

A Pointing Electromagnetic Calorimeter for the SHiP Experiment

Dissertation

zur Erlangung des Grades

Doktor der Naturwissenschaften

am Fachbereich Physik, Mathematik und Informatik
der Johannes Gutenberg-Universität
in Mainz



JOHANNES GUTENBERG
UNIVERSITÄT MAINZ

Sebastian Ritter

geboren in Frankfurt am Main

Mainz den 29.5.2026

Dieses Werk ist lizenziert unter einer Creative Commons Namensnennung 4.0 International Lizenz (CC BY 4.0).

Den Lizenztext finden Sie unter: <https://creativecommons.org/licenses/by/4.0/>

Erstgutachter: aus datenschutzrechtlichen Gründen entfernt

Zweitgutachter: aus datenschutzrechtlichen Gründen entfernt

Tag der mündlichen Prüfung: 9. Juli 2026

Abstract

Feebly interacting particles (FIPs) are a broad class of beyond-Standard-Model candidates for dark matter and the origin of neutrino masses. Their defining feature is a very weak coupling to Standard Model fields, which suppresses their production cross-sections and results in macroscopic decay lengths, making high-luminosity beam-dump experiments at the CERN SPS a great discovery environment. Axion-like particles (ALPs) are a minimal, model-independent subclass of FIPs arising from the spontaneous breaking of approximate global symmetries, experimentally accessible via an $ALP \rightarrow \gamma\gamma$ signature. Both experiments studied in this thesis, SHADOWS and SHiP, target neutral final states with displaced vertices, requiring an electromagnetic calorimeter (ECAL) that reconstructs the energy as well as the direction of photon showers. An evaluation of the ALP decay kinematics in the boosted regime relevant at the SPS shows that the propagated uncertainty of the invariant mass is dominated by the angular resolution, making pointing resolution the primary figure of merit throughout this work.

For SHADOWS, the off-axis geometry and moderate final-state energies set granularity as the key design parameter. A channel pitch of 1 cm removes any granularity-induced reconstruction bias at small angles for showers above 2 GeV. Diphoton vertex and invariant-mass reconstruction are demonstrated for two ALP benchmark scenarios, and the energy resolution of $\sim 11\%/\sqrt{E/\text{GeV}}$ meets the requirements. For the 2024 approved on-axis experiment SHiP, the harder and more forward final states impose stricter angular resolution requirements. A longitudinal air gap is introduced as the central design element, increasing the pointing lever arm without adding active channels or absorber.

A prototype was built to validate the pointing concept experimentally and to provide the data needed to tune the simulation framework. The prototype was tested at the CERN SPS at energies from 3 to 288 GeV, establishing that the plastic scintillator strip-based concept performs calorimetric pointing under realistic conditions. A simulation tuned to the measured detector response reproduces all test-beam observables and is extended to a full detector module for the SHiP ECAL optimisation.

Systematic studies of the gap depth and position with single photons identify the optimal detector geometry and channel structure: a gap with a depth of 100 cm at $3X_0$ combined with an alternating strip layer configuration. The conservative energy resolution estimate of $14.4\%/\sqrt{E/\text{GeV}}$ meets the SHiP requirement. Beyond the optimal configuration, the study establishes a systematic understanding of which design parameters control which performance observables, providing the foundation for well-informed decisions in the subsequent design stages and for studies with more sophisticated figures of merit.

Preface

This thesis was written independently by the author. It has not been submitted elsewhere for examination purposes. Any publications listed below are based on results first developed and presented in this thesis.

Parts of the work presented in this thesis are intended for publication in peer-reviewed journals. In particular, the detector design studies, prototype test beam results, and simulation-based optimisation of the electromagnetic calorimeter for SHiP are being prepared for submission.

The prototype test beam campaign carried out at the CERN SPS, the tuning of the simulation framework to the measured detector response, and the systematic optimisation studies for the SHiP ECAL module are original contributions of the author within the SHiP collaboration. The SHiP collaboration involves scientists and engineers from multiple institutes across Europe and beyond. This implies the usage of common software and infrastructure developed collectively within the collaboration. The collection of the test beam data and the development of the simulation tools used to analyse and interpret the results benefit from the effort of many current and former collaboration members.

Two conference contributions based on the simulation work and ECAL design studies presented in Chapter 5 have already been published as proceedings:

C. C. Delogu et al., “The SHADOWS calorimeter — A pointing calorimeter for future beam dump experiments”, *Nucl. Instrum. Meth. A* **1069**, 169873 (2024). [1]

S. Ritter et al., “A Pointing Electromagnetic Calorimeter for FIP Experiments leveraging $X \rightarrow \gamma\gamma$ Decays”, *EPJ Web Conf.* **320**, 00011 (2025). [2]

The author of this thesis is the main author and speaker of the second contribution and a co-author of the first.

In addition, the studies contributed to the following technical proposal submitted to CERN:

SHADOWS Collaboration, “SHADOWS Technical Proposal”, CERN-SPSC-2023-029, SPSC-P-367 (2023). [3]

The author of this thesis is author or co-author of all three publications and wrote sections describing the SHADOWS ECAL design and simulation studies. As a result, the descriptions in Chapter 5 and in these publications are similar in structure and wording.

A further contribution based on the work presented in this thesis has been accepted for presentation at the *21st International Conference on Calorimetry in Particle Physics* (CALOR 2026) and will be published in the conference proceedings under the title:

“Design, Simulation, and Performance of a High-Granularity Pointing ECAL for the SHiP Experiment at CERN”

The author of this thesis is the main author and speaker of this contribution.

Contents

Introduction	3
1 Theoretical Framework	4
1.1 The Standard Model of Particle Physics	4
1.1.1 Open Questions and the Limits of the Standard Model	4
1.2 The Gauge System and the SM Lagrangian	6
1.3 Symmetries in the Standard Model	6
1.3.1 Spontaneous Symmetry Breaking and Goldstone Bosons	6
1.4 Axion-Like Particles as Pseudo-Nambu-Goldstone Bosons	7
1.4.1 The Peccei-Quinn Axion	8
1.4.2 The ALP Effective Lagrangian	8
1.4.3 The ALP Parameter Space	9
2 Signal Kinematics in Beam-Dump Experiments	11
2.1 From the Energy Frontier to the Intensity Frontier	11
2.2 The Beam Dump Concept and the CERN Beam Dump Facility	12
2.3 ALP Production at Proton Beam Dumps	12
2.3.1 Primakoff Production	13
2.3.2 Light Meson Decays	13
2.3.3 Heavy-Flavor Meson Decays	13
2.3.4 Proton Bremsstrahlung	14
2.3.5 Beam Induced Backgrounds	14
2.4 Signal Kinematics	14
3 Particle Interactions and Detection Principles	16
3.1 Particle Interactions with Matter	16
3.1.1 Ionisation	16
3.1.2 Photon Interactions in Matter	16
3.1.3 Hadronic Interactions	19
3.2 Principles of Calorimetry	19
3.2.1 Energy Resolution	20
3.2.2 Electromagnetic Showers	20
3.2.3 Hadronic Showers	23
3.3 Calorimeter Technologies	25
3.3.1 Homogeneous Calorimeters	25
3.3.2 Sampling Calorimeters	26
3.3.3 Particle Identification from Combined Detector Response	26
3.4 Instrumentation Technologies	27
3.4.1 Plastic Scintillators	27
3.4.2 Silicon Photomultipliers	28
3.4.3 Wavelength-Shifting Fibres	29
3.4.4 Double Sided Readout	30
4 Experimental Options at the CERN SPS	32
4.1 Signal Acceptance and Detector Geometry	32
4.2 SHADOWS: The Off-Axis Concept	33

4.3	SHiP: The On-Axis Experiment	34
4.4	Beam-Dump Experiment Blueprint	35
4.5	Calorimetry for Beam Dump Experiments	39
4.5.1	Diphoton Opening Angle	39
4.5.2	Dominance of Angular Uncertainty	40
4.5.3	Pointing Reconstruction and Calorimeter Segmentation	43
4.5.4	The Design Challenge	43
5	Studies towards the SHADOWS Electromagnetic Calorimeter	44
5.1	Boundary Conditions and Performance Requirements	44
5.2	Reference Design: Tile Calorimeter	44
5.3	Optimisation Studies	45
5.3.1	Simulation Setup	45
5.3.2	Granularity Optimisation	47
5.3.3	Channel Count Optimisation	48
5.4	Baseline Design	49
5.5	Performance	49
5.6	Benchmark Scenarios	49
5.6.1	Invariant Mass Resolution	50
5.6.2	Vertex Reconstruction and Background Rejection	51
5.7	Conclusion	52
6	The SHiP ECAL Design	53
6.1	Boundary Conditions and Design Drivers	53
6.1.1	Longitudinal Gap	53
6.1.2	Strip Width and Channel Count	54
6.2	Baseline Design	55
7	Conceptual Prototype Testing at Test Beams	57
7.1	The Electromagnetic Calorimeter Prototype	57
7.1.1	Thin Strip Layers	57
7.1.2	Wide Strip Layers	59
7.1.3	The Prototype Assembly	59
7.2	The Hadronic Calorimeter Prototype	60
7.3	Readout	64
7.4	Trigger	65
7.5	Test Beams	67
7.6	Physics Program at the SPS Test Beam	68
8	Analysis of Test Beam Data	70
8.1	Data Preparation	70
8.1.1	Event Building	70
8.1.2	Pedestal Correction	70
8.1.3	Intercalibration	71
8.1.4	MIP Calibration	72
8.1.5	SiPM Saturation Compensation	75
8.1.6	Particle Identification	76

8.2	Shower Direction Reconstruction	78
8.2.1	ADC Saturation	79
8.2.2	Transverse Profile Model	80
8.2.3	Cost Function	80
8.2.4	Minimisation	81
8.3	Discussion of Results	82
8.3.1	Intercept Reconstruction	82
8.3.2	Angular Resolution	84
9	Simulation of Test Beam Setup	87
9.1	Particle Gun	87
9.2	Digitisation	87
9.3	Simulation Validation	89
9.3.1	Per-Layer Energy Deposits	89
9.3.2	Shower Maximum	93
9.3.3	Total Shower Energy	94
9.3.4	Comparison of Intercept and Angle Reconstruction	94
10	Optimisation of the SHiP ECAL	98
10.1	From Prototype to Full Detector Section	98
10.2	Simulation	98
10.2.1	Particle Generation	99
10.2.2	Reconstruction	99
10.3	Air Gap Optimisation	102
10.3.1	Gap Depth Optimisation	103
10.3.2	Gap Position Optimisation	107
10.4	Sampling Optimisation	110
10.5	Proposed Setup	111
10.6	Future Studies and Potential Improvements	111
	Conclusion	117
A	Appendix	119
A.1	Comparator schematic	119
A.2	FERS Daisy Chain	120
A.3	SHADOWS ALP Benchmark Scenarios	121
A.4	Additional Plots for Chapter 10	122
A.4.1	Shower Profiles	122
A.4.2	Optimisation plots with statistical sigma	125
A.4.3	Sampling with layers removed in the back	130
A.4.4	Energy reconstruction linearity	131
	List of Figures	132
	List of Tables	136
	References	137

Introduction

The Standard Model leaves several fundamental questions unanswered. A broad class of extensions predicts the existence of new particles with couplings far below the weak scale. These *Feebly Interacting Particles* (FIPs) have small production cross-sections, and their suppressed decay rates produce macroscopic decay lengths in the laboratory frame. *Axion-like particles* (ALPs) are a minimal, model-independent subclass, arising from the spontaneous breaking of approximate global symmetries. This theoretical background is discussed in Chapter 1. Beam-dump experiments with dense targets and long decay volumes are the ideal environment to search for such particles and are discussed in Chapter 2. At the CERN SPS, two experiments have been proposed for this purpose. SHADOWS is a concept for an off-axis detector targeting heavy-flavour production in the lateral halo [3]. SHiP is a large on-axis experiment covering the full forward phase space [4], approved by CERN management in 2024 [5].

Both experiments must be able to reconstruct neutral long-lived particles decaying inside the fiducial volume. This requires an *electromagnetic calorimeter* (ECAL) that measures both the energy and the direction of the incident showers to reconstruct among others the possible $ALP \rightarrow \gamma\gamma$ decay. The required detector physics is discussed in Chapter 3. In the highly boosted regime relevant at the CERN SPS, a derivation of decay kinematics shows that the propagated mass uncertainty is dominated by the angular reconstruction rather than the energy resolution, defining a clear optimisation path for the ECAL design.

