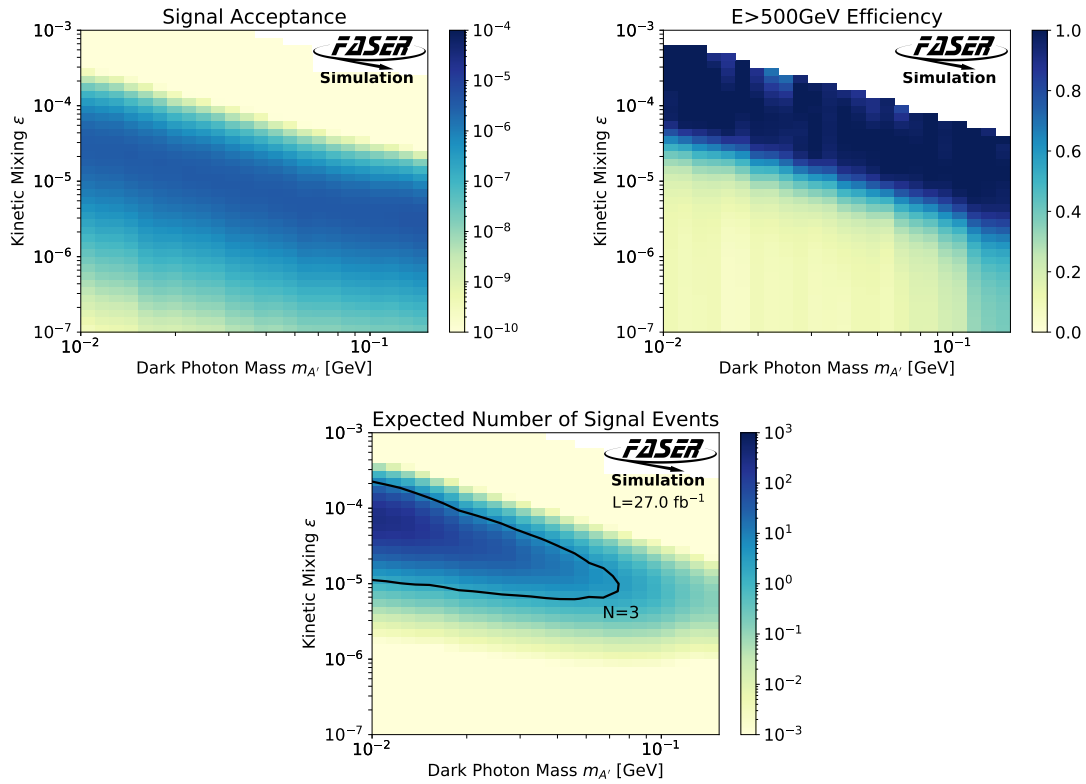


FIGURE 3.6. Top left: The acceptance for dark photons to decay inside the FASER decay volume. Top right: The fraction of dark photons decaying inside the FASER decay volume that have energy greater than 500 GeV. Bottom: The expected number of dark photon events in FASER for 27.0 fb $^{-1}$ of data, assuming a 50% signal efficiency, on top of the requirement that the A' energy is greater than 500 GeV.




FIGURE 3.7. FASER’s expected sensitivity, with luminosities ranging from 1 to 3000 fb^{-1} , is shown as yellow-shaded areas in the dark photon’s parameter space. The gray-shaded regions are excluded by past experiments, and the projected future sensitivities of other experiments are shown as colored contours.

CHAPTER IV

FASER DETECTOR

The FASER detector, described in detail in Ref. [48], and depicted in Fig. 4.1, was primarily designed to be sensitive to the dark photon signal. With dark photons from the ATLAS IP producing a very distinctive signal of a non-interacting particle suddenly decaying to two highly energetic ($\sim \text{TeV}$) e^+e^- particles, the main motivations behind FASER's design were to:

- Efficiently veto background events caused by incoming charged particles (primarily high-energy muons)
- Separate and identify the $\sim \text{TeV}$ e^+e^- tracks
- Measure the energy of the e^+e^- pair

There were also constraints due to the cost, minimizing the time to construct, and having limited space in the TI12 tunnel that all impacted the final design of the detector.

In addition to the primary FASER detector, there is also the FASER ν tungsten/emulsion detector just upstream of FASER. FASER ν is a passive stand-alone detector that is dedicated to neutrino measurements, and is not used in this analysis, but it does contain eight interaction lengths of tungsten that help suppress potential backgrounds.

In the remainder of this chapter, the detector components of interest in this analysis will be discussed further; including the tracking spectrometer, scintillators, electromagnetic calorimeter, and the data acquisition system.




FIGURE 4.1. A depiction of the FASER detector, where each color represents a different sub-detector system. The coordinate system used for FASER is also shown in the bottom left, where the particles coming from the ATLAS IP are travelling in the $+\hat{z}$ direction through FASER.

4.1 TRACKING SPECTROMETER

The tracking spectrometer [54] consists of three tracking stations interleaved with three 0.57 T dipole magnets. Each tracking station has three layers of Semi-Conductor Tracker (SCT) barrel modules from the ATLAS experiment [55], which have a hit position resolution of about $20 \mu\text{m}$ in the precision coordinate. The magnets bend charged tracks in the \hat{y} direction, corresponding to the precision coordinate of the tracker. A rigid piece of aluminum, shown in green in Fig. 4.1, is used to mount and roughly align the three tracking stations to one another. In order to obtain optimal tracking performance, an iterative local alignment procedure is performed that minimizes the χ^2 of reconstructed tracks. The remainder of this section will further describe the magnets, SCT modules, tracking stations, and tracking alignment procedure.

4.1.1 Magnets

The three magnets in FASER are identical except for their lengths, where the most upstream magnet is 1.5 m long and the other two are 1 m long. For this analysis, the central empty volume of the 1.5 m long magnet is where we require the dark photon to decay, and is thus referred to as the decay volume. The active transverse area of the detector is defined by the circular magnet aperture with a radius of 10 cm.

All three magnets are composed of Samarium-Cobalt (Sm_2Co_{17}) permanent magnetic blocks. The blocks are arranged in a Halbach array, as seen in Fig. 4.2, of 16 blocks per ring and 12 (18) rings for the 1 m (1.5 m) long magnets. The result is a uniform dipole magnetic field throughout the central empty volume of the cylinder, depicted in Fig. 4.3. To assemble and hold the magnetic blocks in place, aluminum guides were used. The outer frame of the magnets use construction steel grade S355JR, which provides stability during assembly as the magnetic blocks are attracted to the outer frame until the Halbach array is complete.

To measure the magnetic fields of the magnets, the CERN magnet team used both a single-stretched wire (SSW) and a 3D Hall probe [48]. The SSW method utilizes a conducting wire that is tight under tension and suspended through the length of the magnet, where translations of the wire in directions perpendicular to the magnetic field induce a current in the wire via the Lorentz force that is proportional to the perpendicular component of the magnetic field. Performing the SSW measurement in steps of x and y produces a 2D measurement of the integrated magnetic field along the wire, but fails to give any information on the field dependence of the z position. On the other hand, the 3D Hall probe measurement combines three separate Hall effect sensors that are all perpendicular to one another in order to measure to all




FIGURE 4.2. A cross section of a cylindrical dipole magnet in FASER. The arrows indicate the direction of the magnetic field in each of the permanent magnetic blocks. Note that the final orientation of the inner magnetic field is in the $-\hat{x}$ direction, which would correspond with the vertical direction in this figure.

components of the magnetic field at any position in 3D space. The results of the 3D Hall probe measurement produces a 3D magnetic field map that is used in the simulation, and whose field strengths at the center of one of the 1 m long magnets can be seen in Fig. 4.4. Both methods agree in their measurements of a homogeneous magnetic field of 0.57 T in the $-\hat{x}$ direction for all three magnets.

4.1.2 SCT Module

The ATLAS SCT barrel modules, used to construct the FASER tracker, and depicted in Fig. 4.5, consist of four identical single-sided silicon microstrip sensors. Each sensor is composed of a $285 \mu\text{m}$ thick n-type silicon substrate that has p-type strips implanted across its surface. With a rectangular geometry of $64 \times 63.6 \text{ mm}^2$, and a constant pitch of $80 \mu\text{m}$ between strips, each sensor has 768 readout strips. The four sensors in an SCT module are aligned in pairs along the length of the strips and bonded end-to-end to create ~ 12.8 cm long readout strips on both sides of a base




FIGURE 4.3. The FEMM [15] simulated 2D magnetic field distribution is shown in a quadrant of the dipole magnet. Note that for this figure, the \hat{x} direction is vertical and the \hat{y} direction is horizontal.

board. Thus, there are 1536 total readout channels in a module, with 768 on each side. There is a 40 mrad stereo angle between the front and back strips, providing an intrinsic $17 \mu\text{m}$ spatial resolution in the precision \hat{y} coordinate and an intrinsic $580 \mu\text{m}$ spatial resolution in \hat{x} .

A flex polyimide/copper hybrid readout with front-end ASICs (ABCD3TA chips [56]) is attached near the center on one side of the module. A module has 12 ABCD3TA chips, which readout 128 channels each, and provide a preamplifier, pulse-shaper, and discriminator for each channel. A nominal bias voltage of 150 V




FIGURE 4.4. The magnetic field components, measured with the Hall probe method, are shown against z while at the center of one of the 1 m long dipole magnets.




FIGURE 4.5. A photograph of a SCT barrel module inside an aluminum test-box.

and a common threshold of 1 fC is applied across all channels. In order to account for differences between channels and obtain a uniform activation response, a 4-bit threshold offset parameter is tuned for each channel.

4.1.3 Tracking Station

Eight SCT modules are installed in a 24×24 cm² transverse plane within an aluminum frame, as seen in Fig. 4.6. The modules are displaced in their placement in z , such that they can overlap one another and minimize the gaps in the transverse plane. Each module has a flexible printed circuit board connecting the ASICs to a patch-panel on either side of the tracker plane, which is used to interface with the power supplies and DAQ.

In order to keep the modules cool, such that the epoxy glue used for the module assembly does not melt, a 5 mm diameter water cooling channel is integrated into the frame of the tracker plane, which is thermally coupled to the baseboards of all eight modules.

A tracking station, shown in Fig. 4.7, is composed of three tracking planes that are stacked in the \hat{z} direction. Each plane in a station is staggered in the \hat{y} direction by ± 5 mm with respect to the first plane in order to cover tracking dead zones caused by the module overlaps not fully covering the 1 mm gaps between the sensor cut edge and the sensitive sensor region. The staggering of the planes ensures that there are at least two tracker modules hit for a track travelling through this region of the tracking station. One dead region that was not accounted for, and amounts to 1.6% of the active area of the tracker, was the vertical gaps in between the sensors in the middle of a module. The gaps and staggering of the planes can be clearly seen




FIGURE 4.6. A photograph of a tracker plane, where 8 SCT modules are installed in an aluminum frame.




FIGURE 4.7. A tracker station is depicted as an expanded CAD view on the LEFT, and as a photograph of the final assembly on the RIGHT.

in Fig. 4.8, which shows the number of hits on track across the transverse area of the most upstream tracking station.

Quality tests and calibration runs were routinely performed on the SCT modules throughout the construction of the tracking spectrometer and after its installation,




FIGURE 4.8. The number of hits on track, per mm^2 , is shown across the transverse area of the most upstream tracking station. Both the vertical gaps between the sensors of a module and the staggered horizontal gaps between the modules can be seen. Less hits on track are seen for positions with $r \geq 100$ mm because that region is where the magnetic material lies, which scatters tracks and complicates the track reconstruction.

as described in [54]. At the very start of construction, the spare SCT barrel modules from ATLAS were tested and selected for FASER based upon being operational up to a bias voltage of 300 V and also for having the least amount of dead strips. After their final installation, only about 0.5% of all strips, in all SCT modules used, failed to respond to an injected charge and were deemed to be dead. As for noisy strips, the occupancy of strips that have noise above the 1 fC threshold was measured to be 6.61×10^{-5} . Finally, the detection efficiency for a high energy charged particle passing through the sensitive region of the SCT sensors was found to be above 99.8%, which was measured by reconstructing tracks, as described in Chapter VI, without using the hits from a given SCT module and then seeing if the module had a hit where the reconstructed track says it should.

4.1.4 Tracker Metrology and Alignment

Accurately measuring where each tracking module is positioned with respect to the others is done throughout both detector assembly and regular data-taking. During assembly, metrology and survey scans are performed to ensure the tracking spectrometer is being built to the designed specifications. Whereas, throughout data-taking, events with through-going particles are used to perform an iterative local alignment that minimizes the χ^2 of reconstructed tracks and optimizes the tracking performance.

A Mitutoyo CRYSTA-Apex S CNC coordinate measuring machine was used to perform metrology scans on each assembled tracking layer and station, as seen in Fig. 4.9. Equipped with a touch-probe and an optical camera, the metrology machine can perform precise measurements ($\sim 5 \mu\text{m}$ precision in the plane and $10 - 15 \mu\text{m}$ out of the plane) with respect to visual targets located on the corners of the tracking planes. Comparing these measurements with the perfectly aligned CAD geometry, showed that the SCT modules had a maximal deviation of $100 \mu\text{m}$ from their designated positions. Throughout the assembly of the full tracking spectrometer, the Cern Survey team used a Leica laser tracker to map out, with mm-level precision, all of the magnets and tracking stations with respect to the LHC reference system.

The in situ alignment of the tracking spectrometer using reconstructed tracks, which are described in Chapter VI, was done with $\sim 596 \text{ pb}^{-1}$ of data collected in late July of 2022. An iterative local χ^2 alignment method was used that incrementally shifts the positions of the tracking modules, in the software, in the direction that minimizes the disagreement between the hit positions in that module and the expected hit positions from the reconstructed tracks. Only the most sensitive




FIGURE 4.9. A picture of the metrology scans being performed on a tracker plane.

distortions, translations in the precision tracker coordinate (\hat{y}) and rotations around the longitudinal axis (\hat{z}), at both the individual module and tracking layer level were considered. One must also deal with weak modes in the tracker alignment, which are distortions in the tracking spectrometer that do not affect the track residuals. In order to constrain the weak modes in FASER, which include a common translation or rotation in all of the tracking layers and a sheering of the tracking station and layer positions in \hat{y} , the positions of the modules in the middle layer at stations 1 and 3 are fixed. The results of the alignment after 20 iterations are shown in Fig. 4.10, where a clear improvement in the hit residual distributions are seen when comparing unaligned data, aligned data, and perfectly aligned MC. The remaining disagreement between the aligned data and the perfectly aligned MC can be attributed to the misalignment in the two fixed layers and also the misalignment in the less sensitive translation and rotation directions that were not considered.

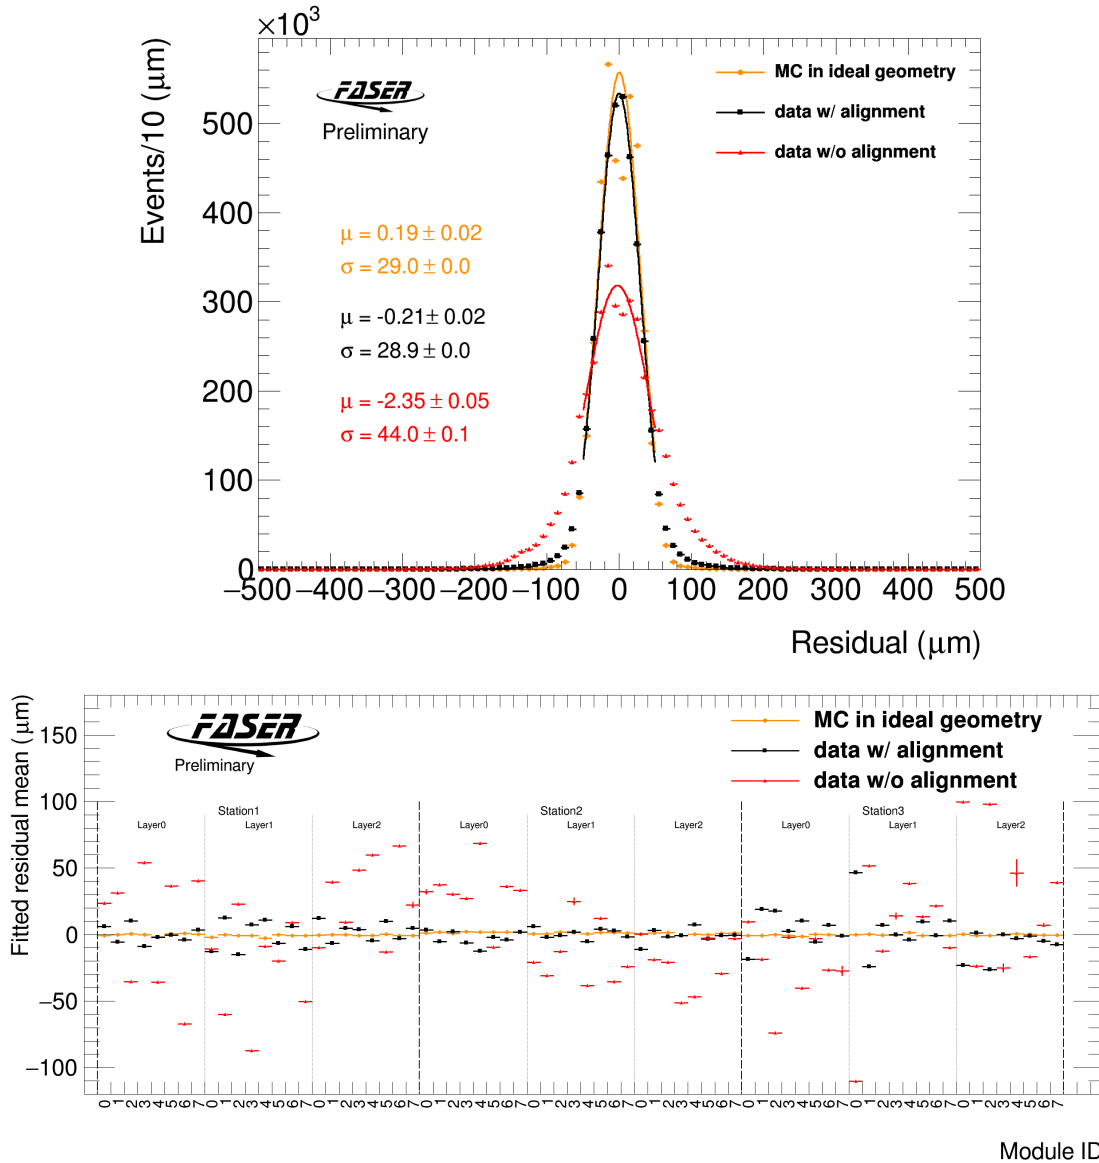


FIGURE 4.10. Comparisons of the reconstructed track residuals are shown for tracks reconstructed from data in an unaligned geometry (red), data in an aligned geometry (black), and from MC simulation with an ideal geometry (orange). In the TOP, the track residual distributions are shown when combining the residuals from all tracker modules. The distributions are normalized to the same number of entries and fit with a Gaussian. The BOTTOM plot shows the means of the track residual distributions for each tracker module. For both plots, the data used has an integrated luminosity of 596 pb^{-1} and only the tracks with momentum greater than 300 GeV and the number of clusters greater than 14 are used.

4.2 SCINTILLATORS

The scintillator system is composed of four stations, each consisting of multiple scintillating modules. A scintillator module is defined as the combination of a scintillating plane, a light guide or wavelength-shifting (WLS) rod, and a Photo-Multiplier-Tube (PMT). The EJ-200 plastic scintillator material [57] was used for all scintillating planes, providing a light output of about 22,000 photons/cm for minimum ionizing particles [58]. Efficient charged particle detection is achieved by using thick scintillating layers of either 1 or 2 cm, corresponding to 2.5–5% of a radiation length. Light generated in each scintillating layer is guided to a PMT, where the light is converted into an amplified electronic signal. In order to prevent light leaks, and to improve fire safety, all scintillating layers and light guides are wrapped in 0.5 mm-thick aluminum tape. The different scintillator stations can be seen individually in Fig. 4.11, and as part of the full detector in Fig. 4.1.

Starting at the front of the detector, nearest the ATLAS IP, is the VetoNu scintillator station. There are two back-to-back scintillator modules that make up the station, each with a scintillating layer of 30×35 cm² transverse area and 2 cm depth. Due to space limitations at the front of the FASER detector, where the area around the FASER ν sub-detector is enclosed, a compact light collection mechanism and PMT are used only for the VetoNu modules. A 1.5 cm \times 1.5 cm \times 37.5 cm EJ-280 plastic WLS rod [57] is coupled with optical grease to both the vertical side of the scintillating layer and also to a compact Hamamatsu H11934-300 PMT [59]. The WLS rod provides a light emission spectrum that peaks around 500 nm, which matches the light wavelength that maximizes the sensitivity of the Hamamatsu H11934-300 PMT.




FIGURE 4.11. The VetoNu (TOP-LEFT), Veto (TOP-RIGHT), Timing/Trigger (BOTTOM-LEFT), and Pre-Shower (BOTTOM-RIGHT) scintillator stations are shown.

Further downstream is the Veto station, constructed from four scintillator modules, and placed just in front of the decay volume. A single module encompasses a $30\text{ cm} \times 30\text{ cm} \times 2\text{ cm}$ scintillator, a triangular light guide, and a Hamamatsu H6410 PMT [60]. The PMT assembly is spring-loaded such that it maintains good optical contact with the light guide. To minimize the effect that the magnets have on

the PMT response, there are two layers of permalloy tube protection surrounding the PMT. The triangular light guide is mounted on top of the scintillator and at an angle from vertical of ± 140 mrad, allowing for two modules to be placed back-to-back without their PMTs clashing. Two pairs of back-to-back modules are placed on either side of a 10 cm-thick lead absorber. The lead is to ensure that incoming photons are either completely absorbed or generate an EM shower that can be detected by the pair of downstream scintillator modules.

