

UC Irvine

UC Irvine Electronic Theses and Dissertations

Title

First Long-Lived Particle Results and Commissioning of the ForwArd Search ExpeRiment

Permalink

<https://escholarship.org/uc/item/2782f2nm>

Author

Shively, Savannah Rose

Publication Date

2024

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NonCommercial-NoDerivatives License, available at

<https://creativecommons.org/licenses/by-nc-nd/4.0/>

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA,
IRVINE

First Long-Lived Particle Results and Commissioning of the ForwArd Search ExpeRiment

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Physics

by

Savannah Rose Shively-Piccardo

Dissertation Committee:
Associate Professor David Casper, Co-Chair
Distinguished Professor Jonathan Feng, Co-Chair
Distinguished Professor Andrew J. Lankford
Professor Steven W. Barwick
Research Professor Dr. William Parker

2024

DEDICATION

Para mi madre, Gabriela Piccardo, who encouraged my curiosity and supported me every step of the way through everything and anything.

-

Para Mema, for teaching me how to use a computer.

-

Para Lolo, for showing me how designs on paper can become real.

-

To my dad, Sean P. Shively, who helps his kids however he can, and my 2nd mom, Melissa “Missy” Rae-Shively, for taking care of me and generally being inspirational.

-

To my little siblings and sorta-sibling, Kyle, Joel, Tatiana, Nicole and the rest of my extended family for cheering me on.

-

To my classmates and faculty of UCI, for their solidarity.

-

Para José Flores Velázquez, who will be remembered for the joy and science he brought us.
Rest in peace, my friend.

TABLE OF CONTENTS

	Page
LIST OF FIGURES	vi
LIST OF TABLES	viii
ACKNOWLEDGMENTS	ix
VITA	x
ABSTRACT OF THE DISSERTATION	xi
1 Introduction	1
1.1 CERN and its Facilities	3
1.2 FASER’s History	5
2 FASER Physics: The Standard Model and Beyond	7
2.1 The Standard Model	7
2.1.1 General Features	8
2.1.2 Mathematical Description	9
2.1.3 Successes and Limitations	10
2.2 Dark matter and the Dark Sector	11
2.2.1 Dark Photon	12
2.2.2 B-L Gauge Boson	13
2.3 Long-Lived Particle Production at the LHC	14
2.3.1 Dark Photon Production and Decay	14
2.3.2 B-L Gauge Boson Production and Decay	16
2.4 Sensitivity	18
3 Experimental Set-Up	20
3.1 Location	21
3.2 Beam Production	22
3.3 Tracking	26
3.3.1 Hardware	26
3.3.2 Calibration	30
3.4 Scintillators	35
3.4.1 Hardware	35

3.4.2	Calibration	40
3.5	Calorimeters	40
3.5.1	Hardware	40
3.5.2	Calibration	44
3.6	Magnets	46
3.7	Emulsion Detector	47
3.8	Mechanical Support Structures	49
3.9	Software	50
3.9.1	Data Acquisition	50
3.9.2	Data Processing	55
3.10	Control, Monitoring and Safety	56
4	Simulation	59
4.1	Event Generation	61
4.1.1	Particle Guns	61
4.1.2	Proton-proton Events	61
4.1.3	Beyond Standard Model Events	62
4.1.4	Neutrino Events	64
4.1.5	FLUKA	66
4.2	Detector Simulation	66
4.3	Digitization	68
5	Reconstruction	71
5.1	Track Reconstruction	71
5.1.1	Clusterization	71
5.1.2	Track Finding	72
5.2	PMT Data	73
5.3	Detector Alignment	74
6	Analysis	77
6.1	Event Selection	77
6.2	Background	80
6.2.1	Veto inefficiency	80
6.2.2	Neutral Hadrons	81
6.2.3	Neutrinos	83
6.2.4	Large-Angle Muons	85
6.2.5	Non-collision	85
6.2.6	Summary	87
6.3	Systematic Uncertainty	87
6.3.1	Signal Systematics	88
6.3.2	Scintillator and Calorimeter Contributions	89
6.3.3	Tracking Contributions	89

7	Results	95
7.1	Statistical Inference	95
7.2	Frequentist Concepts	96
7.3	Statistical Analysis Framework	98
	7.3.1 Model-Dependent Fit	99
	7.3.2 Model-Independent Fit	99
7.4	Sensitivity Tests	100
7.5	Unblinded Results	100
8	Conclusion	104
	Bibliography	106
	Appendix A Tracker Commissioning	112
A.1	Prototype Plane	112
A.2	Module Commissioning	113
A.3	Plane and Station Procedure	113
A.4	Magnetic Field Test	116
A.5	Surface Cosmics	117
A.6	Post-installation and Summary	119

LIST OF FIGURES

	Page
1.1 The Lamppost Landscape	2
1.2 Diagram of LHC facilities as of January 2022, including SPS. Not to scale.	4
2.1 IP meson spectrum for pions.	15
2.2 Dark photon (A') production.	15
2.3 Dark photon decay length, branching fractions, and sensitivity.	17
2.4 $B - L$ gauge boson decay length, branching fractions, and sensitivity.	18
3.1 A model of FASER showing the location of different sub-systems and its standard coordinate system.	21
3.2 CAD sketch overview of integrated supporting systems in TI12.	23
3.3 TI12, before and after civil engineering intervention.	23
3.4 Schematic view of the LHC structures near ATLAS and FASER.	25
3.5 An ATLAS SCT barrel module.	27
3.6 Differences in temperature across all strip detectors in a single tracker plane for different LV configurations.	28
3.7 Sketch of a tracker plane.	29
3.8 Photograph of an assembled tracking station.	30
3.9 Illustration showing stages of a threshold scan.	32
3.10 Response curve of a single strip.	36
3.11 CAD sketches of the veto stations.	38
3.12 CAD sketches of the timing and preshower scintillator stations.	39
3.13 PMT response to cosmic ray muons.	41
3.14 Sketch of FASER calorimeter.	42
3.15 Photograph of assembled FASER calorimeter.	43
3.16 Comparison of Calorimeter resolution between LHCb data and FASER simulation.	45
3.17 Cross-section of FASER magnet design.	47
3.18 One quadrant of FASER magnet cross-section simulated in FEMM.	48
3.19 Sketch of structural and mobility systems for the magnets and spectrometer.	49
3.20 Overview of TDAQ hardware network.	52
3.21 Overview of TDAQ software.	53
3.22 Overview of the tracker control system.	54
3.23 Sketch of the interlock system.	57

4.1	Expected dark photon spectra for dark photons within FASER’s reach. . . .	63
4.2	Comparison of e^+e^- spectra from the simulated decay of A' by FORESEE and ISG.	64
4.3	Neutrino flux and energy estimated by RIVET and BDSIM using Sybill 2.3d hadron generation for muon and electron neutrinos.	65
4.4	FLUKA-simulated muon rates at 409 m from IP1 entering a $1\text{ m} \times 1\text{ m}$ area centered on the LOS as a function of energy.	67
4.5	Fluka-generated muon flux and spectra in TI18, normalized to $2 \times 10^{34}\text{ cm}^{-2}\text{ s}^{-1}$ luminosity.	68
4.6	A comparison of normalized waveforms derived from generated and real events.	70
5.1	Full detector event display with a single-muon event, top and side views.	73
5.2	Crystal ball function fit to PMT pulses.	75
5.3	The effects of alignment on track residuals.	76
6.1	Sketch showing the response of the FASER detector to a neutral LLP decay.	78
6.2	Simulated neutrino background spectrum.	84
6.3	ECAL MC geometric muon energy prior to event selection.	86
6.4	Gaussian-fitted peak position of the E/p from non-muon photon conversion tracks in data and MC simulation as a function of the momentum of photon conversion candidates.	90
6.5	Single track efficiency vs. truth momentum and truth track angle.	91
6.6	The efficiency of overlaid MC and Data as a function of track separation and the ratio of MC and data efficiencies.	93
6.7	500 GeV photon conversion peak with Gaussian fits for momentum resolution and momentum scale.	94
7.1	Spectrum of events with different cuts applied.	101
7.2	New exclusion limits for dark photons and $B - L$ gauge boson.	103

LIST OF TABLES

	Page
3.1 Nominal values for scintillator operation.	39
6.1 Selection criteria for A' and $A'B - L$ events and signal efficiency for one example MC dark photon model.	79
6.2 Cutflow for one example MC dark photon model.	80

ACKNOWLEDGMENTS

I would like to thank my advisors for their support and guidance. Dave Casper introduced me to FASER's calypso software and many academic opportunities. Jonathan Feng has been a source of academic guidance since my undergraduate career; he introduced me to FASER. Thank you Steve Barwick for his guidance during my undergraduate research with ARIANNA. Thank you Bill Parker for being one of my academic advisors. Finally, thank you Andrew Lankford for his bureaucratic assistance.

Also from UCI, I would like to thank Sophia Nasr who has been my peer mentor since I was an undergraduate student. I would also like to thank Mu-Chun Chen for encouraging my participation in diversity, equity and inclusion activities that have brought additional purpose to my work.

I would also like to thank the FASER Collaboration, particularly those based at CERN, for their support, patience, and guidance: Jamie Boyd for his leadership and versatility, Hide Otono for encouraging my participation in tracker activities, Susanne Kuehn for checking on me and helping me settle in among many other things, Claire Antel for her advice, feedback, being a brilliant coworker and one of my first friends at CERN, Michaela Queitsch-Maitland for her many resources and academic advice, Felix Kling for his assistance with FORESEE and theory topics, and finally Lottie Cavanagh and Daniella Kock for being inspiring peers.

My research was primarily funded by the generous contributions of Nancy and Corwin Evans. The final year of my program was funded by the University of California UC President's Pre-Professoriate Fellowship (PPPF). Financial support was also provided by the University of California, Irvine. FASER hardware and my funding were supported by the following grants: Heising Foundation grants 2018-1135, 2019-1179, and 2020-1840; Simons Foundation grant 623683; and the U.S. National Science Foundation grant PHY-2111427.

I would like to thank CERN and FASER for their permission to include copyrighted, public plots as part of my thesis.

VITA

Savannah Rose Shively-Piccardo

- 2015 Undergraduate Internship, NASA Armstrong Flight Research Center
- 2016-18 Undergraduate Student Researcher, Antarctic Ross Ice Shelf Antenna Neutrino Array
- 2018 B.S. in Computational Physics, University of California, Irvine
- 2018-19 Teaching Assistant, University of California, Irvine
- 2020 M.S. in Particle Physics, University of California, Irvine
- 2019-24 Graduate Student Researcher, ForwArd Search ExpeRiment (FASER)
- 2024 Ph.D. in Physics, University of California, Irvine

FIELD OF STUDY

High-Energy Experimental Particle Physics

PUBLICATIONS

“Radio detection of air showers with the ARIANNA experiment on the Ross Ice Shelf”. *Astroparticle Physics* 90, 2017.

“First neutrino interaction candidates at the LHC”. *Phys. Rev. D* 104, 2021.

“The trigger and data acquisition system of the FASER experiment”. *JINST* 16, 2021.

“The tracking detector of the FASER experiment”. *Nucl. Instrum. Methods Phys. Res., A* 1034, 2022.

“FASER Tracker Detector - Commissioning, Installation, and Functionality”. *PoS ICHEP2022*, 2022.

“First Direct Observation of Collider Neutrinos with FASER at the LHC”. *Phys.Rev.Lett.* 131, 2023.

“The FASER detector”. *JINST* 19, 2024.

“Search for Dark Photons with the FASER detector at the LHC”. *Phys. Lett. B* 848, 2024.

ABSTRACT OF THE DISSERTATION

First Long-Lived Particle Results and Commissioning of the ForwArd Search ExpeRiment

By

Savannah Rose Shively-Piccardo

Doctor of Philosophy in Physics

University of California, Irvine, 2024

Associate Professor David Casper, Co-Chair

Distinguished Professor Jonathan Feng, Co-chair

The ForwArd Search ExpeRiment (FASER) is collecting far-forward proton-proton collision data during the Large Hadron Collider’s third running period (2022-2025). This thesis describes two of the first-proposed searches: for dark photon (A') and the $B - L$ gauge boson (A_{B-L}). These searches are motivated by the proposed dark sector — an extension to the Standard Model of particle physics that includes dark matter. These long-lived particles would mediate interactions between the Standard Model and the dark sector.

Although larger experiments like ATLAS and CMS are well-suited to detect new, heavy particles, there may be low-mass particles produced in abundance along the axis of collision, in the far forward region. FASER, a 5m-long detector comprised of a magnetic spectrometer, scintillators, and calorimeters, is designed to study high-energy charged particles in this region. The experiment is installed in the repurposed SPS maintenance tunnel, TI12, 480m from the ATLAS interaction point and oriented along the line-of-sight. FASER’s highly efficient veto system rejects all incoming muon events. Tracker readout is triggered by a timing scintillator or calorimeter signal.

Dark sector models predict A' and A_{B-L} from the decay of SM mesons, which are copiously

produced in the far forward direction at ATLAS. They share a striking experimental signature, decaying into energetic electron-positron pairs. Event selection requires no entering (veto) signal, two good reconstructed tracks, and a calorimeter deposit of at least 500 GeV. The expected background for the search is $(2.3 \pm 2.3 \times 10^{-3})$ events for 27 fb^{-1} luminosity and is dominated by neutrino interactions with the detector materials.

FASER collected its first year of data in 2022, with results first presented at the March 2023 Moriond conferences. No candidates were observed. FASER excluded new regions of parameter space for A' and A_{B-L} at 90% confidence level.

Chapter 1

Introduction

Ever since observation of the Higgs boson in 2012 [1, 2], particle physicists have sought new ways to probe beyond the Standard Model (SM). Although the SM has seen decades of success, there are many phenomena that it cannot describe. For example, over 95% of the mass and energy in the universe is attributed to dark matter and dark energy [3], neither of which are described adequately by the SM. Such cosmological clues motivate us to continue the search for evidence of new particles with new strategies. The Forward Search Experiment (FASER) is part of this movement.

Different strategies for a beyond-standard-model (BSM) search are possible, depending on the properties of hypothesized new particles; see Figure 1.1. Particles with low mass and high interaction strength should already be discovered. Collider experiments, such as ATLAS and CMS, search for heavy particles that can only be made with high energy collisions but nevertheless are as interactive as the low-mass common species. However, searching for heavier particles requires ever more energy; if a new particle is both heavy and feebly interacting, we may never generate and detect enough in collisions to declare discovery. Experiments like FASER probe the lower-

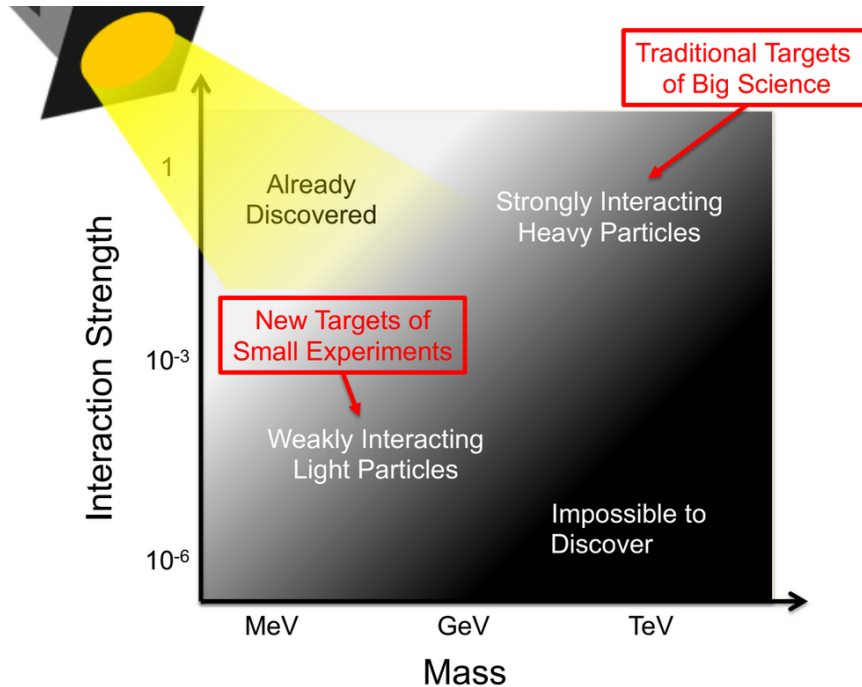


Figure 1.1: The Lamppost Landscape, broadly illustrating different particle searches [4].

left of the diagram, where particles have lower mass but interact feebly. Like neutrinos, such particles could be produced in great abundance but only rarely detected.

Low-mass particles, due to conservation of energy and momentum, are mostly produced along the axis of a collision where ATLAS and similar large collider experiments are the least sensitive. Until recently, this direction, called the “Far Forward” region, was a blind spot for collider physics. Near ATLAS, the far forward region is occupied by essential infrastructure such as the beam line and shielding, preventing the installation of any detectors along the line of sight (LOS) at that distance.

Fortunately, basic geometry and defunct maintenance tunnels provide a solution. The LHC, being generally circular, curves away from the ATLAS axis of collision. This LOS axis eventually intersects with the TI12 tunnel formerly used to transfer electrons and positrons from the Super Proton Synchrotron (SPS), an accelerator, to the now-decommissioned Large Electron-Positron Collider (LEP). This is the perfect spot for catching high-energy, feebly-

interacting particles produced along the ATLAS collision axis.

Between TI12 and ATLAS is approximately 100 m of earth and other dense material. This is a shield against almost everything “already discovered,” with only high energy muons and neutrinos capable of penetrating to TI12. Hypothesized feebly interacting particles such as A' or A_{B-L} could also reach FASER.

The unique opportunities offered by the TI12 site motivated the design and installation of FASER. This thesis covers FASER’s first search and results for A' and A_{B-L} .

1.1 CERN and its Facilities

CERN (Conseil Européen pour la Recherche Nucléaire) is an international organization founded in 1951 for studies in nuclear science. The organization is supported by contributions from 23 countries and provides the infrastructure to run high-energy physics experiments in Europe. As of writing, there are 12,000 “Unpaid Scientific Associates” (a.k.a. “users”) using CERN’s resources [5]. CERN facilitated the development of several accelerator and collider apparatuses that would eventually lead to the construction of the location repurposed for FASER.

CERN hosts many underground accelerator and detector facilities approximately 100 m under the Franco-Swiss border. Tunnels excavated in 1985 for LEP housed the Large Hadron Collider (LHC) from 2000 on [6]. The construction of these machines and their supporting infrastructure included many maintenance tunnels, one of which (TI12) would become FASER’s home. Figure 1.2 illustrates the modern network of LHC facilities.

Operation of the LHC is divided into two types of eras: “Run” or “Long Shutdown.” A “Run” is a period of several months or years during which LHC beam is active and exper-

The CERN accelerator complex *Complexe des accélérateurs du CERN*

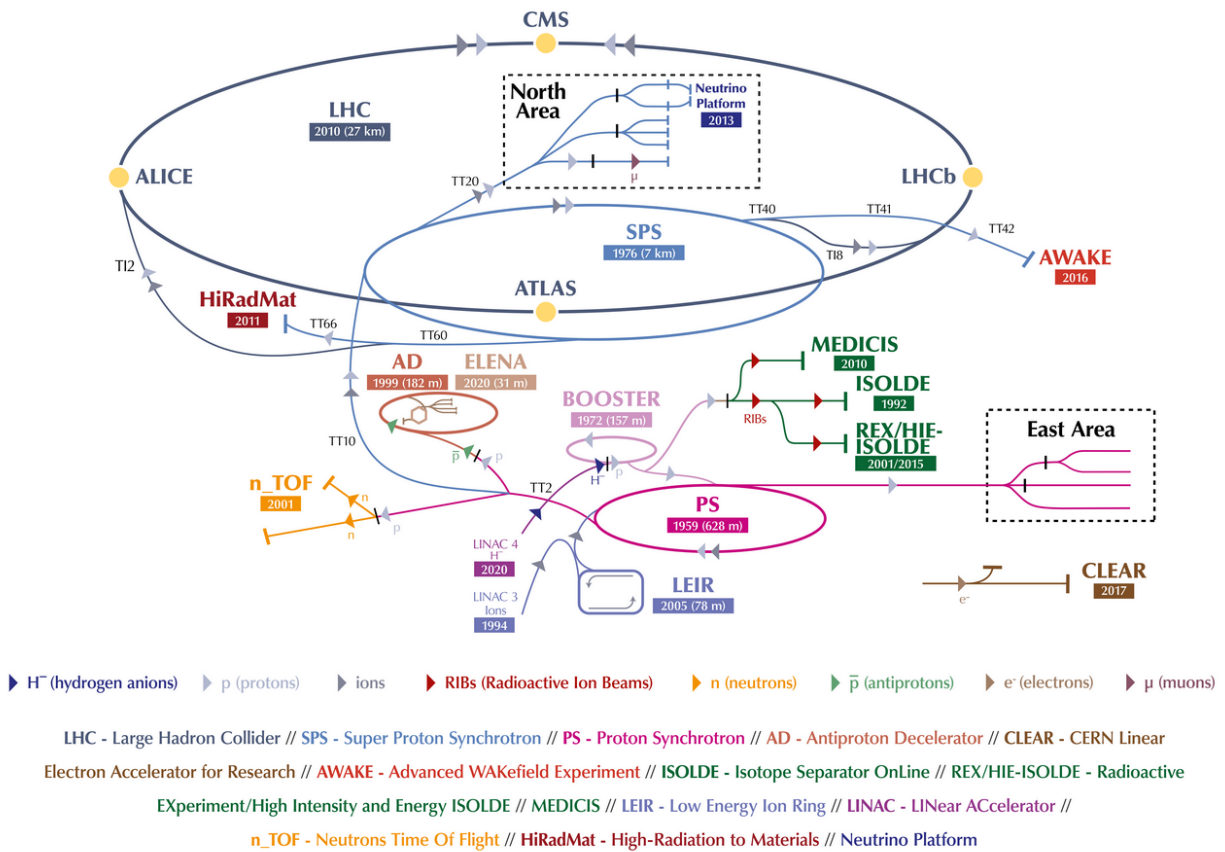


Figure 1.2: Diagram of LHC facilities as of January 2022, including SPS. Not to scale [7].

iments installed along its ring collect data. A Run can have several “technical stops” in which the beam is shutdown and the tunnels are accessed for small-scale activities. Periods between Runs are “Long Shutdowns.” During a Long Shutdown, various integral parts of the LHC, such as super-cooled magnets, are powered down and there is no beam for several years. These periods are for large-scale projects, such as the installation of new experiments. FASER was installed just before the start of Run 3, during the second Long Shutdown of the LHC’s lifetime (LS2).

Along with accelerator infrastructure, the LHC hosts several large particle physics experiments, including ATLAS. ATLAS is a general-purpose detector designed to make measurements on the high-energy LHC collisions. ATLAS is sensitive in all directions except where the colliding beams enter the detector.

1.2 FASER’s History

University of California, Irvine physicists first described FASER in their 2017 paper [8]. FASER’s founding collaborators published a letter of intent [9] and technical design report [10] a year later, introducing FASER as a cost-efficient alternative to heavy particle searches by looking in the far-forward direction yet-unprobed by collider experiments.

The cost-efficiency of FASER’s proposal expedited its approval in 2019. The original schedule expected commissioning and installation to be complete within a year.

COVID-19 shut down CERN for months. Work on FASER was almost exclusively software-related and remote for over a month. The earliest and most strict measures included determining “essential” activities and restricted access to electronics labs used for commissioning. For months after, FASER followed guidelines to prevent infection, such as cataloging visitors to different labs and limiting its occupants. CERN postponed Run 3 to 2022.