Two experimental concepts employing pointing calorimetry based on plastic scintillator strip sampling ECALs are studied in detail, with their respective geometries and the common requirements on the ECAL discussed in Chapter 4. For SHADOWS, a systematic simulation study determines the optimal channel pitch and defines a design which successfully reconstructs benchmark scenarios (Chapter 5). For SHiP, the on-axis geometry produces higher energy, more forward final states with stricter angular resolution requirements. A longitudinal air gap is introduced as the central design element increasing the lever arm within the existing ECAL stack. The SHiP baseline design is described in Chapter 6.

The central experimental part of this work is the construction and validation of a conceptual prototype to provide experimental data as the basis for the SHiP ECAL simulation framework. The prototype design and the test beam campaign carried out at the CERN SPS are described in Chapter 7. A scaled-down but functionally complete prototype was built to demonstrate calorimetric pointing under realistic conditions. The analysis of the test beam data and the shower direction reconstruction are presented in Chapter 8. A simulation of the prototype, validated against the test beam data, is described in Chapter 9 and then extended to a full detector module for the optimisation of the SHiP ECAL design in Chapter 10. In this optimisation, the depth and position of the longitudinal air gap, the sampling strategy, and the associated acceptance and containment effects are studied systematically. The result is the identification of the optimal detector geometry and channel structure within the given boundary conditions, together with a systematic understanding of how each design parameter controls each performance observable, providing a design reference for the next stages of the SHiP ECAL development.

1 Theoretical Framework

The *Standard Model* (SM) of particle physics is a highly successful description of fundamental interactions based on gauge symmetries and their associated fields. Its structure is governed by the symmetry group $SU(3)_C \times SU(2)_L \times U(1)_Y$, with particle masses generated through spontaneous electroweak symmetry breaking via the Higgs mechanism. Despite its experimental success, the SM is regarded as an effective theory, as it does not address several conceptual questions, including the origin of neutrino masses and the role of global symmetries.

One class of extensions arises from the spontaneous breaking of approximate global symmetries. Such theories generically predict the existence of light pseudo-Goldstone bosons, whose properties are determined by the symmetry structure rather than the details of the underlying high-energy theory. *Axion-like particles* (ALPs) provide a minimal and model-independent realisation of this framework, where the mass and couplings are treated as free parameters of the effective theory. The description of SM formalisms in this chapter is based on the textbooks of Griffiths [6] and Thomson [7]. Throughout this thesis natural units ($c = \hbar = 1$) are used.

1.1 The Standard Model of Particle Physics

The SM of particle physics provides the foundation for our current understanding of fundamental interactions. It describes the fundamental components of matter and three of the four known forces within a unified gauge-theoretic framework. Within this framework, matter is organised into three generations of quarks and leptons, with each generation containing an up-type quark, a down-type quark, a charged lepton, and an associated neutrino. The first generation, comprising the up and down quarks, the electron, and the electron neutrino, is the building block of all stable baryonic matter, while the second and third generations are heavier copies observed in high-energy interactions. Spin-1 gauge bosons mediate the forces between these fermions. Eight massless gluons carry the strong force, the massless photon carries the electromagnetic force, and the massive $W^\pm$ and Z bosons carry the weak force. Completing the known particle content is a single electrically neutral scalar boson: the Higgs boson with mass $m_h \approx 125 \text{ GeV}$, responsible for electroweak symmetry breaking and the generation of mass for the $W^\pm$, Z , and the charged fermions [8]. A visualisation of the full SM particle content is shown in Figure 1.

From precision measurements of the anomalous magnetic moment of the electron [10] to the discovery of the Higgs boson at the *Large Hadron Collider* (LHC) in 2012 [11], the SM has predicted the existence of particles and the outcomes of high-energy collisions with unmatched accuracy. No statistically significant deviation from its predictions has been established in any high-energy collider search to date [12], but the SM framework continues to be constrained by high-precision measurements like the anomalous magnetic moment of the muon [13, 14].

1.1.1 Open Questions and the Limits of the Standard Model

Despite its success, the SM is known to be fundamentally incomplete. Several physical observations are not explained by the SM, suggesting the existence of new physics beyond its framework. These include:




Figure 1: The particles of the Standard Model [9].

Dark Matter: Cosmological and astrophysical observations provide compelling evidence for a non-baryonic form of matter accounting for approximately 27% of the energy density of the universe. The SM contains no viable candidate for such a particle [15].

Neutrino Masses: The minimal SM predicts neutrinos to be massless, yet the observation of neutrino oscillations establishes that they carry non-zero masses and mixing angles, constituting the most direct laboratory evidence for physics beyond the SM. The mechanism responsible for generating these masses remains unknown [16].

Matter–Antimatter Asymmetry: The observed excess of matter over antimatter in the universe cannot be accounted for by the CP-violating phase in the CKM matrix alone, requiring additional sources of CP violation not present in the SM [17].

Quantum Gravity: The SM unifies three of the four fundamental forces within a quantum field-theoretic framework but leaves General Relativity entirely unaccounted for.

Strong CP Problem: The QCD Lagrangian admits a CP-violating term whose coefficient is experimentally constrained to very small values with no SM explanation for this suppression [18].

Each of these shortcomings points toward physics beyond the SM, but any such extension must be built on the SM's gauge structure and symmetry-breaking mechanism.

1.2 The Gauge System and the SM Lagrangian

The dynamics of the Standard Model are condensed in the Lagrangian density $\mathcal{L}_{\text{SM}}$, which is required to be invariant under local transformations of the gauge group $G_{\text{SM}} = SU(3)_c \times SU(2)_L \times U(1)_Y$. The total Lagrangian is partitioned into gauge kinetic, fermionic kinetic, scalar, and Yukawa sectors:

$$\mathcal{L}_{\text{SM}} = -\frac{1}{4}G_{\mu\nu}^a G^{a\mu\nu} - \frac{1}{4}W_{\mu\nu}^i W^{i\mu\nu} - \frac{1}{4}B_{\mu\nu} B^{\mu\nu} + \sum_{\psi} \bar{\psi} i \not{D} \psi + |D_{\mu} H|^2 - V(H) + \mathcal{L}_{\text{Yukawa}}, \quad (1)$$

where $G_{\mu\nu}^a$, $W_{\mu\nu}^i$, and $B_{\mu\nu}$ are the field strength tensors of the $SU(3)_c$, $SU(2)_L$, and $U(1)_Y$ gauge fields respectively, and D_{μ} denotes the gauge-covariant derivative. The requirement of local gauge invariance determines the form of all interaction vertices, leaving the fermion masses and mixing angles as the only free parameters not fixed by symmetry. This rigid structure is what makes the SM so predictive and, at the same time, prevents it from explaining the observations of Section 1.1.1. The key to understanding both the internal structure of the SM and its possible extensions lies in the symmetry properties of the Lagrangian.

Each term in the Lagrangian has a direct physical interpretation. The gauge kinetic terms encode the dynamics and self-interactions of the force carriers, fixed by the gauge symmetry. The fermionic kinetic terms determine how matter fields interact with gauge bosons. The Yukawa sector introduces fermion masses and flavour structure after electroweak symmetry breaking, while the scalar sector is responsible for the non-trivial vacuum structure of the theory.

1.3 Symmetries in the Standard Model

The SM Lagrangian admits two distinct classes of symmetries. Local gauge symmetries are imposed as fundamental construction principles and are exact to all orders in perturbation theory. Global symmetries, namely baryon number B and the lepton numbers $L_{e,\mu,\tau}$, emerge at the renormalisable level without being explicitly required and carry no fundamental protection against violation by *Beyond Standard Model* (BSM) effects at high energies.

Many well-motivated extensions introduce additional global symmetries associated with new physics at energy scales far above those explored by current collider experiments. If such a symmetry is not exact¹, the theory predicts the existence of a new light, weakly coupled pseudo-scalar particle whose mass and interactions are both suppressed by the large scale at which the new symmetry is defined.

1.3.1 Spontaneous Symmetry Breaking and Goldstone Bosons

A symmetry is spontaneously broken when the Lagrangian respects it, but is broken in the ground state. In quantum field theory, this situation arises when the potential energy of a field has its minimum not at zero field value but on a continuous manifold

¹A symmetry is said to be exact if every term in the Lagrangian is invariant under the corresponding transformation. In practice, many symmetries hold for the dominant contributions to the dynamics but are violated by subdominant terms, such as small mass parameters or higher-dimensional operators suppressed by a large energy scale.

of degenerate ground states. A well-known example is a complex scalar field ϕ with a potential of the form

$$V(\phi) = -\mu^2|\phi|^2 + \lambda|\phi|^4, \quad (2)$$

where $\mu^2 > 0$ and $\lambda > 0$. This potential, often called the Mexican-hat potential, has a local maximum at $\phi = 0$ and a continuous circle of minima at $|\phi| = \mu/\sqrt{2\lambda}$. The Lagrangian is invariant under global $U(1)$ rotations $\phi \rightarrow e^{i\alpha}\phi$, but the vacuum must settle on a point on this circle, breaking the $U(1)$ symmetry spontaneously. The physical consequence of this symmetry breaking is captured by *Goldstone's theorem* [19], which states that for every generator of a continuous global symmetry that is spontaneously broken, a massless boson appears in the low-energy spectrum. These massless particles are called *Nambu-Goldstone bosons* (NGBs), which physically correspond to excitations of the field along the circle of degenerate minima. This mechanism is already part of the SM. The Higgs field spontaneously breaks the electroweak symmetry at the scale $v \approx 246$ GeV, and the Higgs boson is the massive radial excitation of the Higgs field around the symmetry-breaking minimum. The two excitation modes of the Mexican-hat potential are illustrated in Figure 2. The massive radial mode corresponds to the Higgs boson. The massless angular modes along the circle of minima are the would-be NGBs which give the $W^\pm$ and Z bosons their mass.




Figure 2: Visual representation of the Mexican-hat potential $V(\phi) = -\mu^2|\phi|^2 + \lambda|\phi|^4$ and its degenerate ground states. The Higgs mode (radial excitation) and the Goldstone mode (angular excitation along the circle of minima) are indicated by red and blue arrows, respectively. Illustration taken from [20].

1.4 Axion-Like Particles as Pseudo-Nambu-Goldstone Bosons

The arguments of the last section lead directly to the concrete prediction that any extension of the SM that introduces a new complex scalar field charged under a global $U(1)$ symmetry, with a potential that drives spontaneous symmetry breaking at a scale f_a , will generically predict a light pseudo-scalar called *pseudo-Nambu-Goldstone boson* (pNGB). If the $U(1)$ symmetry is only approximate, the would-be

NGB acquires a mass parametrically smaller than f_a , set by the scale of the explicit breaking, while retaining couplings to SM fields suppressed by $1/f_a$. This light pseudo-scalar is generally called an *Axion-Like Particle* (ALP).

1.4.1 The Peccei-Quinn Axion

The term *axion-like* refers to the historical axion proposed by Peccei and Quinn [21] to resolve the strong CP problem. The SM permits an additional CP-violating term in the *quantum chromodynamics* (QCD) Lagrangian, $\mathcal{L}_\theta = \bar{\theta} (g_s^2/32\pi^2) G_{\mu\nu}^a \tilde{G}^{a\mu\nu}$, whose coefficient $\bar{\theta}$ is experimentally constrained to $|\bar{\theta}| < 10^{-10}$ [22] with no symmetry reason to be small. Peccei and Quinn resolved this by introducing a new global $U(1)_{\text{PQ}}$ symmetry, spontaneously broken at an assumed large scale f_a , whose associated pNGB (the QCD axion) dynamically relaxes $\bar{\theta}$ to zero. The axion mass is not a free parameter but is fixed by the QCD potential:

$$m_a \simeq \frac{m_\pi f_\pi}{f_a} \approx 5.7 \mu\text{eV} \left(\frac{10^{12} \text{ GeV}}{f_a} \right), \quad (3)$$

constraining the QCD axion to a single line in the $(m_a, g_{a\gamma})$ plane [23], where $g_{a\gamma}$ is the axion-photon coupling constant scaling as $g_{a\gamma} \propto 1/f_a$. An ALP is the generalisation of this construction to an arbitrary broken global symmetry with no equivalent of the QCD potential to fix m_a . The mass and couplings become independent free parameters, opening the full two-dimensional parameter space.