Both the VetoNu and Veto stations have scintillators with a transverse size significantly larger than the active region of the detector, which allows for the efficient detection and rejection of muons, even if they enter the detector at an angle with respect to the LOS.

The next downstream scintillator station is the Trigger/Timing station, which is installed just after the decay volume and before the first tracking station. Covering a large transverse area of 40×40 cm², the Trigger/Timing station is used to detect any charged particles traversing the FASER tracking spectrometer and to trigger the recording of all sub-detector signals for that event. This station also provides sub-nanosecond signal timing, which is used to discriminate between particles coming from pp collisions at the ATLAS IP and non-collision backgrounds. To minimize the time it takes for light generated in the scintillator to reach the PMT, the scintillating plane is separated into two planes that are stacked vertically on top of one another, where each plane has two PMT's, with one mounted on each side. The PMTs and light guides used in the Trigger/Timing station are the same as those used in the Veto station, except there is an additional 90-degree bend in the triangular light guide that is needed to keep the Hamamatsu H6410 PMTs from clashing with the sides of the trench. In order to minimize the amount of material in the volume of

the tracking spectrometer, the Trigger/Timing station only uses a single scintillator plane with a thickness of 1 cm.

Finally, the Pre-shower station is near the end of the detector, positioned after the last tracking station and just before the calorimeter. The Pre-shower consists of two scintillator modules that are identical to the Veto scintillator modules, interleaved with two 3 mm thick tungsten plates and three 5 mm thick porous graphite blocks. Both the tungsten plates and graphite blocks have the same $30 \times 30 \text{ cm}^2$ transverse area as the scintillators. The graphite blocks are used to stop low-energy backward-travelling particles that originate from the electromagnetic shower in the calorimeter and could cause additional hits in the last tracking station. With one radiation length of tungsten absorber directly preceding each scintillator module, the primary purpose of the Pre-shower is to cause the electron-positron pair, coming from the dark photon decay, to start their electromagnetic shower and to record the longitudinal development of the shower. As seen in Sec. 4.3, the calorimeter does not provide any longitudinal information about the electromagnetic shower, and thus the Pre-shower is critical for particle identification. The ratio of the two preshower modules can be used to discern between quickly evolving electromagnetic showers (e.g. photons and electrons), slowly evolving hadronic showers (e.g. pions), and non-showering minimum-ionizing particles (e.g. muons). The Pre-shower also improves the energy resolution of the calorimeter for highly energetic ($\sim \text{TeV}$) electromagnetic showers, which can penetrate the entire length of the calorimeter, as the extra radiation lengths of the preshower allow for more of the shower to be captured in the calorimeter.

4.3 CALORIMETER

The electromagnetic energy of particles is measured by an electromagnetic calorimeter, the most downstream component of the detector. The calorimeter is constructed from four outer ECAL modules from the LHCb experiment [61], and is shown in Fig. 4.12. Each module, depicted in Fig. 4.13, is 12×12 cm² in the transverse plane, with 66 layers of interleaved 4 mm thick plastic scintillator and 2 mm thick lead plates, corresponding to a total of 25 radiation lengths. A module has 64 wavelength-shifting fibers that penetrate the length of the module and guide the light generated in the scintillating layers to the PMT at the back of the module. The Hamamatsu R7899-20 PMT [62] is used to detect the light from the calorimeter, which is the same PMT that LHCb uses. The PMT assembly is shown in Fig. 4.14, and includes both a light mixer and light filter upstream of the PMT, which are all covered by a permalloy tube in order to reduce the effect of magnetic fields. A 32 mm long and 8 mm wide rectangular polystyrene light mixer is placed just before the PMT to reduce the positional non-uniformity of the PMT response. A light filter with 10% transmission efficiency is installed between the calorimeter's wavelength-shifting fibers and the light mixer in order to run the PMT at high gain, where its response is most linear, whilst also not saturating the PMTs during \sim TeV EM showers. With this default setup, the readout of the PMTs only saturates for large pulses corresponding to energy deposits above 3 TeV.

The calorimeter energy resolution has been measured with high energy electrons in a test-beam, discussed in Appendix B.1, to be $\mathcal{O}(1\%)$ in the high-energy range most relevant for this analysis.




FIGURE 4.12. A picture of the FASER calorimeter is shown, where four LHCb outer ECal modules are placed 2×2 , and the PMTs are coming out of the back of the calorimeter.




FIGURE 4.13. The calorimeter module design is depicted.




FIGURE 4.14. The calorimeter PMT assembly is shown.

4.4 TRIGGER AND DATA ACQUISITION

The FASER trigger and data acquisition (TDAQ) system is designed to capture the raw data from all scintillator, calorimeter, and tracking modules whenever a highly energetic charged particle passes through the detector. Readout is triggered by signals from the scintillators or calorimeter system, with a typical trigger rate of ~ 1 kHz during an instantaneous luminosity of $2 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$, and is dominated by high energy muons from IP1. Upon readout, the digitized raw data from all sub-detectors is sent via optical fibers to a DAQ PC at the surface, where the data is combined to form an event and then saved to storage.

The remainder of this section will describe both the main TDAQ hardware and the trigger logic used when collecting data.

4.4.1 TDAQ Hardware

The central component of the TDAQ hardware is the Trigger Logic Board (TLB), which receives trigger inputs from the waveform digitizer and then decides when to distribute a global Level-1 accept (L1A) signal to the Tracker Readout Board (TRB) and the digitizer in order to initiate the readout of an event. A diagram of the TDAQ hardware system can be seen in Fig. 4.15, and a detailed description of the the entire system can be found in Ref. [16]. Just a brief overview of the TLB, TRB, and digitizer will be discussed here.

The TLB is a custom general-purpose input/output (GPIO) board that utilizes a CYCLONE V A7 FPGA and is driven by the LHC clock at a frequency of 40.08 MHz. The trigger logic required for the TLB to send out an L1A signal is discussed in Sec. 4.4.2. To track which proton-proton bunch collision we are triggering on, a

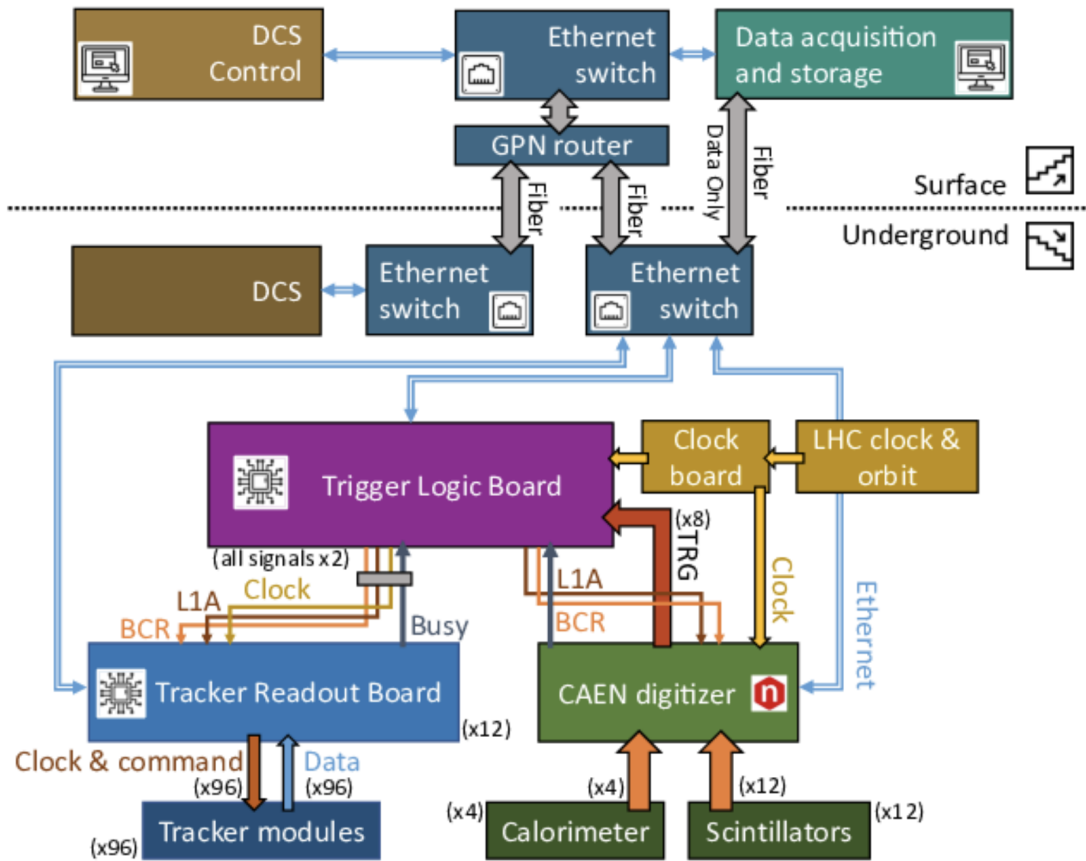


FIGURE 4.15. A diagram depicting the FASER TDAQ architecture. Each component is represented by a colorful box and the connections between the components are shown with arrows. The numbers in parentheses indicate the number of channels. The figure is from Ref [16].

bunch counter in the TLB increments by one for each LHC clock cycle and resets with the bunch counter reset (BCR) signal that comes each time the LHC beam completes an orbit. In order to be sure that the data from the TRB and digitizer are readout with the correct event, the TLB has adjustable delays for sending out both the L1A and BCR signals. The TLB also has an on-board rate limiter, set to a maximum trigger rate of 2.2 kHz in order to protect against possible noise bursts. The time in which triggered events are ignored and not read out is referred to as dead-time. This dead-time does not only come from the rate limiter, but also from

the busy signals of both the TRB and digitizer that veto the L1A signal whilst they are reading out data from a previous trigger.

A TRB is composed of an identical GPIO board to the TLB, also driven by the LHC clock, and an adapter card that allows for connection to the tracker modules through the tracking station patch panels. The TRB firmware translates commands from the host PC, into a bit-stream that the SCT modules can understand. Thus the TRB provides a way to communicate with, configure, and readout the tracker. A single TRB reads out and operates eight SCT modules, corresponding to an entire tracker plane. Thus, there are three TRBs per tracking station and 12 total TRBs when including the IFT tracking station in addition to the three-station tracking spectrometer. Where the TLB can provide a delayed L1A to all of the TRBs, similarly can each TLB delay the readout of the SCT modules such that the data from each module arrives in time with one another.

During testing and calibrations, both the TLB and TRB can be driven by 40 MHz internal clocks rather than the LHC clock.

A 14-bit CAEN VX1730 digitizer board receives the analog PMT signals from the scintillator and calorimeter modules, digitizes the waveforms, and provides trigger input signals to the TLB. The LHC clock is also input into the digitizer such that the relative timing between the PMT waveforms and the LHC clock can be measured during data reconstruction. With one LHC clock and 16 scintillator/calorimeter modules needing to be input into the 16-channel digitizer, there was one too many inputs for the digitizer. It was decided that the most upstream layer in the Veto would not be digitized, as there are still three Veto layers that provide sufficient redundancy. For the remaining scintillator/calorimeter modules and the LHC clock, the waveforms from each digitizer channel are sampled at a rate of 500 MHz and

stored in a 1200 ns circular buffer that is read out upon receiving an event trigger from the TLB. All digitizer channels have a 2 V range and an adjustable baseline offset. A baseline offset of 0.85 V was implemented for the data used in this analysis, such that a PMT signal whose baseline-to-peak amplitude was larger than 1.85 V would produce a saturated digitized waveform. The digitizer also monitors the waveforms at all times and sends trigger signals to the TLB whenever the waveform peak is greater than a predefined threshold for each channel, as described in Sec. 4.4.2.

4.4.2 *Trigger Logic*

The CAEN digitizer produces eight trigger outputs, which the TLB receives and logically combines into four physics triggers that cause the readout and storage of an event.

The eight triggers produced from the digitizer are logical “and/or” combinations of paired up calorimeter and scintillator modules as follows:

1. CaloBottom: either of the two calorimeter modules on the bottom
2. CaloTop: either of the two calorimeter modules on the top
3. Preshower: both of the pre-shower layers
4. TimingBottom: both of the PMTs on the bottom trigger/timing layer
5. TimingTop: both of the PMTs on the top trigger/timing layer
6. 1stVeto: the single digitized Veto layer upstream of the lead shielding
 - the other Veto layer upstream of the lead is not connected to the digitizer
7. 2ndVeto: both of the Veto layers downstream of the lead shielding

8. VetoNu: both of the VetoNu layers

where a given sub-detector is triggered only if its signal crosses a predefined digitizer threshold. The thresholds in the scintillators were tuned to be well below that seen from a single minimum-ionizing particle, while the calorimeter thresholds were set to the amplitude expected from a ~ 20 GeV electromagnetic shower. All digitizer thresholds are at least $10\times$ larger than the rms of the waveform noise, which is dominated by the electronic noise of the digitizer.

The TLB further combines the eight input trigger lines from the digitizer into four trigger logic items that generate an L1A signal, and are defined as follows:

1. CaloBottom OR CaloTop trigger
2. VetoNu OR 1stVeto OR 2ndVeto OR Preshower trigger
3. TimingBottom OR TimingTop trigger
4. A coincidence between (VetoNu OR 1stVeto OR 2ndVeto) AND Preshower trigger

You can see the different FASER triggers and their recorded trigger rates, along with the total rate, for a common LHC fill in Fig. 4.16. The total trigger rate is dominated by the Trigger/Timing scintillator station as it has the largest sensitive transverse area. All of the TLB triggers, except for the coincidence trigger, should fire due to a dark photon signal, providing redundancy that is necessary when trying to capture such a rare event.




FIGURE 4.16. The recorded trigger rates for each of the four FASER TLB triggers are shown during LHC Fill 8143 on August 19, 2022. The triggers are further described in Sec. 4.4.2. The structure and decay in rate over time closely follows the LHC luminosity at IP1.

CHAPTER V

DATASET AND SIMULATION SAMPLES

This search uses 27.0 fb^{-1} of Run 3 collision data collected by FASER between September and November 2022. From July to August 2022, the calorimeter light filters were removed for commissioning purposes, causing saturation for electromagnetic showers with energies greater than 300 GeV; this data is excluded from the dark photon search. A total of 64 different FASER runs were chosen for this analysis, which all met the requirements that the delivered luminosity was greater than 10 pb^{-1} and the calorimeters were setup to measure high energy ($\sim \text{TeV}$) electromagnetic showers. The luminosity of the dataset is provided by the ATLAS experiment [47, 63, 64] and corrected for the FASER detector dead-time. The average dead-time for the runs used in this analysis was 1.3%. The cumulative delivered luminosity and recorded luminosity throughout 2022 and 2023 can be seen in Fig. 5.1.

In order to check that there were no unseen issues with any of the 64 runs used in this analysis, both the number of reconstructed tracks and the number of events with energy in the calorimeter greater than 100 GeV were normalized by the luminosity recorded during that run and then compared to the other runs, as seen in Fig. 5.2. The luminosity normalized detector response is seen to be stable across all runs, providing confidence that the detector conditions are constant for the data used in this analysis.

Monte Carlo (MC) simulation samples are used to evaluate the signal efficiency, in the estimation of background yields, and to calculate the systematic uncertainties. All samples are simulated using **GEANT4** [65] with a perfectly aligned and detailed description of the detector geometry, including passive material. The simulation




FIGURE 5.1. The cumulative luminosity versus time delivered to (yellow) and recorded by FASER (blue) during stable beams are shown.

properly aligns the detector with respect to the IP, where there is a 6.5 cm vertical offset from the longitudinal LOS, as discussed in Sec. 3.2. The simulated response from the scintillator and calorimeter modules are digitized to mimic the waveforms from the PMT output in data, the process of which is described in detail in Appendix A.2. The MC samples include a realistic level of detector noise, and are reconstructed in the same way as the data, as described in Chapter VI.

Dark photon signal events are generated using FORESEE [66] with the EPOS-LHC [13] generator to model very forward π^0 and η meson production in the LHC collisions. The production of dark photons via dark bremsstrahlung is also included, which is modelled using the Fermi-Weizsacker-Williams approximation following Ref. [67] with the additional requirement on the dark photon's transverse




FIGURE 5.2. The luminosity-normalized detector response is shown for all 64 runs used in this analysis. The total number of reconstructed tracks, as well as the number of positively or negatively curved tracks, are shown in the TOP plot; whereas the BOTTOM plot shows the total number of events with energy in the calorimeter that is greater than that of a 100 GeV electromagnetic shower.

momentum, $p_T(A') < 1$ GeV, to ensure the validity of the calculation. Signal samples for the A' model are generated covering the relevant ranges in both coupling and mass, as shown in Fig. 5.3.

A high-statistics high-energy muon sample with 2×10^8 events entering FASER from IP1 is used for several background and systematic uncertainty studies. The




FIGURE 5.3. The dark photon samples generated for this analysis are shown as red dots in the phase space of the A' mass and coupling (ϵ). The yellow regions show the expected sensitivity of FASER for varying amounts of delivered luminosity. Regions previously excluded by other experiments are shown in grey, and regions that could be excluded in the future by other experiments are shown as colorful dotted lines.

sample uses the expected energy and angle of the muons as estimated by FLUKA [52, 53, 68] simulations of incoming muons from IP1. The samples include a detailed description of the LHC components and infrastructure between IP1 and FASER. A similar sample of 8×10^5 large-angle (15-60 mrad with respect to the longitudinal axis of FASER) muon events generated slightly upstream of the VetoNu scintillators, and with a radius spanning 15-30 cm covering the edge region of the scintillators, is produced and used to study the background from large-angle muons that miss the veto system.

Neutrino interactions in FASER [69] are simulated by the **GENIE** [70, 71] generator, following the fluence, energy spectrum and flavour composition obtained in Ref. [72]; which presents a fast simulation of the high-energy neutrino flux in the far-forward direction of proton-proton collisions at the LHC by only accounting for neutrinos produced in the vacuum of the beam pipe, and not in downstream material interactions, which was found to be subdominant. The sample used for the neutrino background study corresponds to 300 ab^{-1} of data, and only includes neutrino interactions in FASER that are either upstream of the Veto scintillators or in the active detector volume.

CHAPTER VI

EVENT RECONSTRUCTION

Event reconstruction is done to prepare the raw data for analysis, and is performed using FASER's Calypso [73] offline software system, based on the open-source Athena framework [74, 75] from the ATLAS experiment.

The reconstruction of charged particle tracks in the tracking spectrometer is done in steps that take the raw data from SCT strips and forms clusters, then spacepoints, then track segments, and finally full tracks. Clusters give the total charge and charge-weighted position of groups of SCT strip neighbors that see a signal above threshold. Separate clusters are formed on either side of an SCT sensor, and thus the cluster position is only sensitive to the precision local coordinate (\hat{y}) of the sensor. To get the global 3D position of where the track interacted with (hit) the SCT sensor, two clusters on the front and back of an SCT sensor are combined into spacepoints. The local \hat{x} position of the hit is calculated using the two clusters and the stereo-angle between the strips on the front and back of an SCT module. The global position, which is stored as a spacepoint, is then calculated by combining the local hit position in the SCT module with the global position of the SCT module in the aligned detector geometry. All combinations of spacepoints in separate tracking layers of a given tracking station are then linearly fit to form track segments. If the x-angle of a track segment is ≥ 0.08 radians, which is the maximum angle that a straight track can traverse the length of the tracking spectrometer, then the track segment is discarded. Also, if multiple track segments share more than 60% of clusters, then only the segment with the most clusters are kept, where the smallest χ^2 is used as a tie breaker. Finally, full track reconstruction is performed

by using the track segments as seeds in the combinatorial Kalman filter from the ACTS library [76], which takes into account material interactions and the measured magnetic field maps. When reconstructing multiple tracks, it is required that they do not share more than 6 clusters; if the number of shared hits exceeds this threshold, then the track with the higher χ^2 is discarded.

Extracting desired event information from the scintillator and calorimeter modules is done through a series of waveform reconstruction steps, and is described in detail in Appendix A.1. The only waveform measurement used directly in this analysis is that of the integrated PMT charge over a small 120 ns window on the waveform, which is where the signal is expected to arrive from a relativistic charged particle originating from the IP. After pedestal subtraction and selecting this waveform window, the PMT charge is calculated by simply summing the digitised waveform values. But, if the waveform is saturated, then the PMT charge is calculated by using the integral of a Crystal Ball fit on the signal, which allows for an estimation of the signal charge if it were not saturated, as demonstrated in Fig. A.4 of Appendix A.1.

The reconstructed track and PMT waveforms are shown in Fig. 6.1 for an event where a single muon traverses the FASER detector.