Despite delays, FASER was ready for data-taking before the start of Run 3. The months ahead of schedule were used to refine software and collect data from cosmic events. FASER collected data for all beam time in 2022. The collaboration presented FASER's first results — including those covered in this thesis — on March 29, 2023 at the 57th Rencontres de Moriond 2023 [11].

Chapter 2

FASEER Physics: The Standard Model and Beyond

This chapter is a brief review and discussion of contemporary physics models, dark matter, and relevant expansions on established theory.

2.1 The Standard Model

The Standard Model (SM) of particle physics is a resilient theory that describes the behavior and properties of matter. The observation of the final SM-predicted elementary particle, the Higgs boson, and its validation of the SM, is one of the greatest triumphs of science.

Although there is unprecedented experimental evidence in support of the SM, it remains an incomplete description of the universe. To account for all observed phenomena, we must go Beyond the SM (BSM), but not before understanding the SM itself.

The following section is supported by these works: [12, 13, 14, 15, 16].

2.1.1 General Features

Of the four fundamental forces, the SM describes all but gravity. The SM categorizes elementary particles as force-carriers (gauge bosons), matter particles (fermions) and one scalar boson (the Higgs boson). Every elementary particle species has a corresponding “charge” for each force and a mass.

Fermions are further classified as either quarks or leptons. Among other, more-nuanced differences, quarks only exist within composite particles (such as protons and neutrons), while leptons (such as electrons, muons, and neutrinos) *never* form composite particles.

Electromagnetism (EM) is mediated by the photon, the only massless boson. The EM force occurs when electrically-charged particles exchange photons, as in most atomic, molecular, and chemical phenomena.

The nuclear strong force (SF) is mediated by eight different gluons. SF charge, “color,” has six possible values: red, green, blue, anti-red, anti-green, and anti-blue. The SF unifies quarks in composite particles through gluon exchange; this is also the force that binds nucleons together. Each type of gluon corresponds to a permitted color exchange. Constituent particles must be “color neutral”: they must have equal amounts of red, green and blue, or equal amounts of a color and its corresponding anti-color.

The nuclear weak force (WF) is mediated by three bosons: W^+ , W^- , and Z . The W^\pm bosons, which also have electric charge, facilitate nuclear decay via a charged current interaction. The neutral Z boson facilitates neutrino scattering through a neutral current interaction.

The Higgs mechanism and its corresponding Higgs boson are associated with the Higgs force. This mechanism gives elementary particles, including the Higgs boson itself, their rest mass.

2.1.2 Mathematical Description

Quantum Field Theory (QFT) is the foundation of the SM. QFT is derived from the principles of quantum mechanics and special relativity; it asserts that each fundamental particle is an excitation of a corresponding field in spacetime.

Group theory, through symmetry operation matrices, is the mathematical mechanism that upholds conservation laws and their corresponding forces to all SM interactions. The SM is represented by the irreducible group: $SU(3) \times SU(2) \times U(1)$. The terms in this direct product are associated with different forces. $SU(3)$ models the SF. This form arose to support quark theory to explain and anticipate the variety of composite particles observed in experiments. EM and WF arise from Glashow–Weinberg–Salam Theory (GWS) as $SU(2) \times U(1)$. The “electroweak” forces reorganize into EM and WF when the electroweak symmetry is broken with the Higgs mechanism.

QFT uses Lagrangians to describe the state of a system, given a set of parameters related to the energy and momentum. The compact form of the SM Lagrangian is

$$\begin{aligned}\mathcal{L} = & -\frac{1}{4}F^{\mu\nu}F_{\mu\nu} \\ & + \bar{\Psi}(i\gamma^\mu D_\mu)\Psi + h.c. \\ & + \bar{\Psi}_i y_{ij}\Psi_j\Phi \\ & + |D_\mu\Phi|^2 - V(\Phi).\end{aligned}\tag{2.1}$$

The following is a surface-level description of what each term represents. The indices μ and ν refer to elements of the Minkowski metric. The *h.c.* terms refer to the Hermitian conjugate of the preceding term.

The first line represents the kinetic energy of the force-carrying bosons. $F_{\mu\nu}F^{\mu\nu}$ is the sum of the field strength tensors for all SM forces.

The second line represents fermion kinetic energy (Ψ) and fermion-force interactions with the covariant derivative D_μ . D_μ contains coupling strength constants for each force, which determine how readily the fermions interact with force-carriers, and the mechanisms that generate all fermions and gauge boson fields. γ^μ are Dirac matrices that represent the Lorentz boosting factor.

The third line describes fermion-Higgs interactions. The indices i and j iterate over all fermions and y_{ij} is the Yukawa coupling matrix, representing the Higgs mechanism that gives non-neutrino fermions their mass.

The last line describes boson-Higgs interactions and the Higgs self-interaction. Φ is the complex scalar field representation of the Higgs field. The $|D_\mu\Phi|^2$ term gives all other bosons their rest mass. Finally, $V(\Phi)$ represents the Higgs scalar potential, which includes the self-interaction mechanism that gives the Higgs boson its mass and enforces a non-zero vacuum expectation value (VEV).

The SM has 19 free parameters, including particle masses, coupling constants, and the Higgs VEV. The values of these parameters were determined experimentally.

2.1.3 Successes and Limitations

The SM successfully models particle behavior, culminating in the Higgs observation fifty years following its prediction. Despite its successes, the model has room for improvement. Its free parameters are entirely arbitrary, constrained only by observation rather than explained by theory. The SM further does not account for neutrino masses or how they oscillate between flavors. Most notably for this work, the SM does not explain dark matter's presence and effects. The following sections address the concept of dark matter and potential amendments to the SM that would accommodate it.

2.2 Dark matter and the Dark Sector

“Dark matter” (DM) is a broad term referring specifically to matter which is only known to interact through gravity. The need for such a term arose from significant discrepancies observed in cosmological data. In 1933, Fritz Zwicky published a study in which he used the virial theorem to estimate a galaxy cluster’s mass with the velocity of its galaxies and compared it to an estimation that used the galaxies’ luminosity [17]. The presence of additional, invisible matter best explained the discrepancy. Although its existence is supported by contemporary cosmological observations and there are many theories that can explain the data, there is no consensus on which theory is correct [18].

Cosmological studies suggest several properties for extant DM. Zwicky’s work supports that DM has little to no interaction with EM while behaving gravitationally similar to massive SM matter. Cosmic Microwave Background (CMB) studies inform estimates of the universal quantity of DM. The latest results reported by the Particle Data Group attribute 84.4% of the universe’s total mass density to DM and support that DM’s lifetime is long compared to cosmological timescales [19]. Gravitational lensing maps of the Bullet Cluster, in which galaxy clusters merged at relativistic speeds, suggest that DM does not readily interact with SM matter or itself [20].

DM’s production mechanisms are unknown. This work considers models in which DM is generated by thermal production — that is, high energy collisions. Such collisions were common in the hot, high-density, early universe following the Big Bang. As the universe continued to expand and cool, interactions between particles became less frequent. The DM production rate would decrease until production interactions were so rare that the amount of DM in the universe became effectively constant; this event is called the “freeze out.”

Hypothesized, thermally-produced DM candidates leftover after freeze out are “thermal relics.” “Relic density” is the amount of DM remaining after freeze out. The equation

for relic density is a solution to the Boltzmann equation¹:

$$\Omega h^2 = 2.755 \times 10^8 \frac{m}{\text{GeV}} Y_0(\langle \sigma v \rangle), \quad (2.2)$$

where Ωh^2 is the relic density of a particle, m is the mass of the particle, and Y_0 is the present-day ratio of the particle density to entropy density as a function of $\langle \sigma v \rangle$. $\langle \sigma v \rangle$ is the thermal average of the product of the annihilation cross-section σ and the relative velocity v of annihilating particles. The relic density’s dependence on the particle’s cross-section and mass constrains how heavy and interactive a dark matter candidate can be.

Dark matter’s stealthiness makes it difficult to study. Further, dark matter may not be just one new particle; it could be an entirely new family: “the dark sector.” Such models create many avenues for indirect DM observation through particles that mediate interactions between the SM and the proposed dark sector. Such a “mediator” — also called a “portal” — could rarely but visibly interact with one sector and decay into the other. This work considers a dark sector model with the dark photon as a mediator.

2.2.1 Dark Photon

A dark photon A' is an example of a mediating particle, named for how it mixes with the SM photon and a generic dark matter particle χ . Unlike the SM photon, a dark photon has mass and a finite lifetime. As with many hypothesized particles, its mass and interaction strength with SM particles have ranges of possible values. These parameters determine the dark photon’s lifetime and its decay products.

The dark photon is a vector boson. The kinematic relationship between the dark photon, the SM photon, and a new gauge boson is

¹See Chapter 9 of D. Majumdar’s text [18] for the derivation.

$$\mathcal{L} \supset -\frac{\epsilon'}{2}F_{\mu\nu}F'^{\mu\nu} + \frac{1}{2}m_{A'}^2X_{A'}^2, \quad (2.3)$$

where the new gauge boson is represented by the field strength tensor $X_{A'}$. $F_{\mu\nu}$ and $F'^{\mu\nu}$ are the field strength tensors of the SM photon and dark photon, respectively, with ϵ' as the coupling constant between the SM bosons and $X_{A'}$ and m' the mass of $X_{A'}$.

Rearranging this Lagrangian with field re-definition and rotating into the mass basis² gives

$$\mathcal{L} \supset \frac{1}{2}m_{A'}^2A'^2 - \epsilon e \sum_f q_f A'^\mu \bar{f} \gamma_\mu f. \quad (2.4)$$

This form shows directly how the new boson, now A' , couples with charged fermions, f . $m_{A'}$ is the dark photon mass and ϵ is the dark photon kinematic mixing parameter. The sum is over all SM fermions with electric charge q_f .

2.2.2 B-L Gauge Boson

“ $B - L$,” spoken as “ B minus L ,” is the difference between baryon number B and total lepton number L , which are accidental global symmetries of the SM. If only either B or L are promoted to global symmetries, the resulting theory would be anomalous. A “ $B - L$ ” gauge boson A'_{B-L} , and the addition of three sterile neutrinos with the right $B - L$ charges, allows a nonanomalous local U(1) symmetry. B and L are each conserved quantities, but do not have a gauge boson associated with either of them. However, if a yet-unknown high-energy interaction converts baryons into leptons, it could be mediated by A'_{B-L} .

This model introduces a $B - L$ charge, Q_{B-L}^f , for every fermion f :

²See Appendix A in [21].

$$\mathcal{L} \supset \frac{1}{2} m_{A'_{B-L}}^2 A'_{B-L}{}^2 - g_{B-L} \sum_f Q_{B-L}^f A'_{B-L}{}^\mu \bar{f} \gamma_\mu f, \quad (2.5)$$

where A'_{B-L} is associated with a mass $m_{A'_{B-L}}$ and gauge coupling constant g_{B-L} .

2.3 Long-Lived Particle Production at the LHC

A' and A'_{B-L} are classified as “long-lived particles” (LLPs) because their small masses and relativistic speeds allow them to travel significant distances before they decay. This section discusses the dominant production modes for A' and A'_{B-L} from proton-proton collisions at 13.6 TeV, the Run 3 center-of-mass energy.

When protons collide in the LHC, their combined energy and mass reforms into new particles. Low-mass collision products are produced at high velocity and their decays generally retain this forward momentum. Such products form collimated beams along the collision axis. Simulation studies, discussed further in Chapter 4, estimated the SM content of these far-forward beams. Preliminary tests found agreement between simulated and observed fluxes, which supported that low-mass, energetic particles would reach FASER [22]. These far-forward production rates, some shown in Figure 2.1, determined the production rates for LLPs reaching FASER.

2.3.1 Dark Photon Production and Decay

The dark photon has similar production mechanisms to its SM cousin for the given center-of-mass energy, depending on the dark photon’s mass. Meson decay and proton bremsstrahlung, their Feynman diagrams shown in Figure 2.2, are the most relevant production modes to this work. Others are considered subdominant and negligible.




Figure 2.1: IP meson spectrum for pions. θ is the angle with respect to the LOS and p is the momentum [23].




Figure 2.2: Dark photon (A') production via neutral-pion (π^0) meson decay (left), via dark bremsstrahlung (right) [23].

Mesons are produced in abundance in pp collisions; for 27 fb^{-1} , $\mathcal{O}(10^{15})$ are expected to be produced within 0.2 mrad of the beam collision axis, dominated by neutral pions π^0 . In the purely standard model process $\pi^0 \rightarrow \gamma\gamma$, the branching fraction is 0.99; the dark photon equivalent $\pi^0 \rightarrow A'\gamma$, accessible to dark photons with $m_{A'} < m_{\pi^0} \simeq 135 \text{ MeV}$, is suppressed by ϵ^2 relative to the SM process.

In the dark bremsstrahlung process $pp \rightarrow ppA'$, a dark photon is radiated from one of the colliding protons. This mode is accessible for $m_{A'} < \mathcal{O}(2 \text{ GeV})$.

Although dark photons can decay into any kinematically accessible charged states, they primarily decay into e^+e^- for the model parameter space considered in this analysis. The dark photon partial decay width Γ_e for this process, which includes electron charge e and mass m_e , is given by

$$\Gamma_e = \frac{\epsilon^2 e^2 m_{A'}}{12\pi} \left[1 - \left(\frac{2m_e}{m_{A'}} \right)^2 \right]^{1/2} \left[1 + \frac{2m_e^2}{m_{A'}^2} \right]. \quad (2.6)$$

This analysis assumes $m_{A'} < 2m_X$ and that A' decays into visible SM products. If $2m_e < 2m_{A'} < 2m_\mu$, then A' always decays into e^+e^- and thermal production occurs through $\chi\chi \leftrightarrow A' \leftrightarrow f\bar{f}$. For $m_{A'}$ in the MeV-GeV range and $\epsilon = 10^{-6}$ – 10^{-3} , the resulting thermal relic density accounts for a significant fraction of cosmological dark matter, and therefore such values are cosmologically favored [8]. The dark photon decay widths as branching ratios and decay lengths are shown in the leftmost plots in Figure 2.3.

2.3.2 B-L Gauge Boson Production and Decay

For $B - L$ gauge boson mass $m_{A'_{B-L}}$ in the MeV-GeV range, A'_{B-L} shares the dominant production mechanisms of the dark photon: light meson decay, and dark bremsstrahlung.

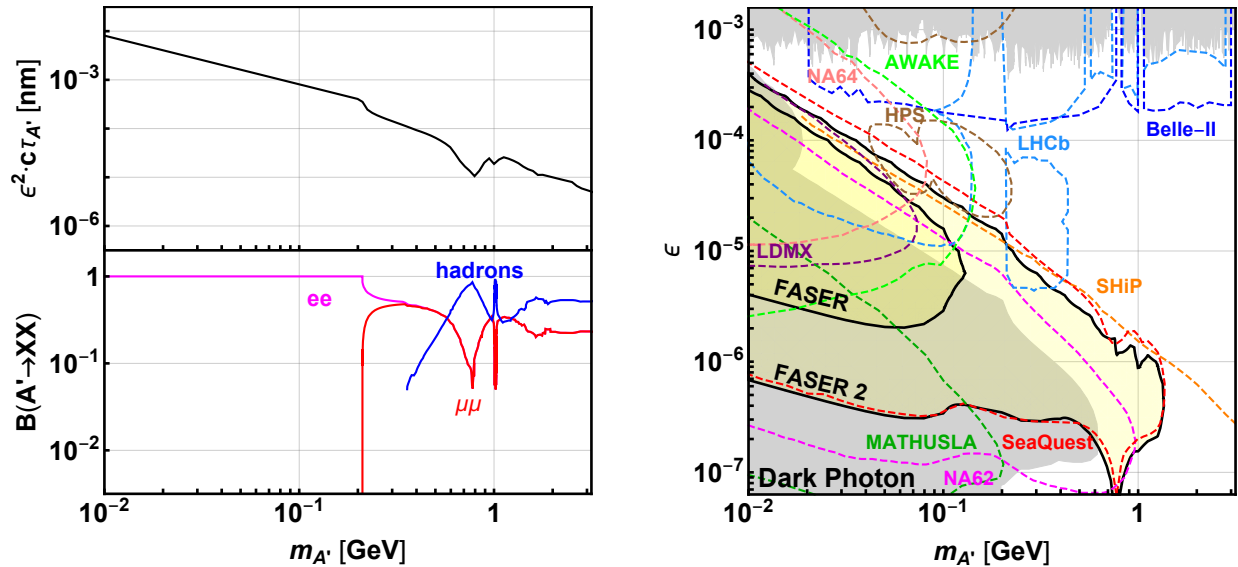


Figure 2.3: Dark photon decay length (top left), branching fractions (bottom left), and sensitivity for 150 fb^{-1} (right) [23].

A'_{B-L} can decay into anything with a $B - L$ charge. Unlike the dark photon, whose interactions mirror EM, A'_{B-L} interacts uniquely and complexly with the SM. Strategies for calculating its decay width are discussed in several, dedicated works [21, 24] which informed FASER's searches.

These studies concluded that for the given experiment parameters, A'_{B-L} primarily decays into leptons, particularly visible e^+e^- pairs and invisible neutrinos. The A'_{B-L} decay widths as branching ratios and decay lengths are shown in the leftmost plots in Figure 2.4.

Although this model requires sterile neutrinos, this study makes no assumptions about sterile neutrino mass. If a sterile neutrino's mass is less than half of A'_{B-L} , its production channels are accessible for the phase space probed in this study, leading to consequences for A'_{B-L} visible yield. A'_{B-L} has less frequent visible decays when sterile neutrino masses are kinematically accessible.

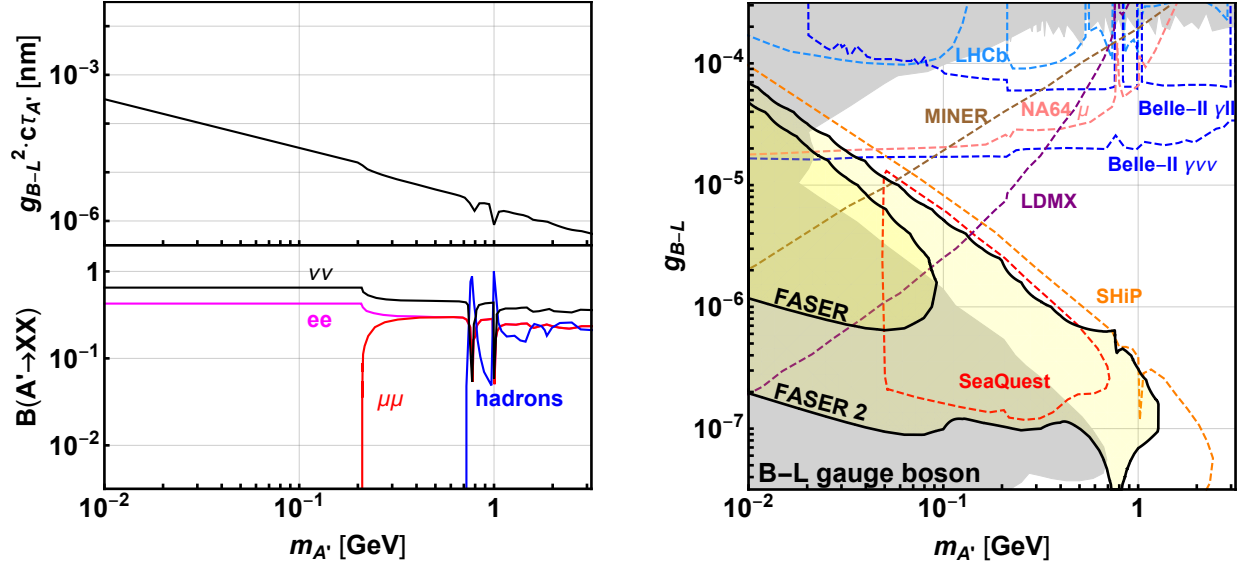


Figure 2.4: $B - L$ gauge boson decay length (top left), branching fractions (bottom left), and sensitivity for 150 fb^{-1} (right) [23].

2.4 Sensitivity

FORESEE, an event generator, used the simulated forward production results, the Lagrangian description for the new particles, and their calculated decay widths to estimate the spectra and sensitivity for those models, given FASER's dimensions and location with respect to the IP. Chapter 4 provides further details about which and how simulations are applied.

First, EPOS-LHC simulated the distribution and amount of far-forward products produced from pp collisions at the given energy. The branching ratios determined how likely A' or A'_{B-L} are produced within 0.2 mrad of the LOS. The spectra given by FORESEE determined the likelihood that A' and A'_{B-L} had energies greater than 500 GeV . The sensitivity studies also assumed a 50% signal efficiency.

The contours illustrating FASER's reach given half of Run 3's total integrated luminosity,

shown in the rightmost plots of Figure 2.3 and Figure 2.4, depend on the detector’s dimensions, location, and the amount of collected data. The top boundary depends on FASER’s distance from the IP. The general downward slope of the reach regions is a consequence of momentum conservation: heavier particles are boosted less. A heavier particle with enough energy to reach FASER needs a proportionally lower interaction coefficient, or else it will decay too quickly to be detected. A lower interaction coefficient also means a particle is produced less often. Collecting more data extends FASER’s reach further, since it will have additional opportunities to observe less frequent new-physics events.

For regions of A' or A'_{B-L} phase space in FASER’s reach, only the e^+e^- channel is visible. A'_{B-L} frequently decays into neutrino pairs as well, though these are not visible in this search. Early estimates demonstrated that FASER explores untouched phase space with as little as 1 fb^{-1} and continues expanding world-leading constraints with more data.

Chapter 3

Experimental Set-Up

FASER’s primary goal is to identify and distinguish high-momentum, neutral particles produced in ATLAS, 480 m away. ATLAS and LHCb donated the tracker modules and electronic calorimeters respectively, while other systems — such as the magnets — were designed and manufactured specifically for FASER. The terms “upstream” and “downstream” refer to the position of detector components as closer or further from the ATLAS interaction point (IP1), respectively.

The general layout of the detector is shown in Figure 3.1. The foremost upstream structure, an emulsion detector, is not relevant to this analysis except as a contributing source of neutrino background. The veto scintillator stations flag entering charged particles, while the downstream stations provide timing and energy loss information. A hypothetical, neutral long-lived particle (LLP) could pass through the veto and decay into energetic charged particles within the approximately-cylindrical empty region inside the most upstream magnet, sometimes referred to as the “decay volume.” Sensitive elements downstream of the decay volume collect measurements on the decay products. The experiment’s design minimizes the amount of material along the LOS, with the central tracker region contributing 2.1% of




Figure 3.1: A model of FASER showing the location of different sub-systems and its standard coordinate system. Particles from IP1 enter from the right side of the drawing [25].

a radiation length [25]. Chapter 6 further describes the new-physics signal, efficiencies and backgrounds.

3.1 Location

LHC houses many high-energy physics experiments, including FASER. Its approximately-circular structure is divided into 16 alternating “intersections” and “arcs.” Arcs are curved tunnel segments comprised of cryostat and superconducting magnet systems. Intersections are 528 m-long straight segments that may host experiments or other utilities. Each segment contains a piece of the beam line. The beam line’s main structures include two parallel beam pipes, also called vacuum chambers, surrounded by magnets; these contain high-energy particle beams when the accelerator is active. The LHC’s circumference is 27 km, making it the largest particle accelerator in the world [6].

There are three LHC vacuum systems. Two systems contribute to magnet temperature control, while the third maintains the vacuum inside the beam pipes. The presence of any gas inside the beam pipe would reduce the lifetime of a beam; therefore, optimal operation requires as complete a vacuum as possible. For example: in Run 1, for a 100-hour beam lifetime, the pressure in the beam pipe must be below 8×10^{-8} Pa at 5 K. The superconducting magnets were cooled to below 2 K for the energies achieved in Run 3.