1.4.2 The ALP Effective Lagrangian

At energies well below f_a , the physics of the ALP is captured by an *effective field theory* (EFT). The residual symmetry of the broken $U(1)$ is a shift symmetry $a \rightarrow a + \text{const}$, which forbids non-derivative mass terms and restricts interactions to derivative couplings at leading order. Including the leading interactions with SM gauge bosons and fermions, the ALP EFT Lagrangian reads:

$$\mathcal{L}_{\text{ALP}} = \frac{1}{2} \partial^\mu a \partial_\mu a - \frac{1}{2} m_a^2 a^2 - \frac{1}{4} g_{a\gamma} a F_{\mu\nu} \tilde{F}^{\mu\nu} - \frac{1}{4} g_{ag} a G_{\mu\nu}^a \tilde{G}^{a\mu\nu} - \frac{\partial_\mu a}{f_a} \sum_\psi c_\psi \bar{\psi} \gamma^\mu \gamma^5 \psi, \quad (4)$$

where $F_{\mu\nu}$ and $G_{\mu\nu}^a$ are the electromagnetic and gluon field strength tensors, $\tilde{F}^{\mu\nu} = \frac{1}{2} \epsilon^{\mu\nu\rho\sigma} F_{\rho\sigma}$ denotes their duals, $g_{a\gamma}$ and g_{ag} are the photon and gluon coupling constants, and c_ψ are dimensionless fermion coupling coefficients.

Because of its clean and model-independent signature, this thesis focuses on the decay channel $a \rightarrow \gamma\gamma$, governed by the photon coupling $g_{a\gamma}$.

The mass m_a and the coupling $g_{a\gamma}$ are independent parameters which open a two-dimensional parameter space that covers a broad landscape of BSM scenarios. Different *ultraviolet* (UV) completions that reduce to the same EFT at low energies populate different regions of this space. Unlike the QCD axion, which predicts a specific target, no single UV completion is theoretically preferred, motivating broad experimental exploration across the full parameter space. The pNGB origin of the ALP restricts its interactions to be suppressed by $1/f_a$ at leading order.

1.4.3 The ALP Parameter Space

The experimentally accessible ALP parameter space spans roughly $m_a \in [10^{-6}, 10^3]$ GeV and $g_{a\gamma} \in [10^{-12}, 10^{-1}]$ GeV $^{-1}$, with large regions already excluded by astrophysical observations, laboratory searches, and collider experiments as shown in Figure 3 with Figure 4 zooming into the mass range accessible to accelerator-based experiments [24]. Stellar cooling bounds, arising from the requirement that ALP emission does not excessively accelerate energy loss in stars, constrain $g_{a\gamma} \lesssim 10^{-10}$ GeV $^{-1}$ for $m_a \lesssim$ keV [25]. At higher masses, bounds from beam-dump experiments, e^+e^- colliders, and rare meson decays carve out large regions of the parameter space. Despite these constraints, significant open windows remain in the MeV–GeV mass range at couplings $g_{a\gamma} \sim 10^{-5}$ – 10^{-3} GeV $^{-1}$, which are inaccessible to both astrophysical probes and high-energy collider experiments.




Figure 3: Current exclusion limits on ALPs with photon coupling in the plane of coupling $g_{a\gamma}$ versus mass m_a , spanning the full experimentally accessible parameter space. Excluded regions from astrophysical observations, laboratory searches, and collider experiments are shown as filled areas. Figure adapted from [24].

The open MeV–GeV window identified above is the direct experimental target of the experiments discussed in this thesis. ALPs in this region are too weakly coupled and long-lived for the reach of collider experiments. Probing them requires a new approach discussed in Chapter 2. There, the ALP production mechanisms and resulting decay kinematics at a proton beam dump are discussed. Chapter 4 then shows how these kinematic features motivate the design of two different experimental concepts.




Figure 4: Zoom into the ALP parameter space in the mass range relevant for accelerator-based experiments (note the different mass units relative to Figure 3). Current exclusion bounds are shown as filled coloured areas. The open window in the MeV–GeV range at couplings $g_{a\gamma} \sim 10^{-5}$ – 10^{-3} GeV^{-1} is the target region of this thesis. Figure adapted from [24].

2 Signal Kinematics in Beam-Dump Experiments

The last chapter developed the theoretical framework for ALPs, the primary physics target of the pointing calorimeter designed in this thesis. However, ALPs are not an isolated prediction but belong to the broader class of *Feebly Interacting Particles* (FIPs), a general category comprising all hypothetical particles whose couplings to SM fields are many orders of magnitude smaller than those of the known weak interactions. This class includes dark photons, heavy neutral leptons, dark scalars and ALPs [26]. Despite their different theoretical origins and quantum numbers, all FIPs share the defining property of a very small coupling constant $g \ll 1$, imposing a common experimental logic: suppressed production cross-sections demand maximising the number of interactions, while the correspondingly suppressed decay rates result in macroscopic lifetimes that allow the particles to travel macroscopic distances before decaying.

This chapter introduces the beam-dump concept, characterises the ALP production mechanisms and backgrounds at the SPS, and derives the kinematic signatures that shape the experimental designs.

2.1 From the Energy Frontier to the Intensity Frontier

The LHC has produced one landmark discovery: the observation of the Higgs boson in 2012 [11], completing the particle content of the SM and validating the Brout–Englert–Higgs mechanism [27–29] as the correct description of electroweak symmetry breaking. Beyond the Higgs boson, no statistically significant signal of new physics has been established. If BSM particles exist and couple to SM fields with order-unity strength, their masses must lie above several TeV [30].

A complementary approach is the search for BSM particles at the *intensity frontier*. Particles that couple feebly to the SM, with coupling constants $g \ll 1$, are kinematically accessible at any collision energy but are produced with cross-sections $\sigma \propto g^2$. Rather than maximising collision energy, one maximises the number of interactions. In fixed-target experiments, the relevant quantity is the number of Protons on Target (PoT) with modern facilities targeting $N_{\text{PoT}} \sim 10^{20}$ [31]. So far, the intensity frontier is utilised mainly for the search for rare or forbidden phenomena in SM processes [32, 33] and precision measurements [13, 34].

The photon coupling $g_{a\gamma}$, which enters the photon interaction term in the ALP-Lagrangian (Eq. (4)), directly governs the dominant decay mode in the photon-portal scenario². In the rest frame of the ALP, this vertex gives rise to the two-photon decay $a \rightarrow \gamma\gamma$, whose partial width follows from a straightforward tree-level calculation:

$$\Gamma(a \rightarrow \gamma\gamma) = \frac{g_{a\gamma}^2 m_a^3}{64\pi}. \quad (5)$$

The cubic dependence on m_a and the quadratic dependence on $g_{a\gamma}$ determine the ALP lifetime $\tau = 1/\Gamma$ and, after boosting to the laboratory frame, its decay length $\lambda = \gamma\beta\tau$ [35]. For the small couplings and moderate masses of phenomenological

²The three terms in Eq. (4) beyond the kinetic and mass terms define the photon, gluon, and fermion portals, respectively. Which of these dominates is determined by the chosen UV completion.

interest, λ reaches macroscopic scales of tens to hundreds of metres in the laboratory frame for GeV-scale ALPs produced at a high-energy beam.

2.2 The Beam Dump Concept and the CERN Beam Dump Facility

In a beam dump, a dense target serves to maximise FIP production from the large secondary meson flux while simultaneously absorbing hadronic shower components that would otherwise overwhelm the downstream detectors. An extended low-density decay volume downstream of the shielding provides the baseline over which long-lived FIPs decay in flight. The sensitivity to a given point in the (m, g) parameter space is determined by the signal yield

$$N_{\text{signal}} \propto N_{\text{PoT}} \cdot \sigma_{\text{prod}}(g) \cdot P_{\text{decay}}, \quad (6)$$

where $\sigma_{\text{prod}} \propto g^2$ is the production cross-section and P_{decay} is the probability for the produced particle to survive the shielding and decay within the instrumented decay volume.

The concrete realisation of this concept takes place at the CERN *Super Proton Synchrotron* (SPS), which delivers a 400 GeV proton beam with an intensity of 4×10^{13} protons per spill to the North Area experimental complex. Throughout this thesis, beam momenta are referred to as beam energies, since all beams discussed are highly relativistic ($E \approx p$). The 400 GeV beam energy is a critical working point since it exceeds the production thresholds for charm and beauty hadrons, and the high beam intensity allows for the accumulation of $\mathcal{O}(10^{19})$ to $\mathcal{O}(10^{20})$ PoT over several years of operation [4]. The focal point for intensity frontier hidden-sector searches is the CERN *Beam Dump Facility* (BDF), set for construction in the ECN3 underground cavern [31]. Its high-density target acts simultaneously as a high-luminosity factory for light mesons and heavy-flavour hadrons and as approximately 30 m of hadron absorber and muon shielding, reducing the overwhelming forward SM flux by several orders of magnitude before the decay volume [36].

2.3 ALP Production at Proton Beam Dumps

Throughout this work, sensitivity is assessed for the decay $a \rightarrow \gamma\gamma$, which is the dominant channel for a photon-coupled ALP and produces the cleanest experimental signature in a calorimetric search. The observable is a pair of photons originating from a common displaced vertex inside the fiducial volume.

When a 400 GeV proton beam strikes a dense target, an intense hadronic and electromagnetic cascade develops within the first interaction lengths of the material, which is the primary factory for ALP production. Four distinct mechanisms contribute, each scaling differently with the ALP mass and coupling, and each imprinting a characteristic kinematic signature on the produced ALP flux. While the discussion here is specialised on the ALP, the same four production topologies recur across all FIP species. Only the relative importance of each channel and the relevant coupling replacement differ [26]. If not stated otherwise, this section is based on [35, 37, 38].

2.3.1 Primakoff Production

The Primakoff process is the coherent conversion of a photon into an ALP in the Coulomb field of a target nucleus: $\gamma + Z \rightarrow a + Z$. High-energy photons are produced abundantly in the hadronic cascade through $\pi^0 \rightarrow \gamma\gamma$ and $\eta \rightarrow \gamma\gamma$ decays, providing an intense source of primary photons throughout the target. For ALPs with dominant photon coupling and $m_a \lesssim \mathcal{O}(1 \text{ GeV})$, Primakoff production is expected to be the dominant channel. The Primakoff cross-section scales as $g_{a\gamma}^2 Z^2$, where the Z^2 scaling makes heavy target materials particularly effective. The process is strongly forward-peaked due to the small momentum transfer to the nucleus. This characteristic makes Primakoff production particularly relevant for long-baseline beam-dump experiments, as a large ALP phase space can be covered with a relatively compact detector. The Feynman diagram for Primakoff production is shown in Fig. 5.




Figure 5: Feynman diagram illustrating the Primakoff production of an ALP in the Coulomb field of a nucleus. The cross-section scales as $g_{a\gamma}^2 Z^2$, motivating the use of heavy target materials. [35]

2.3.2 Light Meson Decays

The hadronic cascade produces a large flux of light neutral mesons, predominantly π^0 and η . For m_a below the meson mass, the ALP can be emitted in radiative decays such as $\pi^0 \rightarrow \gamma a$ and $\eta \rightarrow \gamma a$, or replace the diphoton final state entirely. The branching ratios for these processes are suppressed by $g_{a\gamma}^2/m_{\text{meson}}^2$ relative to the SM decay and are therefore small, but the enormous light meson flux produced at 10^{20} PoT could compensate for this suppression. At 400 GeV, light mesons are produced with a large longitudinal boost and small transverse momentum. The characteristic emission angle of a pion with momentum p_π is

$$\theta_\pi \sim \frac{m_\pi}{p_\pi}, \quad (7)$$

and ALPs produced in their decays consequently follow a distribution that is strongly peaked along the beam axis.

2.3.3 Heavy-Flavor Meson Decays

The SPS proton beam causes a significant production of D and B mesons through strong interactions in the target. For m_a in the range of hundreds of MeV to a few GeV, these heavy-flavour mesons may provide the dominant production channel via

flavour-changing processes such as $B \rightarrow Ka$ and $D \rightarrow \pi a$, as well as through direct radiative decays [39]. The kinematic properties of this channel differ fundamentally from those of Primakoff and light meson production. As Eq. (7) describes, heavy mesons are produced with significantly smaller Lorentz boosts and larger intrinsic transverse momenta, resulting in a substantially broader angular distribution with transverse displacements at the detector that can reach several metres.