The calorimeter charge-to-energy scale calibration is determined using high energy electron and muon beams from the test-beam data described in Appendix B.1. To take into account differences between the detector configurations in the test-beam and in collision data, the most probable calorimeter charge deposited by muons as minimum ionising particles (MIPs) is used as an in-situ normalisation of the energy scale. Special calibration runs are performed at high calorimeter gain to measure the MIP signal. After individually normalising each calorimeter module signal to the

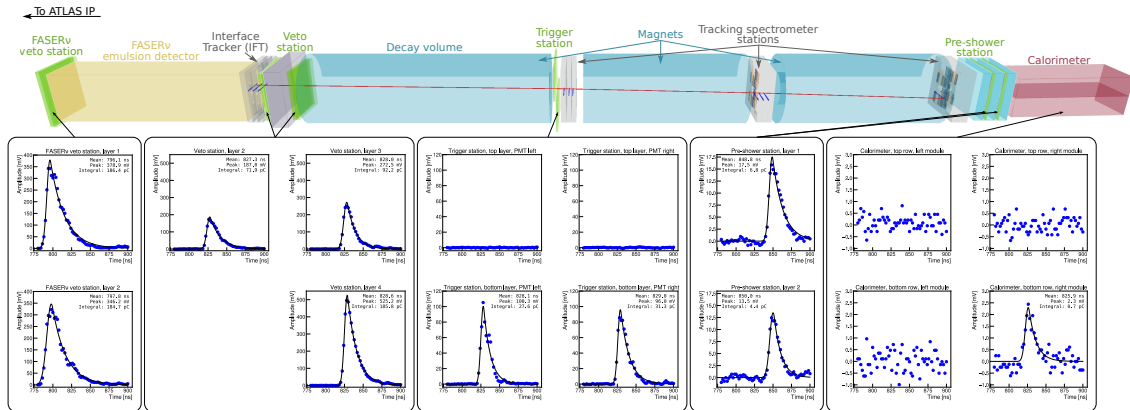


FIGURE 6.1. The reconstructed event display of a muon traversing FASER. The event (run 8336, eventID 1477982) occurred on August 23, 2022 and is in time with a collision event. The measured track momentum is 21.9 GeV. The waveforms are shown for signals in the scintillator and calorimeter modules, and are fit using a Crystal Ball function. All PMT waveforms are consistent with a muon passing through the scintillators and one of the calorimeter modules. The detected hits in the semiconductor tracker modules are shown with blue lines and the reconstructed track is shown with a red line. In the title of the waveform plots, left and right is defined facing the downstream direction.

MIP scale, the test-beam data is used to estimate the initial electromagnetic energy of the particle entering the calorimeter. A detailed account of the calorimeter energy calibration is presented in Appendix B.2.

CHAPTER VII

EVENT SELECTION

The typical A' detector signature, shown in Fig. 7.1, provides a unique signature to investigate. Because the A' is weakly interacting, no signal is expected in the veto scintillator systems. The A' can then decay in the decay volume to a very collimated, high momentum, e^+e^- pair, leaving two closely-spaced oppositely-charged particle tracks in the tracking spectrometer. The e^+e^- then leave signals in both the Timing and Pre-shower scintillators as well as a large energy deposit in the calorimeter. There are no significant SM processes that can mimic this signature, allowing for a practically background-free search.

To avoid unconscious bias affecting the analysis, a blinding procedure is applied to events where there is both no signal in any veto scintillator and the calorimeter energy is above 100 GeV. The event selection, background estimation and systematic uncertainties are then finalised before looking in this signal-dominated region of the data.

The signal region event selection requires the following:




FIGURE 7.1. A sketch presenting a side view of the FASER detector, showing the different detector systems as well as the signature of a dark photon (A') decaying to an electron-positron pair inside the decay volume. The white blobs depict where measurements are taken for the A' signal and the solid red lines represent the reconstructed tracks produced by the e^+e^- pair.

- the trigger from the Timing scintillator fired;
- event time is consistent with a colliding bunch at IP1;
- no signal in any of the five vetoNu and veto scintillators;
 - * required to be less than half that expected from a MIP
- signal in the scintillators that are downstream of the decay volume;
 - * required to be compatible with or larger than expected for two MIPs
- two fiducial reconstructed tracks of good quality;
 - * a good quality track has a track fit $\chi^2/(\text{number of degrees of freedom}) < 25$, at least 12 hits on track, and a momentum > 20 GeV
 - * a fiducial track has an extrapolated position of < 9.5 cm radius at all scintillators and tracking stations
- total calorimeter energy greater than 500 GeV;

Where the exact PMT charge cuts applied to the scintillators are shown in Table 7.1. Before even utilizing the tracker, the cut applied to the Trigger/Timing scintillator allows us to reject most of the single-track events in data whilst keeping the two-track A' signal events, as shown in Fig. 7.2.

The cumulative efficiency of each event selection requirement on a representative signal model in the parameter space where the analysis is most sensitive ($\epsilon = 3 \times 10^{-5}$, $m_{A'} = 25.1$ MeV) is shown in Table 7.2. The total efficiency was found to be about 50% for dark photons that decay in the decay volume (where the probability of the dark photon to decay while within the volume is $\mathcal{O}(10^{-3})$), with the largest inefficiency arising from the two track requirement.

Scintillator	PMT Threshold
Trigger/Timing station: Top OR Bottom layer or Top AND Bottom layer	$> 70 \text{ pC}$
Each Pre-shower layer	$> 2.5 \text{ pC}$
Each Veto layer	$< 40 \text{ pC}$
Each VetoNu layer	$< 40 \text{ pC}$

TABLE 7.1. Summary of the scintillator PMT charge thresholds used in the A' event selection. Because the Trigger/Timing layer is split into a top and bottom plane, the two-MIP minimum cut on the Trigger/Timing PMT charge accommodates for the case if two particles go into the same plane or separate planes.




FIGURE 7.2. The maximum reconstructed charge from the timing station scintillators for data, and MC A' simulated signal. The data is dominated by single muon events. The minimum required charge of 70 pC is shown as a dashed line.

A requirement that the Timing scintillator trigger fired ensures that the trigger efficiency, measured using orthogonal triggers on two-track events, is 100% for the A' phase space of interest.

Selection Criteria	Efficiency
No Veto Signal	99.7%
Timing+Preshower Signal	97.9%
≥ 1 good track	91.6%
= 2 good tracks	57.3% *
Track radius < 95 mm	51.8% *
Calo E > 500 GeV	50.8% *

TABLE 7.2. Cutflow for a representative signal model with $\epsilon = 3 \times 10^{-5}$ $m_{A'}$ = 25.1 MeV, showing the fraction of MC simulated signal events passing the different event selection criteria. The efficiency numbers in the rows marked with a * have been scaled down by 7% following a data driven estimate of the two-track tracking efficiency, as described in Sec. 9.2. The statistical uncertainty is smaller than the last digit shown.

The probability to veto a signal event, due to the presence of an uncorrelated beam-background muon in the same or neighbouring bunch crossing, is estimated to be less than 1 per mille.

CHAPTER VIII

BACKGROUNDS

Several sources of background are considered in the analysis. The dominant background arises from neutrino interactions in the detector. Other processes such as neutral hadrons, or muons that enter the detector volume without firing the veto scintillator systems, either by missing the scintillators or due to scintillator inefficiencies, also contribute to the background. Finally, non-collision backgrounds from cosmic-rays or nearby LHC beam interactions are also considered. The contribution of each of these background sources is described and quantified in the following sub-sections.

8.1 VETO INEFFICIENCY

It is necessary to detect and veto incoming charged particles, dominated by ~ 1 kHz of high energy muons, which can interact in material and produce two-track events with large deposits of energy in the calorimeter. The inefficiency of each of the five planes of veto scintillators is measured independently with data, by selecting events in which there is a single good fiducial reconstructed track and then measuring the fraction of such events in which the scintillator charge is below that of a MIP signal. The single-track PMT charge distributions are shown in Fig. 8.1 for both a module in the VetoNu station and a module in the Veto station. Thanks to the thick scintillators and tight fiducial track requirements, the inefficiencies are at the 10^{-5} level or smaller for all VetoNu and Veto planes, as seen in Table 8.1. Since the planes are independent, this leads to a combined veto inefficiency of smaller than

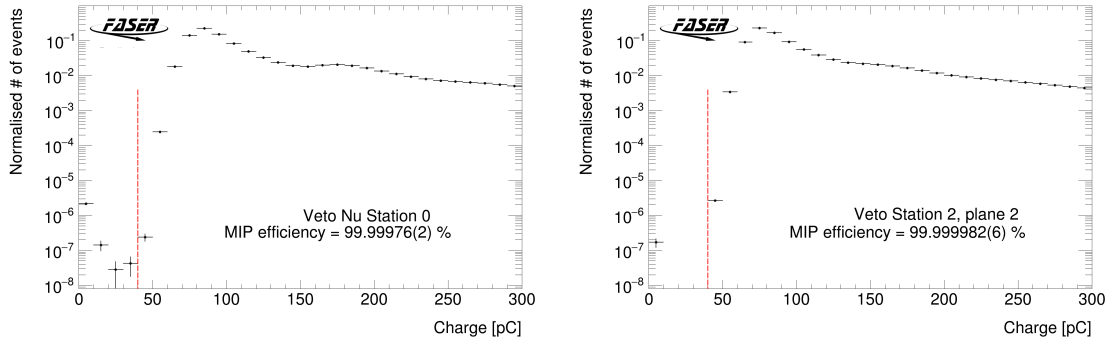


FIGURE 8.1. Example plots of the charge distribution for events with good fiducial tracks for the most downstream VetoNu scintillator (left) and the most upstream Veto scintillator (right). The efficiency is calculated as the ratio of the number of events with charge above the 40 pC threshold (indicated on the figures) to the total number of events.

10^{-20} . As $\mathcal{O}(10^8)$ incoming muons are observed in the 2022 dataset, the background due to the veto inefficiency is taken to be negligible.

Scintillator	Efficiency
VetoNu-0	0.9999976(2)
VetoNu-1	0.9999974(2)
Veto-0	0.9999994(1)
Veto-1	0.99999976(7)
Veto-2	0.99999982(6)

TABLE 8.1. The measured efficiency of the five veto scintillator planes (the number in parenthesis is the statistical uncertainty).

8.2 NEUTRAL HADRONS

Neutral hadrons produced in muon interactions in the rock in front of FASER can be a possible source of background if, when passing through the veto systems undetected and interacting or decaying inside the detector decay volume, they produce exactly two reconstructed charged particle tracks and a calorimeter energy deposit above 500 GeV. This background is heavily suppressed by the need for the

neutral hadron to traverse the full eight interaction lengths of the FASER ν detector, and by the need for the parent muon to scatter to miss the veto scintillators. To determine the fraction of neutral hadron events that deposit at least 500 GeV of energy in the calorimeter, a three-track control region is used, where the parent muon enters the detector and is reconstructed along with the neutral hadron decay products. In these three-track events, the ratio of events with low calorimeter energy ($E < 100$ GeV) to high energy ($E > 500$ GeV) is used to scale the number of events with two reconstructed tracks (in which the parent muon is not present in the detector) at low-energy ($E < 100$ GeV) to estimate the expected background number of two-track events with $E > 500$ GeV. To allow sufficient event counts in the two-track low-energy control region, the veto requirements are relaxed, requiring no signal in the VetoNu scintillators, but with no requirements on the other Veto scintillator signals.

Photon conversion events (with the accompanying parent muon) constitute a significant fraction of the three-track sample defined above and must be removed. This is done by requiring that the invariant mass of the two lowest momentum tracks, where the muon is assumed to be the highest momentum track, is greater than 200 MeV, which was found to be optimal when separating K_S events from photon conversions in MC simulation.

After discarding the photon conversion events, the number of data events in the low- and high-energy three-track regions are 404 and 19 respectively. The ratio of 19/404 is used to extrapolate the one event observed in the low-energy two-track region to the two-track high-energy region, resulting in an estimate of 0.047 expected events. This method provides an estimate of the number of neutral hadron events

that lead to two reconstructed tracks, contain more than 500 GeV of calorimeter energy, and leave no signal in the VetoNu scintillators.

To obtain the final background estimate, the results are corrected to account for the fact that the signal region selection requires no signal in the downstream Veto station as well. The correction is derived by studying the signal recorded in the Veto station using three-track events. With a clear separation in the Veto scintillator signal size for when only one track (the parent muon) traverses the Veto station versus when the other two tracks also leave a signal in the Veto station, the scintillators can be used to select both types of events and the ratio of the number of events in the two cases is used as the correction.

After correcting for the fraction of events that will decay or interact before the second veto system, a final estimate of $(8.4 \pm 11.9) \times 10^{-4}$ events is found; where the 100% statistical uncertainty is driven by the single event observed in the low-energy two-track data region, and an additional 100% systematic uncertainty is applied to account for the assumptions in the method.

In performing this estimation, potential neutrino background to the low-energy two-track data region, predicted to be 3.6 ± 3.8 events from GENIE simulation, is conservatively neglected.

8.3 LARGE-ANGLE MUONS

Another potential background source arises from large-angle muons that miss the veto system and then enter the FASER decay volume. This background is heavily suppressed by the fact that the tracks extrapolated to the front veto scintillators are required to be within the fiducial volume. The MC sample with large-angle muons generated at the edge of the scintillators, described in Chapter V, is used

to study this background. No two-track events are seen in this sample, even before applying the fiducial requirements on the extrapolated tracks or the calorimeter energy requirement, suggesting that this background is negligible in the final analysis.

This was validated via a data-driven method by using events with a signal in the veto scintillators and calculating the ratio of the number of such events with > 500 GeV or < 500 GeV in the calorimeter, which is then used to extrapolate from the number of events with no signal in the veto scintillators and < 500 GeV in the calorimeter to the number of events with no signal in the veto and > 500 GeV in the calorimeter. The results of this validation are consistent with those from the MC estimate, providing confidence that this background is negligible.

8.4 NEUTRINOS

The large flux of high energy neutrinos, whose interaction cross section rises with energy, at the FASER location constitutes an important background, since the neutrinos do not leave any signal in the veto scintillators, and can interact to produce high energy particles. To suppress this background, the detector was designed to minimise the amount of material in the main detector volume.

The expected background from neutrino interactions inside the detector is estimated using the 300 ab^{-1} ($\sim 10000\times$ larger than the data used in this analysis) neutrino MC sample described in Chapter V. The MC simulation shows that 0.0015 neutrino events (0.0012 electron (anti)neutrino events and 0.0003 muon (anti)neutrino events) pass the signal region selection when scaled to 27.0 fb^{-1} of data, with these interactions occurring in the Timing scintillator station or the first tracking station. Fig. 8.2 shows the calorimeter energy distribution for neutrino events that pass the signal region selection when disregarding the requirement on the




FIGURE 8.2. The calorimeter energy in simulated neutrino events passing all signal selection requirements, besides that on the calorimeter energy. GENIE is used to simulate the neutrino interactions. The figure is scaled to a luminosity of 27.0 fb^{-1} .

calorimeter energy. The figure shows that a requirement of $\geq 500 \text{ GeV}$ gives a good suppression of the neutrino background. The uncertainty on the incoming neutrino flux [72] is taken to be 100% for electron neutrinos and 25% for muon neutrinos, and an additional 100% uncertainty is applied to account for the effect of uncertainties in the modelling of neutrino interactions. The total neutrino background estimate when scaled to 27.0 fb^{-1} is $(1.5 \pm 0.5 \text{ (stat.)} \pm 1.9 \text{ (syst.)}) \times 10^{-3}$ events.

8.5 NON-COLLISION EVENTS

Apart from potential background sources that originate from pp collisions at IP1, backgrounds from both cosmic rays and the LHC beam interactions near FASER are also considered.

Cosmic rays are particle showers that occur when high energy particles or nuclei, which originate from our sun or other energetic sources beyond our solar system, interact with the Earth's atmosphere. The vast majority of cosmic ray particles will be absorbed before reaching FASER, which is 100 m underground, but it is possible for weakly interacting particles to reach the detector. These events are studied during 330 hours of data-taking with no beam circulating in the LHC, which corresponds to a similar running time to the full 2022 physics data-taking period. During this time, no event is observed with a calorimeter energy deposit above 100 GeV, and no events are found when requiring at least one good quality track, as shown in Fig. 8.3.

LHC beam interactions with gas or tails of the beam interacting with the beampipe aperture can lead to particles boosted in the direction of FASER. Detector activity is observed in correlation with LHC beam-1, the incoming beam to ATLAS in the FASER location, and it's bunches passing the back of the detector. This beam background is studied by checking the tracker and calorimeter activity in events with the relevant bunch timing, but which do not correspond to colliding bunches at IP1. It is found that beam background events without signal in the veto scintillators do not have a good reconstructed track, and for these events without a track, there are zero events with calorimeter energies above 400 GeV, as seen in Fig. 8.3.

The overall contribution from non-collision backgrounds is therefore considered to be negligible.

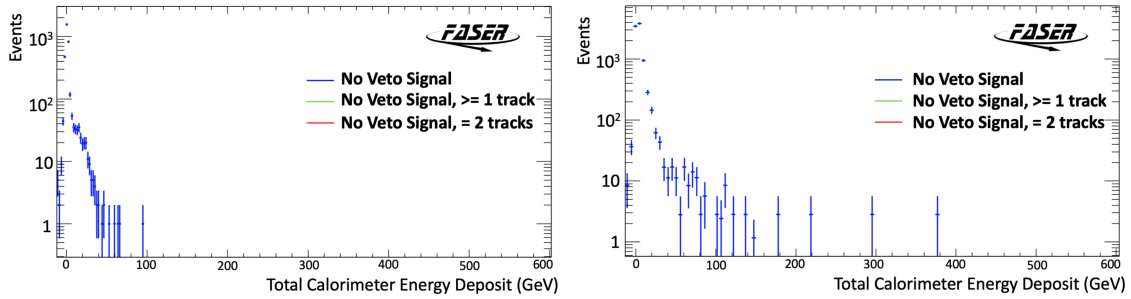


FIGURE 8.3. The calorimeter energy distributions are shown for different track requirements in both the (left) cosmic ray control sample and the (right) beam-1 background control sample.

8.6 SUMMARY OF THE EXPECTED BACKGROUND

The background sources considered and their estimated values are shown in Table 8.2. As background contributions from the veto inefficiency, large-angle muons, and non-collision events are estimated to provide a negligible contribution in the signal region, the total expected background is obtained by combining just the neutrino and neutral hadron estimates, leading to a total background of $(2.3 \pm 2.3) \times 10^{-3}$ events.

Background	Central Value	Error (%)
Veto inefficiency	-	-
Non-collision	-	-
Neutral hadrons	0.8×10^{-3}	1.2×10^{-3} (140%)
Neutrinos	1.5×10^{-3}	2.0×10^{-3} (130%)
Total	2.3×10^{-3}	2.3×10^{-3} (100%)

TABLE 8.2. Summary of the different background estimates. Estimates that are negligible are represented with a “-”.

CHAPTER IX

SYSTEMATIC UNCERTAINTIES ON THE SIGNAL YIELD

Systematic uncertainties on the expected signal yields arise from several sources. The uncertainty in the integrated luminosity is provided by the ATLAS collaboration, and is 2.2% [64], following the methodology discussed in Ref. [47]. The statistical uncertainty from the number of MC simulated signal events is included and ranges from 1 to 3%. There are also spin correlations between dark photon production and decay, but they are not included in the MC simulated signal as their effect on this search is negligible [77]. The remainder of the systematic uncertainties, discussed below, arise from the signal generator and from the modelling of the detector response in the MC simulation.

9.1 SIGNAL GENERATOR

Uncertainties on the number of signal events decaying inside the FASER decay volume are derived by comparing the estimates from using different event generators to model very forward π^0 and η meson production in the LHC collisions. Comparing signal yields from QGSJET II-04 [78] and SIBYLL 2.3d [79] with the central estimate from the EPOS-LHC [13] generator, where these generators have been validated using LHCf's forward photon measurements [80], provides an envelope of estimates as a function of the energy of the signal ($E(A')$), that is parameterized and used as the uncertainty:

$$\frac{\Delta N}{N} = \frac{0.15 + (E(A')/4 \text{ TeV})^3}{1 + (E(A')/4 \text{ TeV})^3} . \quad (9.1)$$

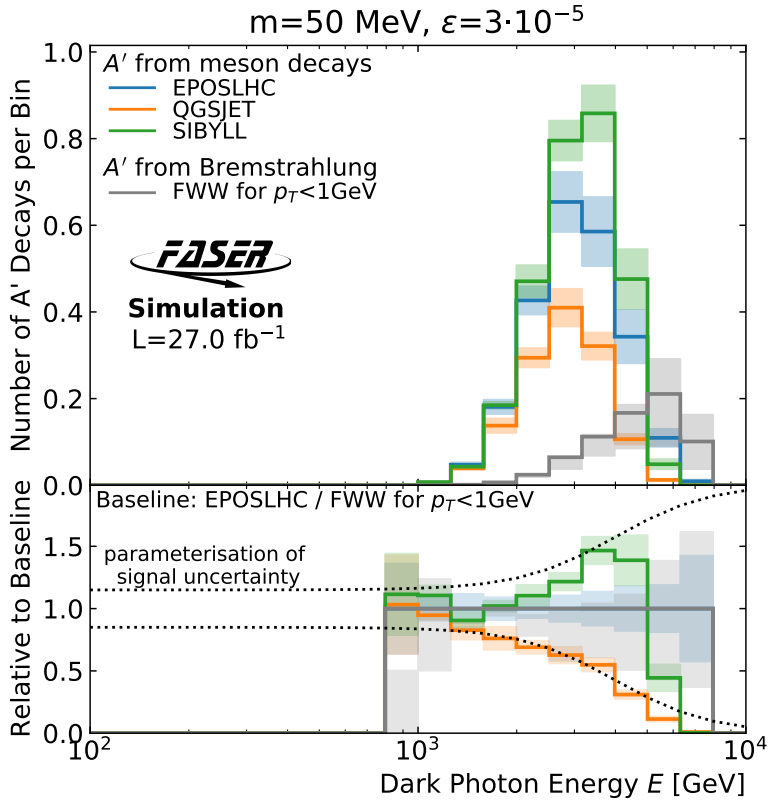


FIGURE 9.1. The energy spectrum of dark photons in FASER produced with meson production modeled by different generators (EPOS-LHC, QGSJET II-04 and SIBYLL 2.3d). Also shown is production from bremsstrahlung with a factor of two variation in the p_T cut off. The bottom panel shows the ratio between the different estimates, and the parameterisation of the uncertainty as a function of energy. A representative signal model (with $m_{A'}=50 \text{ MeV}$ and $\epsilon=3 \times 10^{-5}$) is shown.