Several major LHC structures relate to FASER’s location. ATLAS is installed at the “Point 1” intersection, surrounding IP1 — one of four “interaction points” (IP) where the high-energy beams are crossed to produce collisions. The CERN survey team confirmed that the line-of-sight (LOS) to IP1, also called the axis of collision, intersects maintenance tunnels TI18 and TI12 on either side of ATLAS. TI12 was chosen for FASER’s final location because it could accommodate the largest detector design [10].

Civil engineering and other technical infrastructure teams prepared TI12 and the surrounding infrastructure for FASER’s installation. Notably, TI12’s floor has an upward slope because it connects the LHC to the SPS above it. Civil engineering excavated a flat-bottomed trench such that FASER could be installed along the LOS, shown in Figure 3.3. The other technical infrastructure teams also extended electricity, airflow, and water utilities into TI12 and installed a crane to lift detector components safely over the beam line. These systems are labeled in Figure 3.2.

3.2 Beam Production

Although the LHC also collides heavy ions such as lead, only proton-proton collisions occur at a sufficiently high rate to produce useful data in FASER. Proton production begins in Linear Accelerator 4 (Linac4) with molecular hydrogen gas (H_2). Hydrogen and cesium are




Figure 3.2: CAD sketch overview of integrated supporting systems in TI12 [10].



Figure 3.3: TI12, before (left) and after (right) civil engineering intervention [10].

injected into a plasma chamber surrounded by a radio-frequency (RF) amplifier solenoid. RF-induced discharge ignites the hydrogen-cesium mixture into a plasma [26, 27]. The H-plasma undergoes four stages of RF-induced acceleration to reach 160 MeV before a graphite stripping foil strips its electrons [28]. The remaining nuclei — protons — are subsequently injected into the Proton Synchrotron Booster (PSB). The Linac4 beam is distributed into the PSB’s four, superimposed synchrotron rings; this division of the beam is the first stage of “bunch” formation [29].

“Bunches” are longitudinally-separated packets of particles formed by applying a focusing force, such as from high-voltage RF fields. Harmonic oscillations in the RF fields form the “buckets” that contain and accelerate bunches. The length of a bunch and the separation between them are typically measured in nanoseconds.

The PSB boosts the beam energy to 2 GeV before its injection into the Proton Synchrotron (PS). The PS divides the beam into as much as four batches of 72 bunches of 4 ns length. The PS boosts the bunches to 26 GeV then injects them into the Super Proton Synchrotron (SPS) [30, 31].

For LHC Run 3, the SPS boosts the batches from PS to 450 GeV. At this point in standard operation, each bunch is less than 1.6 ns long and contains approximately 1.6×10^{11} (1.2×10^{11}) protons for a beam with 25 ns bunch spacing [32].

Finally, the beam arrives in the LHC, injected into one of two parallel beam pipes: “Beam 1” begins with TI2 and circulates clockwise; “Beam 2” begins with TI8 and circulates counterclockwise. The dipole magnets require a 1.9 K operational temperature to generate the 8.3 T magnetic field strength to achieve the final 6.8 TeV beam energy, corresponding to 13.6 TeV center-of-mass energy. Once the SPS fills the LHC rings, it takes about 30 minutes for the beams to reach nominal energies [6].

The counter-circulating beams intersect for bunch collisions at the IPs. The beams typically

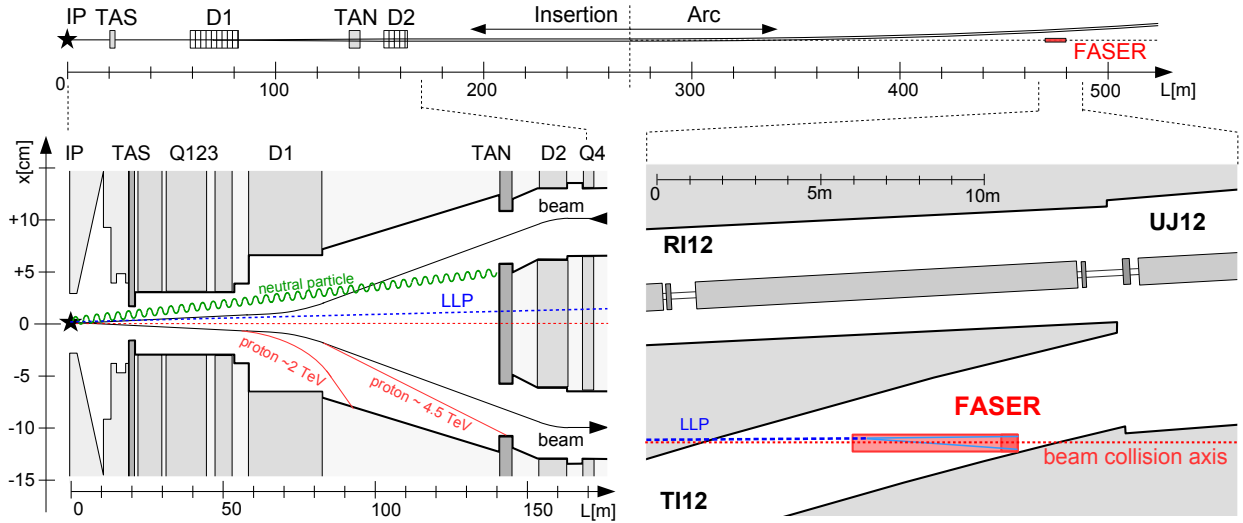


Figure 3.4: Schematic view of the LHC structures near ATLAS and FASER [10].

do not collide head on, but at a “crossing angle.” From Run 2 onward, the crossing angle was $160\mu\text{rad}$ with the direction mirrored yearly to mitigate uneven radiation damage to susceptible instruments. During 2022 data collection, the nonzero crossing angle caused a 7.7 cm downward vertical displacement of the beam collision axis at FASER’s location with respect to the default LOS. Fortunately, the collision axis is still within FASER’s 10 cm-radius aperture.

Although collisions can eject particles in any direction, we are only interested in the far-forward products produced along the LOS. As shown in Figure 3.4, several structures and other materials are between IP1 and TI12, which shield FASER from most far-forward products except high-energy muons and feebly interacting particles such as neutrinos or potentially LLPs. The focusing magnets near ATLAS deflect charged, far-forward products. The Target Absorber Secondaries (TAS) and Target Absorber Neutral (TAN), protect the superconducting magnets and, by extension, FASER from neutral hadron products. Any collision products must also penetrate 10 m of concrete and 90 m of rock to reach TI12 [23].

FASER made measurements on 27 fb^{-1} of integrated luminosity collected during the 2022 LHC Run 3. Run 3 expects 300 fb^{-1} total in proton-proton collisions.

3.3 Tracking

FASER’s tracking system measures positions and momenta of energetic charged particles traversing the detector. Four tracking stations are arranged into two sections: the interface tracker (IFT) and the spectrometer. The most upstream station, the IFT, enables electronic neutrino studies [33]. The remaining three stations downstream of the decay volume, along with the interleaved magnets, form the spectrometer and measure charged particle tracks for BSM searches. Each station has three tracking planes, which each house eight tracker modules [34].

3.3.1 Hardware

Spare ATLAS barrel-module semiconductor trackers (SCTs) [35] are the sensitive components of the tracking system. Figure 3.5 shows a single SCT module. Each module has two parallel-strip detectors glued back-to-back with a 40 mrad stereo angle and 80 μm strip pitch, giving them a spatial resolution of 17 μm in the precision coordinate (\hat{y} , according to Figure 3.1) and 580 μm in the non-precision coordinate (\hat{x}). The parallel-strip detectors each have six application-specific integrated circuit (ASIC) chips, which each read out data from 128 silicon strips. The modules are arranged in each plane to maximally cover the magnet aperture with sensitive hardware.

The modules operate at a 150 V bias voltage and a 3-4 V chip read-out voltage; these are referred to high-voltage (HV) and low-voltage (LV) systems, respectively.

Negative-temperature-coefficient (NTC) thermistors on each side of the module measure its temperature. A module has a maximum operating temperature of 35°C due to the thermal properties of the glue joining the parallel-strip detectors. Applying LV to the ASICs generates heat, therefore the modules must be cooled during operation. FASER’s cooling




Figure 3.5: An ATLAS SCT barrel module [10].




Figure 3.6: Differences in temperature across all strip detectors in a single tracker plane for different LV configurations. During data-taking, LV is on and configured (blue) [34].

system has two air-cooled water chillers that maintain a supply of water at 15°C. One chiller continuously circulates water, while the other is an emergency back-up on stand-by. The cooling power of the chillers is 1.8 kW at 15°C, significantly greater than the 576 W needed to keep all tracker modules sufficiently cold. During operation, there is a 3°C difference between inlet and outlet temperature. Figure 3.6 shows the typical operating temperatures for different LV settings for a single plane. The planes are flushed with dry air to prevent condensation by maintaining a dew point of -40°C or lower.

A tracking plane comprises eight modules affixed to an aluminum frame. An application of thermal paste improves the thermal contact between the modules and the frame. Flexible, printed circuit boards (a.k.a. “pigtailed”) connect the modules to external circuit boards called “patch panels.” The tracker plane is powered and read out by its patch panels. The patch panel also facilitates communication between the plane’s electronics and the Tracker Interlock and Monitoring (TIM) board, which monitors all of a tracking station’s sensors.




Figure 3.7: Sketch of a tracker plane. The yellow circle represents the magnet aperture [34].

Each plane has two patch panels, with each patch panel managing four modules. The patch panels are affixed to the outside of the tracker plane’s aluminum frame. Figure 3.7 shows the patch panel locations and the inlets/outlets for dry air and cooling.

Carbon-fiber plates cover both sides of a tracking station’s combined planes to make a closed-air system and to protect the delicate SCT modules. Since modules within a single plane cannot cover the aperture perfectly, the planes are vertically staggered such that a gap in one plane is covered by a different plane, ensuring that all charged particle tracks have at least four measurements per station [36]. An isolated, assembled tracking station is shown in Figure 3.8.

The tracker’s “backbone” is an aluminum frame that mechanically supports the three spec-



Figure 3.8: Photograph of an assembled tracking station with cooling tubes, patch panels, and carbon-fiber plate shown [34].

trometer tracking stations. In addition to supporting the spectrometer, it serves as a reference for tracker alignment. The backbone is clamped to FASER’s magnets. The backbone is shown alongside the detector’s other mechanical support infrastructure in Figure 3.19.

3.3.2 Calibration

Tracker calibration (*tcalib*) tests measure the response of the SCT modules and configure them to optimize performance. When a charged particle passes through an SCT, it induces a current pulse in the silicon strips. A “hit” is determined by whether the amplitude of the current pulse is above a pre-set threshold voltage. This signal is binary; the analog qualities of the pulse — such as its amplitude — are not recorded. Therefore, chip calibration is essential to maximize hit efficiency. Properly calibrated modules also have low noise and uniform threshold settings. Calibration tests were conducted at each stage of the tracking system’s

development and are regularly done *in situ* during gaps in beam activity. The calibration tests follow a procedure first developed by ATLAS. For more details about commissioning the tracker, see Appendix A.

Each of a module's twelve chips is manufactured with calibration circuitry that can inject charges (also called calibration pulses), bypassing the need for "real" hits. Every strip has an internal 100 fF capacitor whose voltage is controlled by an 8-bit digital-to-analog converter (DAC); an internal chopper circuit applies the changes as voltage steps. The capacitor's voltage setting (0-160 mV) determines the amplitude of the calibration pulses. A calibration pulse may also have as much as 50ns delay from its trigger, which typically comes from a clock signal provided by an external logic board. The DAC also controls the voltage of the discriminator threshold that determines whether an incoming signal is a hit or not; this setting can range from 0-640 mV and is typically set just above the noise level [37].

All *tcalib* tests use one or more variations of a "threshold scan." A threshold scan uses threshold voltage settings and injected charges to characterize a chip's performance. A threshold scan runs in steps that vary the discriminator threshold and inject fixed-amplitude calibration charges.

Each step in a threshold scan measures hit occupancy, the fraction of injected signals that passed the discriminator threshold, for every strip. In an ideal system, the hit occupancy as a function of threshold for a fixed charge is a step function, because the amplitude of the injected charge pulse is sufficient to pass the threshold, or it isn't. However, the pulses are subject to Gaussian electronics noise, which softens the step-function into an s-curve. An s-curve as shown in Figure 3.9 describes the probability that a signal will pass the discriminator threshold. The threshold for which the hit occupancy is 50%, labeled vt_{50} , corresponds to the injected charge amplitude. The width of a Gaussian distribution fitted to the s-curve corresponds to the amount of noise associated with the discriminator.




Figure 3.9: Illustration showing stages of a threshold scan [34].

The first tcalib test, the “mask scan,” identifies dead or noisy channels and masks them to improve the module’s hit efficiency. To identify dead (noisy) channels: the mask scan uses very low (high) thresholds and trigger signals without charge injection. Triggered channels below (above) the threshold are marked dead (noisy). Defective channels are tagged in an in-chip “mask register.” Of the modules lent to FASER from ATLAS, modules with the fewest masked strips were chosen for the tracker planes. The “best” modules were allocated locations where tracking was most important, such as the front of the spectrometer and covering the most magnet aperture. The best modules have at most 0.08% of their channels masked, while the “worst” modules have no more than 0.3%.

A strobe delay is the time between injecting a charge and the clock signal for each chip. The strobe delay test that follows the mask scan determines each chip’s optimum strobe delay. Strobe delay threshold scans maintain a fixed threshold and inject 2 fC charges, while the “strobe delay” setting is varied in discrete steps. The occupancy of each strip on one side of a module is stored in a 2D histogram, where the 2D coordinates are the strip number and the strobe setting. For each side of each module, these histograms are projected onto the strobe setting axis. Strobe delay settings that are too low or too high miss the timing signal to observe the injected charge, resulting in a low occupancy. The optimum strobe delay for each side of the module is calculated from the strobe delay settings at the start and end of a plateau in occupancy plateau:

$$SD = P_{start} + (P_{end} - P_{start}) * 0.25. \tag{3.1}$$

The best strobe delay setting (SD) falls within a finite range in which the occupancy “plateaus” above 95%. This range is found by checking when a projected occupancy is greater than 95% from both directions on the strobe-delay setting axis; P_{start} is the strobe delay setting when the plateau begins, and P_{end} is the setting when it ends. This setting,

optimized for one side of a module, is applied uniformly to each chip on that side.

Next is the “three-point gain” (3PG) test, which is run before and after a “trim scan” test. The first 3PG scan estimates the gain and noise of the signal before amplification (the “preamplification stage”). A single 3PG scan runs several threshold scans injecting different charges: 1.5, 2, 2.5 fC. These scans characterize the modules in several ways. The slope of the linear fit of vt_{50} vs. input charge is the gain in mV/fC and its intercept value is the threshold offset, in mV. Noise associated with the discriminator threshold in mV for 2 fC injected charge, also called the Equivalent Noise Charge (ENC), is calculated as the standard deviation of the threshold scans’ s-curve fits ($\sigma_{vt_{50}}$) divided by the gain. The average gain over all strips is 54 mV/fC which agrees with the 55 mV/fC measured from earlier tests done by ATLAS [35].

The “trim scan” assesses and calibrates the circuitry that corrects variations in the discriminator threshold. A discriminator threshold at its optimal threshold target 1 fC setting may perform differently than expected as it accumulates radiation damage; this variation is called the chip’s threshold dispersion or spread. To correct this offset, maintain uniform efficiency across all strips, and extend the operational lifetime of the module, each strip has a 4-bit DAC (TrimDAC). The TrimDAC applies a correction voltage in one of four ranges; higher-voltage ranges are used for higher threshold spread. Ideally, all chips operate with minimal correction in the lowest range: 0-60 mV. Trim threshold scans inject 1 fC charge for different trim settings. A channel is “trimmable” if an observed offset from the threshold target can be corrected. The optimal trim setting is that which maximizes the number of trimmable channels. Commissioning and post-installation calibration found that most (99.9 %) channels are trimmable with the 2 lowest trim ranges, as expected of undamaged modules; 0.05% of channels required a larger trim range and the remaining 0.05% were not trimmable at any setting.

The final scan that adjusts settings is the “response curve” or “10-point gain” scan. Similarly

to the 3PG scan, it runs threshold scans for ten different injected charge settings: 0.5, 0.75, 1.0, 1.25, 1.5, 2, 3, 4, 6, 8 fC. Response curve results measure gain as a function of injected charge. The dependence of vt_{50} on injected charge is fitted to a first-degree polynomial

$$\tau = p_2 + \frac{p_0}{1 + e^{-q/p_1}} \quad (3.2)$$

that describes the voltage to charge conversion associated with the discriminator threshold setting, where p_0 , p_1 , and p_2 are the fit parameters and q and τ are the threshold values as charge and voltage, respectively. A non-linear response curve suggests the corresponding strip is noisy.

An additional “noise occupancy” test is not part of most routine calibration cycles but is conducted to assess the effects of the preceding calibration steps. This scan measures the dependence of a chip’s noise occupancy on its threshold setting. The noise occupancy of a strip is the probability that it will generate a signal from noise, such as when there is a fluctuation in the discriminator threshold. Noise occupancy threshold scans do not use injected charges; instead, the threshold voltage is varied. As the threshold voltage is increased, we expect fewer noise-induced hits. The scans continue to run until enough data has been collected. 99.7% of strips have less than the expected upper limit 5×10^{-4} at 1 fC and nominal operating temperature.

3.4 Scintillators

3.4.1 Hardware

FASER’s four scintillator stations respond to the passage of charged particles. A single scintillator module has scintillator material, a Photo Multiplier Tube (PMT), and a light

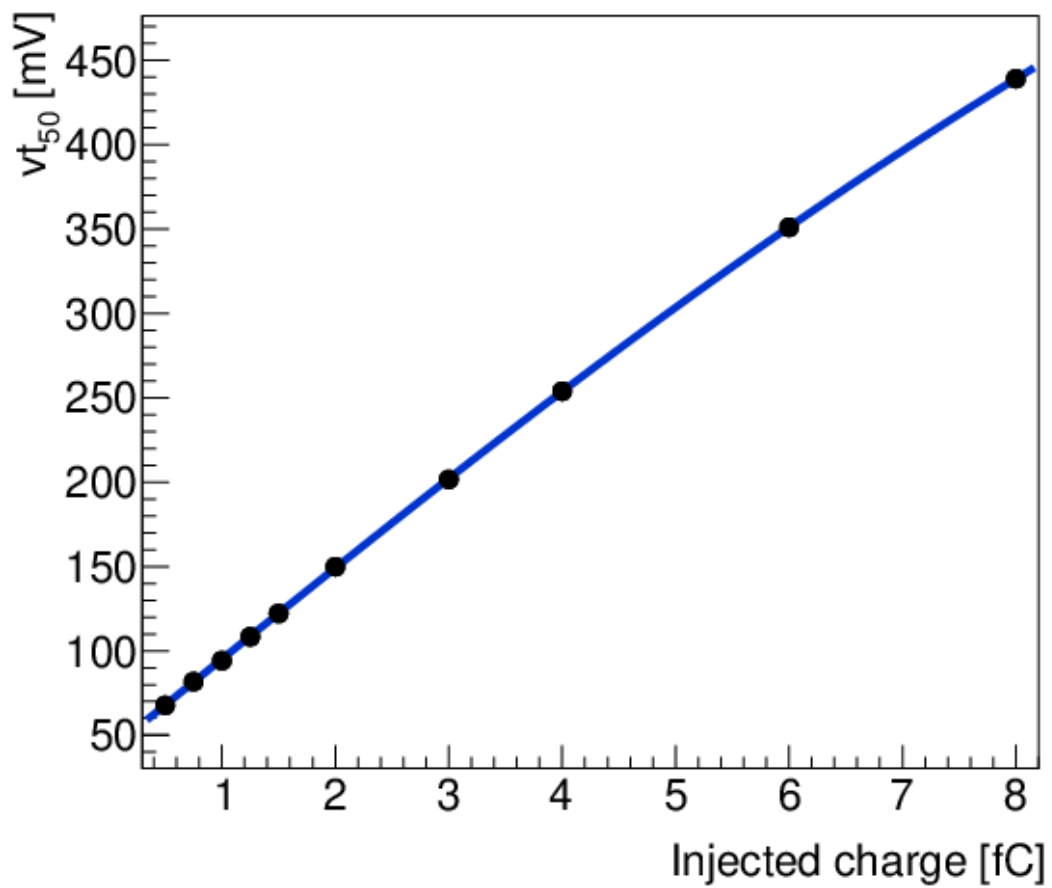


Figure 3.10: Response curve of a single strip, showing ten different injected charges (black dots) and linear fit (blue) [34].

guide or wavelength shifting (WLS) rod. The scintillator material is a rectangular block of 10-20 mm-thick EJ-200 plastic and the active material of the scintillator. Depending on the station, a wavelength shifting (WLS) rod or plastic light guide channels scintillation photons to the module's PMT. The PMT translates scintillation photons into electronic signals. An open-ended optical fiber can inject LED calibration pulses on command. A layer of nickel-iron (permalloy) protects the PMTs from interference caused by FASER's magnets. 0.5 mm-thick aluminum foil covers the modules for light tightness and fire safety. Along with these common factors, each scintillator station has unique characteristics suited to its purpose. Combined, their thickness corresponds to 2.5-5% of a radiation length.

The first two stations, shown in Figure 3.11, are on either side of the emulsion detector. The upstream station, VetoNu, has two identical scintillator modules. The space around the emulsion detector is limited, so VetoNu uses smaller, WLS rods rather than the light guides used by the other stations, and Hamamatsu H11934-300 PMTs instead of Hamamatsu H6410. Notably, while the WLS rods are smaller, their timing measurements are less precise due to their long emission time. VetoNu signals distinguish emulsion detector neutrino events from other upstream background sources. The downstream veto station (Veto) sits between the emulsion detector and the decay volume; it is composed of two pairs of modules separated by a 10 cm-thick lead block. The lead block absorbs or creates detectable showers from photons generated by muon bremsstrahlung. The presence of any veto signal above an analysis-defined threshold classifies an event as background.

The third scintillator station, also called the "Timing" station, is installed between the decay volume and the spectrometer. Two co-planar modules (shown in Figure 3.12 on the left) are arranged to completely cover the magnet aperture with an extended margin, such that any muon events that missed the veto stations could then be detected here. A timing station's signal is the primary trigger that prompts the read-out of all other FASER detector systems. Timing data contributes to the calculation of the arrival time of a FASER event relative




Figure 3.11: CAD sketches of the upstream (left) and downstream (right) veto stations [25].

to the time of the corresponding pp collision in ATLAS. To optimize the timing precision, the signal is transmitted by wave-guide-PMTs on both sides of each layer in a module. The horizontal time-walk is canceled when averaging the arrival time from each pair of PMTs.

The final “Preshower” scintillator station, shown in Figure 3.12 on the right, is installed between the downstream end of the spectrometer and the calorimeter; it consists of two modules, each with a 3 mm-thick tungsten radiator on their upstream side. There is also a 5 cm-thick graphite absorber between the preshower station and the preceding spectrometer. This station provides additional timing information and measurements on any backwards-ejected products (back-splash) from high-energy particles interacting with the calorimeter. The graphite absorber shields the spectrometer from these back-splash particles.

The standard operating voltages and currents for the scintillator PMTs are shown in Table 3.1.




Figure 3.12: CAD sketches of the timing (left) and preshower (right) scintillator stations [25].