2.3.4 Proton Bremsstrahlung

At the highest energies, ALPs can be produced directly through bremsstrahlung-like processes in the primary proton–target interaction, analogous to ordinary photon bremsstrahlung but with the photon replaced by an ALP via the $g_{a\gamma}$ coupling [40]. This channel contributes predominantly at low m_a and small coupling, and produces a forward-peaked ALP flux kinematically similar to Primakoff production. Its contribution is subdominant to Primakoff production for most of the parameter space accessible to the beam dump experiments discussed here.

2.3.5 Beam Induced Backgrounds

The same hadronic cascade that generates the signal also produces the dominant backgrounds. Charged pions and kaons decay in flight via

$$\pi^\pm, K^\pm \rightarrow \mu^\pm + \nu_\mu^{(-)}$$

resulting in an intense flux of muons and neutrinos. Because the parent mesons carry large longitudinal momentum and modest transverse momentum, this flux is strongly forward-peaked and concentrated within a narrow cone along the beam axis.

2.4 Signal Kinematics

Figure 6 shows the solid-angle distributions of various FIPs produced in proton-target collisions at the SPS, computed for a molybdenum target. The top panels show the angular distributions, with the left panel showing FIPs from boosted production channels and the right showing FIPs from heavy-flavour meson decays, illustrating the qualitatively different angular distributions for the two cases. The vertical dashed line indicates the polar angle coverage of the SHiP experiment discussed in Section 4.3.




Figure 6: Kinematics of various FIPs produced in proton–target collisions on a molybdenum target at the SPS. **Top:** Angular distributions for different FIP masses, corresponding to different dominant production channels. The vertical dashed line represents the acceptance of the SHiP experiment discussed in section 4.3. **Bottom:** Energy spectra of the FIPs and their parent mesons (left). For heavy-flavour mesons (B , D), the spectrum is shown separately for $\theta < 0.05$ rad and $\theta > 0.05$ rad, showing the reduced hardness further from the beam axis (right). [41]

3 Particle Interactions and Detection Principles

Charged particles and photons traversing matter lose energy through various processes. The following sections cover the relevant interaction mechanisms, calorimetry principles, and the instrumentation technologies used in this thesis.

3.1 Particle Interactions with Matter

The nature of a particle's interaction with matter depends fundamentally on the particle type. Charged particles interact primarily via ionisation of the traversed medium, while photons undergo a distinct set of processes whose relative importance varies strongly with energy. Both classes of interactions are central to calorimetry and are treated in turn below. Other particles are not of interest for this thesis and are not discussed. If not stated otherwise, this section is based on the textbooks by Grupen [42] and Kolanoski/Wermes [43].

3.1.1 Ionisation

The dominant energy-loss mechanism for charged particles traversing matter is ionisation of the medium. The mean energy loss per unit path length is described by the Bethe–Bloch formula,

$$-\left\langle \frac{dE}{dx} \right\rangle = K z^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\frac{1}{2} \ln \frac{2m_e c^2 \beta^2 \gamma^2 T_{\max}}{I^2} - \beta^2 - \frac{\delta}{2} \right], \quad (8)$$

where $K = 4\pi N_A r_e^2 m_e c^2$ is a proportionality constant (N_A denotes Avogadro's number, r_e the classical electron radius, and m_e the electron mass), z is the charge number of the incident particle, Z and A are the atomic number and mass number of the absorbing material, β and γ are the standard Lorentz kinematic factors, $T_{\max}$ is the maximum kinetic energy transferable to an atomic electron in a single collision, I is the mean excitation energy of the medium, and $\delta(\beta)$ is a density-effect correction term.

Expressed as a function of $\beta\gamma = p/m$, Equation (8) is entirely invariant and exhibits an absolute minimum at $\beta\gamma \approx 3$, followed by a very slow relativistic rise, as illustrated in Figure 7. Particles with $\beta\gamma \gtrsim 3$ are referred to as *Minimum-Ionising Particles* (MIPs). Their energy loss of approximately $1.4 \text{ MeV cm}^2 \text{ g}^{-1}$ is constant for a wide range of materials and energies, which makes MIPs a useful reference for detector calibration.

The energy deposited by a single MIP in a finite absorber does not follow a Gaussian distribution but rather a Landau distribution, reflecting the stochastic nature of rare large-energy-transfer collisions.

3.1.2 Photon Interactions in Matter

Photons interact with matter through three distinct processes whose relative contributions depend strongly on the photon energy and the atomic number of the medium. These are the photoelectric effect, Compton scattering, and pair creation, and the energy and Z domains in which each dominates are illustrated in Figure 8.




Figure 7: The average energy loss per unit path length for positive muons in copper, given as a function of the product of the Lorentz variables. For muon momenta in the range from ~ 5 MeV to GeV, this energy loss is well described by the Bethe-Bloch formula. [44]




Figure 8: The energy domains in which photoelectric effect, Compton scattering and pair production are the most likely processes to occur, as a function of the Z value of the absorber material. [44].

Photoelectric Effect. At low energies, a photon can be absorbed by an atomic structure, ejecting a bound electron from its atomic shell. The kinetic energy of the emitted *photoelectron* (p.e.) is

$$E_{e^-} = h\nu - E_b, \quad (9)$$

where $h\nu$ is the photon energy and E_b is the electron binding energy. In high- Z materials, the photoelectric effect is the dominant process at low energies and is exploited directly in photon detectors such as *Photomultiplier Tubes* (PMTs) and *Silicon Photomultipliers* (SiPMs).

Compton Scattering. At intermediate energies, an incoming photon scatters inelastically off an atomic electron, transferring part of its energy to the recoiling Compton electron. The energy of the scattered photon is

$$E'_\gamma = \frac{E_\gamma}{1 + \varepsilon(1 - \cos \theta)}, \quad (10)$$

where θ is the scattering angle relative to the incident direction and $\varepsilon \equiv E_\gamma/(m_e c^2)$. The Compton scattering probability is proportional to the electron density of the medium. Because the ratio of atomic number to mass number (Z/A) remains approximately constant (~ 0.4 – 0.5) across most elements, the mass attenuation coefficient is largely material-independent. In the intermediate energy regime, the cross-section scales approximately as $1/E_\gamma$ [45].

Pair Creation and Bremsstrahlung. At high energies, electrons lose energy primarily through *bremsstrahlung* ($e^\pm \rightarrow e^\pm \gamma$), emitting photons in the Coulomb field of a nucleus. The mean bremsstrahlung energy loss per unit path length is given by

$$\frac{dE}{dx} = \frac{N_A \rho}{A} \int_0^{E_0} E_\gamma \frac{d\sigma}{dE_\gamma} dE_\gamma, \quad (11)$$

where A is the nuclear mass number, ρ the material density, and E_0 the initial electron energy, which sets the upper integration bound. This loss is parameterised through the *radiation length* X_0 , the material-dependent characteristic scale over which an electron loses all but a fraction $1/e \approx 37\%$ of its energy via bremsstrahlung:

$$\left. \frac{dE}{dx} \right|_{\text{rad}} = -\frac{E}{X_0} \quad \Rightarrow \quad E(x) = E_0 e^{-x/X_0}. \quad (12)$$

For photons, the analogous process is *pair creation* ($\gamma \rightarrow e^+ e^-$), which becomes accessible above the threshold of $2m_e$ and ultimately dominates at high energies. It also occurs in the Coulomb field of a nucleus, and its cross-section can be expressed in terms of X_0 as

$$\sigma_{\text{pair}} \approx \frac{7}{9} \cdot \frac{1}{X_0} \cdot \frac{A}{\rho N_A}, \quad (13)$$

where ρ is the material density, giving a mean free path for pair production of

$$\lambda_{\text{pair}} = \frac{A}{\sigma_{\text{pair}} N_A \rho} = \frac{9}{7} X_0. \quad (14)$$

The probability of pair conversion after traversing one radiation length is therefore $P(\gamma \rightarrow e^+e^-) = 1 - e^{-7/9} \approx 54\%$, somewhat below the corresponding energy loss through bremsstrahlung for electrons.

3.1.3 Hadronic Interactions

In addition to ionisation and bremsstrahlung, hadrons have access to a much wider range of strong-force interactions, making their energy-loss behaviour considerably more complex than that of leptons. Hadronic showers have limited direct relevance to the calorimeter developed in this thesis and are only briefly introduced for completeness.

By analogy with the electromagnetic case, hadronic shower development is characterised by a nuclear absorption length λ_a , defined via

$$N(x) = N_0 e^{-x/\lambda_a}, \quad \lambda_a = \frac{A}{N_A \rho \sigma_{\text{inelas}}}, \quad (15)$$

where N_0 is the number of incident hadrons, σ_{inelas} is the inelastic scattering cross-section, and x is the distance traversed. The radiation length scales as $X_0 \propto A/Z^2$, reflecting the dependence on nuclear charge, whereas the absorption length scales as $\lambda_a \propto A^{1/3}$, corresponding to the geometric nuclear radius. Since $\lambda_a \gg X_0$ for all materials of practical interest, considerably more absorber material is required to contain a hadronic shower rather than an electromagnetic one of comparable energy. This results in measurably different shower topologies for the two cases, a distinction that can be exploited for particle identification.

3.2 Principles of Calorimetry

Calorimetry relies on the initiation of a particle shower, which consists of a stochastic cascade of secondary particles produced through successive interactions in a dense medium. Calorimetric measurements are inherently destructive, with aggregate ionisation, scintillation, or Cherenkov light produced by the resulting shower being collected as the energy signal.

The physical development of these showers is governed by the material properties of the absorber. An *electromagnetic calorimeter* (ECAL) contains showers from electrons and photons, driven by bremsstrahlung and pair-production. The longitudinal scale of this development is defined by X_0 . In contrast, a *hadronic calorimeter* (HCAL) must contain the broader showers initiated by hadronic interactions, governed by the *nuclear absorption length* (λ_a). Since λ_a is typically much larger than X_0 , HCAL systems require significantly more volume for full containment.

By sampling the shower at various depths and lateral positions, the calorimeter may also enable *pointing* capabilities required for neutral decay reconstruction. This allows the incident particle's trajectory to be determined directly from the shower

axis, providing information about the shower direction in searches where the interaction vertex is unconstrained by upstream tracking. Unless stated otherwise, the formalisms in this section follow the textbook by Wigman [44].

3.2.1 Energy Resolution

The energy resolution of a calorimeter is traditionally its most important figure of merit. In the general case, it is described by three terms added in quadrature:

$$\frac{\sigma(E)}{E} = \frac{a}{\sqrt{E/\text{GeV}}} \oplus \frac{b}{E} \oplus c. \quad (16)$$

The three coefficients parameterise distinct physical contributions.

The stochastic term $a/\sqrt{E/\text{GeV}}$ originates from the Poisson statistics governing the number of independent shower particles contributing to the signal. This term represents the dominant resolution component at intermediate energies and is characteristically larger in sampling calorimeters due to sampling fluctuations inherent in incomplete energy collection. Typical values for the coefficient a range from 2–15% for electromagnetic systems and 35–120% for hadronic systems.

The noise term b/E is dominated by electronic noise fluctuations, independent of the incident particle energy. This component dominates the total resolution at low energies, where the signal-to-noise ratio is minimal. Careful optimisation of the front-end electronics and shielding is key to minimising b and extending the dynamic range of the detector toward the low-energy threshold.

The constant term c encapsulates energy-independent systematic effects, including mechanical tolerances, electronic non-uniformities, intercalibration errors, and fluctuations in longitudinal leakage. At high energies, this term dominates the resolution, and high-performance systems are typically designed to achieve $c \lesssim 1\%$.

Additionally, the measured signal must be proportional to the incident energy. In practice, linearity can be compromised by longitudinal shower leakage, which grows logarithmically with energy, or by electronic saturation. Calibration with test beams of known particle species and energy is essential to characterise and correct for such effects.