The parameterization also envelops the uncertainty on the signal predictions due to changing the p_T -cutoff in modelling of the dark bremsstrahlung as described in Chapter V. Fig. 9.1 presents the A' energy distribution as estimated by the different generators for a representative signal model. The parameterization is checked for numerous signal models spanning the relevant phase space of the A' , and is found to be in good agreement with the envelope of the generators.

9.2 MC MODELLING

The remaining uncertainties arise from the modelling of the detector response in the MC simulation, which is used to calculate the signal yield. With the same event selection cuts applied to both data and MC simulation, a mismodelling of the simulation response produces different signal efficiencies between data and MC, which we estimate here for the various event selection requirements defined in Chapter VII.

To measure the uncertainty attributed to the Pre-shower and Trigger/Timing scintillator cuts on the A' signal, we compare the fraction of two track events that are rejected due to the scintillator cuts in both data and MC signal simulation. The scintillator efficiencies for two track events are measured to be $\geq 99.7\%$ in both data and MC signal simulation. Thus, with negligible ($<1\%$) effect on the signal yield, no corresponding uncertainty is assigned to the Trigger/Timing and Pre-shower scintillator cuts.

Similarly for the veto and vetoNu scintillator charge cuts; given that the scintillators are extremely efficient, as summarized in Table 8.1, the corresponding uncertainty on the signal yield is considered to be negligible.

The calorimeter energy scale calibration, as described in Appendix B.2, is applied to both the data and MC simulation identically. The stability of the calorimeter system across the data taking period is tested with regular calibrations using an LED pulse injected into the calorimeter modules [48]. A conservative analysis taking into account all components of the energy calibration is described in Appendix B.3, and leads to a 6% uncertainty on the difference in the calibration of the energy scale between data and MC simulation. This uncertainty is checked in data by comparing the calorimeter response to photon conversion events in both data and the Fluka MC simulation described in Chapter V. Photon conversion events are initiated by

high-energy muons, producing three track events where the muon continues through FASER and the photon decays into an electron-positron pair. After selecting for three-track events, a cut requiring the signal in the downstream preshower layer to be at least twice as large as the upstream preshower layer is used to further isolate photon conversion events from other three track events, as shown in Fig. 9.2. The cut on the ratio of the preshower layers rejects 90% of three-track events that have an $E/P < 0.5$ whilst retaining 77% of three track events with an $E/P > 0.5$; where E is the reconstructed electromagnetic energy in the calorimeter and P is the combined momentum of the two lowest momentum tracks, which are assumed to be the electron-positron pair while the highest momentum track is assumed to be the muon. Thus the E/P ratio should ideally be centered on a value of one for electromagnetic showers if both the calorimeter and track momentum are well calibrated. For the purposes of this analysis, it is not the uncertainty in the absolute energy calibration that is important, but rather the relative difference between the data and MC simulation. The relative difference in the reconstructed E/p peak position in data and MC simulation is consistent and well within the 6% uncertainty across the momentum ranges probed, as shown in Fig. 9.3.

The uncertainty due to the tracking efficiency of single tracks is assessed by comparing the relative efficiency for finding tracks in events with a single track segment in each of the three tracking stations between data and MC simulation. This yields a 1.5% uncertainty per track.

The track finding procedure is more complex when there are two closely-spaced tracks, as is the case in the dark photon signal, especially when the tracks share hits. The uncertainty due to this is assessed by overlaying the raw tracker data from two different events, each of which has a single reconstructed track. The track




FIGURE 9.2. (LEFT) The E/P vs. Invariant Mass distribution is shown for the two lowest momentum tracks in three-track events found in data. Requiring an $E/P > 0.5$ is shown to select photon conversion event candidates, whose invariant mass is zero, apart from most other three-track events that have a non-zero invariant mass. (RIGHT) The ratio of the signal in the downstream vs upstream preshower layer is shown for both photon conversion event candidates (black) and other three track events (red) in data, where a preshower ratio cut of > 2 is used to efficiently select 3-track events in data that have an $E/P > 0.5$.

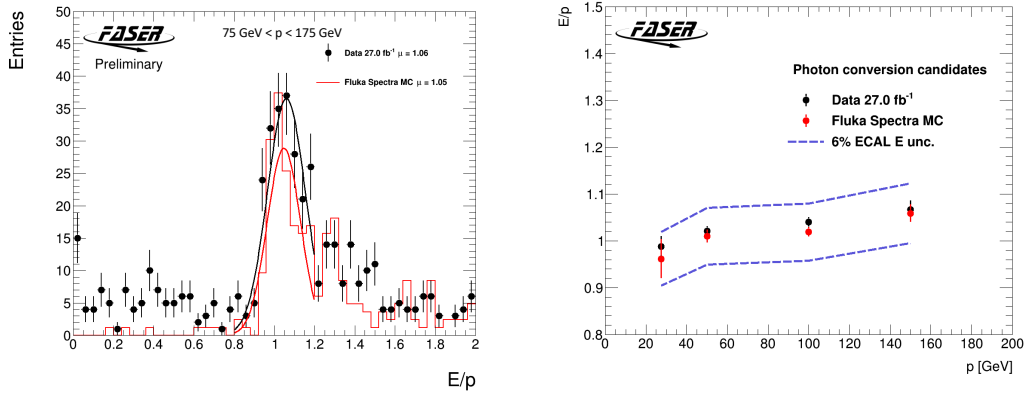


FIGURE 9.3. (LEFT) The E/P distributions are shown for three-track photon conversion events in both data (black) and Fluka MC simulation (red) for a track momentum range of 75-175 GeV. (RIGHT) The Gaussian-fitted peak position of the E/p in data and MC simulation is shown for several different momentum ranges: $20 \text{ GeV} < p < 35 \text{ GeV}$, $35 \text{ GeV} < p < 75 \text{ GeV}$, $75 \text{ GeV} < p < 125 \text{ GeV}$, $125 \text{ GeV} < p < 175 \text{ GeV}$. For all ranges measured, the E/P peak in both data and MC agree within the conservative 6% uncertainty (blue dashed line) used for the calorimeter energy calibration.

reconstruction is then re-run on the combined event, built from the two overlaid events, so that the tracking efficiency can be calculated. This is performed using both data and simulation, shown in Fig. 9.4, where the ratio of the efficiency between the two, as a function of the distance between the two tracks at their first measurements, is used to assess the uncertainty. At track separations comparable to that expected in the A' signals, the efficiency in data is up to 7% less than in MC simulation. Thus a 7% correction to the two-track tracking efficiency is applied, with a corresponding systematic uncertainty that is assumed to be the difference between the nominal and corrected efficiency. This 7% uncertainty in the two-track reconstruction efficiency is also validated using delta-rays, as described in Appendix B.3.5.

The track momentum scale and resolution uncertainty is derived by comparing the mass peak of photon conversion events in data and MC simulation. Upon comparison, both a shift or Gaussian smear of the MC track momentum by 5% were shown to more than account for the difference in the photon conversion mass peak position between data and MC, leading to a conservative uncertainty of 5% on both the track momentum scale and momentum resolution.




FIGURE 9.4. (Top) The two-track reconstruction efficiency versus track separation for overlaid tracks in both data and MC events are shown. The distribution of the separation in e^+e^- tracks of an A' sample is also shown in red with the axis on the right-hand side. (Bottom) The ratio of the overlay tracking efficiencies between MC and data is depicted.

9.3 SUMMARY OF THE SYSTEMATIC UNCERTAINTIES

Table 9.1 summarises the various sources of uncertainty on the signal, showing both the size of the individual uncertainties, and the range of the effect on the overall uncertainty of the signal yield.

Source	Value	Effect on signal yield
Signal Generator	$\frac{0.15+(E_{A'}/4\text{TeV})^3}{1+(E_{A'}/4\text{TeV})^3}$	15-65% (15-45%)
Luminosity	2.2%	2.2%
MC Statistics	$\sqrt{\sum W^2}$	1-3% (1-2%)
Track Momentum Scale	5%	< 0.5%
Track Momentum Resolution	5%	< 0.5%
Single Track Efficiency	3%	3%
Two-track Efficiency	7%	7%
Calorimeter Energy Scale	6%	0-8% (< 1%)

TABLE 9.1. Summary of the systematic uncertainties on the signal yield. For each of the sources of uncertainty, the source and size of the uncertainty is presented. The effect on the signal yield across the full signal parameter space probed is also shown. The numbers in parenthesis indicate the effect on the signals within the parameter space for which this analysis is sensitive.

CHAPTER X

RESULTS

After applying the signal selection described in Chapter VII, zero events are observed in the un-blinded data, which is compatible with the expected background of $(2.3 \pm 2.3) \times 10^{-3}$ events. Fig. 10.1 shows the calorimeter energy distribution for data and three representative signal models at different stages of the signal region selection on the veto scintillator and track information. There are events that have no veto signal and at least one track, but the calorimeter energies are well below the 500 GeV threshold; and there are no events upon further requiring two fiducial tracks.

As no significant excess of events over the background is observed, the results are used to set exclusion limits in the signal scenarios considered. The exclusion limits are made using a profile likelihood approach implemented via the `HistFitter` framework [81], and are set at 90% confidence level to allow for direct comparison with constraints from other experiments. Hypothesis tests are performed using profile likelihood test statistics [82] and the CLs method [83] to test the exclusion of new physics scenarios. For dark photons, the analysis excludes signal models in the range $\epsilon \sim 4 \times 10^{-6} - 2 \times 10^{-4}$ and $m_{A'} \sim 10 \text{ MeV} - 80 \text{ MeV}$, and provides the world-leading exclusion for scenarios in the range $\epsilon \sim 2 \times 10^{-5} - 1 \times 10^{-4}$ and $m_{A'} \sim 17 \text{ MeV} - 70 \text{ MeV}$. Fig. 10.2 shows the A' exclusion limit in the signal parameter space, where the grey regions are already excluded by experimental data from BaBar [84], E141 [85], NA48 [86], NA64 [87], Orsay [88, 89], and NuCal [67, 90], which are adapted from DarkCast [9].

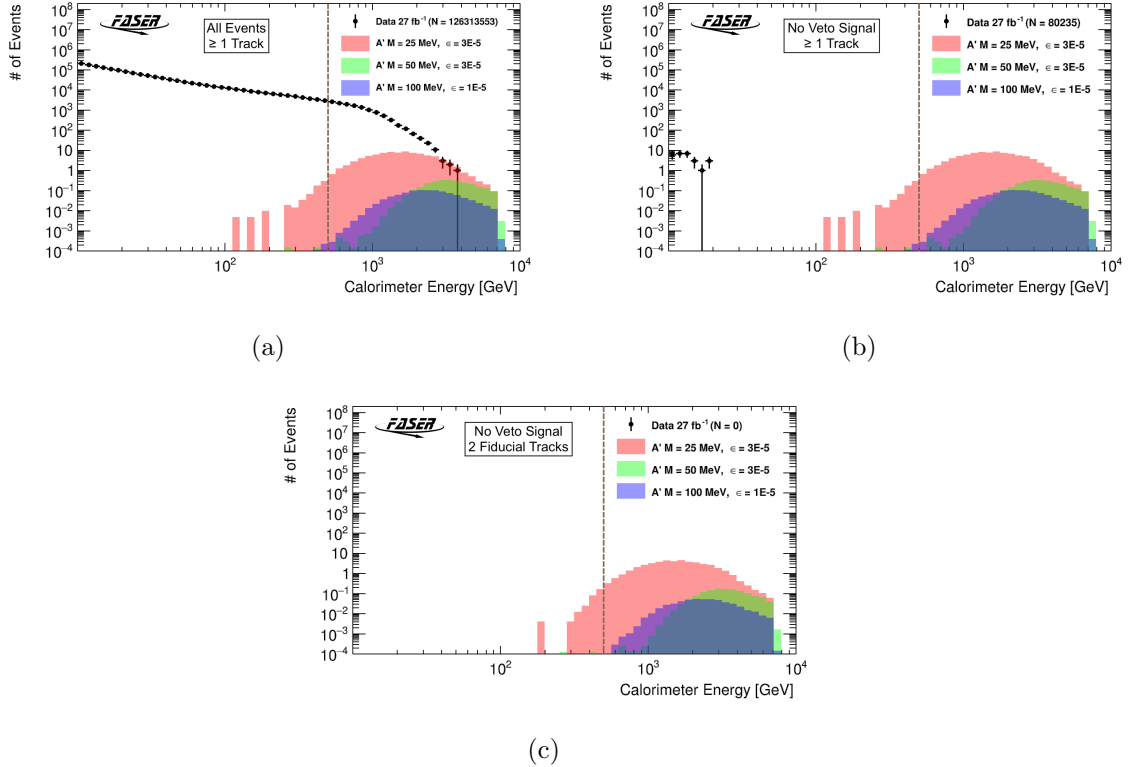


FIGURE 10.1. The calorimeter energy distribution for data and three representative MC simulated signal models are shown for (a) all events with at least one good track, (b) events that have no signal in the veto stations and at least one good track, and (c) events that have no signal in the veto stations and exactly two good fiducial tracks. The distributions and expected events from the MC samples are scaled to 27.0 fb^{-1} .

A key reason for investigating dark photons is their potential as intermediaries between the SM and a dark sector. In particular, they allow for obtaining the correct value of the dark matter relic density, $\Omega_\chi^{\text{total}} h^2 \simeq 0.12$ [91], via the thermal freeze-out mechanism. In Fig. 10.2, an example thermal relic contour is included, obtained for the scenario where the dark photons couple to a light complex scalar dark matter field χ [66]. In particular, this line assumes that the mass ratio between the dark matter candidate and the dark photon is always equal to $m_\chi/m_{A'} = 0.6$ and that the dark photon coupling constant to dark matter has a fixed value of $\alpha_D = 0.1$. This




FIGURE 10.2. 90% confidence level exclusion contours in the dark photon parameter space are shown. Regions excluded by previous experiments are shown in grey (see main text for details). The red line shows the region of parameter space that yields the correct dark matter relic density, with the assumptions discussed in the text.

mass ratio guarantees that the dark photon decays visibly into the SM species and that the dark matter primarily annihilates via $\chi\chi \rightarrow A' \rightarrow ff$. Variations of both the coupling and mass ratio in the dark sector are possible and will lead to a shift of the relic target line. Notably, in the context of this particular dark matter model, the region below the target line would have an over-abundance of dark matter and would be excluded cosmologically: FASER therefore probes a significant fraction of the cosmologically-allowed region of parameter space.

This analysis was also reinterpreted to exclude previously uncovered parameter space of a $U(1)_{B-L}$ model, and is presented in Appendix B.3.5.

CHAPTER XI

CONCLUSION

The first search for dark photons by the FASER experiment has been presented, providing a proof of principle that very low background searches for long-lived particles in the very forward region are possible at the LHC. The search applies an event selection requiring no signal in the veto scintillator systems, two good quality reconstructed charged particle tracks and more than 500 GeV of energy deposited in the calorimeter. No events are observed passing the selection, with an expected background of $(2.3 \pm 2.3) \times 10^{-3}$ events. At the 90% confidence level, FASER excludes the region of $\epsilon \sim 4 \times 10^{-6} - 2 \times 10^{-4}$ and $m_{A'}$ ~ 10 MeV $- 80$ MeV in the dark photon parameter space. This result excludes previously-viable models motivated by dark matter, where the dark photon would provide a portal between the SM and the dark sector and allow for the thermal relic abundance of dark matter to be set via the freeze-out mechanism.

APPENDIX A

PMT WAVEFORM RECONSTRUCTION AND DIGITIZATION

To prepare the calorimeter and scintillator data for analysis, we must first reconstruct the PMT waveforms in order to extract the desired information. We also want to treat the Monte Carlo (MC) simulations in the same way as the raw data, and thus have to digitize the simulation output to mimic the PMT waveforms from real data.

In this Appendix, the details of reconstructing the PMT waveforms and digitizing the MC scintillators and calorimeter modules are presented.

A.1 WAVEFORM RECONSTRUCTION

The raw data of the PMT signals, after passing through the digitizer, as discussed in Sec. 4.4.1, are waveforms 1200 ns long with 2 ns bins and a negative amplitude in ADC counts, as shown in Fig. A.1. The digitizer has a 2 V range with a 14 bit readout resulting in a 0.122 mV/ADC conversion factor.




FIGURE A.1. A typical example of a PMT raw waveform signal coming from the digitizer.

As the first step in reconstructing the waveform, we subtract off the baseline and invert the pulse. To accomplish this, we histogram the waveform ADC values that are before the expected time of the signal and then zoom in on the most common value as shown in Figure A.2. We then fit this distribution with a Gaussian, which has a two sigma fit-range around the most common value such that the occasional out-of-time signal does not bias the baseline measurement. The mean and width of the Gaussian fit are then taken to be the measured baseline mean and rms values respectively. Now we subtract the measured baseline mean from the raw waveform and then invert the signal in order to get a baseline-subtracted positive pulse that we use for the remainder of the reconstruction.

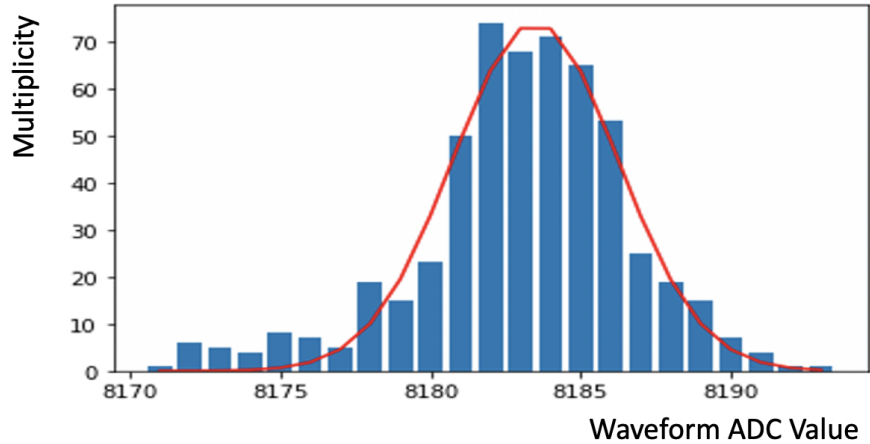


FIGURE A.2. The distribution of ADC values in a given waveform are shown. We fit the distribution around the most common value to obtain the baseline and baseline rms values.

In the next step of the waveform reconstruction chain, we isolate the signal from the rest of the waveform, which is mostly noise. Because the primary signals we are interested in are coming from relativistic particles emanating from the IP, the timing of the signals in the scintillators and calorimeter modules, with respect to the trigger, are roughly constant. Thus, to isolate the primary signal, we select a reconstruction

window that is 40 ns before and 80 ns after the expected primary signal time. Figure A.3 shows an example of this reconstruction window for two different scintillators, where you can see the window fully captures the signals from both scintillators.

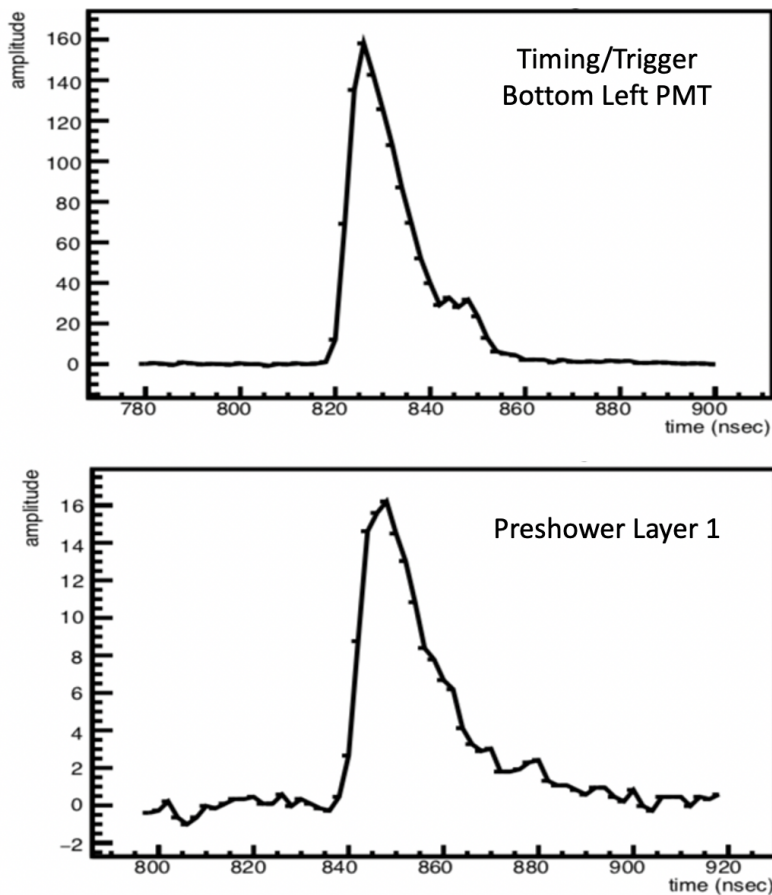


FIGURE A.3. Both figures show a typical waveform after baseline subtraction, inversion, and selection of the reconstruction window. You can see that the reconstruction window differs depending on which sub-detector is being reconstructed, where the (TOP) plot is for a Trigger scintillator and has a reconstruction window from 780 - 900 ns and the (BOTTOM) plot is for a preshower scintillator and has a reconstruction window from 800 - 920 ns.