Table 3.1: Nominal values for scintillator operation [38].

Station	Voltage (V)	Current (μA)
VetoNu	-1100 — -800	270-330
Veto	-1900 — -1400	360-460
Timing	-1900 — -1400	360-460
Preshower	-1100 — -800	270-330

3.4.2 Calibration

All FASER PMTs underwent tests to select the best-performing units to use in the final installation.

Absolute gain characterizes the PMT signal amplification; in Hamamatsu PMTs, the gain is proportional to the operating voltage. The absolute gain of each PMT was measured in a specialized dark room using single photons from a low-intensity LED and scanning different PMT voltage settings. All PMTs tested had similar voltage dependence, within 15% of one another. Each station was given PMTs which had the most-alike voltage dependence between them to improve the effects of voltage calibration.

PMTs record a dark current of random noise, a small electric current that persists passively due to quantum activity in its active material, even when the PMT is not illuminated. Tests in darkness near nominal voltage measured dark rates of 10 Hz or less.

The minimum ionizing particle (MIP) signal was measured with cosmic ray muons in an above-ground lab. Two scintillator modules were suspended over a tracking station upon a table. Each scintillator module measured over 50,000 cosmic muon events. The results for a single veto module are shown in Figure 3.13. After installation, gains for PMTs within the same scintillator station were adjusted such that their MIP response was nearly the same.

3.5 Calorimeters

3.5.1 Hardware

FASER's calorimeter system measures the total electromagnetic energy deposit in four shashlik-type calorimeter modules donated by LHCb [39]. Each module has 66 layers of

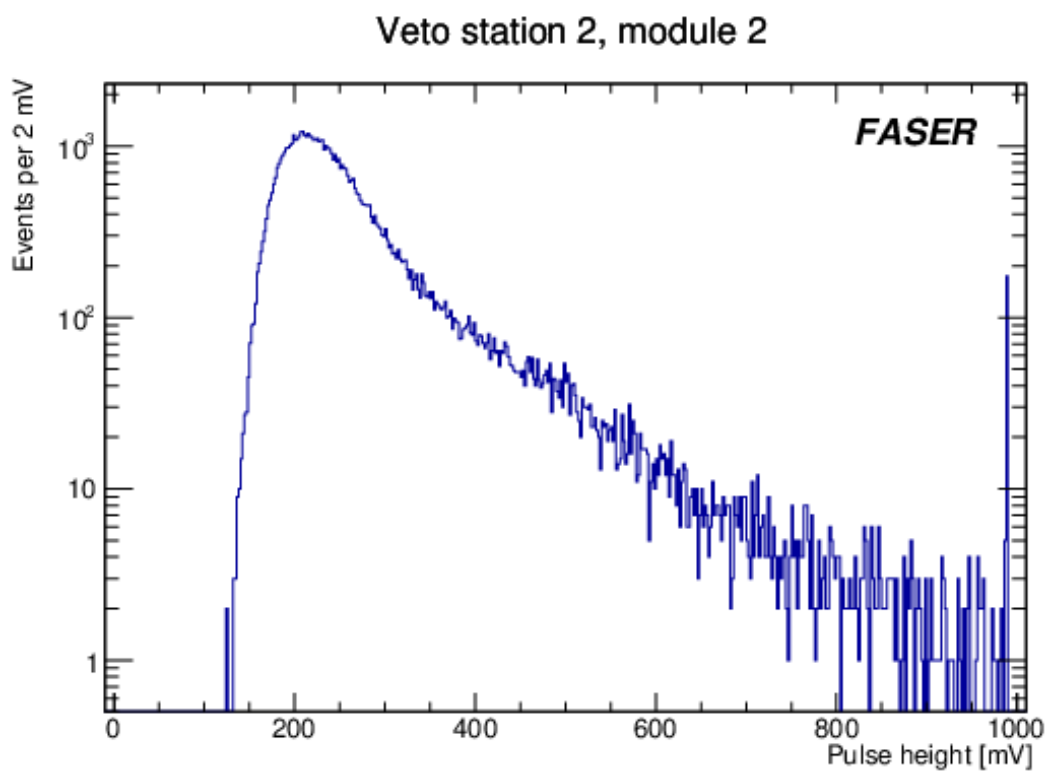


Figure 3.13: PMT response to cosmic ray muons, with saturation at 1000 mV [25].




Figure 3.14: Sketch of FASER calorimeter module showing all fibers leading to a single PMT [25].

interleaved 121.2 mm x 121.2 mm 2 mm-thick lead, 4 mm-thick scintillator plastic, and Tyvek reflective paper. A module's length is 754 mm when including the readout PMT, as shown in Figure 3.14. The layers are penetrated by 64 uniformly-spaced wavelength shifting fibers that transmit light from the scintillating material to a Hamamatsu R7899-20 PMT at the back of the module. Since there is only one PMT per module, there is no longitudinal segmentation and only coarse transverse segmentation.

FASER's magnet aperture is completely covered by the four modules, which are angled horizontally and vertically by 50 mrad with respect to the beam line to minimize insensitive regions along the fibers and between modules. The gaps between modules are additionally covered by aluminum tape, and the entire calorimeter system is protected by 0.5 mm-thick aluminum plates for light tightness. A grounded, aluminum Faraday cage shields the modules from electronic noise generated by a nearby tunnel safety system antenna. Capacitor-driven blue LEDs assist with calibration, and an 8-bit microcontroller controls the entire calibration system. The assembled calorimeter without the Faraday cage is shown in Figure 3.15.

t




Figure 3.15: Photograph of assembled FASER calorimeter with four modules before the Faraday cage installation [25].

3.5.2 Calibration

Since the calorimeter modules were donated from a pre-existing experiment, we can use previous test results to determine the function and quality of the modules in FASER. Figure 3.16 shows just that, where LHCb data is compared to the expected energy resolution in FASER given by

$$\frac{\sigma_E}{E} = \frac{9.2\%}{\sqrt{E}} \oplus 0.2\%. \quad (3.3)$$

This follows the general 3-term description of energy resolution. The first term is the “stochastic” term, which is related to a statistical approximation of how particle showers form in a calorimeter. Second is the “noise” term, referring to an electronic noise the measurement may have; it is expected to be small in photosensitive devices such as PMTs and is omitted from Equation (3.3). Finally, the “constant” term includes all contributions that do not depend on energy. These terms and other calorimetry concepts are described in [40]. Notably, we can compare Equation (3.3) to the LHCb energy resolution equation shown in Figure 3.16. FASER’s energy resolution was estimated with simulation, while the coefficients for the LHCb energy resolution were derived from measurements.

A calorimeter’s energy response is typically calibrated by measuring a frequent, known process. However, the only known particles frequently reaching FASER are muons, which do not deposit all of their energy into the calorimeter¹ and cannot usually be detected at the nominal PMT gains during FASER operation. Calibrating the FASER calorimeters then requires taking muon measurements at higher gains and extrapolating the results into the operating range. Blue LED light injections at equidistant, descending voltage steps ensure the extrapolation is accurate. Measurements from test beams with known particle types and energy also supported the values used for calibration [41].

¹Muons have a most probable energy of 330 MeV.

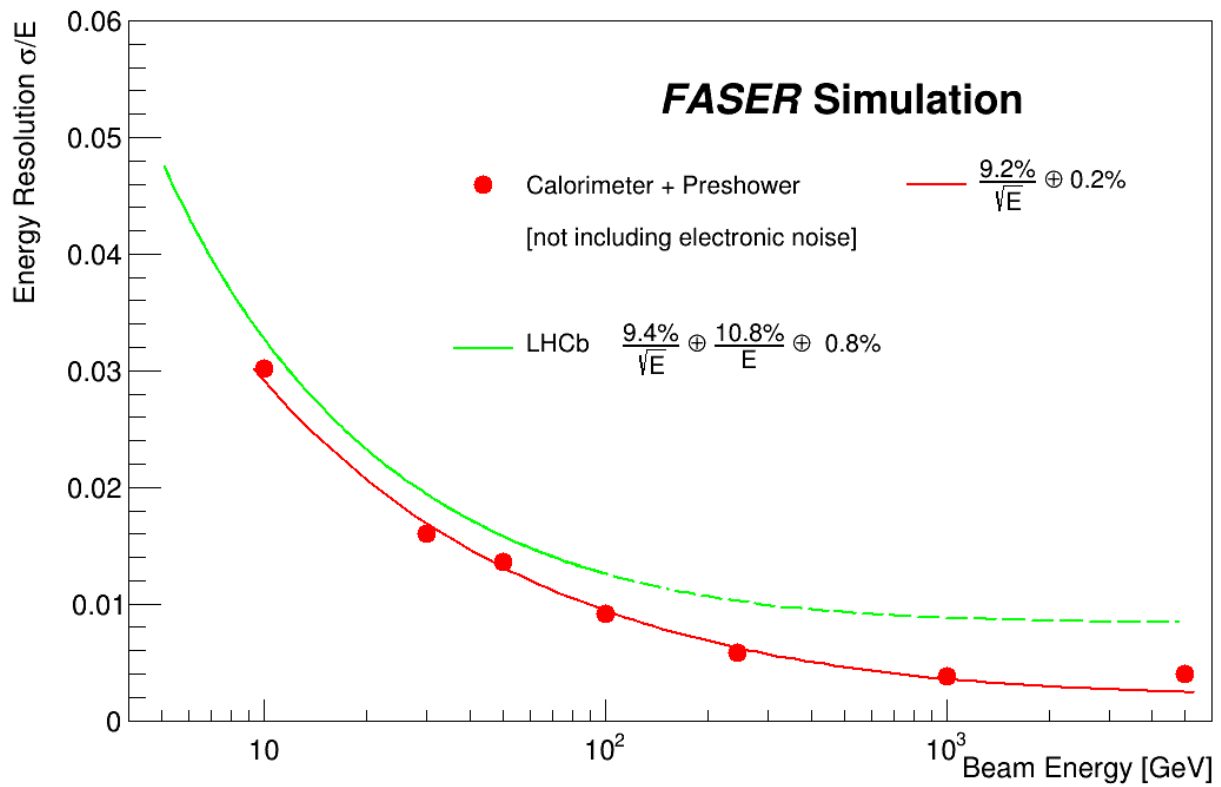


Figure 3.16: Comparison of Calorimeter resolution between LHCb data and FASER simulation. The dashed extension of LHCb data is an extrapolation to higher energies [25].

3.6 Magnets

FASER’s magnet system allows momentum measurements, and charge identification and spatial separation of charged particle tracks essential for identifying new physics signals. One primary purpose of the magnet system is to further separate the charged, highly collimated decay products of LLPs. For momentum less than 1 TeV, track curvature induced by the magnetic field can also tell us the charge and momentum of the track’s particle. Most of FASER’s initial budget was allocated to the design and manufacturing of these magnets. Their dimensions determined FASER’s final size, location and physics reach. Early studies on far-forward LHC particles identified two thresholds that FASER could meet, which depended on the radial spread of each kind of far-forward-produced meson. Most pion-decay products are within 10 cm of the LOS, while the same proportion of B-meson decay products could have a spread as far as 50 cm from the LOS. Since a magnet large enough to contain B-meson products was beyond FASER’s target budget, the pion threshold determined the size of the detector.

The magnet system consists of three cylindrical, Samarium Cobalt $\text{Sm}_2\text{Co}_{17}$ permanent dipole magnets. All FASER magnets have a 43 cm outer diameter and 20 cm aperture. The most upstream magnet, which encompasses the decay volume, is 1.5 m long. The remaining two downstream magnets are each 1 m long. The magnets have a total mass of 3.16 metric tons. Figure 3.17 shows the dimensions and significant features of the cross-section of the FASER dipole magnets.

The magnets were designed and evaluated with several software packages: Finite Element Method Magnetics (FEMM) and Opera 3D/TOSCA. All three magnets have a Halbach array design [42]: 16 trapezoidal magnets arranged in an aluminum frame to generate a homogeneous 0.57 T interior field, shown in Figure 3.18. The stray field falls below 10 mT at 25 cm from an aperture. Optical references installed on the magnet’s exterior surfaces inform





Figure 3.17: Cross-section of FASER magnet design showing approximate internal field direction of each segment [25]. This early drawing and Figure 3.18 are rotated by 90 degrees from their final orientation.

changes to its alignment, facilitated by its supporting structures, which allow adjustment in three directions and a rolling angle. When any detector work is needed, plastic shields cover the magnet apertures to prevent material from entering its interior volume; the covers are removed during data collection to minimize the material in the detector [25].

3.7 Emulsion Detector

The emulsion detector is a 30 cm × 30 cm × 1.35 m aluminum box containing 770 30 cm × 25 cm tungsten plates interleaved with emulsion film, corresponding to 220 radiation lengths and 7.8 hadronic interaction lengths. Its primary purpose is for neutrino studies, with tungsten as the neutrino target and the emulsion film for interaction observation. Although this analysis does not use emulsion detector data, its presence affects background studies, which are described in Chapter 6.






Figure 3.19: Sketch of structural and mobility systems for the magnets and spectrometer [34].

3.8 Mechanical Support Structures

FASER has several mechanical support structures, shown in Figure 3.19, that give it mobility and additional stability.

The detector is seated on baseplates: two layers of machined aluminum plates. These plates, along with a hydraulic jack pushing system, can move the detector as much as 5 cm towards the TI12 tunnel wall. Adjustable supports attach FASER’s magnets directly to the upper baseplate. The position of the LOS and crossing angle could change during routine LHC operation; the baseplates allow FASER to move with them as needed.

The “upper frame” is an aluminum profile structure comprised of various thin beams meant to support non-detecting hardware near the detector. It also serves as the underlying skeleton for other systems, such as cooling, or a place to tie extra lengths of cable.

To keep all metallic components of FASER equipotentially bonded, the entire detector is connected via two copper cables to the main LHC earth.

The detector paper [25] has additional details about FASER’s supporting structures.

3.9 Software

FASER software is divided between “online” and “offline.” “Online” software, described in Section 3.9.1 and Section 3.10, manages operation and monitoring of the detector. “Calypso” is FASER’s offline software framework which performs simulation (Chapter 4) and reconstruction (Chapter 5) independent of the detector’s status. FASER has benefited from pre-existing projects by adapting them to suit its own, often simpler needs. FASER’s software development and usage is documented and stored on GitLab [43] and an internal CERN Twiki [38].

3.9.1 Data Acquisition

During 2022 operations, the trigger and data acquisition system (TDAQ) is designed to accommodate data collection occurring at a trigger rate up to approximately 1500 Hz at luminosity $2 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$, dominated by muons. TDAQ software is implemented through DAQling [44], an open-source C++ framework.

The data acquisition process is represented as hardware in Figure 3.20 and software in Figure 3.21. Data acquisition begins when a scintillator or calorimeter PMT measures a signal exceeding a preset threshold. Such a signal can come from a high-energy particle or a pulse from the calibration system. A 14-bit CAEN VX1730 digitizer board digitizes the PMT pulses and propagates the signal to the Trigger Logic Board (TLB). The TLB issues a

global, single-bit Level-1 Accept (L1A) trigger that prompts each subsystem to readout data from its buffer. The triggered readout boards each send their data fragments via Ethernet to a server running the event builder on the surface. The event builder combines fragments sharing an event identifier (ID) into a single event, as described in Section 3.9.2. Events are stored in a buffer until they can be written to a file. Stored events are further prepared for analysis with processes described in Chapter 5.

Read-out data comes from several sources: the tracker system, the PMTs, and TLB.

A Tracker Readout Board (TRB) controls and manages signals from the SCT modules of a single tracker plane. The spectrometer's nine TRBs and the three IFT TRBs are installed in aluminum mini-crates. Twinax Firefly cables connect each SCT module by its patch panel interface to the plane's TRB. In data-taking mode, the TLB sends the TRB clock signals at 40.08 MHz and a bunch counter reset signal (BCR). The TLB forms a bunch counter ID (BCID) by counting the number of LHC clock cycles passed between the BCR and the L1A signal. A single clock cycle corresponds to one bunch crossing (25 ns). Data read from module chips is temporarily stored in the TRB's memory before being sent by Ethernet to the event builder. The SCT hit readout window spans three clock cycles. The readout delay can be finely tuned to 390 ps. SCT readout is three bits derived from three time bins, one for the bunch crossings before, shared with, and after the L1A signal. The TRB can determine which 3-bit hit patterns are read out [34]. Figure 3.22 shows an overview of this system, including powering and environmental control.

The CAEN digitizer board reads out all PMT channels and continuously digitizes signals from them at 500 MHz. During data collection, digitized signals are stored in a buffer until an L1A signal from the TLB triggers readout. A scintillator PMT's threshold for initializing the read-out process is set to be less than one MIP and the calorimeter threshold is set to respond to energy deposits greater than 20 GeV.

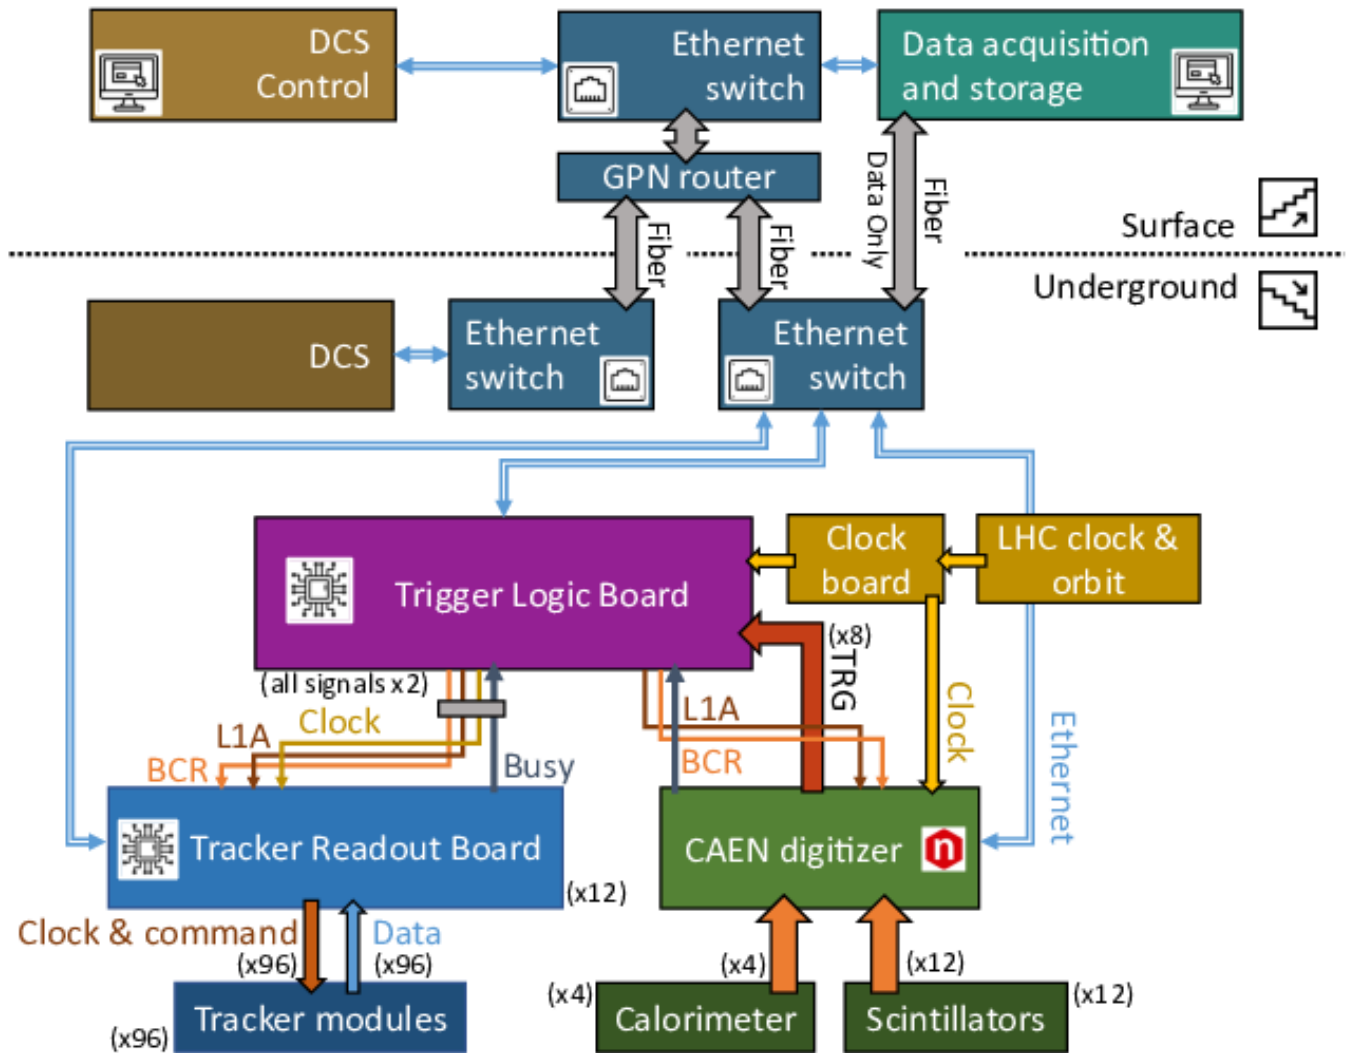


Figure 3.20: Overview of TDAQ hardware network. Numbers in parentheses indicate multiple units/channels [25].

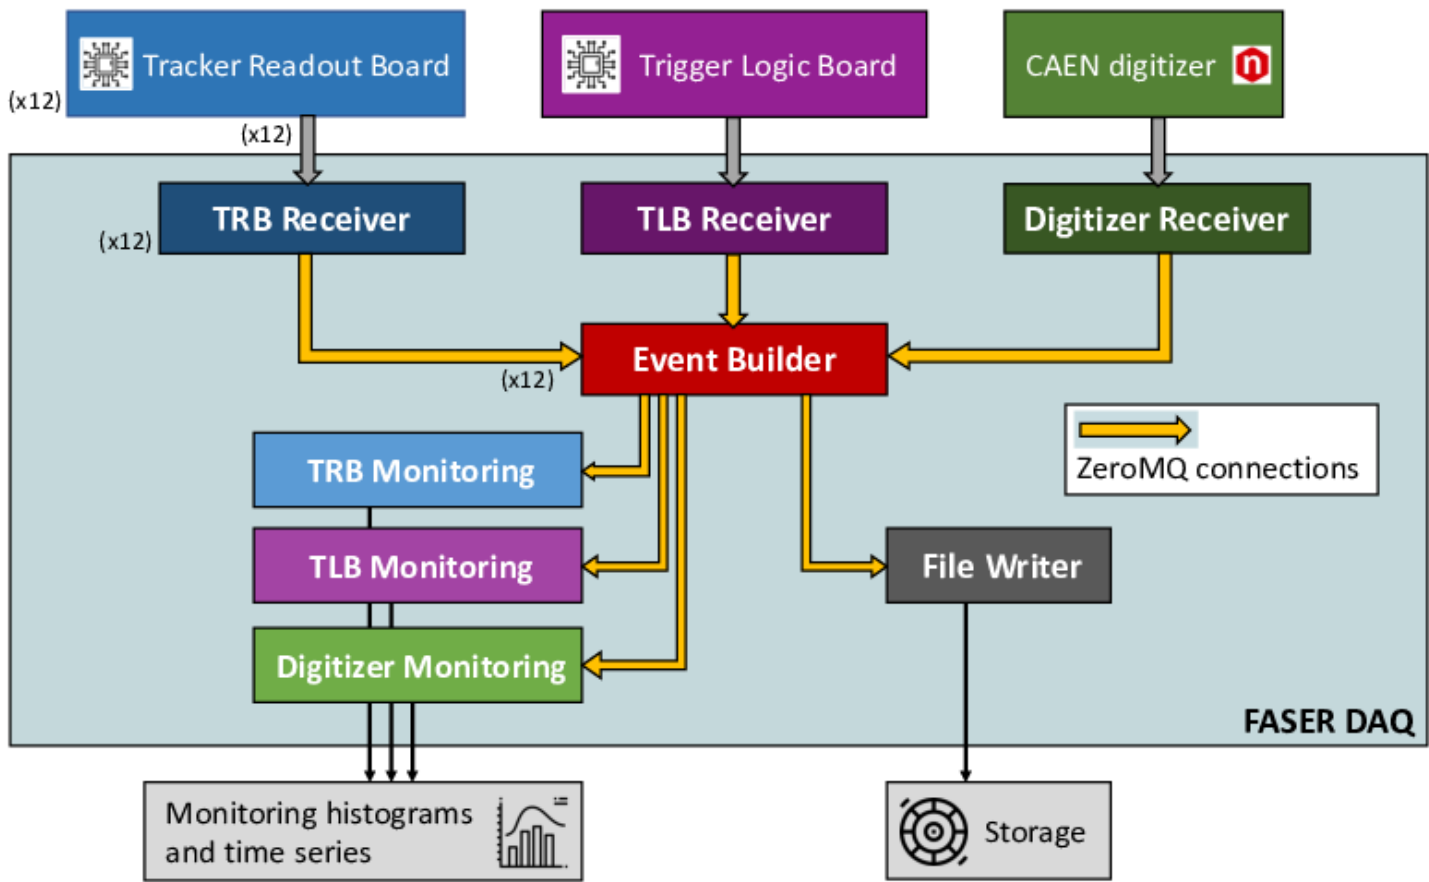


Figure 3.21: Overview of TDAQ software. ZeroMQ is a messaging software library [25].