3.2.2 Electromagnetic Showers

Highly energetic electrons, positrons, and photons traversing a dense medium initiate electromagnetic showers through the repeated alternation of bremsstrahlung and pair creation. Showers induced by $e^\pm$ and γ are similar, with the latter exhibiting a slightly delayed development due to the photon conversion probability discussed in section 3.1.2. These processes occur within the nuclear Coulomb field with cross-sections scaling as $\sigma \propto Z^2$, necessitating high- Z absorbers (e.g., lead or tungsten)

for compact calorimeter designs. Due to the high boost of the primary particle, secondary products are strongly collimated with emission angles $\theta \propto m_e/E$. The longitudinal development of the cascade is well-approximated by the simplified Rossi-Heitler model [46], which operates under the following assumptions:

- Each particle undergoes exactly one interaction per radiation length.
- Each interaction (bremsstrahlung or pair production) splits the parent energy equally between two secondary particles.
- Energy loss via ionisation is neglected above the critical energy E_c , and is treated as the dominant, constant loss mechanism below E_c .

Under these assumptions, the number of particles N and their average energy E at a depth t (in units of X_0) evolve as:

$$N(t) = 2^t, \quad E(t) = \frac{E_0}{2^t}. \quad (17)$$

The shower continues to grow until the average particle energy reaches the critical energy ($E(t) \approx E_c$), at which point the multiplication ceases. This defines the shower maximum:

$$t_{\max} = \frac{\ln(E_0/E_c)}{\ln 2}, \quad N_{\max} \approx \frac{E_0}{E_c}. \quad (18)$$

Consequently, the total path length of all charged particles, $s_{\text{tot}} \approx (E_0/E_c)X_0$, and the particle count at maximum scale linearly with the primary energy E_0 .

Longitudinal Profile. The longitudinal energy distribution of an EM shower is well described by the empirical Longo formula,

$$\frac{dE}{dt} = E_0 \frac{b^a}{\Gamma(a)} t^{a-1} e^{-bt}, \quad (19)$$

where $t = x/X_0$ is the shower depth measured in radiation lengths, Γ is the gamma function, and the parameters a and b depend on the primary energy E_0 and the atomic number Z of the absorber. The shower maximum occurs at

$$t_{\max} = \frac{a-1}{b} \approx \ln \frac{E_0}{E_c} + \begin{cases} -0.5 & \text{(electrons),} \\ +0.5 & \text{(photons),} \end{cases} \quad (20)$$

confirming that $t_{\max}$ scales only logarithmically with E_0 through E_c , which is material dependent. The parameter $b \approx 0.5$ for heavy absorbers with minor energy dependence. [47] The longitudinal development of electromagnetic showers is intrinsically coupled to the atomic number Z of the absorber medium. As illustrated in Figure 9, several differences emerge in the longitudinal profiles across different materials:

- As Z increases, the shower maximum $t_{\max}$ shifts toward greater depths. This is primarily a consequence of the decreasing critical energy ε_c , which lowers the threshold at which ionisation losses begin to dominate over bremsstrahlung and pair production.

- As Z increases, the shower profiles exhibit a more gradual decay beyond the shower maximum. This slower attenuation reflects the increased contribution of low-energy photons and the modified absorption cross-sections in high- Z environments.




Figure 9: Energy deposit as a function of depth, for 10 GeV electron showers developing in aluminium, iron and lead, showing approximate scaling of the longitudinal shower profile, when expressed in units of radiation length X_0 . Results of EGS4 [48] calculations. [44].

Figure 10 demonstrates the characteristic evolution of electron showers across a range of incident energies. Figure 9 shows that standard analytical parameterisations are approximations which do not account for noticeable differences in the longitudinal shower shapes between materials.




Figure 10: The energy deposit as a function of depth, for 1, 10, 100 and 1,000 GeV electron showers developing in a block of copper. To compare the energy deposit profiles, the integrals of these curves have been normalised to the same value [44].

As a practical example, in lead the 98% containment depth for 10 GeV electrons is $18.8 X_0$, rising to $22.2 X_0$ at 100 GeV. Consequently, an ECAL depth of $20 X_0$ is considered standard in contemporary high-energy physics. Longitudinal segmentation of the readout can substantially improve electron–hadron separation and enable shower direction reconstruction by making the shower profile visible layer by layer.

Transverse Profile. At high energies, shower particles are initially emitted in a narrow cone around the beam axis. However, the lateral extent of the shower is eventually determined not by the primary high-energy interactions, but by the displacement of lower-energy components. This transverse spread is driven by two distinct physical mechanisms:

1. **Multiple Coulomb Scattering:** Electrons and positrons undergo successive small-angle deflections, gradually migrating away from the shower axis. This process is the primary driver of lateral spread during the early stages of shower development.
2. **Isotropic Secondary Processes:** Beyond the shower maximum, the lateral profile is dominated by photons and electrons produced in increasingly isotropic processes, such as Compton scattering and the photoelectric effect. Furthermore, bremsstrahlung photons emitted by electrons already travelling at significant angles relative to the axis contribute to a further broadening of the shower mantle.

The dominance of the latter processes is particularly pronounced in high- Z absorber media, where the lower critical energy allows for a more prolonged multiplication phase. The development of the lateral shower profile with increasing X_0 can be seen in Figure 11.

The characteristic transverse scale used to quantify this spread is the Molière radius, defined as:

$$R_M = X_0 \frac{E_s}{\varepsilon_c}, \quad (21)$$

where $E_s \approx 21.2 \text{ MeV}$ is the scale parameter from Molière’s multiple-scattering theory. Approximately 90% of the shower energy is contained within a cylinder of radius R_M , while 95% and 99% are contained within $2 R_M$ and $3.5 R_M$, respectively. Notably, the ratio $R_M/X_0 \propto 1/\varepsilon_c \propto Z$ increases with the atomic number. This implies that while showers in high- Z materials are absolutely more compact due to the shorter X_0 , they appear relatively broader when measured in units of radiation length. The concentrated shower core motivates a lateral segmentation at a scale of R_M or finer.

3.2.3 Hadronic Showers

Hadronic showers develop through inelastic nuclear collisions, governed by the nuclear absorption length $\lambda_a \gg X_0$, requiring substantially deeper calorimeters than their electromagnetic counterparts. A key complication is the electromagnetic sub-component of hadronic showers, which results from approximately one-third of secondary pions produced in each nuclear interaction being neutral, and their prompt decay $\pi^0 \rightarrow \gamma\gamma$ feeding compact EM sub-showers within the hadronic cascade. The




Figure 11: Longitudinal and lateral development of an electron shower (6 GeV) in lead. [42]

resulting fluctuation of the electromagnetic energy fraction from event to event is the dominant source of hadronic energy resolution degradation and introduces non-linearity in the calorimeter response. An additional fraction of the hadronic energy is permanently lost as invisible nuclear binding energy, producing no detectable signal. The distinct shower topologies of hadronic and electromagnetic showers, in particular the greater depth and broader lateral extent of hadronic showers, are exploited for *particle identification* (PID). The difference in shapes and depth of electromagnetic and hadronic showers can be seen in Figure 12. For lack of relevance in this thesis, hadronic calorimetry is not further elaborated.




Figure 12: Simulation of showers in an iron block. The picture compares an electron and a pion shower at the same energy, at 100 GeV. Only ionising particles are plotted. [43]

3.3 Calorimeter Technologies

The transformation of an electromagnetic or hadronic shower into a measurable electronic signal and further into physically meaningful, MIP-calibrated data requires a medium that can both induce particle multiplication and translate the resulting energy deposition into a detectable format. While the longitudinal and transverse profiles discussed in Section 3.2.2 are universal physical phenomena, the precision with which they are reconstructed depends entirely on the choice of calorimeter architecture and the underlying technology. If not stated otherwise, this section is based on [44].

Broadly, calorimetry is divided into two structural categories based on the distribution of the absorber and the active medium:

- **Homogeneous Calorimeters:** Where the entire detector volume consists of a single material that serves simultaneously as the absorber and the active medium.
- **Sampling Calorimeters:** Which use alternating layers of a dense passive absorber and a sensitive active material.

Both concepts are illustrated schematically in Figure 13. The experimental work presented in this thesis relies on the specific subset of sampling calorimetry.




Figure 13: Schematic of a homogenous (left) and a sampling calorimeter (right) with the active medium in blue, absorbers in grey, and areas with deposited energy in orange.

3.3.1 Homogeneous Calorimeters

In a homogeneous calorimeter, the entire shower develops within a single active medium, eliminating sampling fluctuations and providing the best achievable electromagnetic energy resolution. Common implementations use inorganic scintillating crystals, which achieve stochastic terms of $a \approx 2\text{--}5\%$ with constant terms below 1%. Their principal limitations are cost, restricted segmentation flexibility, and radiation hardness requirements at high luminosities. Homogeneous calorimeters are not the technology of choice for the detector developed in this thesis and are introduced here for completeness only.

3.3.2 Sampling Calorimeters

In a sampling calorimeter, dense passive absorber layers are interleaved with thinner active readout layers. The absorber initiates and contains the shower while the active medium provides the detectable signal. Only the fraction of the shower energy deposited in the active layers contributes to the signal. This *sampling fraction*,

$$f_s = \frac{\Delta E_{\text{active}}}{\Delta E_{\text{total}}}, \quad (22)$$

is typically of the order of a few per cent. The incomplete energy collection gives rise to *sampling fluctuations*, as the number of charged shower particles crossing an active layer fluctuates from event to event, contributing to the stochastic term of the energy resolution as

$$\left. \frac{\sigma(E)}{E} \right|_{\text{samp}} \propto \frac{1}{\sqrt{f_s}}. \quad (23)$$

Thinner absorber layers and larger sampling fractions improve the stochastic term at the cost of increased detector depth.

Compared to homogeneous designs, sampling calorimeters achieve worse intrinsic energy resolution but offer substantially reduced cost. A visualisation of two different sampling fractions is shown in Figure 14, which illustrates the effect of the sampling fraction on energy measurement and longitudinal shower position reconstruction.




Figure 14: Schematic of calorimeters with two different sampling fractions, with active medium in blue, absorbers in grey, and areas with deposited energy in orange. The black measuring points indicate the increased information available in the setup with a higher sampling fraction (right).

3.3.3 Particle Identification from Combined Detector Response

The combined response of different detector subsystems provides sufficient discriminating power for particle identification in most experimental contexts. Each particle species produces a characteristic pattern of activity across the tracking, calorimeter, and muon systems, and identification is possible by matching an observed pattern against these signatures.

Electrons are identified by a charged track in the tracking system associated with a compact, high-energy deposit in the electromagnetic calorimeter. The ratio of calorimeter energy to track momentum, $E/p \approx 1$, is a strong discriminant against

charged hadrons, which typically deposit only a small fraction of their energy in the ECAL and produce a hadronic shower further downstream. Photons produce an electromagnetic shower with no associated charged track. The absence of a track upstream of the ECAL cluster is therefore the primary photon identification criterion. Neutral hadrons similarly produce no track but deposit their energy in the hadronic calorimeter rather than the ECAL, providing separation from photons. Charged hadrons leave a track and a hadronic shower, with little or no signal in the ECAL. Muons are identified by a track that traverses the full detector with only MIP-like energy deposition in both calorimeters and a matching signal in the downstream muon system. The combination of several independent observables suppresses misidentification to a level sufficient for most analyses.




Figure 15: Schematic illustration of particle identification in a detector at a high energy accelerator. [43].

3.4 Instrumentation Technologies

The detector developed in this thesis employs a three-stage readout chain of plastic scintillator strip, *wavelength-shifting* (WLS) fibre, and *Silicon Photomultiplier* (SiPM) and is illustrated in Figure 16. The WLS fibre is the key element for signal uniformity: it collects scintillation photons along the full strip length and re-emits them at a wavelength for which the fibre attenuation length far exceeds the strip length. Each component is described below.

3.4.1 Plastic Scintillators

Plastic scintillators are organic scintillators based on aromatic polymer matrices, typically *polyvinyltoluene* (PVT) or *polystyrene* (PS), doped with one or more fluorescent wavelength-shifting agents. An ionising particle traversing the base material excites molecular states, and the subsequent de-excitation produces fluorescence photons on timescales of $\mathcal{O}(\text{ns})$. The base polymer is opaque to its own emission




Figure 16: Conceptual sketch of a readout chain of plastic scintillator, wavelength-shifting fibre and photosensors. [43]

wavelength, so a primary fluor is added whose absorption spectrum overlaps with the base emission and which re-emits at a longer wavelength with high quantum efficiency. A secondary fluor may further shift the emission to better match the spectral sensitivity of the photosensor. A schematic illustration of the *Stokes shift* between absorption and emission spectra of the scintillator is shown in Figure 17.