Now that we have isolated the primary signal, we fit it with a Crystal Ball (CB) [92] function if the waveform is above threshold. The CB function is a Gaussian

with a power law tail, and is defined as:

$$f(x; \alpha, n, \bar{x}, \sigma) = N \cdot \begin{cases} \exp(-\frac{(x-\bar{x})^2}{2\sigma^2}), & \text{for } \frac{x-\bar{x}}{\sigma} > -\alpha \\ A \cdot (B - \frac{x-\bar{x}}{\sigma})^{-n}, & \text{for } \frac{x-\bar{x}}{\sigma} \leq -\alpha \end{cases} \quad (\text{A.1})$$

where $A = (\frac{n}{|\alpha|})^n \cdot \exp(-\frac{|\alpha|^2}{2})$, $B = \frac{n}{|\alpha|} - |\alpha|$, N is the amplitude, \bar{x} is the Gaussian mean, σ is the Gaussian width, n is the power law of the tail, and α is the number of standard deviations away from the mean that the power law tail takes over.

The CB function nicely fits the signals and even lets us estimate the true size of a signal that saturated the digitizer, as shown in Fig. A.4. In order to fit the saturated pulses, we disregard the waveform values which are at the maximum and fit the rest of the waveform with the Crystal Ball function. The resulting fit nicely matches the unsaturated part of the waveform and gives a reasonable estimate for the shape of the true unsaturated signal.

In the primary reconstruction window, the threshold required to try and fit the waveform is a peak ≥ 5 times the baseline rms. In order to improve the fit quality, a Gaussian fit is done first to get estimates of the pulse amplitude, width, and timing position which are then used as initial parameters in the Crystal Ball fit.

We also look for “secondary signal” regions, which are pulses that occur outside of the primary signal region. The primary signal region is always reconstructed, as it is the time range we expect the signal to be in, but secondary signals are out of time and are only reconstructed if they have a peak value greater than a threshold of 10 times the baseline rms. Due to ringing seen after the pulse in test-beam data, the reconstruction code only looks for secondary signals before the primary reconstruction window if a primary signal was seen, but if there is no primary signal




FIGURE A.4. Crystal Ball fits in blue are shown atop data points in red for a regular waveform pulse (TOP) and a saturated waveform pulse (BOTTOM).

above the threshold, then the code will look both before and after the primary signal region for secondary signals. The secondary signals are fit and processed just like a primary signal but are identified as a secondary signal by a dedicated flag in the status word.

Once the waveform is successfully fit with the Crystal Ball function, we can extract and save its parameters. We also integrate the Crystal Ball function over the reconstruction window to give us a waveform integral. Finally, a constant-fraction time measurement of the signal is obtained by extrapolating the fit to the time of the rising edge where the waveform value is 40% that of the peak; where the optimal

threshold of 40% was measured from minimizing the timing resolution in data, as shown in Fig. A.5.




FIGURE A.5. The timing resolution of the first preshower layer is shown against the constant-fraction threshold used to measure the timing, as described in the text. The other scintillator and calorimeter channels have a similar shape, and thus a 40% constant-fraction threshold was chosen to measure the pulse time as it minimizes the timing resolution. The signals used to do this measurement were from muons in TT12 collision data.

In addition to using the fit to extract values from the waveform, we also store the raw peak and raw integral of the reconstruction window, which are obtained from simply finding the maximum waveform value in the window and adding up all of the waveform values in the window, respectively.

Finally, a status word is saved for each reconstructed waveform, which has bit flags that give information about the waveform. The meaning of the status word values are shown in Table A.1.

By the end of the waveform reconstruction, we have obtained all of the parameters shown in Table A.2, which are then saved and ready to use in a physics analysis.

Status Value	Description
0	good hit
1	below threshold
2	secondary signal
4	amplitude overflow
8	find baseline failed
16	Gaussian fit failed
32	CB fit failed
64	invalid clock
128	waveform missing
256	waveform invalid

TABLE A.1. This table gives the values and descriptions of the waveform status word.

Variable	Description
localtime	time of signal measured at 40% of peak
peak	peak of CB fit
width	Gaussian width of CB fit
integral	integral of CB fit over reconstruction window
bcid_time	time difference of localtime and LHC clock edge
raw_peak	max value of raw waveform in reconstruction window
raw_integral	sum of raw waveform bins over reconstruction window
baseline_mean	measured baseline
baseline_rms	measured baseline rms
hit_status	reconstruction status word, see Table A.1
mean	mean of CB fit
alpha	alpha of CB fit
nval	power law of CB fit

TABLE A.2. All reconstructed waveform variables are shown and described.

The most common waveform variable used in the analysis, is that of the PMT charge, which is calculated from the waveform integral. The charge deposited in the PMT can be simply calculated by dividing the waveform integral (in units of mV*ns)

by the 50Ω terminal resistance of the signal cable, resulting in a charge with units of pC.

A.1.1 Quality Checks

After developing a relatively simple and robust waveform reconstruction algorithm, several quality checks were performed to ensure the algorithm was behaving as expected on the physics dataset used in the dark photon analysis.

As a first check, the mean waveform baseline, baseline RMS, and timing measurements were all plotted for each FASER run used in the dark photon analysis. As shown in Fig. A.6, the waveform baseline reconstruction is very stable, and only deviates by $< 0.01\%$ from run-to-run. Similarly, Fig. A.7 shows that the waveform noise, or baseline rms, was also quite constant, with $< 1\%$ variation in the scintillators and $< 6\%$ variation in the calorimeter modules. Although there are jumps in the calorimeter noise that are not fully understood, this is not an issue as the noise is still negligible when compared to signals from large EM showers. The mean timing of signals in the preshower and calorimeter modules show a drift of $+2$ ns in later runs, as depicted in Fig. A.8, but this is suspected to be due to the LHC clock drifting with seasonal temperature changes and is not an issue as the small drift will not effect the signal being well within the 120 ns waveform reconstruction window.

Another check was to use the waveform status word to see how often steps in the reconstruction chain failed. In the entire dataset used for the dark photon analysis, there was never an instance where waveform reconstruction failed in any way besides the occasional failed fit. The ratio of waveforms that had a failed fit, or had overflowed and saturated the digitizer, are shown in Table A.5. The calorimeter modules saw the largest fraction (1-3%) of failed fits, yet this can be attributed to the low calorimeter




FIGURE A.6. The mean reconstructed waveform baseline is plotted per FASER run used in the dark photon analysis. The baseline measurements are shown for each calorimeter and scintillator module, and are all normalized to the first run (8725).

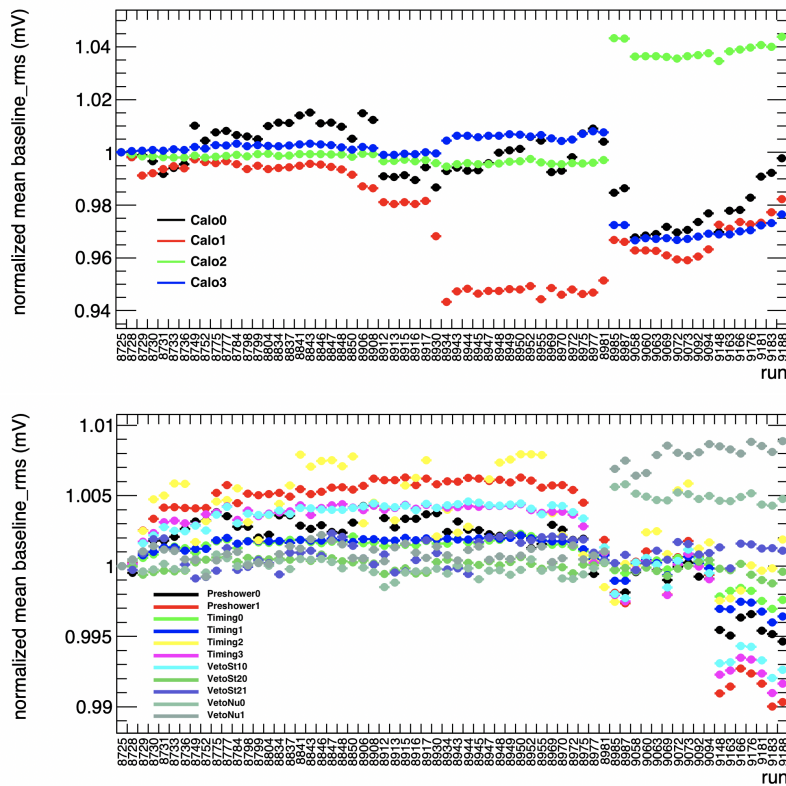


FIGURE A.7. The mean reconstructed waveform baseline rms is plotted per FASER run used in the dark photon analysis. The rms measurements are shown for each calorimeter (TOP) and scintillator (BOTTOM) module, and are all normalized to the first run (8725).




FIGURE A.8. The mean reconstructed waveform time is plotted per FASER run used in the dark photon analysis. The time measurements are shown for each calorimeter and scintillator module, and are all normalized to the first run (8725).

gain and the MIP signals in the calorimeter being dominated by noise. The waveform fits worked better on the larger signals in the scintillators, with a $< 1\%$ fit failure rate.

As a final check, the fitted measurements and the raw measurements of both the waveform peak and integral were compared, as shown in Fig. A.9. As you can see, the fit tends to overestimate the raw peak by $\sim 5\text{-}10\%$, whereas the raw and fitted integral are in much better agreement.

	Failed %	Overflow %	Gaussian Fail %	CB Fail %	Both Fits Fail %
Calo/0	2.6	none	1.5	0.9	0.18
Calo/1	1.4	none	0.71	0.6	0.084
Calo/2	2.3	none	1.4	0.78	0.14
Calo/3	2.4	none	1.4	0.79	0.16
Preshower/0	0.16	0.027	0.14	0.022	0.00011
Preshower/1	0.26	0.026	0.22	0.038	0.00016
Timing/0	0.44	none	0.3	0.11	0.029
Timing/1	0.38	none	0.29	0.064	0.028
Timing/2	0.58	none	0.43	0.11	0.041
Timing/3	0.58	none	0.43	0.11	0.042
Veto/10	0.098	3.1	0.046	0.021	0.0041
Veto/20	0.059	3.9	0.039	0.015	0.0035
Veto/21	0.055	3.4	0.035	0.014	0.003
VetoNu/0	0.06	1.4	0.022	0.017	0.0022
VetoNu/1	0.057	1.3	0.021	0.017	0.002

TABLE A.3. The ratio of waveforms that had a failed fit or an overflow are shown for the 2022 data set.

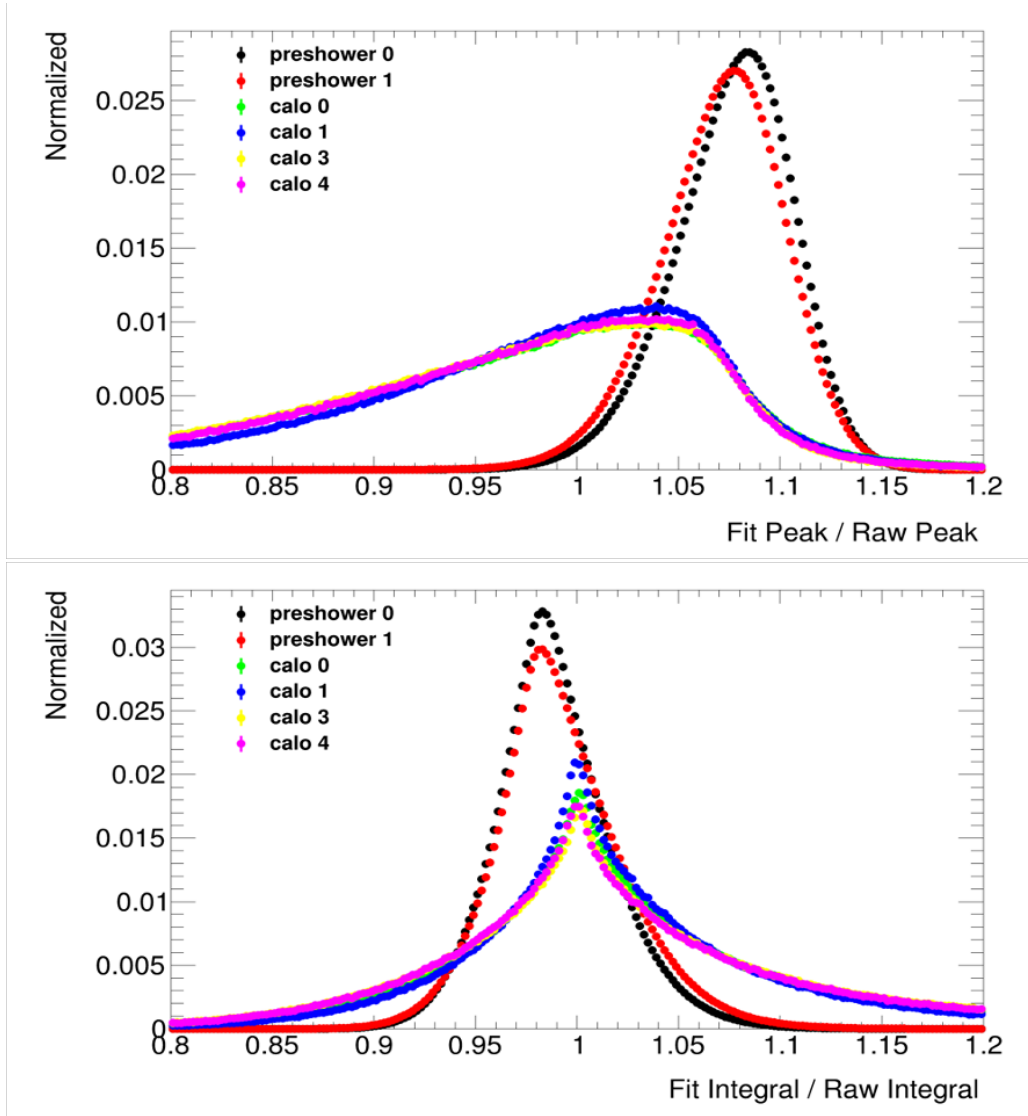


FIGURE A.9. The ratios of the fitted and raw measurements of the waveform peak (TOP) and integral (BOTTOM) are shown for the preshower and calorimeter modules. The entire 2022 data-set was used, where the ratio was calculated for each waveform that had a good fit. Each distribution is normalized such that the sum of it's bins is equal to one.

A.2 MC WAVEFORM DIGITIZATION

MC simulation produces Geant4[65, 93] hits that record the position and amount of energy deposited for each interaction step in a sensitive material, whereas real data from a scintillator module is a waveform pulse showing the PMT charge collected over time. For a direct comparison of the two, we make the simulation mimic the waveform output of a PMT. This is done via a parameterization of the shape of real data waveforms which is then scaled proportionally by the total simulated deposited energy in the scintillator or calorimeter module. As a result of the MC digitization, we can not only mimic electronic noise in the simulation, but also perform the same waveform reconstruction steps on the simulation as we do with the real data.

To find the shape of the waveform signal we want to mimic, we start by measuring an average waveform shape for each detector channel. To get an average waveform from the TI12 FASER data, we first select and normalize the waveforms by performing the following:

- select events with a through-going charged particle by requiring the coincidence trigger between the preshower and the veto layers was fired,
- require the waveform has a signal above $5\times$ the baseline rms,
- shift each waveform by its measured time such that it starts at $t=0$,
- and normalize each waveform by scaling it by the inverse of it's integral.

After averaging these normalized waveforms, we then fit the average waveforms with a CB function in order to obtain the CB parameters needed to reproduce the signal shape in MC.

With a Crystal Ball shape obtained from data, we can start to produce our MC PMT waveform by normalizing the Crystal Ball function and then discretizing it into 2 ns bins across a 1200 ns range. The timing of the digitized MC waveform is constant for a given detector channel as we do not mimic the 12 ns digitizer timing jitter seen in real data. Scaling the discretized CB function by the total deposited energy in a given sub detector maintains the proportionality of the waveform integral to the deposited energy, as depicted in Fig. A.10. Now, to get the MC waveform peak to match the ADC amplitudes seen in data, we have to scale each waveform again by a factor that is obtained by comparing the MIP integral distributions of MC and TI12 data and measuring the ratio between the two.

After this final scaling, if the digitized peak is above the maximum value of the digitizer then we mimic the waveform saturation seen in real data by setting the peak values to the maximum value of the digitizer. An example of a MC digitized waveform which mimics saturation is seen in Fig. A.11.




FIGURE A.10. The digitized MC waveforms of the second preshower layer are shown for varying initial electron energies.




FIGURE A.11. A saturated MC digitized waveform for the Trigger/Trigger layer is shown.

For the calorimeter modules, the deposited energy to ADC scaling parameter changes when we vary the PMT HV or the presence of the optical light filter. This is tracked via a database that stores the four combinations of settings between having the HV at low or high gain and having the light filter in or out. The four calibration calorimeter scaling values can be seen in Table A.4.

Calo Configuration	Calo Digi CB Scale Factor
Low-HV no-Filter	5.1
Low-HV with-Filter	0.51
High-HV no-Filter	153
High-HV with-Filter	15.3

TABLE A.4. The digitization scale factor for the Calorimeter are shown for the four different run configurations.

The next step of the digitization process is to produce a noisy waveform baseline that matches the data. The baseline rms distributions measured from the data show that the noise is very similar for all channels and is around 3 ADC counts. The

noise is dominated by the digitizer noise and thus does not change significantly with different HV settings. The MC baseline mean is also obtained by measuring the real data and is then smeared across a 1200 ns window using a Gaussian whose width is 3 ADC. At this point, we have all of the parameters needed to produce waveforms from each detector type, where the CB waveform and baseline parameters can be seen in Table A.5.

	Calo	VetoNu	Veto	Timing	Preshower
Baseline mean (ADC)	15650	15650	15650	15650	15650
Baseline rms (ADC)	3	3	3	3	3
CB scale factor	Table A.4	7040	6840	4240	400
CB mean	820	815	815	815	846
CB sigma	3.67	5.0	3.35	3.21	4.0
CB alpha	-0.32	-0.28	-3.2	-4.24	3.2
CB n	1000	1000	9.0	6.14	1000

TABLE A.5. The MC digitization parameters for each detector type are shown.

Finally, we subtract the scaled CB signal pulse from the noisy baseline and then convert the waveform values to integers to obtain a discretized negative pulse that mimics the raw waveforms in data. Assigning this digitized waveform its corresponding detector identifier and saving it as a “Raw-Waveform” allows us to send it through the waveform reconstruction chain, just as we would a waveform from real data.

Several quality checks were done to ensure the digitized MC waveforms accurately mimic those seen in FASER physics data. Figures A.12, A.14, and A.13 all show comparison plots overlaying different distributions from reconstructed FASER data and a reconstructed digitized MC muon sample. Although the comparison plots are only shown for a single calorimeter module and preshower layer, the other calorimeter

and scintillator modules all showed similar agreements between digitized MC and data. Figure A.12 shows good agreement between the comparisons of the average reconstructed waveform shapes, where the only significant discrepancy is that the digitized MC waveform has a slightly larger amplitude because the CB fit often overestimates the peak value, as discussed in Appendix A.1.1. Fig. A.13 shows good agreement in the reconstructed baseline noise distributions, where an exact match is not necessary as the noise is small compared to the expected signals. And lastly, Fig. A.14 shows the two distributions of the reconstructed charge being centered right on top of each other, yet the widths of the distributions differ as the MC does not mimic the positional non-uniformities of light collection that the real calorimeter and scintillator modules experience.




FIGURE A.12. Comparisons of the average normalized PMT waveforms are shown between FASER physics data and a digitized MC muon sample for a calorimeter module (LEFT) and a preshower layer (RIGHT).




FIGURE A.13. Comparisons of the reconstructed waveform baseline RMS are shown between FASER physics data and a digitized MC muon sample for a calorimeter module (LEFT) and a preshower layer (RIGHT).




FIGURE A.14. Comparisons of the reconstructed PMT charge distributions are shown between FASER physics data and a digitized MC muon sample for a calorimeter module (TOP) and a preshower layer (BOTTOM).

APPENDIX B

CALORIMETER ENERGY CALIBRATION

The FASER calorimeter consists of four LHCb outer ECAL modules, as discussed in Sec. 4.3. To get an energy measurement from the calorimeter, we use a test-beam to measure the calorimeter signal from electrons with known energies. Normalizing the calorimeter response with that of the MIP signal allows us to then estimate the initial electromagnetic energy of a particle outside of the test-beam, namely in the TI12 tunnel where FASER is installed.

In this Appendix, I give an overview of the test-beam setup, followed by detailed accounts of both the calorimeter energy calibration and the uncertainty measurement for comparing the calibrated calorimeter energies between data and MC.

B.1 TEST-BEAM

The test-beam measurements were carried out in July of 2021 at the H2 beam line [94], located in the North Area experimental hall of the SPS at CERN. This beam-line provided us with 85–150 GeV muons, 200 GeV pions and electrons with an energy between 5 GeV and 300 GeV. The purity of the electron beam varied with energy, but was estimated to be above 99% for most energies. The electron energy spread was about 0.3% for most beam energies, though data was taken with looser beam-collimator settings for energies above 150 GeV to allow for a higher particle rate, but resulted in a larger estimated energy spread of 0.85–0.90%. For the energy calibration used in this analysis, electrons with energies above 150 GeV were not used due to the reduced statistics and larger energy spread.

The detector setup deployed for test-beam measurements used a subset of the detector systems from the FASER experiment, described in Sec. 3.2. The detector setup is illustrated in Fig. B.1 and a picture of the final setup in the H2 beam line is shown in Fig. B.2. The two VetoNu scintillators were placed furthest upstream and were used to trigger events. A three layer FASER tracking station was placed just downstream of the VetoNu scintillators, which was used to measure the position and trajectory of charged particles. The remaining downstream detector systems were the preshower station followed by the calorimeter, where two additional spare calorimeter modules were used such that there were 6 total calorimeter modules. The full setup was located on a large, movable scissor table which could move both vertically and horizontally to accurately change the position of the detector with respect to the beam line and thus illuminate different parts of the setup.




FIGURE B.1. A diagram of test-beam setup, where the beam-line is coming in from the left.