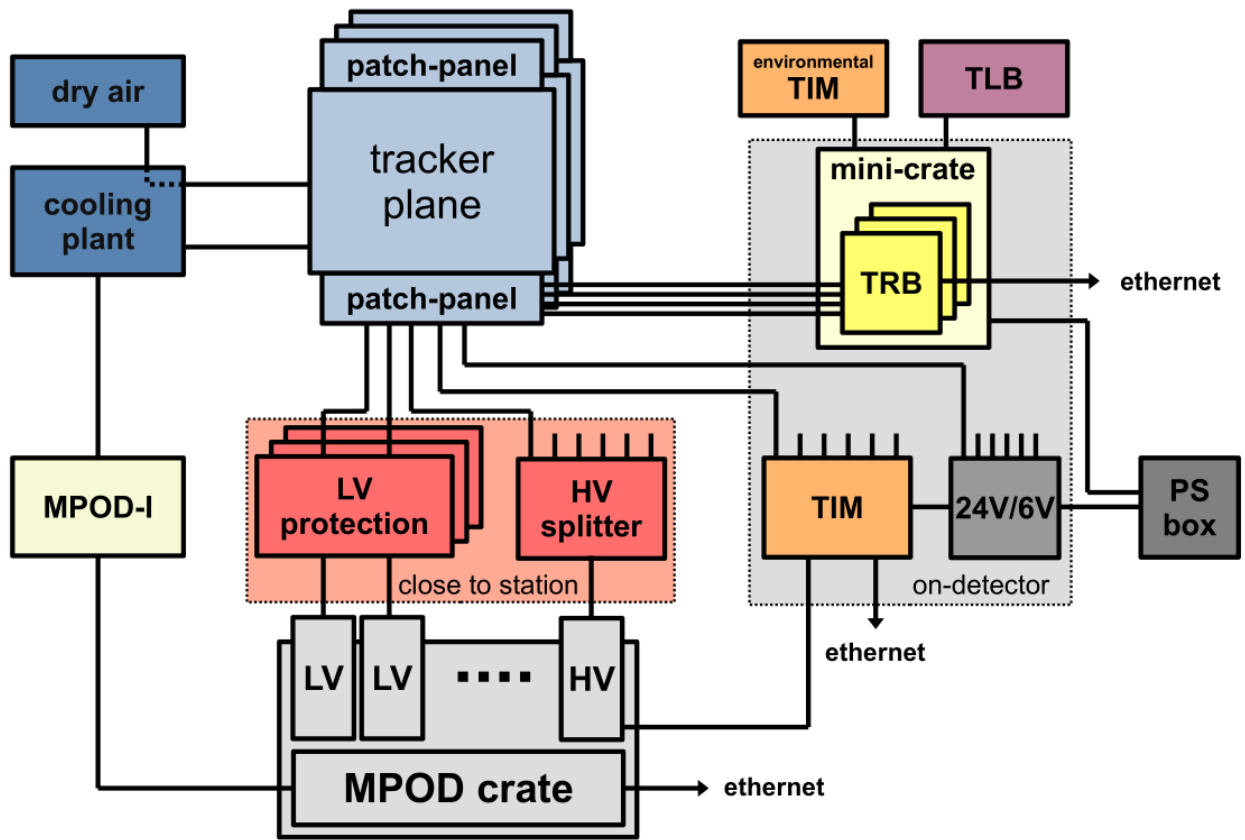


Figure 3.22: Overview of the tracker control system [34].

FASER depends on timing signals sent by the LHC beam synchronous timing (BST) system. FASER’s clock board has a “BST receiver interface for beam observation” (BOBR) [45] system that synchronizes FASER with LHC timing. The BST arrives with expected, non-negligible noise and jitter which must be cleaned to synchronize FASER even when the system is reset or there is a power cycle. The clock board is installed near the TLB in a VME² crate.

3.9.2 Data Processing

Data goes through several stages of software-based processing before it is analyzed. These software structures, typically programming objects, are standardized under an event data model (EDM). A good EDM ensures that the data can be easily accessed and used by different algorithms within the experiment’s software framework. FASER’s software framework, Calypso, uses the xAOD EDM from ATLAS’s Athena framework [46]. The following is an overview of how data is organized at different stages of processing within the event builder.

A data fragment is the basic unit of data for FASER. All data fragments have the same structure and header format. The header contains information about which subdetector measured the data. The body of the data fragment contains read-out data from the system that issued it.

The event builder assembles the data fragments into a raw data object (RDO). An RDO contains all triggered read out data from the TLB, CAEN digitizer, and TRBs as well as an assigned event number and timestamp. An RDO is approximately 25 kB in size, dominated by PMT readout.

Once the data is packaged into RDOs and stored to disk, the event data undergoes reconstruction as described in Chapter 5.

²A box for storing and powering circuit boards in parallel.

All raw and reconstructed data is written to tape with the CERN Tape Archiving system [47].

3.10 Control, Monitoring and Safety

FASER was designed for remote operation. With the appropriate permissions, the detector can be controlled from anywhere in the world. To that end, its systems require a degree of autonomy to respond promptly to any issues that may arise. The Project and Experiment Safety Support (PESS) and the CERN safety group (HSE-OHS) directed the design and installation of safety-related structures at the installation site [10, 48].

The detector control system (DCS) is a high-level, software-based system that responds to monitoring data and user commands to act on the detector accordingly. The DCS is run on a surface computer and monitors approximately 600 parameters at 1 Hz, corresponding to 1.3 kB/s. If the DCS receives monitoring data that indicates that the detector is approaching a damaging situation, the DCS issues commands to shut down endangered systems. The DCS provides a user interface to control and configure settings for HV and LV systems and Power Distribution Units (PDUs). The DCS also generates alerts that notify on-call experts about concerning changes to the detector's state.

FASER handles two kinds of monitoring data: metrics and histogram distributions. Metrics are typically a single variable that varies over time; environment conditions are usually metrics, as are the total number of events recorded in a data-collection run. Histogram distributions are best for noticing extreme or rare outliers in otherwise uniform data; these are used for data quality control. Monitoring data is archived in an Oracle database. FASER's Grafana webpage can retrieve and plot archived monitoring data. The Grafana interface can also report on TDAQ status and other information about ongoing activities.

FASER's tracking system is particularly vulnerable to changes in the detector environment

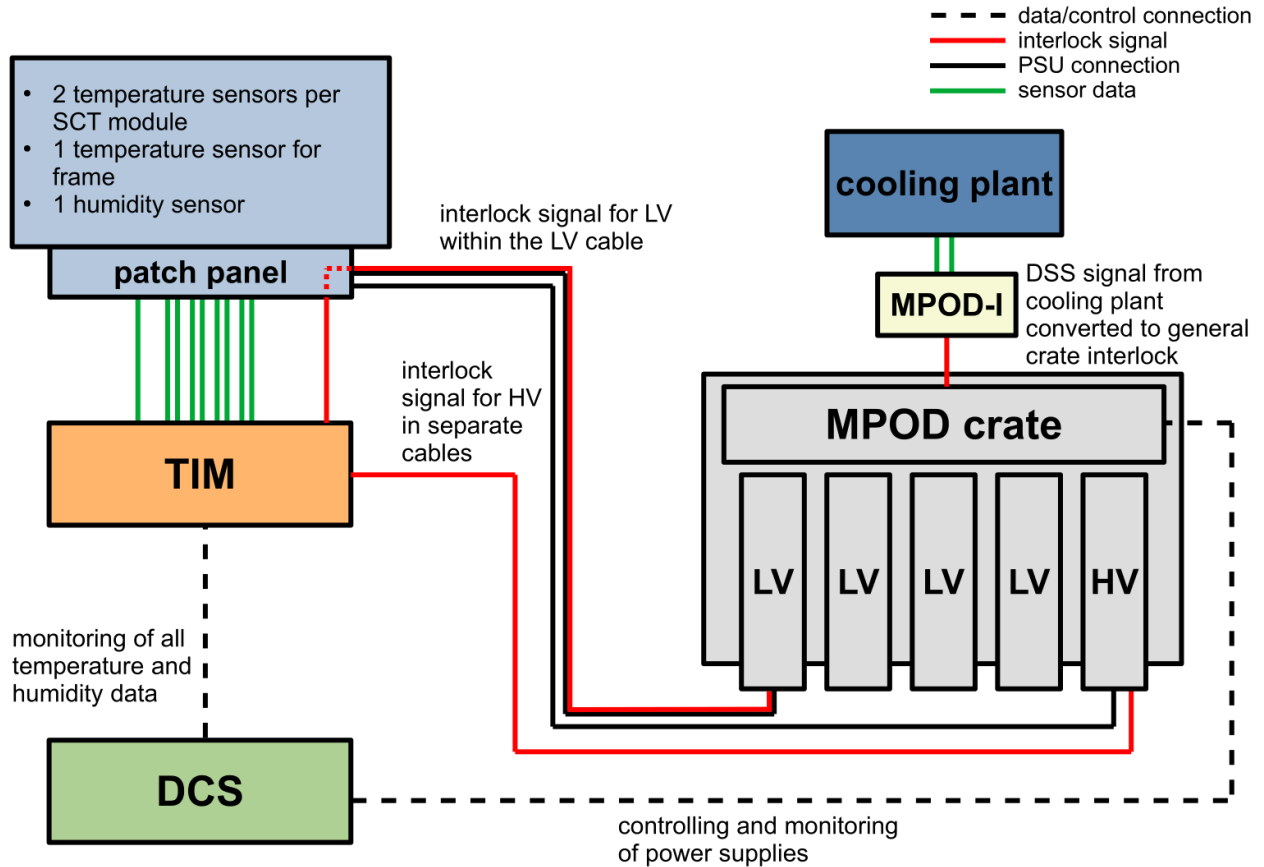


Figure 3.23: Sketch of the interlock system, showing the exchange of signals and data [25].

due to its need for temperature and humidity control. To that end, the Tracker Interlock and Monitoring Board (TIM) provides a low-level, hardware-based response to troubling conditions. There is one TIM per tracking station, each with a microprocessor that digitizes and monitors data from temperature and humidity sensors within the tracking planes; digitized monitoring data is also sent to the DCS. In the event of temperatures or humidity rising beyond acceptable levels, emergency procedures will shut down the tracking detector. This interlock system is illustrated in Figure 3.23. An environmental TIM monitors the ambient temperature and humidity of TI12.

There are also protections related to how much voltage is applied to the tracker modules. The power supply has a supervision module that turns off a channel when it leaves its operation range. There is also a LV protection board ASIC that protects the chips from getting too

much voltage.

Low voltage interlock is propagated through a dedicated signal line in the patch panel. There is also an LV protection board in the LV power path, designed to provide an extra layer of protection against over-voltage. The power supplies of the detector also have their own layers of security, including an MPOD interlock, which acts in the event of a cooling failure to shut off the power supplies of the detector.

FASER's run control system transitions the detector between different operating modes. The run control additionally records and labels the different data-collection periods called "runs." When there is no beam, FASER can run in "Cosmics" mode indefinitely. In "Physics" mode, a new run starts automatically whenever LHC fill number changes. The run control system receives all information about the LHC's status through CERN's Data Interchange Protocol (DIP).

The run control operates alongside a finite state machine (FSM). The FSM directly controls the power of the subdetectors and supporting systems, such as the chiller, and their voltage settings.

Chapter 4

Simulation

The outcome of a high-energy collision is probabilistic. A single measurement is insufficient to observe the full breadth of possible outcomes. Depending on the experiment, millions of measurements may be needed to make a confident statement about the results; the same principle applies to simulations of such experiments. To that end, particle physicists make liberal use of Monte Carlo (MC) generators [49].

Proton-proton (pp) collisions can cascade into myriad outcomes, beginning with all the possible interactions between their constituents. A simulated event created by an event generator is typically a table of values representing the timing, participants, and properties of an interaction. Event generators create these initial interactions whose probabilities are informed by first principles¹. A detector simulation propagates the interaction products through the materials of an experiment's environment. Many modern event generators include a detector simulator in their frameworks, or are otherwise compatible with external projects.

FASER needed many specialized simulators during its development and analysis. Early simulation of the detector environment supported a strong case for far-forward physics studies

¹In other words, calculated from SM-derived Lagrangians, sometimes embellished with new physics as needed.

in TI12 [9]. After FASER’s approval, detector simulations optimized the experiment’s geometry. Early online software development used digitized, simulated data to prepare the TDAQ system [48]. During data analysis, simulations estimated background, refined the new-physics signal criteria, and quantified measurement efficiencies [50]. Each of these tasks required at least one of the projects described in this chapter.

Simulation begins with an MC generator producing simulated particles relevant to a specific study, such as quantifying background or developing selection criteria. The detector simulation propagates the simulated particles using a software description of all the geometries and materials of FASER. The events which interact with sensitive elements are recorded as hits. Since all properties of simulated events can be perfectly known, simulated events have truth information in addition to measured information. Truth information includes what kind of particle deposited energy, and how much, in a sensitive detector, or the location of a decay vertex.

Simulated data is digitized such that its format is compatible with real data from the experiment. Digitization requires an accurate model the detector response and electronics. A digitized, simulated event is an estimation of the detector’s response to a similar, real event and can be reconstructed with the same procedure as real data, as described in Chapter 5.

FASER’s software framework, Calypso, organizes and standardizes the use and output of the many different software projects that contribute to FASER’s development and analysis. Calypso adapts many of its tools from Athena, ATLAS’ framework [51], which is itself derived from LHCb’s Gaudi framework [52]. All rely on the LHC Computing Grid software stack maintained by CERN. FASER generators must use physics tuning based on cosmic ray data, CRMC [53], to model the collimated products of pp collisions near the LOS. Simulations of particle propagation through the detector use GEANT4 [54].

4.1 Event Generation

4.1.1 Particle Guns

A particle gun is a simple generator which creates single particle events along the LOS for a given particle species and energy. Particle guns are a tool for studying the interactions of a single species, independent of their production mode.

FASER’s ParticleGun is adapted from Athena’s generator of the same name. ParticleGun by default generates neutral test-particle (geantino) events. ParticleGun accepts parameters that determine the sample size, the species of particle generated, and in what distribution for momentum, energy, and mass.

The cosmic ray gun, also called the cosmic sampler, uses ParticleGun’s modules and spectral data from a Daya Bay study [55] to generate cosmic ray events. The generator chooses a random angle of incidence and a random energy from the Daya Bay distributions. The minimum energy for the cosmic ray to reach FASER depends on the chosen angle of incidence. If the randomly chosen energy is greater than the required minimum energy and the event passes MC selection, the particle acquires generated coordinates and momentum derived from those parameters. ParticleGun uses the FTFP_BERT physics list.

4.1.2 Proton-proton Events

These generators provided the initial SM states for background and signal studies. FASER used the same event generators as LHCf, which studies forward-produced photons, to estimate the signal yield: QGSJET II-04 [56], Sibyll 2.3d [57], and EPOS-LHC [58]. These generators are primarily designed to model the extensive air showers (EAS) that follow cosmic rays and, more generally, the far-forward spectra of secondary particles from hadronic

collisions. Figure 4.1 illustrates the difference between generators and that EPOS-LHC’s results are central to the envelope defined the Sibyll and QGSJET outputs. These differences contribute to systematic uncertainty, as described in Chapter 6. The difference between generators is due to missing, different, or incomplete validation and tuning data. For example, Sibyll and DPMJET do not include $\Lambda \rightarrow p e \nu$ in their decay tables, which has consequences for any neutrino simulation using those generators [59].

4.1.3 Beyond Standard Model Events

Modelling forward BSM physics begins with pp-collision event generation. EPOS-LHC generated π^0 and η meson production events from LHC pp collisions. This output is provided to FORESEE (FORward Experiment SENSitivity Estimator), a numerical package ² that estimates the sensitivity of far-forward experiments, such as FASER.

FORESEE uses the FTFP_BERT physics list, EPOS-LHC output, and user-defined BSM models and branching ratios to propagate and decay the light meson products. For this analysis, these are the models and branching ratios described in Chapter 2. FORESEE further propagates and decays the BSM products of the meson decays, yielding their spectra. These spectra and the detector geometry determine the sensitivity of the experiment, as shown in Figure 2.3 and Figure 2.4.

FORESEE was verified by comparing its results to an independent signal generator, ISG, written specifically for this purpose [60]. ISG also generated π^0 and η decays, which decayed instantly at the IP into $A' \gamma$. The ISG dark photons decayed in a uniform distribution between 476.23 m and 477.73 m, corresponding to FASER’s decay volume, into $e^+ e^-$. Both generators used CRMCv.2.0.1 [53] to give generated samples their 4-momenta for this validation study. The validation study found excellent agreement between FORESEE and ISG output for

²A “numerical package” is a collection of scripts which perform integration numerically instead of analytically.

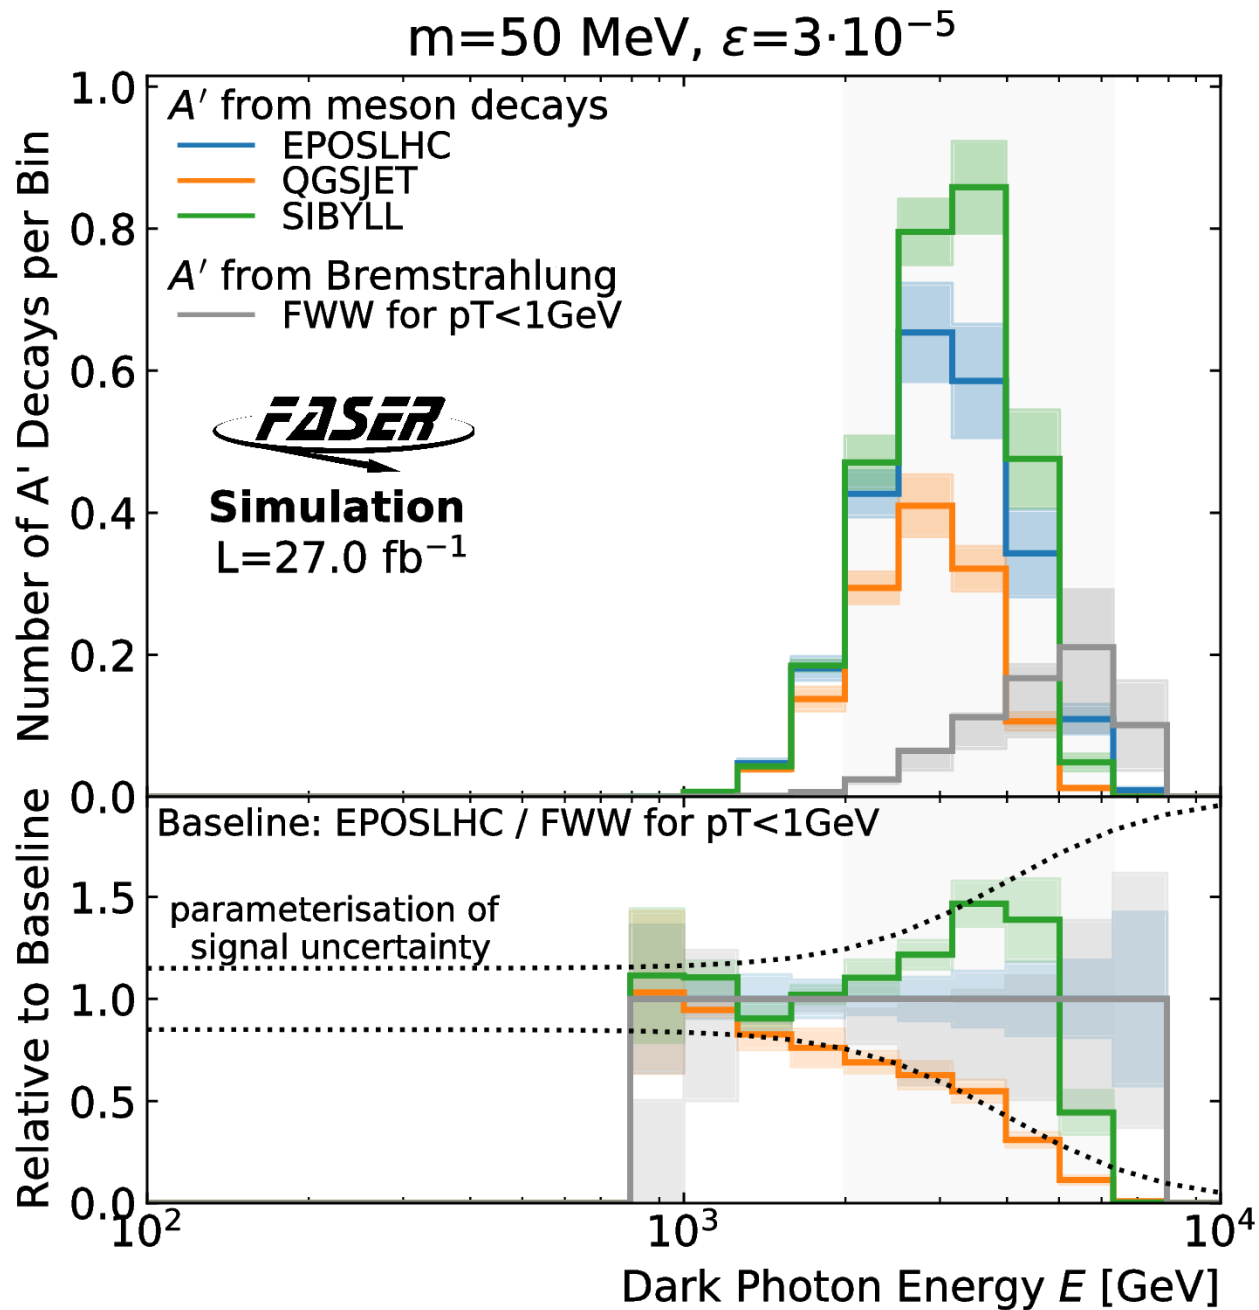


Figure 4.1: Expected dark photon spectra for dark photons within FASER's reach [50].