Figure 17: Absorption and emission spectrum of scintillators (schematic). [43]

The requirement for a large Stokes shift between absorption and emission peaks is essential to minimise self-absorption and achieve attenuation lengths of $\mathcal{O}(\text{m})$ in the bulk material. The light yield L_γ , defined through $\langle N_\gamma \rangle = L_\gamma \times E$, is an important figure of merit and is tunable by adjusting the dopant concentrations. Plastic scintillators combine good performance, low cost, and geometrical flexibility (available as bars, plates, tiles or fibres), making them a widely used active medium in large-scale sampling calorimeters. Their principal limitation is radiation hardness, as light yield degrades measurably above absorbed doses of $\mathcal{O}(\text{kGy})$.

3.4.2 Silicon Photomultipliers

SiPMs, also referred to as *multi-pixel photon counters* (MPPCs), are the standard photosensors for plastic scintillator-based calorimeters in modern experiments. A SiPM consists of a matrix of *avalanche photodiode* (APD) cells operated in Geiger mode, each comprising a *single-photon avalanche diode* (SPAD) and a quenching resistor R_{pix} , as shown in Figure 18.

Typical SiPMs contain pixel counts from hundreds to thousands, with individual pixel spacings of $15\text{--}75\,\mu\text{m}$ and total active areas of 1 to a few mm^2 .




Figure 18: Schematic of the array structure of a SiPM and the electric circuit representation. [49]

When a photon creates an electron–hole pair in a pixel, the Geiger avalanche is triggered, producing a charge pulse $Q_{\text{pix}} = C_{\text{pix}} \Delta V$, where $C_{\text{pix}} \sim 100$ fF is the pixel capacitance and $\Delta V = V_{\text{bias}} - V_{\text{BR}}$ is the overvoltage, typically 2–5 V with $V_{\text{BR}} \sim 50$ V. Each fired pixel produces a charge of approximately 100 fC, corresponding to one p.e. The intrinsic gain of $\mathcal{O}(10^6)$ renders the signal directly processable without external amplification. The current is quenched by the voltage drop across R_{pix} , which reduces the bias below V_{BR} and allows the pixel to recharge with a dead time $\tau = R_{\text{pix}} C_{\text{pix}} \sim \mathcal{O}(50 \text{ ns})$.

The total SiPM signal is the sum over all fired pixels. The *photon detection efficiency* (PDE) governs the probability that an incident photon produces a count,

$$\text{PDE} = Q \times f_A \times P_G, \quad (24)$$

where Q is the quantum efficiency of the silicon, f_A is the geometrical fill factor (fraction of sensitive area), and P_G is the probability of triggering a Geiger avalanche. Typical PDE values reach up to about 50% at the optimal wavelength. The dynamic range is set by the total pixel count: at high photon fluxes, multiple photons may strike the same pixel within its dead time, causing saturation and a non-linear response. The discrete, well-separated peaks in the single-photon spectrum (Figure 19) demonstrate the single-photon resolution of a SiPM.

The principal noise contributions are *dark counts*, thermally induced spontaneous avalanches occurring independently of incident light, *cross-talk*, in which photons generated during one avalanche trigger a neighbouring pixel, and *afterpulses*, arising from carriers trapped during an avalanche and released with a delay of $\mathcal{O}(100 \text{ ns})$. All three contributions scale with overvoltage and, in the case of dark counts, with temperature. SiPMs operate at low reverse-bias voltages and are insensitive to magnetic fields, and are available in compact form factors compatible with surface-mount assembly, making them well suited for the high-channel-count, fine-grained readout required in a pointing calorimeter.

3.4.3 Wavelength-Shifting Fibres

WLS fibres collect scintillation photons from the scintillator and transport them to the photosensor. A WLS fiber consists of a polymer core doped with a wavelength-shifting agent, surrounded by one or more cladding layers of lower refractive index, as shown in Figure 20. Typical WLS fiber diameters range from 0.5 to 2 mm.




Figure 19: Single photon spectrum of a SiPM showcasing the resolution of individual p.e. peaks. [50]

A photon entering the fiber is absorbed by the dopant and re-emitted isotropically at a longer wavelength. Photons emitted within the acceptance cone defined by the critical angle for total internal reflection,

$$\theta_{\text{crit}} = \arcsin \frac{n_{\text{clad}}}{n_{\text{core}}}, \quad (25)$$

are trapped and guided to the fibre end. The fraction of re-emitted photons trapped in this way is the *trapping efficiency*,

$$f_{\text{trap}} = \frac{1}{2} \left(1 - \frac{n_{\text{clad}}}{n_{\text{core}}} \right), \quad (26)$$

which is typically only a few per cent for single-clad fibres. Multi-cladding (adding a second outer layer with a yet lower refractive index) increases the effective acceptance cone and improves f_{trap} . The wavelength shift suppresses self-absorption within the fibre, extending the effective attenuation length to several metres.

3.4.4 Double Sided Readout

In double-sided readout configurations, SiPMs are coupled to both ends of the WLS fiber. This suppresses the residual position dependence of the signal, as the signal falloff towards one end is partially compensated by the larger signal at the other. The time difference of arrival between the two SiPM signals provides a rough longitudinal coordinate along the strip direction, otherwise inaccessible from the strip geometry alone. Requiring a coincident signal in both SiPMs at one scintillator strip also suppresses dark-count noise, improving the single-channel signal-to-noise ratio.




Figure 20: Illustration of the captured light fraction for a typical fibre with two cladding layers. [43]

4 Experimental Options at the CERN SPS

The production mechanisms described in Section 2.3 give the possible FIP flux a distinct kinematic signature, characterised by the decay length λ and the transverse displacement r_T at the detector plane. At a given production and decay cross section, these two observables directly shape the concrete experimental setups: the decay length determines the required shielding and fiducial baseline, while the transverse displacement defines the detector placement and acceptance. The following subsections develop these kinematic quantities and show how they motivate the two experimental concepts of SHADOWS and SHiP. This section is based on the SHADOWS Technical Proposal and EOI [3, 51] as well as several SHiP publications [4, 36, 52, 53].

4.1 Signal Acceptance and Detector Geometry

The number of signal events within the acceptance of a beam dump experiment depends on the signal kinematics discussed in Section 2.3 and a small set of geometric parameters: the longitudinal distance from the target to the start of the decay volume L_1 , the fiducial length $L_2 - L_1$ from L_1 to the end of the fiducial volume L_2 , and the detector acceptance ϵ_{det} derived from the transverse detector cross-section and the transverse displacement of the detector centre from the beam axis. In this discussion, the whole decay volume is considered to be the fiducial volume. The production channel determines the angular and momentum distributions of the FIP flux, which directly impacts which decay products fall within the geometric acceptance of a given detector layout.

In the long-lifetime regime $\lambda \gg L_2$, the geometric acceptance can be described as

$$\epsilon_{\text{geom}} \simeq \epsilon_{\text{FIP}} \times \frac{(L_2 - L_1)^{\text{eff}}}{L_2 - L_1} \times \epsilon_{\text{dec}}, \quad (27)$$

where ϵ_{FIP} is the fraction of produced FIPs pointing towards the detector cross-section, $(L_2 - L_1)^{\text{eff}}$ is the effective traversed length (see Figure 21), and ϵ_{dec} is the probability that all required decay products reach the detector acceptance. The daughter particle acceptance ϵ_{dec} introduces a non-trivial dependence on FIP mass and boost. If the opening angle of the decay products becomes comparable to the angular size of the detector as seen from the decay point, ϵ_{dec} decreases, effectively limiting the useful fiducial length for heavy or slowly moving FIPs (see also [41]). The signal yield is then given by the product of the angular distributions, the geometric acceptance and the probability of FIPs decaying within the fiducial volume

$$P_{\text{decay}} = \exp\left(-\frac{L_1^{\text{eff}}}{\lambda}\right) - \exp\left(-\frac{L_2^{\text{eff}}}{\lambda}\right), \quad (28)$$

where L_1^{eff} and L_2^{eff} are the projected longitudinal positions of the entry and exit points of the fiducial volume, and λ is the mean decay length of the hypothetical FIP.




Figure 21: Illustration of the different contributions to the geometric acceptance of beam dump experiments. The FIPs produced in the target have to point to the detector (red dashed arrow). The fraction of such events is given by ϵ_{FIP} . The effective length inside the decay volume passed by decaying FIPs (dashed blue line) may differ significantly from the nominal decay volume length L_{fid} . The decay products of FIPs (green arrows) also have to be in the detector acceptance (ϵ_{dec}). Figure adapted from [41].

4.2 SHADOWS: The Off-Axis Concept

SHADOWS (Search for Hidden And Dark Objects with the SPS) was a fixed-target experiment proposed for installation at the CERN North Area ECN3 experimental complex. It targets FIPs in the MeV-to-GeV mass range, with particular sensitivity to higher mass (1-5 GeV) candidates produced in the decays of charmed and beauty hadrons. SHADOWS was proposed to operate in parallel with the HIKE experiment [54], which was intended to occupy the on-axis position downstream of the target and focus on kaon physics. The two experiments were designed to exploit the same beam operation from complementary geometric positions: HIKE would capture the strongly forward-peaked flux on-axis, while SHADOWS would intercept the heavy-flavour halo that spreads laterally at large angles. A floorplan view of the SHADOWS layout in Fig. 22 shows the off-axis placement and the sequential arrangement of detector subsystems, which will be discussed in more detail in Section 4.4. The dimensions of SHADOWS are constrained to a $2.5 \times 2.5 \text{ m}^2$ cross-section and about 20 m in length space by availability.

Off-Axis Kinematic Motivation The defining experimental choice of SHADOWS is the off-axis placement of the decay volume about 10 m downstream of the target and displaced laterally by approximately 3 m from the beam axis. The length of the fiducial volume is 19 m, extending to about 30 m from the target. Displacing the detector off-axis has a negative effect on the signal yield, which decreases with the displacement for all FIP models [41]. The off-axis configuration has essentially no sensitivity to FIPs produced via proton bremsstrahlung or Primakoff production. The sensitivity to heavy-meson production is partially preserved, as can be seen in Figure 23 showing the illumination from $\pi^+\mu^-$ pairs of *heavy neutral lepton* (HNL)




Figure 22: Visualisation of the SHADOWS detector system at its off-axis position next to the beam line. HIKE would occupy the on-axis further down the beamline to the left. The beam comes from the right. Figure adapted from [3]

decays at the first SHADOWS tracking layer.

In exchange, the steep transverse fall-off of the beam-induced background fluxes provides a large geometric suppression factor before any active rejection is applied. Active downstream veto and sweeping systems are therefore much less critical in SHADOWS compared to an on-axis position.




Figure 23: Illumination at the first SHADOWS tracking layer (≈ 50 m from the dump) of the $\pi^+\mu^-$ pair from the decay of a HNL of mass 1.5 GeV from charm (left) and from beauty (right) hadron decays. [51]

4.3 SHiP: The On-Axis Experiment

SHiP (Search for Hidden Particles) is a fixed-target beam-dump experiment approved by the CERN management in March 2024, together with a BDF [5]. The experiment is designed to perform an extensive, largely background-free search for FIPs with masses in the MeV-to-GeV range, targeting the broadest possible set of hidden-sector candidates accessible at a proton beam dump. Compared to SHADOWS, SHiP focuses on a wider mass range and different production channels, specifically towards smaller masses. This also motivates the extent of the fiducial volume, which is 50 m long and extends to an aperture of 4×6 m at its end about 90 m from

the target. SHiP is foreseen to begin data taking in 2032, targeting an integrated exposure of 4×10^{19} PoT per year and 6×10^{20} PoT over fifteen years of operation.

On-Axis Geometry and Decay Volume Length For an on-axis detector, ϵ_{FIP} scales with the solid angle of the detector, while the dependence of the signal yield on L_1 is small (reducing L_1 from 38 m to 8 m increases the yield by less than a factor of two for most FIP models). The dependence on the fiducial length is non-trivial. A longer decay volume increases the decay probability but reduces the effective solid angle of the detector, which suppresses ϵ_{dec} for heavy FIPs. The acceptance reaches a maximum near $L_2 - L_1 \approx 50$ m for a 6×4 m detector aperture, directly motivating the SHiP decay volume length. For on-axis configurations, geometry optimisation is therefore driven by background suppression and cost rather than by signal acceptance. This is discussed in more detail in [41].