The test-beam data was recorded using the same FASER digitizer, TRB, and TLB that were described in Sec. 4.4.1. Similarly, the reconstruction of the data was also done in the same way as described in Sec. 4.4.2, although the single tracking station




FIGURE B.2. A photo of test-beam setup in the H2 beam line. During operation, the detector apparatus would be raised such that it is at the same height as the beam-line.

in the test-beam setup only provides track segments and not fully reconstructed multi-station tracks.

Apart from the calorimeter energy calibration, the test-beam data was also used to characterize the response of the different sub-detectors and compared with MC simulation in order to improve the simulated geometry description [95].

B.2 ENERGY CALIBRATION

The first step in the calibration of a calorimeter module is to compare the signal PMT charge Q_{signal} to the most probable value (MPV) of the MIP PMT charge distribution Q_{μ} . Unfortunately, while we are collecting physics data, we can not see the small MIP signal as we install a light filter and use a low-HV for the PMT such

that we are not saturated by EM showers in the TeV range. In order to know how large our signal charge is with respect to the MIP signal, N_{MIP} , one has to measure the MIP signal at a higher PMT HV setting and then apply a correction, labeled *GainRatio*, to extrapolate the PMT charge of the physics signal from the low-HV to the high-HV settings of the PMT:

$$N_{MIP} = \frac{Q_{signal} \times GainRatio}{Q_{\mu}} \quad (\text{B.1})$$

The process is similar for calibrating both the TI12 data and the MC samples, with the main difference being that the *GainRatio* for the TI12 data is obtained by comparing the PMT HV gains, whereas the *GainRatio* for the MC is obtained by comparing the different digitization scaling parameters between the sample being analyzed and the muon sample used to measure Q_{μ} .

Once we know how large the signal charge is with respect to that from a MIP, we can use 2021 test-beam data to estimate the initial EM energy (E_{EM}) of the particle entering the calorimeter via:

$$E_{EM} = N_{MIP} \times \frac{E_{TB}}{\bar{N}_{MIP}^{TB}}, \quad (\text{B.2})$$

where E_{TB} is the energy of an electron test-beam and \bar{N}_{MIP}^{TB} is the average size of the signals from that test-beam with respect to the MIP signal.

The remainder of this section will give a detailed description of the steps in the FASER calorimeter energy calibration.

B.2.1 TI12 MIP Normalization

To measure Q_μ in TI12 FASER data, we start by selecting clean single muon events in each calorimeter module. This is accomplished by requiring that the veto and preshower triggers fired in coincidence, there is one and only one reconstructed track and it has measurements in all three tracking stations, and that the track is within distinct regions of the calorimeter module area when extrapolated to the face of the calorimeter. The regions of the calorimeter module area used to measure Q_μ are shown for a given calorimeter module in Fig. B.3. Comparing the MIP distributions from the three different cuts show that the MPVs do not change by more than 1%. As you can see, there is a hot spot in the calorimeter where the MIP passes through the PMT and leaves extra charge, which is why we exclude this region in our track position cuts. The track position cuts also include some area of the magnet material at $R \geq 100$ mm, but this does not seem to have a large effect on the measured MPV as the different position cuts still agree to within 1%.

After using the event selection to get a clean sample of MIPs, we then fit the MIP charge distribution in order to measure the MPV of the distribution.

Because the waveform noise can be significant enough to smear the MIP charge distribution in the calorimeter and bias the MPV, correcting for the noise is important. For this reason, the raw integral of the waveform is used for producing the MIP charge distribution, as correcting for the noise in the raw integral measurement is much more straight forward than correcting for the noise in the fitted integral measurement. An example of a uniform distribution smearing an asymmetric distribution and changing the MPV can be seen in Ref. [96]. Where, in our case, it is the uniform Gaussian noise distribution that smears the asymmetrical Landau MIP charge distribution. In order to account for this effect, we fit the distributions with

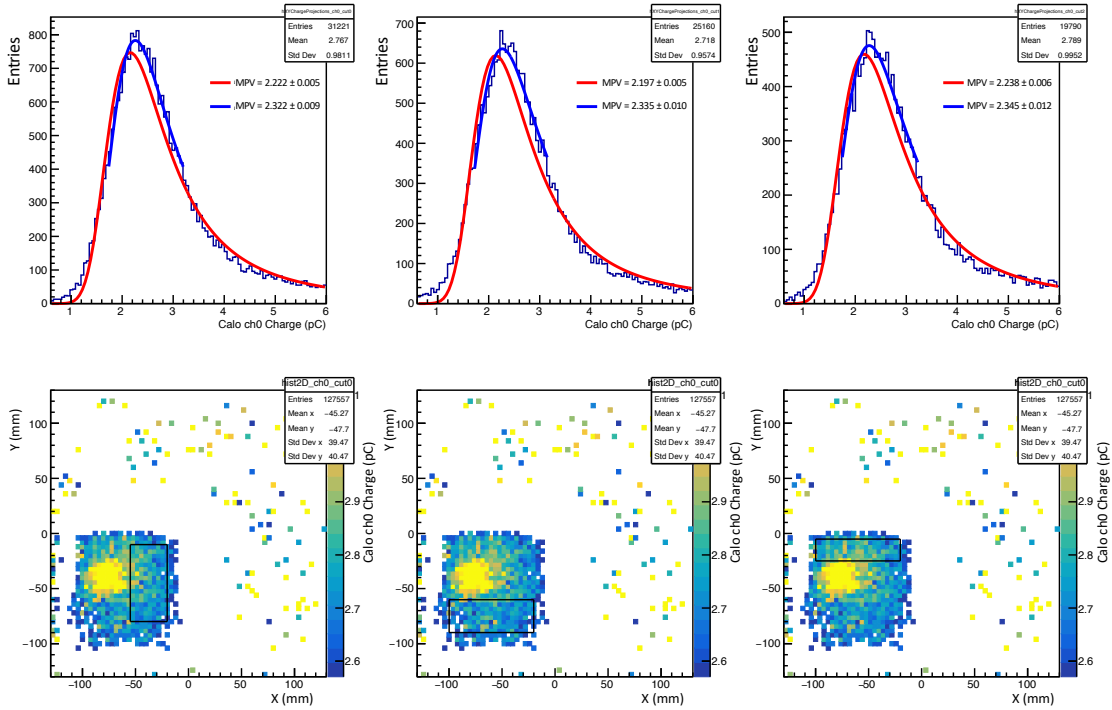


FIGURE B.3. The above plots show the calorimeter channel-0 MIP charge distributions (TOP) and the charge vs position (BOTTOM) for three different track position cuts. The cuts are depicted by the black rectangle in the bottom scatter plots. The distribution above each scatter plot corresponds to the charge distribution using that cut. The bright yellow spot in the bottom plots depict the area where the MIP goes through the PMT and produces a larger signal. The data is from FASER runs 8715 and 8717, where the calorimeter is at high gain and has the light filters installed. The charge distributions are fit with a Landau+Gaussian convolution function, described in the text, over the whole histogram range (Red) and also over -0.5 to $+1.0$ standard deviations around the MPV (Blue).

a Landau convoluted with a Gaussian whose width is equal to the waveform noise rms measured from charge distributions of randomly triggered events. Because the noise in all of the PMT waveforms is dominated by the digitizer noise, the measured noise in the charge distributions was similar for all calorimeter modules, with values ~ 0.14 pC. We then use the Landau MPV fit parameter of the convoluted fit to back out the true Landau MPV without the noise smearing effect. This was tested on a toy model, where a landau distribution was smeared with a Gaussian of varying

width and the resulting distribution was fit with both a regular landau function and a landau+Gaussian convolution function. The MPV measurements from the two fits on the toy model were compared to the MPV of the original landau distribution, as shown in Fig. B.4, where the landau+Gaussian convolution fit was able to retrieve the MPV of the original landau distribution.




FIGURE B.4. The results are shown for comparing the original MPV of a toy model with the MPV from two different fits. The toy model is a landau distribution with a MPV = 1.45 pC and width = 0.27 pC, that is then smeared by a Gaussian of varying width. The simple landau fit shows the bias in the MPV due to the Gaussian noise, whereas the landau+Gaussian convolution fit described in the text allows us to measure the MPV of the original Landau distribution regardless of the noise.

To improve the convergence of the fit, the charge distributions are first fit with a pure Landau function, whose parameters are then used as initial parameters of the Landau+Gaussian convolution function fit over the full range of the histogram. We finally zoom in on the MPV and fit a region of -0.5 to +1.0 standard deviations around the MPV with a Landau+Gaussian convolution function again, where the MPV of this final fit is the MIP charge MPV (Q_μ) that is used in the energy calibration.

The MIP charge distribution and the fits used to obtain Q_μ for each calorimeter module are shown in Fig. B.5.




FIGURE B.5. The above four plots show the MIP distributions used for the calibration of each TI12 calorimeter module. The data is from FASER runs 8715 and 8717, where the calorimeter is at high gain and has the light filters installed. The distributions are fit with a Landau+Gaussian convolution function over the whole histogram range (Red) and also over -0.5 to $+1.0$ standard deviations around the MPV (Blue). The Gaussian width used in the fit function was measured

B.2.2 PMT HV Gain Curves

Because Q_μ can only be measured in special runs with high-gain PMT settings, PMT HV gain curves are used to extrapolate the signals seen during low-gain physics running to the high-gain settings that were used to measure Q_μ .

The PMT HV gain curves used in the calorimeter calibration were obtained from analyzing 2021 test-beam data, where calorimeter PMT charge distributions resulting from a 75 GeV electron beam were measured with varying calorimeter PMT HV settings. To select clean electron events in a given calorimeter module, the following event selection was applied to the test-beam data:

- triggered by coincidence of the two VetoNu trigger layers,
- a signal greater than 10x the noise rms is seen in both preshower layers,
- and one and only one track segment exists which is linearly extrapolated to the calorimeter and is well within the transverse area of the module.

The PMT charge distributions are then fit with a CB function and the mean of the fit is plotted against the PMT HV. For easier comparison of the HV gain curves, the gain is scaled by $10^{-6} \times (1700/V)^{6.6}$ to remove the expected first-order HV dependence that comes from the electron multiplication in the PMT's dynodes. The curves are then fit with a 5 degree polynomial as seen in Fig. B.6 and the fits are stored in a database. The plots in Fig. B.6 also show gain curves measured with LED signals, where LED pulses were periodically flashing into the calorimeter modules throughout the time that the 75 GeV electron beam data was being collected. A designated LED trigger was implemented that allowed for the LED events to be easily selected apart from the electron test-beam events. The PMT charge distributions from the LED

events were then fit with a Gaussian, whose mean was plotted against the HV in order to create the HV gain curve. As you can see, the two ways of measuring the HV curves show very little difference at the high-HV value of 1400 V and also at the low-HV value of 850 V which are used during physics running.

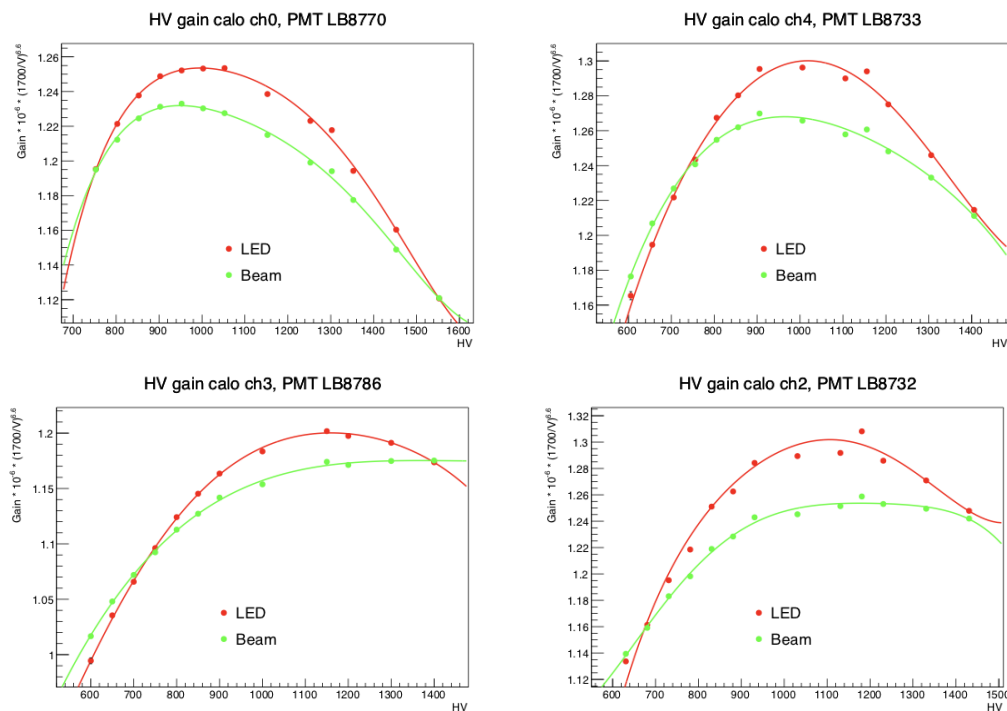


FIGURE B.6. PMT HV gain curves, described in the text, are shown for the four PMT's used in the TI12 Calorimeter. The HV gain curves were measured with a 75 GeV electron beam (Green) and an LED signal (Red) during the 2021 test-beam period. The data points show the mean PMT charge measurement vs PMT HV and the curves show a 5-degree polynomial fit of the data.

To utilize these curves and obtain a ratio for extrapolating a signal to the HV settings that were in place when recording the MIP PMT charge distribution, we use:

$$GainRatio_{data} = \frac{f(HV_{\mu})}{f(HV_{signal})} \left(\frac{HV_{\mu}}{HV_{signal}} \right)^{6.6} \quad (B.3)$$

where $f(V)$ is the HV gain curve fit, HV_μ is the PMT HV setting used when measuring Q_μ , and HV_{signal} is the PMT HV setting used when recording the data that is being calibrated.

B.2.3 MC MIP Normalization

Similarly to the calorimeter calibration in the TI12 physics data, we also calibrate the MC samples by comparing the signals to the MIP scale, Q_μ , which is measured from a 100 GeV MC muon sample.

For measuring the MC MIP charge distribution, we use the same event selection and cuts as used for the TI12 data, except we do not require the veto and preshower coincidence trigger as we don't have trigger information in the MC samples. After the event selection, the MIP charge distribution is fit in the same way as the data in order to extract the MPV. The MC MIP MPV of the charge was found to be 2.14 pC for each calorimeter module, as they are identical in the simulation.

The MC muon sample used for the measurement of Q_μ was digitized such that it mimics the signal size of TI12 MIP data when the PMT HV is at high gain. All other MC samples in this analysis, such as the dark photon signal samples, were digitized to mimic the data taken with a low-HV PMT setting. In order to calibrate these low-HV MC samples to the MIP scale, we do not need to use HV gain curves to extrapolate from low-HV to high-HV PMT settings as we did in data, but rather we simply scale the MC PMT charge by a factor of 30, which is the ratio of the waveform scale parameters used in the digitization of the low-HV and high-HV PMT samples.

B.2.4 EM Energy Calibration

At this point, the calibration steps tell us how large the signal is compared to a MIP in both MC and TI12 data via Eq. (B.1), and we can now simply scale this value to get a rough estimate of the initial EM energy following Eq. (B.2). For this, we turn to the 2021 test-beam data to compare calorimeter signals from electrons of known energies to signals from muons. First, we calibrate the test-beam data following the same calibration steps described previously, such that the calorimeter signals are normalized to the MIP scale and can be expressed as N_{MIP}^{TB} . We fit the N_{MIP}^{TB} distributions for each electron test-beam energy with a CB function, and take the mean of the CB fit to be \bar{N}_{MIP}^{TB} . Plotting the electron test-beam energy divided by \bar{N}_{MIP}^{TB} for various electron test-beam energies, as depicted in Fig. B.7, shows that there is an energy dependence to $E_{TB}/\bar{N}_{MIP}^{TB}$. If there were no energy dependence, then the value of $E_{TB}/\bar{N}_{MIP}^{TB}$ would tell us the initial energy of an electron that produces a signal in the calorimeter of equivalent size to the signal from a MIP. This is not the case though, because the ratio of the EM shower lost in the preshower is dependent upon the electrons initial energy due to the depth of an EM shower scaling with $\log(E)$. There could also be an energy dependence due to punch-through at high energies causing losses out the back of the calorimeter, but simulations have only shown a small, 1-2%, effect when the energies of the electron are in the TeV range. The current calibration does not try to correct for either the preshower or punch-through effects and instead simply uses a constant value of $E_{TB}/\bar{N}_{MIP}^{TB} = 330$ MeV, which is the value LHCb measured in a test-beam with no material in front of the calorimeter modules [97]. In this case, the initial EM energy can be

estimated in both data and MC via:

$$E_{EM} = N_{MIP} * 330 \text{ MeV}. \quad (\text{B.4})$$

This rough estimate of E_{EM} could be improved, but an accurate calibration to the absolute energy scale is not important for this dark photon analysis, and rather we only care about the agreement between the energy calibrations of the MC and data.




FIGURE B.7. The initial test-beam energy (E_{TB}) divided by the average MIP-normalized signal size \bar{N}_{MIP}^{TB} in calorimeter module 0 is shown across electron beam energies ranging from 5 to 300 GeV. The dashed blue line shows the 0.33 GeV value used in the simple estimate of E_{EM} .

B.2.5 Correction to MC Calibrated Energy

In performing the same calibration procedure as described above on the 2021 test-beam data and MC, it was noticed that the calibrated energy of the MC was consistently less than the calibrated energy of the test-beam data across all beam

energies studied. There are several reasons to think that the MC is not perfect, for example we know we mis-modelled the density of the thin light-reflecting Tyvek layers in the Calorimeter. To quickly correct the MC, we can scale it's calibrated energy up to match the calibrated energies as seen in the test-beam data.

A MC correction factor of +8.8% was obtained from comparing the average calibrated energy in each of the six test-beam calorimeter modules with the MC calibrated energy at 100 GeV. After scaling up the MC calibrated energy by this correction factor, there is much better agreement between the data and MC, as shown in Fig. B.8. Yet, the dependence of the beam energy upon the calibrated energy still differs between the MC and data.

Because the MC geometry of the detector systems is the same for the test-beam and TI12 simulations, we apply the same 8.8% correction to the TI12 MC calibrated energies as well.

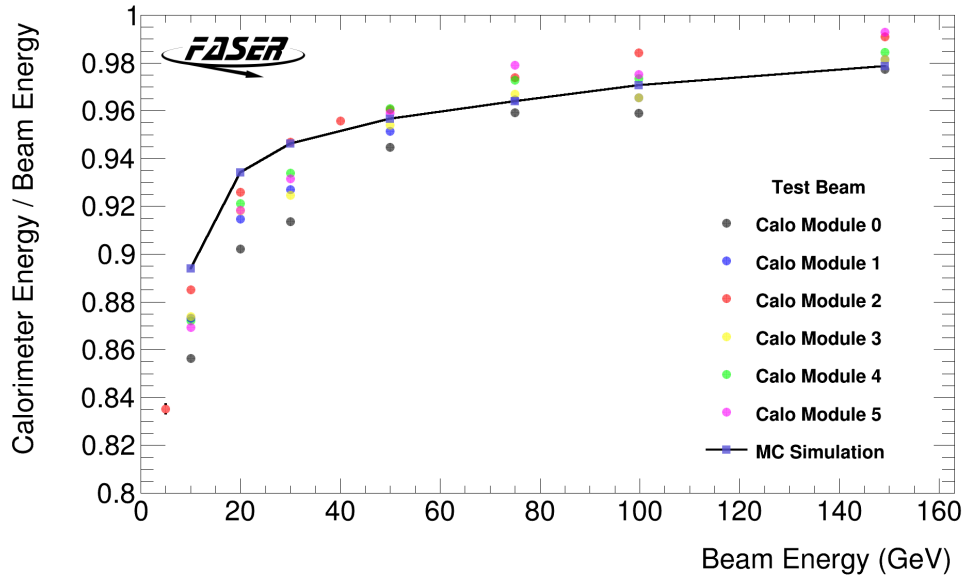


FIGURE B.8. The normalized calibrated energies for the 2021 test-beam data and MC are shown. Each color of the dots represents the calibrated energy at one of the six calorimeter modules in the test-beam setup, whereas the squares with the line represent the corrected MC calibrated energy.

B.3 UNCERTAINTY IN COMPARING DATA AND MC ENERGIES

For a physics analysis like the dark photon search, a calorimeter energy cut is applied to separate the large signal EM showers from the low energy backgrounds. The effect of that cut on the sensitivity of the analysis is determined in MC by applying the same cut to the MC calibrated energy of the signal sample. Thus, the uncertainty due to the calorimeter energy cut comes from the relative uncertainty of how well the calibrated energies agree between MC and data, and it is not the uncertainty on the absolute energy calibration.

In measuring this uncertainty for comparing the TI12 data and MC energy calibrations, I have split the uncertainties into four main groups:

1. Uncertainties in TI12 data calibration
 - MIP MPV fit uncertainty
 - HV gain extrapolation uncertainty
 - PMT drift over time uncertainty
2. Uncertainties in TI12 MC calibration
 - MIP MPV fit uncertainty
3. Uncertainties in correction on MC calibrated energy
 - Test-beam data calibration uncertainty
 - Test-beam MC calibration uncertainty
 - Uncertainty due to different energy dependence in test-beam MC and data
4. Uncertainties from other differences in MC and data

- Uncertainty from differences in edge loss

Where we are comparing the calibration for data and MC, so the overall uncertainty includes the individual uncertainty in the calibration of the data and the individual uncertainty in the calibration of the MC. Furthermore, since we correct the MC calibration using the test-beam, there is also an uncertainty associated with that correction. Finally, non-calibration related differences in the MC and data also contribute an uncertainty to the overall comparison of the two.