Figure 4.2: Comparison of e^+e^- spectra from the simulated decay of A' by FORESEE and ISG [60].

total number events, total events for each investigated decay channel, electron spectra, and electron distributions in the transverse plane. Figure 4.2 compares the spectral output of FORESEE and ISG.

4.1.4 Neutrino Events

Neutrinos appear in high-energy and abundance in FASER and contribute most to background events, as further described in Chapter 6. The expected number of neutrino events at FASER was estimated with GENIE (Generates Events for Neutrino Interaction Experiments) [61, 62], a neutrino MC event generator, and results from a dedicated LHC forward




Figure 4.3: Neutrino flux and energy estimated by RIVET and BDSIM using Sybill 2.3d hadron generation for muon and electron neutrinos, through a 20 cm \times 20 cm (left) and 40 cm \times 40 cm (right) cross-sectional area centered on the LOS, 480 m from IP1 and 150 fb $^{-1}$. [59].

neutrino study [59].

The study used RIVET (Robust Independent Validation for Experiment and Theory) to get a leading-order estimate of the neutrino flux. RIVET uses output from a pp event generator and propagates long-lived hadrons through a simplified LHC model containing only the beam pipe and magnets. The neutrino flux estimation is derived from the hadron decay products and stored in histograms.

EPOS-LHC, QGSJET II-04, DPMJET III.2017.1, Sibyll 2.3d, and Monash-tuned Pythia 8.2 were used to estimate light hadron production. DPMJET III.2017.1, Sibyll 2.3d and Pythia 8.2 estimated charm production. RIVET was validated against output from the full BDSIM LHC simulation, shown in Figure 4.3. The difference between simulations is attributed to RIVET omitting the contributions from subdominant neutrino production channels.

Cross-section splines for neutrino energies between 10 GeV and 7 TeV were generated for 100 target nuclei present in FASER. GEANT4 built the mass distribution, which was given to GENIE, along with the flux, energy spectra, and flavor composition estimates derived from each generator. From these, GENIE generated the positions of the neutrino interactions.

This study used the GENIE physics tune G18_02a, which specifies the models contributing to the cross-section calculations. At GeV energies and higher, we must use structure functions to adequately model nuclear effects. This physics tune uses the Bodek-Ritchie Fermi Gas model [63] for the initial nuclear state, with a di-nucleon tail. Deep inelastic scattering is described by the Bodek-Yang model [64], with default parameter values based on GRV98 LO parton distributions [65]. The AGKY 2018 model is used to simulate hadronization [66].

GEANT4 built the target mass distribution from the detector geometry, which GENIE imported to generate the neutrino interaction positions. GENIE propagated the neutrinos through the detector materials; only interactions with materials downstream of the veto scintillators were considered for the neutrino background study (see Section 6.2.3).

4.1.5 FLUKA

FLUKA is a MC generator and detector simulator built to handle complex detector geometries and high-energy beams [67].

The muon products of 80 mb pp collisions were propagated 409 m through 75 m of rock, on average losing 60 GeV, through a 1 m^2 area. Their directions were smeared across 1 mrad [68]. This sample yielded an estimate of muon flux (Figure 4.4) and spectra (Figure 4.5).

4.2 Detector Simulation

The FASER detector, including non-sensitive materials, is modeled in a “perfectly aligned and detailed description” [50] using ATLAS’s GeoModel library [54, 69]. The simulated detector is aligned with respect to the IP along the LOS with a 6.5 cm vertical offset to account for the crossing angle. This model also includes models for each subdetector and a




Figure 4.4: FLUKA-simulated muon rates at 409 m from IP1 entering a 1 m × 1 m area centered on the LOS as a function of energy [22].




Figure 4.5: Fluka-generated muon flux and spectra in TI18, normalized to $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ luminosity [10].

simplified, 0.6 T uniform dipole magnetic field [10].

4.3 Digitization

Simulation samples must be processed through electronic response models such that they can be studied with the same analysis tools as real data. In addition to simulating the detector response, digitized simulated data includes truth information. Truth tracking information includes the energy contribution of each passing particle to the detector element that measured it. Simulated particles are tracked through the detector geometry while recording hits in sensitive elements [25].

Simulated tracking information is digitized using the same techniques as ATLAS [70]. A simulated particle’s path through the SCT sensitive detector element has several steps between

entering and exiting the wafer, with the total deposited charge divided equally among them. Surface charges and their drift times are input to a model of SCT front-end electronics. RDOs are created and stored for each strip with a model response above a fixed threshold. This model includes detector noise and crosstalk in neighboring channels.

Simulated calorimeter and scintillator data is the position and total energy deposited in a module that the simulated particle passed; the process is almost the reverse of PMT waveform reconstruction. The following is a summary of the MC digitization process described in Appendix A.2 of [16]. This information must be translated into PMT waveforms to replicate how the event would be measured. TI12 data with veto and preshower scintillator measurements yielded the average waveform for each calorimeter module. PMT waveforms are parametrized by the Crystal Ball (CB) [16] function

$$f(x; \alpha, n, \bar{x}, \sigma) = N \cdot \begin{cases} \exp(-\frac{(x-\bar{x})^2}{2\sigma^2}), & \frac{x-\bar{x}}{\sigma} > -\alpha \\ A \cdot (B - \frac{x-\bar{x}}{\sigma})^{-n}, & \frac{x-\bar{x}}{\sigma} \leq -\alpha \end{cases} \quad (4.1)$$

where $A = (\frac{n}{|\alpha|})^n \cdot \exp(-\frac{|\alpha|^2}{2})$, $B = \frac{n}{|\alpha|} - |\alpha|$, N is the waveform amplitude, \bar{x} is the Gaussian mean, σ is the Gaussian width, n corresponds to the signal decay time, and α is the number of standard deviations away from the mean. This function is also used to reconstruct the PMT waveforms in Chapter 5. The parameterized waveform is discretized into 2 ns across a 1200 ns range, mimicking the PMT readout window. The MC-derived waveform does not include the 12 ns jitter seen in real data; the timing is therefore constant for a given module. The discretized CB function is scaled by the total deposited energy in that module. As with tracking data, the simulated waveform is Gaussian-smearred. Finally, the CB signal pulse is abstracted from a noisy baseline and inverted to resemble a raw waveform. Figure 4.6 compares the digitized simulated event to real data. The differences are attributed to the simulation not modelling nonuniform light detection.




Figure 4.6: A comparison of normalized waveforms derived from generated and real events for a sample calorimeter module (left) and scintillator module (right).

Chapter 5

Reconstruction

Event reconstruction is the process of extracting and inferring physical information about particles from raw detector data. Measured and simulated data both enter reconstruction as Raw Data Objects (RDOs). Calypso manages all reconstruction processes within its framework. This chapter will cover FASER's reconstruction workflow.

5.1 Track Reconstruction

5.1.1 Clusterization

A hit is the output measured from in a single SCT strip. A cluster is the combination of adjacent hit strips on the same side of a module. A cluster's location is relative to the local precision coordinate of its module. Since a cluster measures only one coordinate in the plane of a module, it provides only a two-dimensional point on the trajectory of the associated track. Full three-dimensional information is recovered using the stereo angle between strips on different sides of a module. FASER uses the SCT clustering algorithm developed by

ATLAS.

5.1.2 Track Finding

Track finding attempts to determine the trajectories of all charged particles using clusterized tracker data.

A linear fit algorithm organizes hits from each tracking plane within a tracking station into track segments. The number of hits used to make a segment and the χ^2 value associated with the fit determine its quality. A smaller χ^2 is associated with a better fit. Track seeds, rough estimates of track parameters, are derived from the combination of track segments from different stations that have the best overall fit. Since we are interested in long tracks traversing the length of the detector, seeds at high angle to the detector planes are discarded. If multiple track segments share more than 60% of their clusters, only the segment with most clusters kept; if more than one has the same number of clusters, the segment with the smallest χ^2 is kept [71, 50].

The ACT's Kalman Filter algorithm [72] uses the track seeds to resolve ambiguities in multi-track events. A Kalman Filter is an iterative process to find the best fitting track from predictions and measurements. The first prediction is derived from the given track seeds. Kinematics extrapolate the first estimate to the next tracking layer, yielding a new set of estimated track parameters. The prediction is updated by comparing the new estimate to the track seed in that layer. The updated prediction becomes the new estimate for that layer. This process continues through all tracking layers, with multiple iterations. The final, updated predictions of one iteration become the initial estimates for the next iteration. In multi-track events, tracks cannot share more than six clusters in neighboring strips on either side of the module. If they do, the track with the highest χ^2 value is discarded. A fully reconstructed track is shown in Figure 5.1.




Figure 5.1: Full detector event display with a single-muon event, top and side views. Blue lines show hit strips, and the red line in each view is the reconstructed track [73].

Track reconstruction efficiency and its contributions to uncertainty are discussed in Chapter 6. Their results show that reconstruction is approximately as efficient on data as it is on digitized simulated samples, within 2%.

5.2 PMT Data

The PMT waveform reconstruction procedure is found in Appendix A of D. Feller’s thesis [16]. RDO data from the calorimeters and scintillators are digitized waveforms. The PMT-signal parameter of interest in this analysis, the integrated waveform, is derived from fitting a Crystal Ball function to the raw waveform.

The raw waveform is 1200 ns long with a negative amplitude in Analog-to-Digital Converter (ADC) counts and a global positive ADC counts offset called the “baseline.” The waveform is initially processed as a histogram of ADC counts with 2 ns bins. Reconstruction begins with finding and subtracting the baseline and inverting the counts’ amplitude, yielding a positive pulse waveform.

Next, the signal must be isolated from the rest of the waveform. Events of interest to this analysis are relativistic and therefore have approximately the same arrival time in all scintillators and the calorimeter. These constant timing values select a 120 ns “reconstruction window” from the 1200 ns waveform containing the signal.

We now fit the CB function (Equation (4.1)) to the isolated signal. The CB can fit all waveforms regardless of PMT saturation. If a PMT is saturated, the fitted CB function yields an estimate of the saturated signal’s true size, as shown in Figure 5.2. Finally, integrating the fitted CB function over the reconstruction window yields the waveform integral, also called the “charge” of the measurement.

5.3 Detector Alignment

Track reconstruction uses detector simulation to incorporate the precise positions of the spectrometer’s components when calculating track parameters. The detector is already perfectly aligned for MC-generated data. Alignment begins with metrology studies, which check that the design and placement of the planes is consistent with their CAD models. Early metrology studies on individual tracker planes found that the tracker modules were at most 100 μm from their designated positions.

After the spectrometer was fully assembled in TI12, the CERN survey team used a laser tracker to map, down to millimeter precision, the positions of all magnets and tracking stations with respect to the LHC reference system.

During data-taking, an iterative method adjusted the detector model using high-energy muon events traversing the entire detector. Each iteration incrementally moved the position of the detector model’s modules in the direction that minimized the difference between the unadjusted hit position and the position expected by the reconstructed track. This method only considered the precision coordinate, \hat{y} , and the rotation of a module around \hat{z} . The method holds the positions of the central layers of the first and last spectrometer stations fixed. Figure 5.3 shows how alignment improved the distribution of track residuals when compared to unaligned data and MC. The remaining difference between the results using





Figure 5.2: Crystal ball function (blue) fit to data (red) for unsaturated (top) and saturated (bottom) PMT pulses [16].

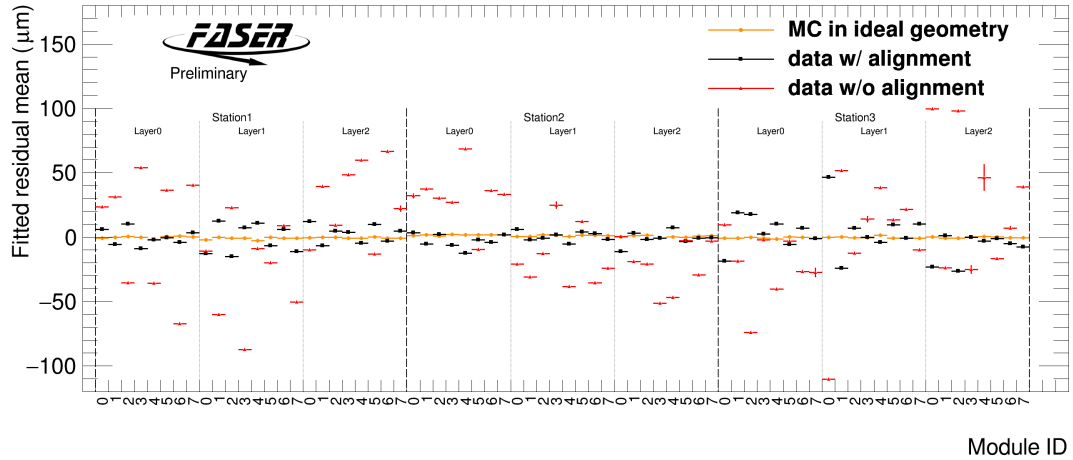


Figure 5.3: The effects of alignment on track residuals [73].

an iteratively-aligned detector model and the simulated results with the perfectly aligned model are attributed to misalignment in the fixed layers and other distortion directions not considered in this method [71, 16].

Chapter 6

Analysis

This analysis uses data collected from September to November 2022 and corresponds to an integrated luminosity of 27 fb^{-1} at 13.6 TeV [50]. The analysis procedure was developed and tested on blinded data, a subset of the full data set which included only events with at least one veto scintillator signal and less than 100 GeV total calorimeter energy. The final results used the same procedure on the full, unblinded data set.

6.1 Event Selection

A' and A'_{B-L} simulations informed event selection criteria (a.k.a. “cuts”). Both models expect the same products given FASER’s sensitivity and are therefore subject to the same cuts. As described in Chapter 2, the models are distinguished by the number of new-physics signal events observed. A cut’s efficiency is the number of simulated signal events passing the cut divided by the total number of simulated signal events in the fiducial region. Section 6.1 summarizes all selection cuts with their exact values when relevant.

A' and A'_{B-L} decay into e^+e^- within the FASER decay volume, as shown in Figure 6.1. The




Figure 6.1: Sketch showing the response of the FASER detector to a neutral LLP decay. White circles indicate individual measurements [50].

highly collimated charged pair is separated by the magnets, curving their paths apart. The charged pair leaves a signal in each penetrated scintillator layer and most tracking layers. They continue until they reach the calorimeter, where they deposit their remaining energy.

A new-physics event does not have signal in any veto scintillator. In the timing and preshower scintillators, it must generate a signal consistent with the passage of two MIPs. The scintillators must not be over-saturated, as from a shower.

Since we expect a highly energetic charged pair from the signal, and low-energy backgrounds, a calorimeter energy cut is stable and is compatible with most models within FASER’s reach without a loss of efficiency. Since some models have a broader energy spectra, the ideal calorimeter cut is as low as possible while allowing much less than one background event. The total reconstructed calorimeter energy is therefore required to be greater than 500 GeV, informed by the neutrino background study shown in Figure 6.2; this cut eliminates most of the dominant neutrino background [71].

Reconstructed tracks can also be used to efficiently reject backgrounds. First, any tracks must be “long,” meaning they have clusters in all downstream stations. A new-physics candidate must have exactly two good tracks. A good track has momentum >20 GeV, at least 12 hits total across at least 7 of the 9 spectrometer tracking planes, and a good fit. A track has a good fit when the ratio of the track’s χ^2 and its degrees of freedom (NDOF)¹ is

¹Degrees of Freedom in this context is the number of clusters minus 5. 5 corresponds to the number of spatial coordinates (x, y, ϕ, θ) and the ratio of charge to momentum (q/p) .

Table 6.1: Selection criteria for A' and $A'B - L$ events and signal efficiency for one example MC dark photon model, $m_{A'} = 50.1$ MeV, $\epsilon = 3 \times 10^{-5}$ [50].

Description	Value	Single Cut Eff.
Ntuple Production		
Stable beams	-	
Good run; no issues	-	
At least one cluster in every station 1-3	-	
Trigger and Data Quality		
Colliding BCID	-	
Timing scint. trigger fired	-	
No timing scint. saturation	-	
Scintillators		
Timing: (Ch0 + Ch1) OR (Ch2 + Ch3) OR (Ch0 + Ch1) AND (Ch2 + Ch3)	> 70 pC > 70 pC > 30 pC	98.327%
Each Veto channel	< 40 pC	99.665%
Each Preshower channel	> 2.5pC	99.719%
Tracking		
Fiducial track		92.057%
Number of "good" tracks	2	61.424%
"Good" Track Criteria: Momentum χ^2/NDOF Tracking layers with hits Number of hits	> 20 GeV < 25 >= 7 >= 12	
Calorimeter		
Sum of all channels	> 500 GeV	100.000%

<25. The track must also be entirely within FASER's fiducial volume, including when the track is extrapolated to the Veto station. FASER's fiducial volume is a cylinder of radius 95 mm whose length extends between the Veto scintillator and the preshower scintillator, coaxial with the magnets. The fiducial cut excludes events that miss the veto scintillators.

An ideal cut would select only and all signal events. Simulation studies use truth data to determine the efficiency of the cuts. "Real" data contains only detector read-out. Cuts that exclude all background may also exclude some desirable events. The signal efficiency is the

Table 6.2: Cutflow for one example MC dark photon model, $m_{A'}$ =50.1 MeV, $\epsilon = 3 \times 10^{-5}$, showing the cumulative efficiency of each criteria [50].

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

number of simulated signal events that pass the selection cuts divided by all simulated signal events. The total efficiency including contributions from all cuts is 50.8%. Section 6.1 shows the cumulative affect of each selection criteria on signal efficiency.

As part of FASER’s first public results, this analysis’ priority was establishing a procedure to find an unambiguous signal if one was present. The cuts chosen for this first analysis are considered conservative; signal efficiency is lower, but any observed new-physics event is associated with less uncertainty. In particular, requiring two good tracks has the largest effect on signal efficiency. The cuts on detector readouts are highly stable—that is, the signal efficiency does not change significantly when substituting other, nearby cuts.

6.2 Background

6.2.1 Veto inefficiency

FASER’s five veto scintillators were individually evaluated for their efficiency. This study identified the circumstances in which muons went undetected. A scintillator’s inefficiency is the fraction of all particles passing through the sensitive material which are not measured.

FLUKA simulations estimated $O(10^8)$ muons in this data set, which was confirmed with in-situ measurements [71]. When comparing the expected number of muon events with the measured number of muon events, the inefficiency of a single veto scintillator is on the order of 10^{-5} . The combined inefficiency of the independent scintillators is less than 10^{-20} , corresponding to a negligible fraction of the total muons expected for the data set. Effectively, there are no fiducial muon events without a veto signal.

6.2.2 Neutral Hadrons

When a muon interacts with the rock surrounding the FASER environment, it may produce a neutral hadron whose decay products mimic the new physics signal. Additionally, a muon may undergo photon conversion, in which the muon produces a photon that converts into $e+e-$ as it passes through material near the veto scintillators.

Such background is heavily suppressed for several reasons. First, the parent muon must stop or scatter before reaching the VetoNu station, or else be tagged as a muon event. A neutral hadron must also be produced before the VetoNu station. Additionally, the neutral hadron must survive 8 interaction lengths of tungsten in the emulsion detector and decay downstream of the veto scintillators. Finally, the neutral hadron decay products must pass event selection, including having at least 500 GeV total calorimeter energy.

Studies on control regions — partially unblinded data that excludes events that could qualify as signal — inform background estimations. This estimation assumes that the energy spectrum of the neutral hadrons are the same whether the event includes a muon.

The first control region included all 3-track events with energy below 100 GeV or above 500 GeV. For this selection, the highest momentum track is assumed to be the muon that caused the neutral hadron or photon conversion. Simulation of K_S and photon conversion events

showed that photon conversion events could be excluded by requiring the invariant mass of the two lowest-momentum tracks to be greater than 200 MeV.

The second control region selects all 2-track events with < 100 GeV and no signal in the VetoNu scintillators. The requirement of having no signal in any veto was relaxed for this selection to allow a sufficient number of events for the study. The “high energy” counterpart of this selection would remain blinded, as it represents the signal region.

The ratio of low-energy (< 100 GeV) to high-energy (> 500 GeV) events in the first selection is used to estimate the number of expected background events for high-energy 2-track events. There were 404 and 19 events for low and high energy respectively in the first control region. Only one event was observed in the second control region, therefore 0.047 events from neutral hadrons were expected for 2-track events with no VetoNu signal and 500 GeV calorimeter energy.

This result is corrected to account for the new-physics signal, requiring no signal in any veto scintillator. The first control region was used to study the veto scintillator readout for 3-track events. We looked for events that had three reconstructed tracks, whose pre-veto data resembled a single-track event. The criteria for such an event required the number of IFT clusters to be fewer than 8, and for all the veto channels to read out less than 200 pC. The ratio of 3-track events with a single-MIP scintillator charge to all 3-track events becomes the correction. In total, 1.7% of 3-track events fit these criteria, making the final estimate for neutral hadron background $(8.4 \pm 11.9) \times 10^{-4}$ events. The statistical uncertainty arises from only observing a single event in the low-energy 2-track control region, with an additional 100% uncertainty applied to account for the assumptions made during calculation.

6.2.3 Neutrinos

A neutrino that interacts with any material downstream of the veto scintillators and before the end of the first tracking station can also produce high-energy charged particles that mimic the new physics signal. To quantify this effect, 300 ab^{-1} of neutrino events fitting these criteria (corresponding to about 10,000 times more neutrinos than expected for this data set) were simulated with GENIE. Only neutrino events that interacted within a 100 mm radius were kept. Before additional cuts were applied, the simulated neutrinos were observed to interact at the same rate in all three downstream tracking stations, and most of those interactions were near the lateral sides of the magnet aperture. Most neutrinos that would have interacted near the axis of the detector would have done so in the FASERnu tungsten plates.

After event selection cuts, the expected yield is 11 events for 300 ab^{-1} or 0.0015 events when scaled to 27 fb^{-1} . The 500 GeV calorimeter cut is the most effective at excluding most neutrino background events. 80% of this background is attributed to electron neutrinos and antineutrinos, with muon (anti)neutrinos contributing to the remaining 20%. Due to mis-modelling of material near the timing scintillator, there is some uncertainty attributed to the difference in approaches to correcting the number of background events; a conservative 100% uncertainty is applied to account for this. An additional 100% uncertainty is applied in consideration for the uncertainty in electron neutrino flux, and similarly another 25% for muon neutrino flux. Another 100% uncertainty to account for the differences and uncertainties when modelling neutrino interactions. Finally, the total neutrino background for 27 fb^{-1} is $(1.5 \pm 0.5(\text{stat.}) \pm 1.9(\text{syst.})) \times 10^{-3}$ events. The spectrum of neutrino events is shown in Figure 6.2.




Figure 6.2: Simulated neutrino background spectrum. The dotted line represents the cut on calorimeter energy [50].

6.2.4 Large-Angle Muons

This type of background, also called a geometric background, refers to the circumstance that a muon misses the veto system entirely but still traverses the decay volume. The fiducial extrapolation requirement suppresses this background, since such a track requires missing the veto scintillator.

To estimate this contribution, FLUKA simulated 800k muons. The generated muons began before the VetoNu scintillator and had an angle of approach between 15-60mrad, which encompasses the 40 mrad angle that allows the muons to miss all veto stations. The simulation did not consider any changes in trajectory due to the presence of magnetic fields or interaction with materials. Before applying any other cuts, none of the simulated events had 2 tracks; events with a single track did not have energies higher than 50 GeV.