Beam and Target Setup SHiP will use the 400 GeV primary proton beam of the CERN SPS delivered to the BDF at up to 4×10^{13} protons per spill with a 1.2 s spill duration at nominal intensity. The production target is a composite structure of a molybdenum alloy and tungsten, maximising radiation hardness, charm and beauty production. A *Magnetised Hadron Absorber* follows immediately downstream, serving simultaneously as the final hadron absorber and as the first stage of the muon shield sequence, providing an initial magnetic sweeping field to the emerging muon flux.

The Muon Shield and the SND As discussed above, the same interactions that generate a possible signal also produce an overwhelming secondary flux of approximately 2×10^{10} muons and neutrinos per spill. The active muon shield is specifically designed to reduce the muon flux hitting the detectors by six orders of magnitude, using a sequence of alternate-polarity electromagnets to sweep muons of both charge signs laterally out of the detector acceptance. The SHiP *Scattering and Neutrino Detector* (SND), placed immediately behind the shield, exploits the intense neutrino flux for a dedicated programme of neutrino cross-section measurements.

A rendering of the full BDF and SHiP setup is shown in Fig. 24. The individual components are discussed in the next section.

4.4 Beam-Dump Experiment Blueprint

The descriptions of SHADOWS and SHiP reveal two very different experiments in placement and scale. SHADOWS is a comparably compact, off-axis detector relying on geometric displacement to suppress its dominant backgrounds. SHiP is a dedicated on-axis facility with a 50 m decay volume and an active muon shield. Yet the functional problem both experiments must solve is identical: suppress beam-induced backgrounds to below one event over the full physics run while maintaining acceptance for a long-lived neutral particle that decays into photons or charged tracks inside a fiducial volume downstream of the target. Every beam-dump experiment of this type must incorporate the same chain of functional elements, each motivated by a specific background class or signal-reconstruction requirement identified in the




Figure 24: Visualisation of the SHiP detector system at its on-axis position in the ECN3 cavern. [4]

preceding chapters. The following subsections trace this common architecture sub-system by subsystem, identifying in each case the physical driver and showing how the two experiments arrive at different solutions.

Muon Sweeping System After the passive absorber, a large flux of high-energy muons remains, which must be reduced as much as possible before reaching the fiducial decay volume. A large-aperture magnetic field deflects positive and negative muons in opposite transverse directions, clearing the geometric aperture of the downstream instrumentation. The required integrated bending power is set by the maximal muon momentum to be swept and by the transverse size of the decay volume. The total volume of active iron scales with the ambient muon flux, which differs by many orders of magnitude between the on-axis and off-axis configurations. In SHADOWS, three stages of passive *Magnetised Iron Blocks* (MIB) are placed off-axis immediately upstream of the decay vessel perform this function. Multiple stages push both charge populations further off-axis until neither falls within the aperture of the decay vessel, achieving a reduction of almost two orders of magnitude in the muon background rate at the first tracking station. Figure 25 shows the effect of the MIBs on the muon flux.

With approximately 2×10^{10} muons per spill exiting the target complex on-axis, SHiP faces a much more substantial muon flux, which must be suppressed by six orders of magnitude before reaching the decay volume. The active muon shield deployed is a highly optimised, more than 20m long, series of electromagnets of varying strength and polarity designed to deviate muons of a wide range of energies and direction. There are designs for a normal and a superconducting version. The result is a low-background central region flanked by two prominent lateral muon flames, as illustrated in Figure 26. These flares are also the reason for the unregular 2:3 aperture aspect ratio of the SHiP experiment³.

³The muon flares are diverted sideways for radiation protection reasons, resulting in the unintuitive upright detector positioning.




Figure 25: Simulated muon flux distribution at the first tracking station (red rectangle) of the SHADOWS detector system with (right) and without muon sweeping by the MIBs (left) viewed from downstream of the detector. [3]




Figure 26: Muon flux distribution after the SHiP active muon shield. (a) Muon flux map at the beginning of the SHiP decay vessel (rate per $5 \times 5 \text{ cm}^2$ area). [55]. (b) Expected prompt dose (dominated by muons) along an older design of the SHiP experiment (top view) for illustrative purposes. [56]

Fiducial Decay Volume The decay volume is an enclosed region providing the fiducial volume over which long-lived particles can decay in flight. Its length $L_2 - L_1$ is chosen to maximise the integral of Eq. (28) over the target coupling and mass range of interest, typically resulting in baselines of order tens of metres. The precise reconstruction of the decay vertex position within this volume is central to the pointing-based background suppression strategy.

To suppress neutrino-induced backgrounds, which scale with the material density traversed, the volume must be evacuated or filled with a medium of very low interaction probability. Both SHADOWS and SHiP originally planned a vacuum vessel for this purpose, but both designs converged on a helium fill as the preferred solution, since helium at atmospheric pressure reduces the secondary interaction cross-section relative to air by more than an order of magnitude while eliminating the structural demands of a large vacuum vessel, and therefore reducing secondary interactions from the vessel walls.

SHADOWS realises this as a modular rectangular vessel with a constant inner aper-

ture of $2.5 \times 2.5 \text{ m}^2$ and a longitudinal extent of approximately 19 m.

SHiP employs a decay volume shaped as a pyramidal frustum with a longitudinal extent of 50 m and transverse dimensions growing from $1.0 \times 2.7 \text{ m}^2$ at the upstream face to $4 \times 6 \text{ m}^2$ at the downstream face.

Veto Detectors Any charged particle entering the fiducial decay volume from outside, whether a residual muon that survives the sweeping system or a secondary particle produced in upstream material, must be tagged and the corresponding event rejected. A single missed muon may produce a reconstructed track pair that mimics a two-body feebly-interacting particle decay. The veto system completes the hermetic boundary of the fiducial region.

In SHADOWS, the front face and lateral walls of the decay vessel facing the beamline are instrumented to tag residual charged particles. SHiP similarly employs a two-component system of the *Upstream Background Tagger* based on plastic scintillators and the *Surrounding Background Tagger* enclosing the whole decay volume with liquid scintillator cells to detect secondaries from muon and neutrino interactions on the helium fill.

Spectrometer Tracker and Dipole Magnet For final states involving charged particles, a downstream magnetic spectrometer reconstructs individual track momenta and charge signs. In the beam-dump context, the spectrometer additionally provides a displaced-vertex constraint, rejecting tracks that do not converge to a common upstream point within the fiducial volume. Both SHADOWS and SHiP realise this function with a straw drift-tube tracker paired with a normal-conducting and a superconducting dipole magnet, respectively. Only the transverse apertures and bending strengths differ substantially, given the respective decay volume geometries and expected particle momenta.

Timing Detector In both experimental setups, high-precision timing detectors are placed between the tracker and the calorimeter section, which associate reconstructed calorimeter clusters and tracker hits with a common time stamp. This coincidence suppresses combinatorial backgrounds from uncorrelated particles from the same beam spill, which are otherwise kinematically indistinguishable from genuine two-body decays. Time resolutions have to be in the range 100–300 ps, set by the requirement that the two signal photons originate from the same decay vertex.

Calorimetry and Particle Identification The two experiments arrive at different calorimeter architectures driven by their respective background environments. SHiP faces a substantially higher flux of residual hadrons and muons, which poses a significant challenge for the PID. This is accounted for by a full HCAL behind the ECAL. SHADOWS benefits from the geometric suppression of the forward hadron flux so that an ECAL is sufficient for the experimental program, with residual hadronic contamination addressed by a muon identification system.

Common to both experiments is the requirement that the ECAL, in addition to energy measurement and E/p -based particle identification, must reconstruct the directions of individual photon showers to enable invariant mass reconstruction of neutral final states such as $a \rightarrow \gamma\gamma$. The momentum vectors of all reconstructed

signal particles have to extrapolate back to a common vertex inside the fiducial volumes. The following section discusses the specific requirements for the ECAL to be able to reconstruct $\text{ALP} \rightarrow \gamma\gamma$ decays, which is a key driver for the design of the calorimeter system in both experiments.

4.5 Calorimetry for Beam Dump Experiments

4.5.1 Diphoton Opening Angle

The two photons produced in $a \rightarrow \gamma\gamma$ are emitted back-to-back in the ALP rest frame. In the laboratory frame, the opening angle between them is

$$\theta_{\gamma\gamma} \approx \frac{2m_a}{E_a}, \quad (29)$$

valid in the highly boosted limit $E_a \gg m_a$. This angle decreases with increasing boost and sets the angular resolution required to resolve the two photons as separate showers. In SHiP and SHADOWS, the distance between the end of the decay volume and the front face of the ECAL is about 10 m, which sets the minimum distance for a valid decay vertex. Everything discussed regarding the reconstruction of an $a \rightarrow \gamma\gamma$ decay naturally also holds for any other $X \rightarrow \gamma\gamma$ decay.

For an ALP with $m_a = 0.5 \text{ GeV}$ and $E_a = 50 \text{ GeV}$, the opening angle is $\theta_{\gamma\gamma} \sim 20 \text{ mrad}$, corresponding to a transverse separation of approximately 20 cm at the calorimeter front face, while at higher energies or lower masses, this separation decreases rapidly. For $E_a = 50 \text{ GeV}$, a $m_a = 0.1 \text{ GeV}$ ALP produces $\theta_{\gamma\gamma} \sim 4 \text{ mrad}$ and a transverse separation of only 4 cm.

Photon shower separation is crucial for kinematic reconstruction. Two electromagnetic showers can only be resolved as distinct objects if their transverse separation in the calorimeter exceeds the intrinsic lateral shower size. Taking two Molière radii in lead ($2R_M^{\text{Pb}} \approx 3.2 \text{ cm}$) as a conservative resolution threshold, and a lever arm of $d_{\text{lever}} = 10 \text{ m}$ between the end of the decay vessel and the ECAL front face, the minimum reconstructable opening angle is

$$\theta_{\min} \approx \frac{2R_M^{\text{Pb}}}{d_{\text{lever}}} = \frac{3.2 \text{ cm}}{10 \text{ m}} = 3.2 \text{ mrad}. \quad (30)$$

Inverting Eq. (29), this defines a minimum reconstructable mass at a given forward boost E_a :

$$m_a^{\min} \approx \frac{\theta_{\min} E_a}{2}. \quad (31)$$

For a representative forward boost of $E_a = 70 \text{ GeV}$, this gives $m_a^{\min} \approx 0.11 \text{ GeV}$, confirming that the two-photon final state can be reconstructed down to 100 MeV for boost values typical for production mechanisms at the energy scale of the SPS⁴. This calculation assumes the ALP decay to happen right at the end of the decay volume, not considering the possibility of the ALP decaying earlier, leveraging the longer lever arm as illustrated in Figure 27. For all masses and boosts of phenomenological interest, the two showers can be resolved without issue.

⁴The π^0 mass is $\approx 135 \text{ MeV}$ making it possible to reconstruct $\pi^0 \rightarrow \gamma\gamma$ decays, possibly providing a valuable calibration method.




Figure 27: Effect of decay vertex position on the transverse separation of the two-photon showers at the ECAL face for $X \rightarrow \gamma\gamma$ decays. A particle decaying late in the decay volume (left) leaves only a short baseline between the decay vertex and the calorimeter face, reducing the transverse separation. A particle decaying early (right) benefits from the full length of the decay volume as a lever arm, maximising the separation. Representations are not to scale. Figure taken from [57]

In [57], a dedicated EventCalc [58] simulation of Primakoff-produced ALPs in the SHiP geometry confirms that the majority of signal events sit well above the conservative $2R_M$ threshold across the accessible parameter space. Figure 28 shows the photon separation at the ECAL front face for representative mass and lifetime combinations. The same conclusion holds for SHADOWS, where the lower boosts of heavier ALPs further increase the opening angle and therefore the photon separation at the calorimeter face.