All of the relevant uncertainties and how they were conservatively estimated will be described in the remainder of this section.

B.3.1 Uncertainties in TI12 MIP Normalization

There are several factors that go into normalizing the TI12 data to the MIP scale and that can contribute to the uncertainty in the calibration.

The first uncertainty we will mention is that in fitting and extracting the MPV of the MIP PMT charge distributions, as described in section B.2.1. To measure this uncertainty, we can use the fit uncertainty of the MPV measurement in each calorimeter MIP PMT charge distribution as shown in Figure B.5. The uncertainties are small at 0.39%, 0.34%, 0.38%, and 0.37% for calorimeter modules 0, 1, 2, and 3 respectively. Adding the errors in quadrature gives a total uncertainty of 0.74% due to the statistical uncertainty of measuring the MPV. Another way to show that this uncertainty is indeed small is to compare the MIP MPV measurements from three different regions of each calorimeter module, as depicted in Figure B.3. Taking the standard deviation of the MPV measurements from the three different regions gives uncertainties of 0.49%, 0.93%, 0.11%, and 1.6% for calorimeter modules 0, 1, 2, and 3 respectively. Adding the errors in quadrature gives a total uncertainty of 1.90%

due to the MIP MPV measurement varying for different regions of the calorimeter module. This second way of measuring the uncertainty in the MPV PMT charge measurement is likely an overestimate as it is dominated by a few of the areas being very near the edges of the calorimeter, where there could be significant loss out the side which biases the MIP PMT charge distributions. Regardless, we take the conservative approach and use the larger uncertainty measurement of 1.90% in the total uncertainty calculation.

Another area that contributes to the uncertainty of the TI12 data energy calibration is that of the HV gain extrapolation, described in section B.2.2. For this uncertainty measurement, we compared the gain extrapolations from high-HV to low-HV (approximately 1430V to 850V, module dependent) for the two different ways of measuring the HV gain curves, as shown in Fig. B.6. Taking the difference in the gain extrapolation from the two curves gives uncertainty measurements of 0.16%, 1.00%, 2.36%, and 2.18% for calorimeter modules 0, 1, 2, and 3 respectively. Thus a combined uncertainty of 3.37% is attributed to the HV gain extrapolation.

The last significant uncertainty contributing to the TI12 data energy calibration is the possibility of the PMT response changing with time. Applying steady LED light signals to the PMTs over the course of two months, as depicted in Fig. B.9, shows that the PMTs are quite stable and the response of the PMTs mostly change in unison, due to the LED intensity shifting with temperature. The spread of the normalized LED PMT charge over time depicts the drift of the PMT response, which is $<1\%$. Another way of measuring the PMT drift over time is to compare the MIP MPV charge measurements obtained from two different time periods. For this, we use the two different times in 2022 that we ran at high gain in order to see the MIP signal, which were on September 30th (runs 8715 and 8717) and October 25th (runs

8932 and 8933), and calculate the difference in the MIP MPV to get uncertainties of 0.66%, 1.09%, 0.11%, and 0.68% for calorimeter modules 0, 1, 2, and 3 respectively. The combined uncertainty of 1.45% is used even though it is likely an overestimate as it includes the uncertainties of fits in measuring the MIP MPVs.




FIGURE B.9. The mean calorimeter PMT charge from LED pulses are shown normalized to the first measurement and spanning a two month period. The trend of each PMT response seems to vary in unison with the other calorimeter modules, depicting the LED intensity changing with temperature. The spread of the PMT responses over time depicts the PMT drift, which is $< 1\%$ over this time period.

B.3.2 Uncertainties in MC MIP Normalization

The uncertainty attributed to the normalization of the MC to the MIP scale, described in section B.2.3, is much less than that of data because the only measurement that goes into the MC normalization is that of the MIP MPV.

The uncertainty in fitting and extracting the MPV of the MC MIP PMT charge distributions is estimated via differences in the MPV measurements from the four calorimeter modules. In MC, the four calorimeter modules are identical and thus we expect them to have the same MIP PMT charge distributions. An uncertainty of 1.16% in the MPV extraction was measured from the standard deviation of the four separate MPV measurements for each calorimeter module. The MC MIP charge distributions had relatively low statistics and it is expected that this uncertainty should go down considerably if the calibration was repeated on a larger statistics sample.

B.3.3 Uncertainties in Correcting MC Calibrated Energy

The correction on the MC calibrated energy, as described in section B.2.5, was obtained from comparing the calibrated energies of a 100 GeV electron beam in the test-beam data and MC. Thus the uncertainties associated with this correction include the uncertainties in the calibration of both the test-beam data and the test-beam MC. In addition, because the test-beam data and MC showed different energy dependencies, there is an additional uncertainty in the correction if we are looking at energies other than 100 GeV.

For the uncertainty in the calibration of the test-beam data, we do not have to add up the individual uncertainties that go into the calibration like we did for the calibration of the TI12 data. Instead, in the test-beam, we have the measured calibrated energy from distinct energy electron samples in each of the six test-beam calorimeter modules which can be compared to one another in order to achieve an uncertainty on the average calibrated energy, which is what was used when calculating the MC correction factor. The standard deviation of the calibrated

energy in the six different calorimeter modules at an EM energy of 100 GeV gives an uncertainty of 0.74%.

As for the uncertainty in the calibration of the test-beam MC, we have to take into account that the calibration was done on old MC samples that were not digitized to look like real raw data waveforms. Instead, the test-beam MC samples used in measuring the correction factor were simply calibrated to the MIP scale by comparing the Geant4 deposited energies of the electron samples to the MPV of a Landau fit on the Geant4 deposited energy distribution of a 150 GeV muon sample. Thus there is an uncertainty coming from applying the correction factor, obtained from comparing un-digitized MC to data in the test-beam, to the digitized TI12 MC. To get an estimate of this uncertainty due to the test-beam MC not being digitized, we take the difference in the calibrated energies from a 100 GeV electron sample in both the old test-beam MC and the digitized TI12 MC, which results in a 2.35% uncertainty.

Lastly, since there is disagreement in the energy dependence of the test-beam data and MC calibrated energies, and we measured the correction at 100 GeV but will be using a calibrated energy cut of 500 GeV in the dark photon analysis, we must account for the uncertainty due to the disagreement in the calibrated energies between test-beam data and MC at 500 GeV. We only have test-beam data with tight beam-collimator settings (low beam energy spread) up to electron energies of 150 GeV for all six calorimeter modules, and thus need to extrapolate the calibrated energies of data and MC to 500 GeV in order to compare the two. Instead of extrapolating the MC, we could of course simulate a 500 GeV electron sample, but this has yet too be done. To extrapolate the calibrated energies of data and MC to 500 GeV, we linearly fit both the corrected MC test-beam calibrated energies and the average test-beam data calibrated energies, obtained from averaging the calibrated energy of

each of the six test-beam calorimeter modules, as seen in Fig. B.10. The simple linear extrapolation to 500 GeV provides a rough overestimate of the difference between MC and data as both should be approaching an asymptote and their rate of divergence is expected to decrease with higher energies. Using this conservative approximation of the difference in test-beam data and MC at 500 GeV gives an uncertainty of 2.46%.




FIGURE B.10. This figure depicts the extrapolation done on the test-beam data and MC calibrated energies to estimate the difference at 500 GeV. The average data points were obtained from averaging the calibrated energy of each of the six test-beam calorimeter modules. The MC calibrated energies here include the 8.8% correction factor used to make the MC and data agree at an EM energy of 100 GeV. Both data and MC were linearly fit using the range from 50 to 150 GeV, then this fit was used to extrapolate the calibrated energies to 500 GeV, resulting in an approximated difference of 2.46% at 500 GeV.

B.3.4 Uncertainties from other differences in MC and data

Apart from the factors that go into the energy calibration, other differences in MC and data can also contribute to the uncertainty when comparing the two. One such case is the lateral shower development in the calorimeter and the loss out the

edges. These positional dependencies can be compared in test-beam data and MC as shown in Fig. B.11, where there is little difference in data and MC up until 20 mm from the edge, at which point small inaccuracies in the MC calorimeter positions can produce large discrepancies in MC and data when right at the edge.

For the dark photon analysis, we require the track to be within the magnetic aperture and thus the particles should never get closer than 20 mm to the edge of a calorimeter module. To measure an uncertainty due to the different positional dependencies in MC and data, we use the most extreme differences measured in the test-beam when further than 20mm from the edge, and thus obtain an uncertainty of 2.5%.

B.3.5 Total Uncertainty

The previously mentioned uncertainties are summarized in Table B.1, where the combined uncertainty in comparing the calibrated energies in data and MC at 500 GeV is measured to be 6%.

This conservative uncertainty measurement was validated by comparing the E/p distributions of photon conversion events in both data and MC, as discussed in Sec. 9.2.

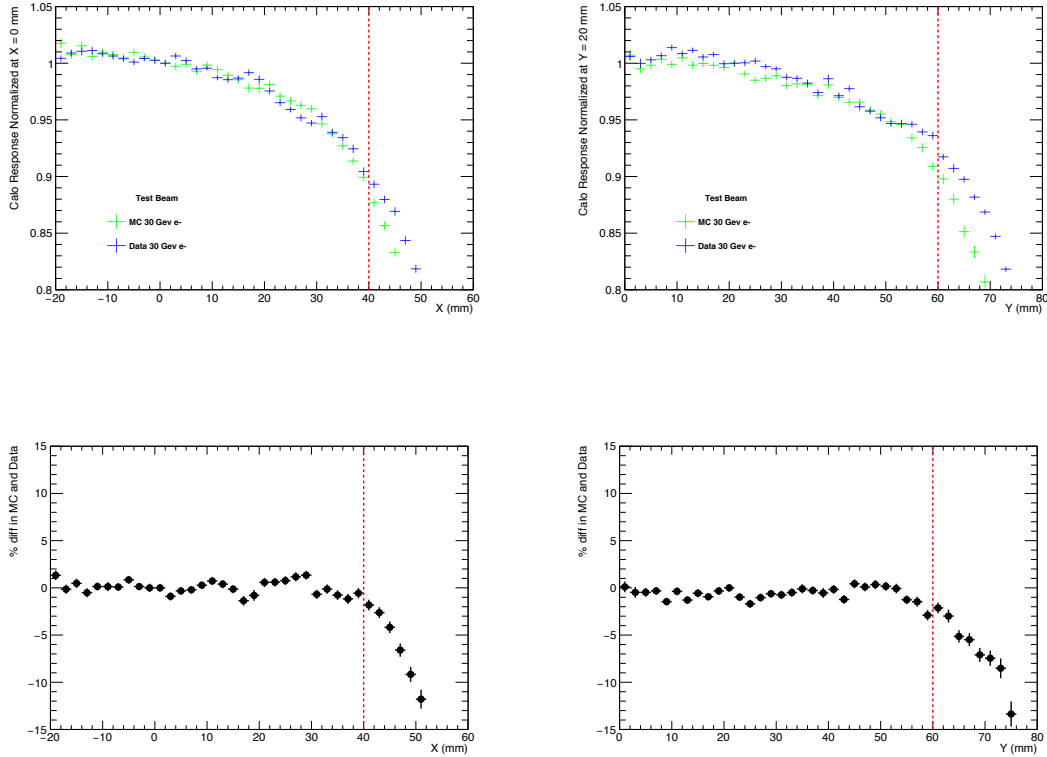


FIGURE B.11. The above plots depict the difference in test-beam MC with regards to lateral shower development and loss out the edges of a calorimeter module that is centered at $X=0$ and $Y=20$ mm. The top plots show the normalized calorimeter responses in both data (Blue) and MC (Green) for 30 GeV electrons traversing the edge of a calorimeter module in X (Left) and Y (Right). The edge in X is adjacent to another calorimeter module that is not being read out, whereas the edge in Y is accompanied by just air. The bottom plots show the % difference in the data and MC vs position. The dashed red line shows the position approximately 20 mm from the edge. The positions are obtained from linearly extrapolating the track segment in the test-beam tracking station to the face of the calorimeter. For the scan across X (left plots) we restrict the Y position to be within -10 and 0 mm, and for the scan across Y (right plots), we restrict the X position to be between 20 and 30 mm.

Uncertainty Description	Measured Uncertainty
TI12 data: MIP MPV fit	1.90%
TI12 data: HV gain extrapolation	3.37%
TI12 data: PMT drift over time	1.45%
TI12 MC: MIP MPV Fit	1.16%
MC energy correction: test-beam data calibration	0.74%
MC energy correction: test-beam MC calibration	2.35%
MC energy correction: Extrapolate difference to 500 GeV	2.46%
Other: Edge loss and positional dependence	2.5%
Total	6.06%

TABLE B.1. This table summarizes all uncertainties and their respective values when comparing the calibrated energies of data and MC at 500 GeV.

APPENDIX C

VALIDATING TWO-TRACK RECONSTRUCTION WITH DELTA RAYS

Delta-rays in FASER are produced from high-energy muons interacting in upstream material and kicking out an electron, thus providing events with two tracks of a common origin that can be used to probe the two-track reconstruction differences in MC and data. Unlike the two high momentum e^+e^- particles from the dark photon, these delta ray events include a high momentum muon and a low momentum electron, whose distribution is shown in Fig. C.1, that will likely be more separated than the \sim TeV dark photon tracks. Nonetheless, delta rays can still be used to validate the results of the track overlay study described in Sec. 9.2, which found a 7% discrepancy in the two-track reconstruction efficiency between data and MC.




FIGURE C.1. The reconstructed momentum of the delta-ray candidate track is shown for both data and the Fluka MC sample described in Chapter V. To normalize the two distributions, both were scaled by the inverse of their total number of entries.

Selecting delta-ray candidate events from both data and the Fluka MC sample is done by first requiring at least one good reconstructed track that looks like an electron. This is accomplished by selecting tracks which have the curvature of a negatively charged particle, and whose reconstructed momentum is similar to the electromagnetic energy measured in the calorimeter. Requiring that the signal of the downstream preshower layer is at least twice as large as the upstream layer is also used to pick out electromagnetic showers. Once we have a reconstructed track from an electron, we enhance the two-track topology of the event selection without explicitly requiring that a second track was reconstructed; which is done by cutting on the charge of the scintillators, number of clusters, and track segment multiplicity to suggest at least two charged particles passed through the detector.

Now that we have selected candidate delta ray events, we can look at the number of reconstructed tracks to see if the accompanying muon track was missed (1 track reconstructed) or if it was successfully reconstructed (2 tracks reconstructed). Comparing the reconstructed track distributions in Fig. C.2, shows that the discrepancy in the two-track reconstruction efficiency between data and MC agrees, within error, to the 7% value found in the track overlay study.

Also, when looking at the upstream track separation of delta ray events that have two successfully reconstructed tracks, and zooming in to the smallest separations, as depicted in Fig. C.3, we see successful reconstruction of tracks that are only separated at the tens of microns level; giving assurance that we can reconstruct the finely separated tracks from the dark photon decay.




FIGURE C.2. The reconstructed track multiplicity in the delta-ray control sample described in the text is shown. The distributions for both data and MC are shown on the TOP. The ratio of MC to data is shown on the BOTTOM plot, where the MC was scaled to match the total number of candidate delta ray events in data.




FIGURE C.3. The separation between the two reconstructed tracks at the first tracking stations, in the delta-ray sample in data and MC simulation. The MC was scaled to match the total number of candidate delta ray events in data.

APPENDIX D

B-L GAUGE BOSON

The dark photon analysis presented in this paper was also used to probe the parameter space of the BSM $B - L$ gauge boson A'_{B-L} , and is discussed in this appendix.

The accidental conservation of baryon number B and (total) lepton number L in the SM suggests that these conserved quantities may be linked to local gauge symmetries, such as the $U(1)_{B-L}$ [98, 99]. This local gauge symmetry is not only conserved classically, but is also free of quantum anomalies when three sterile (right-handed) neutrinos are introduced. Much like the BSM dark photon extension, the $U(1)_{B-L}$ kinetically mixes with the SM $U(1)_Y$ gauge symmetry and predicts a new BSM particle, the $B - L$ gauge boson A'_{B-L} , that can interact with the SM. The properties of A'_{B-L} are determined by the Lagrangian terms [14]

$$\mathcal{L} \supset \frac{1}{2} m_{A'_{B-L}}^2 A_{B-L}'^2 - g_{B-L} \sum_f Q_{B-L}^f A_{B-L}'^\mu \bar{f} \gamma_\mu f, \quad (\text{D.1})$$

where Q_{B-L}^f is the $B - L$ charge of fermion f . The parameter space of this model is defined by the $B - L$ gauge boson's mass $m_{A'_{B-L}}$ and the $B - L$ gauge coupling g_{B-L} .

The A'_{B-L} gauge boson models that FASER is sensitive to are produced in a similar manner to the dark photon, discussed in Sec. 2.3.2, with light meson decays and dark bremsstrahlung the dominant production mechanisms. The production rates are proportional to g_{B-L}^2 , compared to ϵ^2 as in the dark photon model. For masses in the MeV to GeV range and small $B - L$ gauge couplings of $\sim 10^{-5}$, up to

10^8 A'_{B-L} gauge bosons may be produced in $\sim 300fb^{-1}$ of proton-proton collisions at the LHC and travel invisibly hundreds of meters towards FASER [14, 100].

The boson can decay to all kinematically accessible states that possess $B - L$ charge. In this analysis, the region of phase space which FASER is sensitive to is confined to the mass range $2m_e < m_{A'_{B-L}} < 2m_\mu \simeq 211$ MeV, where the possible decays are to electrons, SM neutrinos, and possibly sterile neutrinos. It is assumed that sterile neutrinos have masses $> \frac{1}{2}m_{A'_{B-L}}$, and so decays to sterile neutrinos are kinematically inaccessible. The visible signal from decays to electrons therefore has a branching fraction of $B(A'_{B-L} \rightarrow e^+e^-) \approx 40\%$. If decays to sterile neutrinos are allowed, the visible branching fraction could be as low as $B(A'_{B-L} \rightarrow e^+e^-) \approx 25\%$. This sudden decay into high energy e^+e^- particles in FASER produces a signal identical to the dark photon scenario.

With an identical signal, the dark photon analysis can be directly reinterpreted to probe the A'_{B-L} gauge boson. The dark photon event selection, described in Chapter VII, is applied to MC $B - L$ signal samples, which cover the relevant ranges in both $m_{A'_{B-L}}$ and g_{B-L} , in order to measure the sensitivity of the analysis on the A'_{B-L} parameter space. The $B - L$ signal events are generated using FORESEE [66] with the EPOS-LHC [13] generator, just like the MC dark photon signal samples as discussed in Chapter V. Even the parameterization of the uncertainties on the number of dark photon signal events decaying inside the FASER decay volume, discussed in Sec. 9.1, can be directly applied to the A'_{B-L} analysis because the production mechanisms are so similar. To be sure, the parameterization of the signal generator uncertainty was checked for numerous A'_{B-L} signal models spanning the relevant parameter space, and was found to be in good agreement with the envelope of the QGSJET II-04 [78] and SIBYLL 2.3d [79] generators.




FIGURE D.1. 90% confidence level exclusion contours in the $B - L$ gauge boson parameter space are shown. Regions excluded by previous experiments are shown in grey.

The null results, discussed in Chapter X, of the analysis provides exclusion contours for the $B - L$ gauge boson, as shown in Fig. D.1. In grey are the regions already excluded by experimental data from Orsay [88, 89] and NuCal [67, 90] as adapted from DarkCast [9], as well as from a dedicated search for invisible final states by NA64 [101]. FASER provides the first exclusion for models in the range $g_{B-L} \sim 5 \times 10^{-6} - 2 \times 10^{-5}$ and $m_{A'_{B-L}} \sim 15 \text{ MeV} - 40 \text{ MeV}$, with a total region between $g_{B-L} \sim 3 \times 10^{-6} - 4 \times 10^{-5}$ and $m_{A'_{B-L}} \sim 10 \text{ MeV} - 50 \text{ MeV}$ excluded.

REFERENCES CITED

- [1] Quantum Diaries. The standard model: A beautiful but flawed theory. URL <http://www.quantumdiaries.org/2014/03/14/the-standard-model-a-beautiful-but-flawed-theory/>.
- [2] Mario De Leo. Own work, cc by-sa 4.0. URL <https://commons.wikimedia.org/w/index.php?curid=74398525>.
- [3] Edvige Corbelli and Paolo Salucci. The extended rotation curve and the dark matter halo of M33. *Monthly Notices of the Royal Astronomical Society*, 311(2): 441–447, 01 2000. ISSN 0035-8711. doi: 10.1046/j.1365-8711.2000.03075.x. URL <https://doi.org/10.1046/j.1365-8711.2000.03075.x>.
- [4] Douglas Clowe, Maruša Bradač, Anthony H Gonzalez, Maxim Markevitch, Scott W Randall, Christine Jones, and Dennis Zaritsky. A direct empirical proof of the existence of dark matter. *The Astrophysical Journal*, 648(2):L109, 2006.
- [5] M. Tristram, A. J. Banday, M. Douspis, X. Garrido, K. M. Górski, S. Henrot-Versillé, L. T. Hergt, S. Ilić, R. Keskitalo, and G. Lagache. Cosmological parameters derived from the final Planck data release (PR4). *A&A*, 682:A37, 2024. doi: 10.1051/0004-6361/202348015. URL <https://doi.org/10.1051/0004-6361/202348015>.
- [6] Tongyan Lin. Dark matter models and direct detection. *PoS*, 333:009, 2019. doi: 10.22323/1.333.0009.
- [7] Particle Data Group. Review of Particle Physics. *Progress of Theoretical and Experimental Physics*, 2022(8):083C01, 08 2022. ISSN 2050-3911. doi: 10.1093/ptep/ptac097. URL <https://doi.org/10.1093/ptep/ptac097>.
- [8] Jia Liu, Neal Weiner, and Wei Xue. Signals of a light dark force in the galactic center. *Journal of High Energy Physics*, 2015(8), August 2015. ISSN 1029-8479. doi: 10.1007/jhep08(2015)050. URL [http://dx.doi.org/10.1007/JHEP08\(2015\)050](http://dx.doi.org/10.1007/JHEP08(2015)050).
- [9] Philip Ilten, Yotam Soreq, Mike Williams, and Wei Xue. Serendipity in dark photon searches. *JHEP*, 06:004, 2018. doi: 10.1007/JHEP06(2018)004.
- [10] Renilde Vanden Broeck. The cern accelerator complex. complexe des accélérateurs du cern, 2019. URL <https://cds.cern.ch/record/2693837>.
- [11] AC Team. Diagram of an lhc dipole magnet. schéma d’un aimant dipôle du lhc, 1999. URL <https://cds.cern.ch/record/40524>.