Comparison to real data events with veto signal validated the simulation. The ratio of number of events with energy below 500 GeV and number of events above 500 GeV was used to estimate the number of events without a veto signal and >500 GeV. This result was consistent with the MC estimate shown in Figure 6.3.

These results show that the contribution from large-angle muons is negligible.

6.2.5 Non-collision

Data collected during beam-down time, a total of 330 hours, showed that cosmics and electronics noise can still cause the detector to read out data at a rate of 35 Hz. This rate is dominated by Timing PMT noise, with the calorimeter trigger contributing less than 0.1 Hz and cosmics-triggered events 0.01 Hz. None of the events in this beam-down data had energies above 100 GeV and none had at least one good track.




Figure 6.3: ECAL MC geometric muon energy prior to event selection [71].

Although FASER is 480 m from ATLAS, its entrance is less than 10 m from the LHC beam line. The beam traveling towards ATLAS—Beam 1—interacts with materials in and around the beam line, such as the LHC Q12 magnet. These secondary particles may be ejected towards FASER. They are studied by observing detector activity during times when the beam has good bunch timing but no colliding bunches in IP1. Fortunately, these events have well-defined characteristics, such as a 12.5 ns signal delay with respect to the collision signal and BCID. The events without a veto signal have no good tracks. No events were observed for calorimeter energy higher than 400 GeV.

Since neither source would have any events in the signal region, the non-collision contribution to background is negligible.

6.2.6 Summary

The dominant contribution to background in this analysis comes from neutrinos interacting with material downstream of the veto scintillators. Neutral hadrons contribute an order of magnitude less. Other sources—non-collision events and large-angle muons—are negligible. The total background is then $2.3 \pm 2.3 \times 10^{-3}$ events.

6.3 Systematic Uncertainty

The expected signal yield is subject to systematic uncertainties. The integrated luminosity provided by the ATLAS collaboration is uncertain to 2.2%. There is also a statistical uncertainty attributed to the number of MC simulated signal events that ranges from 1-3% depending on the model. Other sources include the physics of signal generation and modelling of the detector response.

6.3.1 Signal Systematics

The uncertainty for the number of signal events decaying in the FASER decay volume is determined by comparing different simulations of their primary production mode, meson decays. Little published work has explored forward particle production at colliders. There are no models that perfectly fit multi-TeV photon data measured by LHCf, an experiment that measures forward particles near the beam line. QGSJET II-04, SIBYLL 2.3d, and EPOS-LHC came the closest and were used for FASER’s signal yield estimates. Since these generators only provide a “central” estimate and have no uncertainty bands, the spread in the different predictions (also referred to as an envelope in publications) was used as an estimation of the signal yield’s systematic uncertainty. All models showed that the yield had a dependence on the parent particle’s energy. The signal yield uncertainty was parameterized in Equation (6.1).

$$\frac{\Delta N}{N} = \frac{0.15 + (E_{A'}/4TeV)^3}{1 + (E_{A'}/4TeV)^3} \quad (6.1)$$

Where ΔN is the uncertainty, $E_{A'}$ is the energy of the parent particle, and N is the number of events. This parameterization was checked using FORESEE with multiple models for both dark photons and $B - L$ gauge bosons, and found to be within the “envelope” established by the generators. This energy-dependent uncertainty ranges between 15% at low energies and as high as 85% at $E_{A'} = 6.8$ TeV. Subdominant production modes were not found to contribute significantly to signal yield uncertainty [60].

Any difference between real and simulated detector response could and its simulated response contribute to the systematic uncertainty.

6.3.2 Scintillator and Calorimeter Contributions

Some uncertainty could be attributed to the scintillator system if the MC and data showed a difference in number of rejected events. As demonstrated in background studies, FASER's scintillators are highly efficient. A comparison between real and simulated 2-track events measured the effects of the scintillator-related event selection. The study found that the scintillator charge cuts rejected less than 1% of signal events. Scintillator charge distributions for dark photon MC agreed with 2-track scintillator charge distributions in data. Since the overall uncertainty is less than 1%, it is negligible.

The calorimeter-dedicated internal note [41] has a detailed overview of calorimeter uncertainty and efficiency. The dominant calorimeter contributions to uncertainty is from calibration calculations and the uncertainties when measuring the MPV. A conservative analysis found a total 6.06% uncertainty on the difference between the calibrated values of data and MC across momentum ranges probed by this study. This was validated by checking the energy/momentum (E/p) distribution of 3-track events in data and MC. Such events typically arise from photon conversion, so the highest-energy track is assumed to be a muon, with the remaining two used to calculate E/p . The results, shown in Figure 3.16 were found to be in agreement within 6% of simulated E/p .

6.3.3 Tracking Contributions

Track reconstruction also contributes to uncertainty. The single-track efficiency determines how likely a single charged particle traversing the detector is reconstructed. The study required the following event selection: charge in all veto scintillators > 40 pC, a total timing scintillator charge > 20 pC, charge in both preshower scintillators > 2.5 pC and at least one good track segment in each of the spectrometer tracking stations in the fiducial region. A good track segment has $\text{DoF} \geq 1$, $\chi^2 < 25$, and distance from the LOS < 95 mm. The




Figure 6.4: Gaussian-fitted peak position of the E/p from non-muon photon conversion tracks in data and MC simulation as a function of the momentum of photon conversion candidates [50].




Figure 6.5: Single track efficiency vs. truth momentum p_{truth} and truth track angle θ_{truth} [71].

fraction of events passing this first set of cuts and have a good track as defined in Section 6.1 is the single-track efficiency. This selection was applied to the FLUKA-simulated muon sample and real data, finding efficiencies 96.9% and 98.4% respectively. The difference between data and MC yields a 1.5% systematic uncertainty for a single track.

The next step is understanding two-track efficiency. Unlike single-track events, which could be directly studied using muons, there are no high-rate SM processes that produce energetic two-track topologies in FASER. To resemble the signal, the two tracks must be from

closely-spaced, oppositely-charged particles. Instead, raw track data from different single-track events were combined to make artificial closely-spaced 2-track events. First, raw data events were paired according to their proximity; then, the combined events underwent reconstruction. This process can be done with data or simulation; the ratio between the two as a function of track separation provides a measure of the 2-track systematic uncertainty. This difference in efficiency, shown in Figure 6.6, is as much as 7% and quantifies the effect on signal yield.

Finally, we consider the uncertainties of track momentum scale and resolution by examining the photon conversion mass peak: the peak in calorimeter data associated with photon conversion. The location and width of this peak is different when considering simulation or data, quantified by fitting a Gaussian curve to each. For the momentum scale, the position of the peak varied by 5%. For the momentum resolution, the ratio of the widths of these fitted curves was considered; this also varied by 5%, as shown in Figure 6.7. This small difference has less than half a percent effect on signal yield [50, 71].


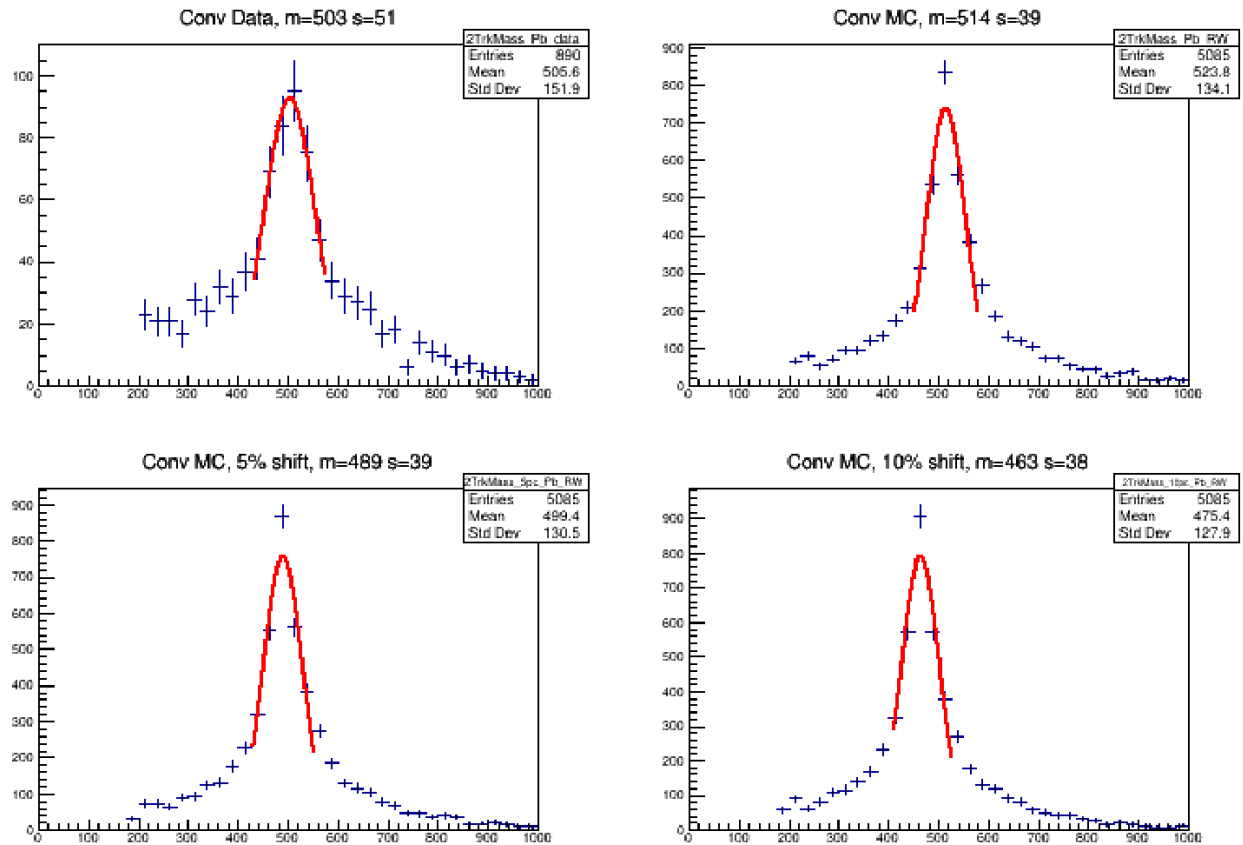


Figure 6.6: The upper plot shows the efficiency of overlaid MC and Data as a function of track separation, with the distribution of track separation of dark photon products in red for a given model. The lower plot shows the ratio of MC and data efficiencies for each track separation bin [50]



Chapter 7

Results

In anticipation of unblinding data, the analysis used the HistFitter software framework to consider the possibilities of discovery and exclusion. HistFitter is used for analyzing large data sets collected from LHC collisions and is a standard statistical tool for ATLAS [74]. This analysis uses a “profile-likelihood fit” approach. To understand what that is and its use, we must understand how and why high-energy physics uses statistical inferences.

7.1 Statistical Inference

We can broadly describe science as the observation of repeatable, natural phenomena. In other words: no matter how many times you repeat an experiment following the same procedure, it will produce the same results. Repeatability alludes to an underlying model for a phenomenon which can be used to predict and control nature. High-energy physics at the LHC is an experimental science in which nigh-uncountable observations are made within a highly controlled environment. The overwhelming amount of data demands a statistical approach [19].

Strategies for handling statistics are typically categorized as either “frequentist” or “Bayesian.” A frequentist approach interprets probability as a limit of how often an outcome is observed in a repeatable experiment. Frequentist tools include parameter estimation, statistical tests, and confidence intervals. Bayesian statistics attempt to quantify the “degree of belief” in a hypothesis. The Bayesian approach includes the concept of a probability density function (pdf) to measure how much we know about the value of a parameter. Each is a fundamentally different approach to address different questions asked of a data set. Most of this statistical analysis will use frequentist strategies.

7.2 Frequentist Concepts

Frequentist hypothesis tests provide a standard for accepting or rejecting a hypothesis depending on the outcome of a measurement. This includes significance tests. Finding the “p-value” of an observation is a frequentist significance test that quantifies how likely it is to make another observation that is as or more incompatible with the null hypothesis as the base observation; it assumes the null hypothesis is true. A low p-value suggests that an alternative hypothesis to the null hypothesis is true.

Likelihood is the probability for an experiment to have have outcome values \mathbf{x} given a hypothesis characterized by parameters θ . Likelihood is designed to be open-ended, such that it can consider new inferences about the observations it is derived from [19].

Sprott [75] asserts that the purpose of an experiment is to make inferences about some parameter of interest, conditional on a model. The phenomenon to be studied and the experiment’s design are intrinsic to specifying this model. When trying to describe what parameters affect the observation of a phenomenon, there are those we are interested in and those we are not, but have to account for anyway. These “uninteresting” parameters are

called “nuisance parameters.”

Calibration measurements and environment control mitigate the impact of nuisance parameter(s) ν by constraining their values. Suppose that a set of measurements \mathbf{x} that determines the values of parameters $\boldsymbol{\theta}$ is described by the data probability model $P(\mathbf{x}|\boldsymbol{\theta})$. Control or calibration measurements, \mathbf{y} , which are statistically independent of \mathbf{x} , are then described by $P(\mathbf{y}, \nu)$. If the \mathbf{y} are independent of \mathbf{x} , then the likelihood is the product of their probability models, as in

$$L(\boldsymbol{\theta}, \boldsymbol{\nu}) = P(\mathbf{x}|\boldsymbol{\theta}, \boldsymbol{\nu})P(\mathbf{y}|\boldsymbol{\nu}). \quad (7.1)$$

When using MC, \mathbf{x} and \mathbf{y} are generated for each repetition, while $\boldsymbol{\theta}$ and $\boldsymbol{\nu}$ are fixed.

Profile likelihood is used to find statistical errors in θ and is

$$L_p(\theta) = L(\theta, \hat{\boldsymbol{\nu}}(\theta)), \quad (7.2)$$

where $\hat{\boldsymbol{\nu}}(\theta)$ are the profiled nuisance parameter values that maximize the likelihood for specified θ .

Suppose we want to test the significance of (i.e. get a p-value for) hypothetical values for θ with a model containing nuisance parameters ν . A “test statistic” q_θ is defined such that it is proportional to the incompatibility between the data and a hypothesis. A hypothetical value for θ is only rejected at some chosen confidence level if the p-value associated with the hypothetical value is less than the probability of rejecting a true, null hypothesis (also called α or “test size”) for all possible values of the nuisance parameters.

At this point, we would still really like to ignore the nuisance parameters as much as possible by defining q_θ such that its distribution is independent of them. A “profile likelihood ratio”

(PLR) is the ratio of the profile likelihood divided by the value of the likelihood at its maximum:

$$\lambda_p(\theta) = \frac{L(\theta, \hat{\nu}(\theta))}{L(\hat{\theta}, \hat{\nu})}. \quad (7.3)$$

Equation (7.3) yields correct p-values as long as the true value of the nuisance parameters are the same as the profiled values; otherwise, it is only an estimate.

7.3 Statistical Analysis Framework

In this analysis, the likelihood function is the product of independent distributions and determined with HistFitter's profile-likelihood fit function. The likelihood function is

$$L(\mathbf{n}, \theta^0 | \mu_{sig}, b, \theta) = \mathcal{P}_{SR} \times \mathcal{C}_{syst}, \quad (7.4)$$

where n_i are the observed number of events in each signal and control region bin i , θ^0 is the central value around which the nuisance parameters θ can be varied, \mathcal{P}_{SR} is the Poisson probability of the number of (simulated and observed) events in a signal region (SR), and \mathcal{C}_{syst} is the Gaussian probability that accounts for systematic uncertainties [74].

This likelihood function is used in both of the following statistical interpretation strategies, with all details found in the internal analysis note [71].

7.3.1 Model-Dependent Fit

Also known as the exclusion fit, this method studies specific signal models and depends on the presence and significance of an excess of events in the SR. The normalization parameter (a.k.a. signal strength) for a given model is calculated for each simulated sample in FASER's reach.

When there is no excess in the SR, the fit determines exclusion limits for a specific model. A p-value in this situation represents the probability of rejecting the background-only hypothesis. If the p-value for a given model is greater than 0.05, then the specified model is excluded at 95% CL. An exclusion contour interpolates the phase space between combined p-value results for all sampled models.

If there is an excess of events, then the same process is used to quantify the signal strength of a model.

7.3.2 Model-Independent Fit

Also called the discovery fit, this approach establishes an upper limit to the number of excess events expected in the SR. This method considers any potentially-contributing generic model as an explanation of the observed excess. If there is evidence for observation in the data, along with a discovery fit, there must then also be an estimate of the production process' cross-section that would explain it.

7.4 Sensitivity Tests

Result calculations were repeated with modified uncertainty and number of neutrino background events to test their robustness. These sensitivity tests assumed no new physics signal. The modified estimated background varied between 0.01 and 0.1 events and its uncertainty was varied in scale by 0%, 50%, 100%, or 200%. Other signal uncertainties as described in Chapter Chapter 6 were applied in an uncorrelated manner as a flat percent on the nominal yield and assumed to be Gaussian. The sensitivity changed very little — well within 1σ uncertainty band of the original, expected contour — with variation of the background. Variation of the background uncertainty also had negligible effect. Thus, the observed limits are insensitive to the modeling of the background, and its systematic uncertainties.

7.5 Unblinded Results

After validating the modeling of background with data for calorimeter energy deposits less than 100 GeV, the signal region was unblinded. A signal event would have, among other cuts described in 6, no veto signal, 2 tracks, and total calorimeter energy greater than 500 GeV. No events were observed in the signal region for the 27 fb^{-1} data collected, which is compatible with the expected background of $(2.3 \pm 2.3 \times 10^{-3})$ events. Further, as shown in Figure 7.1, there were no veto-less events with at least one track and a total calorimeter energy larger than 100 GeV.

Since no evidence of a signal was observed, the model-dependent fit determined an exclusion contour in $(m_{A'}, \epsilon)$ model parameter space. The contour was generated with a 90% CL for comparison with other experiments probing a similar phase space. Figure 7.2 shows the new exclusion limits established by this analysis alongside prior results from other experiments. The gray exclusion regions are adapted from results for other models by the Darkcast





Figure 7.1: Spectrum of events before event selection (top) and after (bottom) [50].

project [24] using data from BaBar [76], E141 [77], NA48 [78], NA64 [79, 80], Orsay [81, 82], and NuCal [83, 84].

This analysis excludes dark photons signals within $4 \times 10^{-6} < \epsilon < 2 \times 10^{-4}$ and $10 \text{ MeV} < m_{A'} < 80 \text{ MeV}$, including world-leading exclusion for $2 \times 10^{-5} < \epsilon < 1 \times 10^{-4}$ and $17 \text{ MeV} < m_{A'} < 70 \text{ MeV}$. The red line in the dark matter exclusion plot distinguishes a scenario in which dark photons couple to a low-mass, complex scalar dark matter field, χ . The contour assumes $m_\chi/m_{A'} = 0.6$ and that the dark photon coupling constant to dark matter, α_D , is 0.1. These assumptions guarantee that the dark photon decays visibly into SM species and that dark matter primarily annihilates via $\chi\chi \rightarrow A' \rightarrow ff$. Different assumptions would shift the relic target line. In this scenario, the region below the target line represents models that would predict an over-abundance of dark matter in disagreement with cosmological observations. Therefore, FASER probes a significant portion of cosmologically-allowed parameter space.

The analysis also constrains $B - L$ gauge boson models by excluding models within $3 \times 10^{-6} < g_{B-L} < 4 \times 10^{-5}$ and $10 \text{ MeV} < m_{A'_{B-L}} < 50 \text{ MeV}$, with the first exclusion within $5 \times 10^{-6} < g_{B-L} < 2 \times 10^{-5}$ and $15 \text{ MeV} < m_{A'_{B-L}} < 40 \text{ MeV}$. These bounds also contain cosmologically relevant regions for the scenario that $0.5m_{A'_{B-L}} < m_\chi < m_{A'_{B-L}}$.





Figure 7.2: New exclusion limits for dark photons (top) and $B - L$ gauge boson (bottom). The gray regions are already excluded from prior experiments [50].

Chapter 8

Conclusion

The development of this analysis procedure establishes a foundation for future FASER analyses. This work demonstrates that FASER is capable of far-forward, long-lived particle searches with low background at the LHC. The event selection for both A' and A'_{B-L} required no signal in the veto scintillators, two “good” reconstructed tracks, and at least 500 GeV calorimeter energy. No events passed this selection; this result is within the expected background of $(2.3 \pm 2.3 \times 10^{-3})$ events. With no evidence of a signal, this analysis set new exclusion boundaries for A' and A'_{B-L} at 90% confidence level. These results establish constraints on cosmologically-motivated models for dark photons and $B - L$ gauge bosons.

These results were first presented at Rencontres de Moriond 2023 [11, 85] alongside FASER’s electronic neutrino analysis that marked the first observation of collider-generated neutrinos. FASER will collect data for the remainder of Run 3. In early 2024, CERN approved the continued operation of FASER during Run 4. FASER’s preshower system will be upgraded with a pixel detector during the 2024 year-end technical stop; the upgrade facilitates multi-photon searches, such as for axion-like particles (ALPs). Future analysis will extend this first result, with less conservative cuts informed by more detailed studies of the calorimetry.

As a research group and experiment, FASER was developed with exceptional speed, even when faced with a global pandemic. This swift success is owed to its simplicity, compact size, donated hardware, and the considerable efforts of its collaborators. FASER benefited directly from decades of work for other experiments, two of which (LHCb and ATLAS) loaned hardware, software, and administrative support.

FASER's development has pioneered strategies for approaching forward physics at collider experiments and has informed experiments currently in-development. The experiment has brought attention to the need for more far-forward physics searches. Many of FASER's collaborators are also contributing to the development of the Forward Physics Facility (FPF), a proposed addition to the LHC tunnel system. FPF would house several forward experiments, including FASER's successor, FASER 2. Together, these projects form a robust far-forward program that extends into the remainder of the LHC's lifetime [86, 87].

Bibliography

- [1] **ATLAS** Collaboration, 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) 1–29, [arXiv:1207.7214 \[hep-ex\]](#).
- [2] **CMS** Collaboration, 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** (2012) 30–61, [arXiv:1207.7235 \[hep-ex\]](#).
- [3] **Planck** Collaboration, N. Aghanim *et al.*, “Planck 2018 results. VI. Cosmological parameters,” *Astron. Astrophys.* **641** (2020) A6, [arXiv:1807.06209 \[astro-ph.CO\]](#). [Erratum: *Astron. Astrophys.* 652, C4 (2021)].
- [4] J. Feng, “ForwArd Search ExpeRiment at the LHC,” in *Annual Theory Meeting National Center for Theoretical Sciences, Taiwan*. 2018. https://drive.google.com/file/d/1_a632CajerxDv7feUr_sdYWkifL70D1e/view.
- [5] CERN Human Resources Department, “CERN Personnel Statistics 2021,” CERN Document Server:2809746.
- [6] “LHC Machine,” *JINST* **3** (2008) S08001.
- [7] E. Lopienska, “The CERN accelerator complex, layout in 2022. Complexe des accélérateurs du CERN en janvier 2022,” CERN Document Server:2800984. General Photo.
- [8] J. L. Feng, I. Galon, F. Kling, and S. Trojanowski, “ForwArd Search ExpeRiment at the LHC,” *Phys. Rev. D* **97** no. 3, (2018) 035001, [arXiv:1708.09389 \[hep-ph\]](#).
- [9] **FASER** Collaboration, A. Ariga *et al.*, “Letter of Intent for FASER: ForwArd Search ExpeRiment at the LHC,” [arXiv:1811.10243 \[physics.ins-det\]](#).
- [10] **FASER** Collaboration, A. Ariga *et al.*, “Technical Proposal for FASER: ForwArd Search ExpeRiment at the LHC,” [arXiv:1812.09139 \[physics.ins-det\]](#).
- [11] **FASER** Collaboration, B. Petersen, “First Physics Results from the FASER Experiment,” in *57th Rencontres de Moriond on Electroweak Interactions and Unified Theories*. 5, 2023. [arXiv:2305.08665 \[hep-ex\]](#).