- [12] Xabier Cid Vidal and Ramon Cid Manzano. Magnetic multipoles, 2022. URL https://www.lhc-closer.es/taking_a_closer_look_at_lhc/0.magnetic_multipoles.
- [13] T. Pierog, Iu. Karpenko, J. M. Katzy, E. Yatsenko, and K. Werner. EPOS LHC: Test of collective hadronization with data measured at the CERN Large Hadron Collider. *Phys. Rev. C*, 92(3):034906, 2015. doi: 10.1103/PhysRevC.92.034906.
- [14] Akitaka Ariga et al. FASER’s physics reach for long-lived particles. *Phys. Rev. D*, 99(9):095011, 2019. doi: 10.1103/PhysRevD.99.095011.
- [15] Finite element method magnetics. URL www.femm.info.
- [16] The trigger and data acquisition system of the FASER experiment. *JINST*, 16(12):P12028, 2021. doi: 10.1088/1748-0221/16/12/P12028.
- [17] FASER Collaboration. Search for dark photons with the faser detector at the lhc. *Physics Letters B*, 848:138378, 2024. ISSN 0370-2693. doi: <https://doi.org/10.1016/j.physletb.2023.138378>. URL <https://www.sciencedirect.com/science/article/pii/S0370269323007128>.
- [18] Gianfranco Bertone, Dan Hooper, and Joseph Silk. Particle dark matter: Evidence, candidates and constraints. *Phys. Rept.*, 405:279–390, 2005. doi: 10.1016/j.physrep.2004.08.031.
- [19] Jonathan L. Feng. Dark Matter Candidates from Particle Physics and Methods of Detection. *Ann. Rev. Astron. Astrophys.*, 48:495–545, 2010. doi: 10.1146/annurev-astro-082708-101659.
- [20] Peter W. Higgs. Broken symmetries and the masses of gauge bosons. *Phys. Rev. Lett.*, 13:508–509, Oct 1964. doi: 10.1103/PhysRevLett.13.508. URL <https://link.aps.org/doi/10.1103/PhysRevLett.13.508>.
- [21] Michael E. Peskin and Daniel V. Schroeder. *An Introduction to Quantum Field Theory*. Addison-Wesley, Reading, USA, 1995. ISBN 978-0-201-50397-5.
- [22] Mark Thomson. *Modern Particle Physics*. Cambridge University Press, 2013. ISBN 978-1-107-29254-3.
- [23] B. D’Urso B. Odom, D. Hanneke and G. Gabrielse. New measurement of the electron magnetic moment using a one-electron quantum cyclotron. *Phys. Rev. Lett.*, 97:030801, 2006.
- [24] ATLAS Collaboration and 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. *Phys. Lett. B*, 716, 2012.

- [25] CMS Collaboration and S. Chatrchyan et al. Observation of a new boson at a mass of 125 gev with the cms experiment at the lhc. *Phys. Lett. B*, 716:30, 2012.
- [26] Gianfranco Bertone and Dan Hooper. History of dark matter. *Rev. Mod. Phys.*, 90:045002, Oct 2018. doi: 10.1103/RevModPhys.90.045002. URL <https://link.aps.org/doi/10.1103/RevModPhys.90.045002>.
- [27] R. L. Workman and Others. Review of particle physics. *PTEP*, 2022:083C01, 2022. doi: 10.1093/ptep/ptac097.
- [28] Gary Steigman and Michael S. Turner. Cosmological constraints on the properties of weakly interacting massive particles. *Nuclear Physics B*, 253: 375–386, 1985. ISSN 0550-3213. doi: [https://doi.org/10.1016/0550-3213\(85\)90537-1](https://doi.org/10.1016/0550-3213(85)90537-1). URL <https://www.sciencedirect.com/science/article/pii/0550321385905371>.
- [29] Gerard Jungman, Marc Kamionkowski, and Kim Griest. Supersymmetric dark matter. *Physics Reports*, 267(5):195–373, 1996. ISSN 0370-1573. doi: [https://doi.org/10.1016/0370-1573\(95\)00058-5](https://doi.org/10.1016/0370-1573(95)00058-5). URL <https://www.sciencedirect.com/science/article/pii/0370157395000585>.
- [30] Bob Holdom. Two U(1)’s and Epsilon Charge Shifts. *Phys. Lett.*, B166:196, 1986. doi: 10.1016/0370-2693(86)91377-8.
- [31] Marco Fabbrichesi, Emidio Gabrielli, and Gaia Lanfranchi. *The Physics of the Dark Photon: A Primer*. Springer International Publishing, 2021. ISBN 9783030625191. doi: 10.1007/978-3-030-62519-1. URL <http://dx.doi.org/10.1007/978-3-030-62519-1>.
- [32] HENRI RUEGG and MARTÍ RUIZ-ALTABA. The stueckelberg field. *International Journal of Modern Physics A*, 19(20):3265–3347, August 2004. ISSN 1793-656X. doi: 10.1142/s0217751x04019755. URL <http://dx.doi.org/10.1142/S0217751X04019755>.
- [33] Haipeng An, Maxim Pospelov, Josef Pradler, and Adam Ritz. Direct detection constraints on dark photon dark matter. *Physics Letters B*, 747:331–338, July 2015. ISSN 0370-2693. doi: 10.1016/j.physletb.2015.06.018. URL <http://dx.doi.org/10.1016/j.physletb.2015.06.018>.
- [34] C. Boehm and Pierre Fayet. Scalar dark matter candidates. *Nucl. Phys. B*, 683: 219–263, 2004. doi: 10.1016/j.nuclphysb.2004.01.015.
- [35] Maxim Pospelov, Adam Ritz, and Mikhail B. Voloshin. Secluded WIMP Dark Matter. *Phys. Lett. B*, 662:53–61, 2008. doi: 10.1016/j.physletb.2008.02.052.

- [36] Jonathan L. Feng and Jason Kumar. The WIMPless Miracle: Dark-Matter Particles without Weak-Scale Masses or Weak Interactions. *Phys. Rev. Lett.*, 101:231301, 2008. doi: 10.1103/PhysRevLett.101.231301.
- [37] P. A. Zyla et al. Review of Particle Physics. *PTEP*, 2020(8):083C01, 2020. doi: 10.1093/ptep/ptaa104.
- [38] Jonathan L. Feng, Iftah Galon, Felix Kling, and Sebastian Trojanowski. ForwArD Search ExpeRiment at the LHC. *Phys. Rev. D*, 97(3):035001, 2018. doi: 10.1103/PhysRevD.97.035001.
- [39] Motoi Endo, Koichi Hamaguchi, and Go Mishima. Constraints on hidden photon models from electron g-2 and hydrogen spectroscopy. *Physical Review D*, 86(9), November 2012. ISSN 1550-2368. doi: 10.1103/physrevd.86.095029. URL <http://dx.doi.org/10.1103/PhysRevD.86.095029>.
- [40] R. et al Aaij. Search for dark photons produced in 13 tev pp collisions. *Physical Review Letters*, 120(6), February 2018. ISSN 1079-7114. doi: 10.1103/physrevlett.120.061801. URL <http://dx.doi.org/10.1103/PhysRevLett.120.061801>.
- [41] Philip Ilten, Jesse Thaler, Mike Williams, and Wei Xue. Dark photons from charm mesons at lhcb. *Physical Review D*, 92(11), December 2015. ISSN 1550-2368. doi: 10.1103/physrevd.92.115017. URL <http://dx.doi.org/10.1103/PhysRevD.92.115017>.
- [42] Communications CERN Education and September 2021 CERN-Brochure-2021-004-Eng Outreach Group. Lhc guide faq. URL <https://cds.cern.ch/record/2809109/files/CERN-Brochure-2021-004-Eng.pdf>.
- [43] Xabier Cid Vidal and Ramon Cid Manzano. Linac4 in very short, 2022. URL https://lhc-closer.es/taking_a_closer_look_at_lhc/0.linac4.
- [44] Hannes Bartosik and Giovanni Rumolo. Performance of the lhc injector chain after the upgrade and potential development, 2022.
- [45] H. Frischholz, W.R. Fowkes, and C. Pearson. Design and construction of a 500 kw cw, 400 mhz klystron to be used as rf power source for lhc/rf component tests. In *Proceedings of the Particle Accelerator Conference (PAC)*, Vancouver, 1997. CERN/SL/97-39 RF.
- [46] Oliver Sim Brüning, Paul Collier, P Lebrun, Stephen Myers, Ranko Ostojic, John Poole, and Paul Proudlock. *LHC Design Report*. CERN Yellow Reports: Monographs. CERN, Geneva, 2004. doi: 10.5170/CERN-2004-003-V-1. URL <https://cds.cern.ch/record/782076>.

- [47] Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC. 2022. doi: 10.48550/ARXIV.2212.09379.
- [48] Henso Abreu et al. The FASER Detector. July 2022.
- [49] N.V. Mokhov and I.L. Rakhno. Protecting lhc components against radiation resulting from colliding beam interactions. In *Proceedings of the Particle Accelerator Conference (PAC'01)*, Chicago, June 2001. Fermilab-Conf - 01/131.
- [50] Lyndon Evans and Philip Bryant. LHC Machine. *JINST*, 3:S08001, 2008. doi: 10.1088/1748-0221/3/08/S08001.
- [51] Henso et al Abreu. First neutrino interaction candidates at the lhc. *Phys. Rev. D*, 104:L091101, Nov 2021. doi: 10.1103/PhysRevD.104.L091101. URL <https://link.aps.org/doi/10.1103/PhysRevD.104.L091101>.
- [52] A Ferrari, Paola R Sala, A Fasso, and Johannes Ranft. *FLUKA: A Multi-particle Transport Code (Program Version 2005)*. CERN Yellow Reports: Monographs. CERN, Geneva, 2005. URL <http://cds.cern.ch/record/898301>.
- [53] T. T. Böhlen, F. Cerutti, M. P. W. Chin, A. FassÃ², A. Ferrari, P. G. Ortega, A. Mairani, P. R. Sala, G. Smirnov, and V. Vlachoudis. The FLUKA Code: Developments and Challenges for High Energy and Medical Applications. *Nucl. Data Sheets*, 120:211–214, 2014. doi: 10.1016/j.nds.2014.07.049.
- [54] The tracking detector of the FASER experiment. *Nucl. Instrum. Meth. A*, 1034:166825, 2022. doi: 10.1016/j.nima.2022.166825.
- [55] A. Abdesselam et al. The barrel modules of the ATLAS semiconductor tracker. *Nucl. Instrum. Meth. A*, 568:642–671, 2006. doi: 10.1016/j.nima.2006.08.036.
- [56] F. Campabadal et al. Design and performance of the abcd3ta asic for readout of silicon strip detectors in the atlas semiconductor tracker. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 552(3):292–328, 2005. ISSN 0168-9002. doi: <https://doi.org/10.1016/j.nima.2005.07.002>. URL <https://www.sciencedirect.com/science/article/pii/S0168900205013926>.
- [57] Eljen Technology. General Purpose EJ-200, EJ-204, EJ-208, EJ-212. <https://eljentechnology.com/products/plastic-scintillators/ej-200-ej-204-ej-208-ej-212>.

- [58] A. Tkaczyk et al. Characterization of ej-200 plastic scintillators as active background shield for cosmogenic radiation. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 882:96–104, 2018. URL <https://www.sciencedirect.com/science/article/pii/S0168900217311774>.
- [59] Hamamatsu. Photomultiplier tube assembly h11934. <https://www.hamamatsu.com/eu/en/product/optical-sensors/pmt/pmt-assembly/head-on-type/H11934-100.html>, .
- [60] Hamamatsu. Photomultiplier tube assembly h6410. <https://www.hamamatsu.com/eu/en/product/optical-sensors/pmt/pmt-assembly/head-on-type/H6410.html>, .
- [61] *LHCb Calorimeters: Technical Design Report*. CERN, Geneva, 2000. URL <http://cds.cern.ch/record/494264>.
- [62] Hamamatsu. Photomultiplier tube r7899. <https://www.hamamatsu.com/eu/en/product/type/R7899/index.html>, .
- [63] G. Avoni et al. The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS. *JINST*, 13(07):P07017, 2018. doi: 10.1088/1748-0221/13/07/P07017.
- [64] Preliminary analysis of the luminosity calibration of the ATLAS 13.6 TeV data recorded in 2022. ATL-DAPR-PUB-2023-001, 2023. URL <https://cds.cern.ch/record/2853525>.
- [65] GEANT4 – a simulation toolkit. *Nucl. Instrum. Meth. A*, 506:250, 2003. doi: 10.1016/S0168-9002(03)01368-8.
- [66] Felix Kling and Sebastian Trojanowski. Forward experiment sensitivity estimator for the LHC and future hadron colliders. *Phys. Rev. D*, 104(3): 035012, 2021. doi: 10.1103/PhysRevD.104.035012.
- [67] Johannes Blümlein and Jürgen Brunner. New Exclusion Limits on Dark Gauge Forces from Proton Bremsstrahlung in Beam-Dump Data. *Phys. Lett. B*, 731: 320–326, 2014. doi: 10.1016/j.physletb.2014.02.029.
- [68] Giuseppe Battistoni, Till Boehlen, Francesco Cerutti, Pik Wai Chin, Luigi Salvatore Esposito, Alberto Fassò, Alfredo Ferrari, Anton Lechner, Anton Empl, Andrea Mairani, Alessio Mereghetti, Pablo Garcia Ortega, Johannes Ranft, Stefan Roesler, Paola R. Sala, Vasilis Vlachoudis, and George Smirnov. Overview of the fluka code. *Annals of Nuclear Energy*, 82:10–18, 2015. ISSN 0306-4549. doi: <https://doi.org/10.1016/j.anucene.2014.11.007>.

- [69] Henso Abreu et al. First Direct Observation of Collider Neutrinos with FASER at the LHC. *Phys. Rev. Lett.*, 131(3):031801, 2023. doi: 10.1103/PhysRevLett.131.031801.
- [70] C. Andreopoulos et al. The GENIE Neutrino Monte Carlo Generator. *Nucl. Instrum. Meth.*, A614:87–104, 2010. doi: 10.1016/j.nima.2009.12.009.
- [71] Costas Andreopoulos, Christopher Barry, Steve Dytman, Hugh Gallagher, Tomasz Golan, Robert Hatcher, Gabriel Perdue, and Julia Yarba. The GENIE Neutrino Monte Carlo Generator: Physics and User Manual. 2015.
- [72] Felix Kling and Laurence J. Nevay. Forward neutrino fluxes at the LHC. *Phys. Rev. D*, 104(11):113008, 2021. doi: 10.1103/PhysRevD.104.113008.
- [73] **FASER** Collaboration, CALYPSO software framework <https://gitlab.cern.ch/faser/calypso>.
- [74] The ATLAS Collaboration Software and Firmware. ATL-SOFT-PUB-2021-001, 2021. URL <https://cds.cern.ch/record/2767187>.
- [75] **ATLAS** Collaboration. “Athena” [software], release-22, <https://zenodo.org/record/2641997#.XhWRDC2ZPyI>.
- [76] Xiacong Ai et al. A Common Tracking Software Project. *Comput. Softw. Big Sci.*, 6(1):8, 2022. doi: 10.1007/s41781-021-00078-8.
- [77] J. L. Feng, M. Toman, and E. Welch. 2023. In preparation.
- [78] Sergey Ostapchenko. Monte Carlo treatment of hadronic interactions in enhanced Pomeron scheme: I. QGSJET-II model. *Phys. Rev. D*, 83:014018, 2011. doi: 10.1103/PhysRevD.83.014018.
- [79] Felix Riehn, Ralph Engel, Anatoli Fedynitch, Thomas K. Gaisser, and Todor Stanev. Hadronic interaction model Sibyll 2.3d and extensive air showers. *Phys. Rev. D*, 102(6):063002, 2020. doi: 10.1103/PhysRevD.102.063002.
- [80] O. Adriani et al. Measurement of forward photon production cross-section in proton-proton collisions at $\sqrt{s} = 13$ TeV with the LHCf detector. *Phys. Lett. B*, 780:233–239, 2018. doi: 10.1016/j.physletb.2017.12.050.
- [81] M. Baak, G. J. Besjes, D. Côte, A. Koutsman, J. Lorenz, and D. Short. HistFitter software framework for statistical data analysis. *Eur. Phys. J. C*, 75: 153, 2015. doi: 10.1140/epjc/s10052-015-3327-7.
- [82] Glen Cowan, Kyle Cranmer, Eilam Gross, and Ofer Vitells. Asymptotic formulae for likelihood-based tests of new physics. *Eur. Phys. J. C*, 71:1554, 2011. doi: 10.1140/epjc/s10052-011-1554-0. [Erratum: *Eur.Phys.J.C* 73, 2501 (2013)].

- [83] Alexander L. Read. Presentation of search results: The CL_s technique. *J. Phys. G*, 28:2693–2704, 2002. doi: 10.1088/0954-3899/28/10/313.
- [84] J. P. Lees et al. Search for a Dark Photon in e^+e^- Collisions at BaBar. *Phys. Rev. Lett.*, 113(20):201801, 2014. doi: 10.1103/PhysRevLett.113.201801.
- [85] E. M. Riordan et al. A Search for Short Lived Axions in an Electron Beam Dump Experiment. *Phys. Rev. Lett.*, 59:755, 1987. doi: 10.1103/PhysRevLett.59.755.
- [86] J. R. Batley et al. Search for the dark photon in π^0 decays. *Phys. Lett. B*, 746:178–185, 2015. doi: 10.1016/j.physletb.2015.04.068.
- [87] D. Banerjee et al. Improved limits on a hypothetical X(16.7) boson and a dark photon decaying into e^+e^- pairs. *Phys. Rev. D*, 101(7):071101, 2020. doi: 10.1103/PhysRevD.101.071101.
- [88] M. Davier and H. Nguyen Ngoc. An Unambiguous Search for a Light Higgs Boson. *Phys. Lett. B*, 229:150–155, 1989. doi: 10.1016/0370-2693(89)90174-3.
- [89] Sarah Andreas, Carsten Niebuhr, and Andreas Ringwald. New Limits on Hidden Photons from Past Electron Beam Dumps. *Phys. Rev. D*, 86:095019, 2012. doi: 10.1103/PhysRevD.86.095019.
- [90] J. Blumlein et al. Limits on the mass of light (pseudo)scalar particles from Bethe-Heitler e^+e^- and $\mu^+\mu^-$ pair production in a proton - iron beam dump experiment. *Int. J. Mod. Phys. A*, 7:3835–3850, 1992. doi: 10.1142/S0217751X9200171X.
- [91] N. Aghanim et al. Planck 2018 results. VI. Cosmological parameters. *Astron. Astrophys.*, 641:A6, 2020. doi: 10.1051/0004-6361/201833910. [Erratum: *Astron. Astrophys.* 652, C4 (2021)].
- [92] M. J. Oreglia. A Study of the Reactions $\psi' \rightarrow \gamma\gamma\psi$. *Ph.D. Thesis, SLAC-R-236 (1980), Appendix D*.
- [93] GEANT4: A Simulation toolkit. *Nucl. Instrum. Meth.*, A506:250–303, 2003. doi: 10.1016/S0168-9002(03)01368-8.
- [94] H2 Beam Line.
http://sba.web.cern.ch/sba/BeamsAndAreas/H2/H2_presentation.html.
- [95] FASER Calorimeter and Tracker Test Beam Characterization. 2023. In preparation.
- [96] URL https://en.wikipedia.org/wiki/Convolution#/media/File:Convolution_of_spiky_function_with_box2.gif.

- [97] Irina Machikhiliyan and LHCb calorimeter group. The lhcb electromagnetic calorimeter. *Journal of Physics: Conference Series*, 160(1):012047, apr 2009. doi: 10.1088/1742-6596/160/1/012047. URL <https://dx.doi.org/10.1088/1742-6596/160/1/012047>.
- [98] Aharon Davidson. $B - L$ as the fourth color within an $SU(2)_L \times U(1)_R \times U(1)$ model. *Phys. Rev. D*, 20:776, 1979. doi: 10.1103/PhysRevD.20.776.
- [99] R. E. Marshak and Rabindra N. Mohapatra. Quark–Lepton Symmetry and $B - L$ as the $U(1)$ Generator of the Electroweak Symmetry Group. *Phys. Lett. B*, 91:222–224, 1980. doi: 10.1016/0370-2693(80)90436-0.
- [100] Martin Bauer, Patrick Foldenauer, and Joerg Jaeckel. Hunting All the Hidden Photons. *JHEP*, 07:094, 2018. doi: 10.1007/JHEP07(2018)094.
- [101] Yu. M. Andreev et al. Search for a New B-L Z' Gauge Boson with the NA64 Experiment at CERN. *Phys. Rev. Lett.*, 129(16):161801, 2022. doi: 10.1103/PhysRevLett.129.161801.