- [12] D. Griffiths, *Introduction to elementary particles*. Wiley, 2008.
- [13] M. E. Peskin and D. V. Schroeder, *An Introduction to quantum field theory*. Addison-Wesley, Reading, USA, 1995.
- [14] D. Tong, “David Tong: Lectures on Theoretical Physics.” 2021. <https://www.damtp.cam.ac.uk/user/tong/teaching.html>.
- [15] W. Buchmuller and C. Ludeling, “Field Theory and Standard Model,” in *2005 European School of High-Energy Physics*. 9, 2006. arXiv:0609174 [hep-ph].
- [16] **FASER** Collaboration, D. E. Fellers, “A Search for Dark Photons with the FASER Detector at the LHC,” **CERN Document Server:2851248**.
- [17] F. Zwicky, “Die Rotverschiebung von extragalaktischen Nebeln,” *Helv. Phys. Acta* **6** (1933) 110–127.
- [18] D. Majumdar, *Dark Matter: An Introduction*. Taylor & Francis Group, Boca Raton, 8, 2014.
- [19] **Particle Data Group** Collaboration, R. L. Workman and Others, “Review of Particle Physics,” *PTEP* **2022** (2022) 083C01.
- [20] D. Clowe, M. Bradac, A. H. Gonzalez, M. Markevitch, S. W. Randall, C. Jones, and D. Zaritsky, “A direct empirical proof of the existence of dark matter,” *Astrophys. J. Lett.* **648** (2006) L109–L113, arXiv:astro-ph/0608407.
- [21] M. Bauer, P. Foldenauer, and J. Jaeckel, “Hunting All the Hidden Photons,” *JHEP* **07** (2018) 094, arXiv:1803.05466 [hep-ph].
- [22] **FASER** Collaboration, H. Abreu *et al.*, “First neutrino interaction candidates at the LHC,” *Phys. Rev. D* **104** no. 9, (2021) L091101, arXiv:2105.06197 [hep-ex].
- [23] **FASER** Collaboration, A. Ariga *et al.*, “FASER’s physics reach for long-lived particles,” *Phys. Rev. D* **99** no. 9, (2019) 095011, arXiv:1811.12522 [hep-ph].
- [24] P. Ilten, Y. Soreq, M. Williams, and W. Xue, “Serendipity in dark photon searches,” *JHEP* **06** (2018) 004, arXiv:1801.04847 [hep-ph].
- [25] **FASER** Collaboration, H. Abreu *et al.*, “The FASER detector,” *JINST* **19** no. 05, (2024) P05066, arXiv:2207.11427 [physics.ins-det].
- [26] R. Zorat and D. Vender, “Global model for an rf hydrogen inductive plasma discharge in the deuterium negative ion source experiment including negative ions,” *Journal of Physics D: Applied Physics* **33** no. 14, (7, 2000) 1728.
- [27] K. N. Leung, G. J. DeVries, W. B. Kunkel, L. T. Perkins, D. S. Pickard, K. Saadatmand, A. B. Wengrow, and M. D. Williams, “Recent development on RF driven multicusp H⁻ ion sources: Guidelines for submission of proposals,” *Conf. Proc. C* **960610** (1996) 1513–1515.

- [28] W. Weterings, C. Bracco, L. Jorat, R. Noulibos, and P. van Trappen, “First experience with carbon stripping foils for the 160 MeV H^- injection into the CERN PSB,” *AIP Conf. Proc.* **1962** no. 1, (2018) 030003.
- [29] V. M. Vollaie, J *et al.*, *Linac4 design report*, vol. 6/2020 of *CERN Yellow Reports: Monographs*. CERN, Geneva, 9, 2020.
- [30] W. Weterings, G. Bellodi, J. Borburgh, T. Fowler, F. Gerigk, B. Goddard, K. Hanke, M. Martini, and L. Sermeus, “160-MeV H^- injection into the CERN PSB,” *Conf. Proc. C* **070625** (2007) 1628.
- [31] A. Newborough, “The Upgrade of the CERN Proton Synchrotron Booster Transfer Line Magnets,” *IEEE Trans. Appl. Supercond.* **30** no. 4, (2020) 4000405.
- [32] “LHC Injectors Upgrade, Technical Design Report: v.1: Protons,”.
- [33] **FASER** Collaboration, H. Abreu *et al.*, “First Direct Observation of Collider Neutrinos with FASER at the LHC,” *Phys. Rev. Lett.* **131** no. 3, (2023) 031801, [arXiv:2303.14185](https://arxiv.org/abs/2303.14185) [hep-ex].
- [34] **FASER** Collaboration, H. Abreu *et al.*, “The tracking detector of the FASER experiment,” *Nucl. Instrum. Meth. A* **1034** (2022) 166825, [arXiv:2112.01116](https://arxiv.org/abs/2112.01116) [physics.ins-det].
- [35] A. Abdesselam *et al.*, “The barrel modules of the ATLAS semiconductor tracker,” *Nucl. Instrum. Meth. A* **568** (2006) 642–671.
- [36] **FASER** Collaboration, S. Shively, “FASER Tracker Detector - Commissioning, Installation, and Functionality,” *PoS ICHEP2022* (2022) 1064, [arXiv:2211.04319](https://arxiv.org/abs/2211.04319) [hep-ex].
- [37] F. Campabadal *et al.*, “Design and performance of the ABCD3TA ASIC for readout of silicon strip detectors in the ATLAS semiconductor tracker,” *Nucl. Instrum. Meth. A* **552** (2005) 292–328.
- [38] FASER Collaboration, “FASER Internal Twiki.” 2024. <https://twiki.cern.ch/twiki/bin/viewauth/FASER/InternalFaserSite>.
- [39] **LHCb** Collaboration, “LHCb calorimeters: Technical design report,”.
- [40] C. W. Fabjan and F. Gianotti, “Calorimetry for particle physics,” *Rev. Mod. Phys.* **75** (2003) 1243–1286.
- [41] FASER Collaboration, “FASER PMT Waveform Reconstruction and Calorimeter Energy Calibration,” CERN Document Server:2851265. Internal.
- [42] K. Halbach, “Design of Permanent Multipole Magnets with Oriented Rare Earth Cobalt Material,” *Nucl. Instrum. Meth.* **169** (1980) 1.

- [43] FASER Collaboration, “FASER GitLab Repository.” 2024.
<https://gitlab.cern.ch/faser/>.
- [44] M. Boretto, W. Brylinski, G. Lehmann Miotto, E. Gamberini, R. Sipos, and V. V. Sonesten, “DAQling: an open-source data acquisition framework,” *EPJ Web Conf.* **245** (2020) 01026.
- [45] J.-J. Savioz, “The Beam Synchronous Timing Receiver Interface for the Beam Observation System,” *Tech. Rep.* (2003) .
<https://edms.cern.ch/ui/file/406137/1.0/LHC-BOBR-ES-0001.pdf>.
- [46] A. Buckley, T. Eifert, M. Elsing, D. Gillberg, K. Koeneke, A. Krasznahorkay, E. Moyse, M. Nowak, S. Snyder, and P. van Gemmeren, “Implementation of the ATLAS Run 2 event data model,” *J. Phys. Conf. Ser.* **664** no. 7, (2015) 072045.
- [47] CERN Computing Centre, “CERN Tape Archive,” <https://cta.web.cern.ch/cta/>.
- [48] **FASER** Collaboration, H. Abreu *et al.*, “The trigger and data acquisition system of the FASER experiment,” *JINST* **16** no. 12, (2021) P12028, [arXiv:2110.15186](https://arxiv.org/abs/2110.15186) [[physics.ins-det](https://arxiv.org/abs/2110.15186)].
- [49] T. Sjostrand, “Monte Carlo Generators,” CERN Document Server:999717.
- [50] **FASER** Collaboration, H. Abreu *et al.*, “Search for dark photons with the FASER detector at the LHC,” *Phys. Lett. B* **848** (2024) 138378, [arXiv:2308.05587](https://arxiv.org/abs/2308.05587) [[hep-ex](https://arxiv.org/abs/2308.05587)].
- [51] M. Nowak, P. van Gemmeren, and J. Cranshaw, “ATLAS Event Store and I/O developments in support for Production and Analysis in Run 3,” *EPJ Web Conf.* **245** (2020) 02031.
- [52] G. Barrand *et al.*, “GAUDI - A software architecture and framework for building HEP data processing applications,” *Comput. Phys. Commun.* **140** (2001) 45–55.
- [53] R. Ulrich, T. Pierog, and C. Baus, “Cosmic ray monte carlo package, crmc.” Aug., 2021. <https://doi.org/10.5281/zenodo.5270381>.
- [54] M. Asai, “Geant4-a simulation toolkit,” *Trans. Amer. Nucl. Soc.* **95** (2006) 757.
- [55] M. Guan, J. Cao, C. Yang, Y. Sun, and K.-B. Luk, “Muon simulation at the Daya Bay site,” .
- [56] S. Ostapchenko, “Monte Carlo treatment of hadronic interactions in enhanced Pomeron scheme: I. QGSJET-II model,” *Phys. Rev. D* **83** (2011) 014018, [arXiv:1010.1869](https://arxiv.org/abs/1010.1869) [[hep-ph](https://arxiv.org/abs/1010.1869)].
- [57] F. Riehn, R. Engel, A. Fedynitch, T. K. Gaisser, and T. Stanev, “Hadronic interaction model Sibyll 2.3d and extensive air showers,” *Phys. Rev. D* **102** no. 6, (2020) 063002, [arXiv:1912.03300](https://arxiv.org/abs/1912.03300) [[hep-ph](https://arxiv.org/abs/1912.03300)].

- [58] T. Pierog, I. 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** no. 3, (2015) 034906, [arXiv:1306.0121 \[hep-ph\]](#).
- [59] F. Kling and L. J. Nevay, “Forward neutrino fluxes at the LHC,” *Phys. Rev. D* **104** no. 11, (2021) 113008, [arXiv:2105.08270 \[hep-ph\]](#).
- [60] **FASER** Collaboration, E. M. Welch, F. Kling, and J. L. Feng, “Dark Photon Signal Validation and Systematic Uncertainties for LHC Run 3 in 2022,” **CERN Document Server:2845232**. Internal.
- [61] C. Andreopoulos *et al.*, “The GENIE Neutrino Monte Carlo Generator,” *Nucl. Instrum. Meth.* **A614** (2010) 87–104, [arXiv:0905.2517 \[hep-ph\]](#).
- [62] C. Andreopoulos, C. Barry, S. Dytman, H. Gallagher, T. Golan, R. Hatcher, G. Perdue, and J. Yarba, “The GENIE Neutrino Monte Carlo Generator: Physics and User Manual,” [arXiv:1510.05494 \[hep-ph\]](#).
- [63] A. Bodek and J. L. Ritchie, “Further Studies of Fermi Motion Effects in Lepton Scattering from Nuclear Targets,” *Phys. Rev. D* **24** (1981) 1400.
- [64] A. Bodek and U. K. Yang, “Higher twist, $\xi(\omega)$ scaling, and effective LO PDFs for lepton scattering in the few GeV region,” *J. Phys. G* **29** (2003) 1899–1906, [arXiv:hep-ex/0210024](#).
- [65] M. Glück, E. Reya, and A. Vogt, “Dynamical parton distributions revisited,” *Eur. Phys. J. C* **5** (1998) 461–470, [arXiv:hep-ph/9806404](#).
- [66] **GENIE** Collaboration, J. Tena-Vidal *et al.*, “Hadronization model tuning in genie v3,” *Phys. Rev. D* **105** no. 1, (2022) 012009, [arXiv:2106.05884 \[hep-ph\]](#).
- [67] A. Ferrari, P. R. Sala, A. Fasso, and J. Ranft, “FLUKA: A multi-particle transport code (Program version 2005),”.
- [68] F. Cerutti. Personal communication, November 2, 2020.
- [69] M. Bandieramonte, R. M. Bianchi, J. Boudreau, A. Dell’Acqua, and V. Tsulaia, “The GeoModel tool suite for detector description,” *EPJ Web Conf.* **251** (2021) 03007.
- [70] S. Gadomski, “Model of the SCT detectors and electronics for the ATLAS simulation using Geant4,”.
- [71] **FASER** Collaboration, J. Boyd, D. Casper, *et al.*, “FASER Dark Photon Analysis Note (Moriond 2023),” **CERN Document Server:2851861**. Internal.
- [72] X. Ai *et al.*, “A Common Tracking Software Project,” *Comput. Softw. Big Sci.* **6** no. 1, (2022) 8, [arXiv:2106.13593 \[physics.ins-det\]](#).
- [73] **FASER** Collaboration, “FASER Public Plots.” 2022. <https://faser-public-plots.app.cern.ch/>.

- [74] 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** (2015) 153, arXiv:1410.1280 [hep-ex].
- [75] D. A. Sprott, *Statistical Inference in Science*. Springer New York, NY, 2000.
- [76] **BaBar** Collaboration, J. P. Lees *et al.*, “Search for a Dark Photon in e^+e^- Collisions at BaBar,” *Phys. Rev. Lett.* **113** no. 20, (2014) 201801, arXiv:1406.2980 [hep-ex].
- [77] E. M. Riordan *et al.*, “A Search for Short Lived Axions in an Electron Beam Dump Experiment,” *Phys. Rev. Lett.* **59** (1987) 755.
- [78] **NA48/2** Collaboration, J. R. Batley *et al.*, “Search for the dark photon in π^0 decays,” *Phys. Lett. B* **746** (2015) 178–185, arXiv:1504.00607 [hep-ex].
- [79] **NA64** Collaboration, 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** no. 7, (2020) 071101, arXiv:1912.11389 [hep-ex].
- [80] **NA64** Collaboration, Y. M. Andreev *et al.*, “Search for a New B-L Z’ Gauge Boson with the NA64 Experiment at CERN,” *Phys. Rev. Lett.* **129** no. 16, (2022) 161801, arXiv:2207.09979 [hep-ex].
- [81] M. Davier and H. Nguyen Ngoc, “An Unambiguous Search for a Light Higgs Boson,” *Phys. Lett. B* **229** (1989) 150–155.
- [82] S. Andreas, C. Niebuhr, and A. Ringwald, “New Limits on Hidden Photons from Past Electron Beam Dumps,” *Phys. Rev. D* **86** (2012) 095019, arXiv:1209.6083 [hep-ph].
- [83] J. Blümlein and J. Brunner, “New Exclusion Limits on Dark Gauge Forces from Proton Bremsstrahlung in Beam-Dump Data,” *Phys. Lett. B* **731** (2014) 320–326, arXiv:1311.3870 [hep-ph].
- [84] 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** (1992) 3835–3850.
- [85] **FASER** Collaboration, C. Gwilliam, “First FASER Physics Results.” March, 2023. 57th Rencontres de Moriond on QCD & High Energy Interactions.
- [86] L. A. Anchordoqui *et al.*, “The Forward Physics Facility: Sites, experiments, and physics potential,” *Phys. Rept.* **968** (2022) 1–50, arXiv:2109.10905 [hep-ph].
- [87] J. L. Feng *et al.*, “The Forward Physics Facility at the High-Luminosity LHC,” *J. Phys. G* **50** no. 3, (2023) 030501, arXiv:2203.05090 [hep-ex].

Appendix A

Tracker Commissioning

FASER's tracker detector was tested at each stage of assembly, beginning with individual SCT modules. Each stage of testing served to validate the next one. Individuals conducting the commissioning were called “shifters.” There were also several “experts,” experienced tracker specialists, who supervised and directed shifter activity.

A.1 Prototype Plane

Along with components that were used in the final detector, there is a prototype plane kept on the surface. The prototype plane uses an early version of the aluminum plane frame and eight of the best SCT modules left after others were selected for the main tracker stations. The prototype plane has dedicated dry air, cooling, power and interlock systems. Almost all components of the plane may be switched out to test specific electronics such as a replacement module or patch panel, after which the commissioning procedure is repeated. The prototype plane additionally serves to test new versions of the DAQ and DCS software before deployment to the main detector.

A.2 Module Commissioning

Each SCT module donated by ATLAS was tested for function and quality at the University of Geneva using a readout system provided by Cambridge University. Individual modules were maintained at temperatures lower than 30 °C in an enclosed, aluminum box during testing. FASER tests generally followed ATLAS procedures for testing the modules' performance. This included studying the dependence of noise on bias voltage. These tests found that 150 V is the minimum bias voltage that achieves maximal noise reduction. Each module was associated with a number of functional strips and their comparison to prior ATLAS results. The best performing modules were selected for the tracker planes, as described in chapter 3.

A.3 Plane and Station Procedure

Experts prepared a spreadsheet from a blank template containing all standard commissioning tests. Shifters filled these spreadsheets with test results. Some form of this procedure is conducted whenever there is a change in software or hardware for any component of the tracking system to record any changes for posterity. Some steps may be skipped with the approval of an expert, depending on how the changes are expected to affect the tracking system. Test results were also documented in a CERN-hosted Electronic Logbook, (ELOG).

Commissioning begins with assessing the state of the tested system; this includes power, position, cooling, grounding, dry air, cabling, and checking for ongoing tests. An expert notes the plane and module ID numbers associated with the tested system, which are verified with visual inspection by the shifter. The shifter notes any changes from the default configuration in the spreadsheet. Any physical modifications require the tested system to be powered off. Early commissioning setups applied Kapton tape to electrical connections to minimize sparking and readout noise.




Figure A.1: Kapton tape applied to the underside of a patch panel, preventing contact between soldered pins and the aluminum frame.

Although the final detector setup reads out only two HV channels per plane, the commissioning setup facilitates studying individual modules with a “splitter” cable that allocates an HV channel to each module in a plane. When the commissioning setup includes the splitter cable, it is referred to as an “8-channel” configuration, in contrast to the operational “2-channel” configuration.

The shifter checks that all cooling and dry air tubes are properly connected before activating the chiller. The tested system is monitored for 10 minutes for any signs of water leakage. The shifter notes the flow rates for water and dry air in the spreadsheet.

After cabling verification, the shifter may start powering the system, beginning with PSUs dedicated to the TRB, patch panels, and TIM. Each physical cable is associated with a label in software that must be confirmed with a mapping check. Each module is briefly toggled through the FSM, and the shifter checks that the appropriate sensors respond in both the MPOD and temperature sensors. Any discrepancy is noted and resolved as needed.

Commissioning also monitored the system’s response to changes in temperature caused by different LV settings. LV is off for at least 3 minutes prior to any measurement. The commissioner recorded the temperatures for each NTC sensor as they appeared in the FSM and the time of the measurement. This was repeated for when LV was on but unconfigured, and when LV was on and configured. This informed the expected range of temperatures for all modules.

Current-Voltage (IV) scans measure the dependence of a module or patch panel’s leakage current on bias voltage (HV). Starting from 0 V HV, the shifter runs scripts that apply different bias voltage steps at fixed time intervals. This test was repeated for the LV-off scenario and the LV-on-configured scenario, each for 2-channel and 8-channel configurations, for a total of four IV scans per plane tested. Figure A.2 shows the IV scan user interface, while fig. A.3 is an example result of a standard set of IV scans for a single plane.

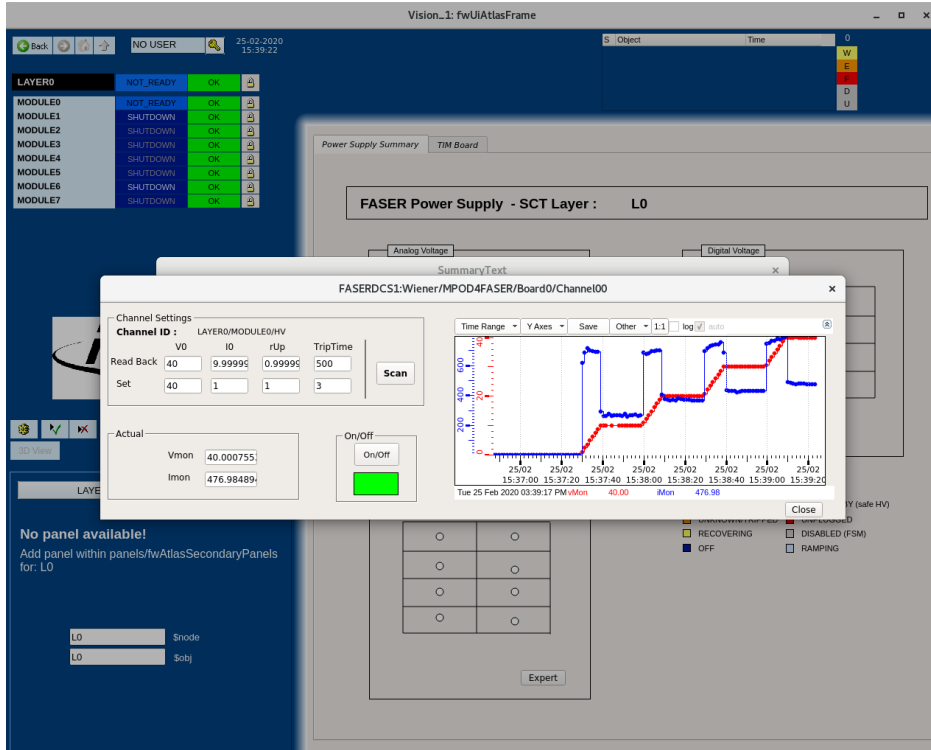


Figure A.2: Example user interface showing operation of IV scan. The leakage current (blue) changes in response to bias voltage (red) which is ramped up in steps.

Plane and station commissioning included a routine tcalib run as described in chapter 3.

After tcalib concludes, the powered system is monitored for 24 hours to observe its long-term behavior.

A.4 Magnetic Field Test

A plane was gradually exposed to one of FASER's permanent magnets to address concerns about how the stray magnetic field could affect its performance. As a control, the plane's performance was assessed 1 m from the magnet aperture, where the stray field was negligible. The plane was moved to 65 mm from the magnet aperture, where the stray field had 60 mT strength. Then, the plane was lowered. at 0.03 m/s until centered at the magnet aperture. Finally, the plane was returned to the initial control position. Any changes in tracker plane




Figure A.3: IV scan results for a single plane comparing the results from 8-channel and 2-channel tests, which are in good agreement.

performance were consistent with its performance before exposure to the magnetic field.

A.5 Surface Cosmics

Cosmic rays provide high-energy particle data without using a collider, which is used to test hardware and software performance. Prior to installation, a test setup on the surface used cosmics to check that the trigger scintillator, DAQ, and a tracking station worked together.

Figure A.4 shows a horizontal track station between scintillators that cover its full acceptance, alongside other supporting systems such as power and cooling. Figure A.5 shows an event collected during this test period, demonstrating that these systems successfully worked together. These tests were the final part of station commissioning prior to installation in T112.




Figure A.4: Photograph of surface cosmics testing setup, including scintillators above and below the tracker station.




Figure A.5: Event display showing a single tracking station with a cosmic ray event. The yellow lines are hit strips; the red dots are clusters; the orange line is a reconstructed track.

A.6 Post-installation and Summary

After installation, the station commissioning procedure repeated to confirm that all systems were operating as expected. Cosmic rays contributed to these tests, but cosmic events that traversed more than one tracker station were exceedingly rare; cosmic rays traversing all three spectrometer stations were expected to occur fewer than one event every 80 days.

Each installed station had good performance comparable to prior commissioning stages. Multi-week cosmic runs observed and confirmed the detector's long-term stability, as well as facilitating the development of the operations shifting procedure.