4.5.2 Dominance of Angular Uncertainty

The ALP mass is reconstructed from the two-photon system via the relation

$$m_a^2 = 2 E_1 E_2 (1 - \cos \theta_{\gamma\gamma}), \quad (32)$$

where E_1 , E_2 are the photon energies and $\theta_{\gamma\gamma}$ is the opening angle between them, measured in the laboratory frame. In the highly boosted limit $\theta_{\gamma\gamma} \ll 1$, using $1 - \cos \theta \approx \frac{1}{2}\theta^2$, this reduces to

$$m_a \approx \sqrt{E_1 E_2} \theta_{\gamma\gamma}, \quad (33)$$

so that the reconstructed mass depends on three independently measured quantities: E_1 , E_2 , and $\theta_{\gamma\gamma}$. Propagating the measurement uncertainties in quadrature gives

$$\left(\frac{\sigma_{m_a}}{m_a}\right)^2 = \frac{1}{4} \left(\frac{\sigma_{E_1}}{E_1}\right)^2 + \frac{1}{4} \left(\frac{\sigma_{E_2}}{E_2}\right)^2 + \left(\frac{\sigma_{\theta_{\gamma\gamma}}}{\theta_{\gamma\gamma}}\right)^2. \quad (34)$$

The angular uncertainty $\sigma_{\theta_{\gamma\gamma}}$ is related to the single-photon angular resolution σ_θ by $\sigma_{\theta_{\gamma\gamma}} \approx \sqrt{2} \sigma_\theta$, assuming independent shower reconstruction. Substituting the boosted-limit expression Eq. (29) for the denominator yields

$$\frac{\sigma_{\theta_{\gamma\gamma}}}{\theta_{\gamma\gamma}} \approx \frac{\sqrt{2} \sigma_\theta}{2m_a/E_a} = \frac{E_a \sigma_\theta}{\sqrt{2} m_a}. \quad (35)$$




Figure 28: Distance between the two photons from $\text{ALP} \rightarrow \gamma\gamma$ decays at the SHiP ECAL front face, for representative mass and lifetime combinations generated via Primakoff production in the SHiP geometry. The yellow vertical line marks the conservative two-Molière-radii separation limit of Eq. (30). In all cases, the majority of signal events are well separated, confirming that photon-pair reconstruction is feasible across the accessible ALP parameter space. [57]

This shows that the relative angular contribution to the mass resolution scales as E_a/m_a , which grows rapidly with boost. The energy terms in Eq. (34) are bound by the sampling resolution $\sigma_E/E \sim \mathcal{O}(10\text{--}15\%/\sqrt{E/\text{GeV}})$. The angular term, therefore, dominates whenever

$$\frac{E_a \sigma_\theta}{\sqrt{2} m_a} \gg \frac{1}{2} \frac{\sigma_E}{E}, \quad (36)$$

which is satisfied for all ALP masses and boost values within the highly boosted regime at the SPS.

This conclusion is confirmed numerically by a systematic model-independent reconstruction study covering 14280 combinations of particle mass, boost, angular resolution, and energy resolution discussed in [57]. As illustrated in Figure 29, degrading the angular resolution by a factor of ten while keeping the energy resolution fixed smears the reconstructed mass peak into the background, making it inseparable, whereas doubling the energy resolution at fixed angular resolution leaves the peak essentially unchanged.

The mass resolution of the full system is consequently dominated by how well the calorimeter can reconstruct the direction of each photon shower. This makes pointing resolution the primary design requirement for the SHiP and SHADOWS ECALs.




Figure 29: Reconstructed invariant mass for $X \rightarrow \gamma\gamma$ decays with $m_X = 800 \text{ MeV}$ and $B = 20 \text{ GeV}$ forward boost under three detector configurations with different energy and angular resolutions. The red and black vertical lines mark the π^0 and η background masses, respectively. Decreasing the energy resolution (a)→(b) broadens the peak only modestly, while degrading angular resolution by a factor of ten (a)→(c) smears it into the background. [57]

4.5.3 Pointing Reconstruction and Calorimeter Segmentation

Resolving the directionality of individual showers requires both longitudinal and lateral segmentation of the calorimeter. By measuring the transverse shower barycentre $\bar{x}$ across at least two longitudinal layers at depths z_1 and z_2 , the incident angle θ is determined as:

$$\tan \theta = \frac{\bar{x}_2 - \bar{x}_1}{z_2 - z_1}. \quad (37)$$

The precision of this estimator is fundamentally limited by the transverse segmentation of each layer. As established in the discussion of the transverse profile (Section 3.2.2), the transversal resolution must be on the scale of the Molière radius or finer to accurately sample the shower core and suppress fluctuations in the shower tail.

Furthermore, sufficient longitudinal sampling is required to capture the development of both low-energy showers, which peak early, and high-energy showers, which penetrate deeper into the detector.

The angular precision of the estimator of Equation (37) follows directly from propagating its uncertainties. For two layers each with barycentre resolution σ_x , the pointing resolution σ_θ for small polar angles θ is given by:

$$\sigma_\theta \approx \frac{\sqrt{2} \sigma_x}{\Delta z}, \quad \Delta z \equiv z_2 - z_1. \quad (38)$$

The single-layer barycentre resolution σ_x is limited by the resolution of individual channels $p/\sqrt{12}$, set by the channel width p . The longitudinal lever arm Δz is the other factor to decrease σ_θ in direct proportion with a larger physical separation between the first and last sampling layer.

4.5.4 The Design Challenge

The calculation and examples in section 4.5.2 establish that pointing resolution is the dominant contribution to the ALP mass resolution in the highly boosted regime. The pointing resolution of a longitudinally segmented calorimeter is governed by the transverse channel pitch and the longitudinal layer separation Δz . A finer channel pitch increases the channel count and the associated readout cost linearly with the number of channels, while a larger Δz decreases the geometrical acceptance of the ECAL.

The energy resolution must remain within $\sigma_E/E \sim 10\text{--}15\%/\sqrt{E/\text{GeV}}$ across the signal energy range, but it does not need to be minimised beyond this threshold. At low energies, a high sampling fraction is advantageous to keep stochastic fluctuations small, while at high energies, a lower sampling fraction is acceptable and necessary to achieve sufficient longitudinal containment within the available depth without high cost.

The design of the calorimeter, therefore, reduces to finding a working point in the space of channel pitch, longitudinal spacing and sampling fraction that optimises the pointing resolution within given geometrical and financial boundary conditions while keeping the energy resolution within the acceptable band across the signal spectrum.

5 Studies towards the SHADOWS Electromagnetic Calorimeter

In this chapter, a systematic simulation study was carried out to identify the optimal calorimeter design for the SHADOWS ECAL. The resulting baseline design and its expected performance in terms of energy resolution, angular resolution, and invariant mass reconstruction are presented and evaluated against two benchmark ALP signal scenarios. The studies described in this chapter were conducted by the author as part of the preparation of the SHADOWS Technical Proposal [3]. Parts of this work have been published in [1], [2], and [3]. As the corresponding sections of the proposal were also written by the author, the descriptions in this chapter and in the proposal are similar in structure and wording. Figures reproduced or adapted from those publications are indicated by the respective citation in the caption.

5.1 Boundary Conditions and Performance Requirements

In the scope of SHADOWS, a sampling calorimeter energy resolution of $\sigma_E/E \sim 10\text{--}15\%/\sqrt{E/\text{GeV}}$ is sufficient for particle identification purposes. As established in Section 4.5.2, the invariant mass resolution of a strongly boosted diphoton decay is dominated by the angular uncertainty on the reconstructed photon directions. Consequently, the pointing resolution is the primary figure of merit, and all design choices described in this chapter are evaluated against this criterion.

Before the optimisation studies, a set of baseline assumptions is established to constrain the design space. The active area of the ECAL is, as the rest of the experiment, constrained to $2.5 \times 2.5 \text{ m}^2$ by the general space availability in the ECN3 hall. The total absorber depth is chosen to be $20 X_0$ ⁵ and longitudinally segmented into 40 sampling layers of about $0.5 X_0$ each. Each layer consists of 9 mm iron absorber plates ($R_m \approx 1.7 \text{ cm}$ [8]) interleaved with 1 cm plastic scintillator layers, starting with an absorber layer. This sampling ensures the required energy resolution as well as constant performance throughout the ECAL for late showering photons.

The active layers are based on plastic scintillator read out by SiPMs, a technology choice motivated by its proven performance, cost-effectiveness, and operational simplicity.

5.2 Reference Design: Tile Calorimeter

The natural starting point for a high-granularity sampling calorimeter with pointing capability is the SiPM-on-tile design, following the original approach of the CALICE AHCAL [63]. In such a design, square scintillator tiles are read out by an individual SiPM each, providing full three-dimensional position information. A picture showing the concept can be found in Figure 30.

⁵Electromagnetic calorimeters in high-energy physics experiments are typically designed with absorber depths around $20 - 30 X_0$ [59–62]. Increasing the depth with identical sampling directly improves the energy resolution by reducing longitudinal shower leakage. At $20 X_0$, shower containment is sufficient to achieve the target energy resolution for the photon energies relevant to SHADOWS.




Figure 30: Picture of the SiPM-on-tile concept showing individual $3 \times 3 \text{ cm}^2$ square scintillator tiles of the CALICE AHCAL read out by a single SiPM each. [64]

The CMS high-granularity endcap based on the same technology employs tiles between $2 \times 2 \text{ cm}^2$ and $6 \times 6 \text{ cm}^2$ depending on radial position [65]. This design achieves excellent spatial resolution and timing performance, but has not previously been evaluated against the specific requirement of shower direction reconstruction in an environment with a displaced vertex.

5.3 Optimisation Studies

Starting from the tile-based reference design, a systematic simulation study was performed to determine the optimal channel granularity and to evaluate whether a strip-based geometry can achieve comparable pointing performance at a drastically reduced channel count. The two key figures of merit are the angular residual as a function of energy and incident angle, and the total number of readout channels required to instrument the full $2.5 \times 2.5 \text{ m}^2$ active area.

5.3.1 Simulation Setup

The detector performance was evaluated using data generated by a GEANT4-based simulation [66] developed for the CALICE AHCAL [67]. To enable efficient exploration of the design space, the simulation is set up such that the absorber and scintillator volumes are implemented as continuous volumes in GEANT4, with detector segmentation being applied virtually during reconstruction by assigning energy deposits to channels according to their true spatial coordinates. This approach allows a single simulated dataset to be analysed under multiple segmentation geometries without repeating the full simulation, making systematic granularity scans computationally less demanding. This method captures only the shower-shape limitations of each geometry and does not account for detector effects such as SiPM saturation, or signal attenuation within the scintillator. The ECAL is represented in GEANT4 with alternating layers of 9 mm iron absorber, 10 mm scintillator (PS), 0.8 mm PCB (FR-4 fiberglass), and a 2 mm air gap. The conversion of deposited energy to p.e. is consistent with experimental results obtained by the CALICE collaboration [68]. Poisson statistics are applied to the estimated number of photoelectrons with no

further digitisation implemented. A rendering of the simulated ECAL geometry can be found in Figure 31.




Figure 31: Render of the SHADOWS ECAL baseline design simulation showing the outer dimensions and the longitudinal and lateral layer segmentation.

Clustering. The initial clustering of the showers is performed on a segmentation of the active layers into square channels with the side length equivalent to the channel pitch intended to study. The shower reconstruction is implemented by Antoine Laudrain and proceeds in three stages [3]:

1. **2D Clustering:** Within each layer, adjacent cells are grouped into planar clusters around seed cells that exceed a seed energy threshold. Cells above a lower secondary threshold adjacent to the growing cluster are iteratively added. The cluster is considered complete when no further cells above the secondary threshold remain adjacent. This procedure is executed independently for each seed cell in each layer.
2. **3D Clustering:** Planar clusters from consecutive layers are combined into 3D shower objects. The algorithm starts from the middle layer of the ECAL and identifies unprocessed planar clusters. For each such cluster, adjacent planar clusters in the neighbouring layers are located using position and distance criteria. If compatible adjacent clusters are found, they are merged into a new 3D cluster. The process continues iteratively until all planar clusters have been assigned.

3. **Directional Fit:** The direction of each 3D cluster is determined by computing the energy-weighted barycentre of each planar cluster and performing a linear fit.

In the case of strip-based geometries, only the coordinate along the strip orientation is used for the fit. This method is chosen for its computational efficiency. In the background, spatial information is available in each dimension for each layer, but depending on the chosen layout, only a fraction of this information is used for the direction reconstruction. This allows for fast evaluation of the pointing performance across a wide range of design parameters.

5.3.2 Granularity Optimisation

Systematic granularity scans are performed using electrons rather than photons. Electrons shower reliably in the first active layer, removing the stochastic variation of the photon conversion and thereby isolating the effect of channel granularity on the pointing and energy resolution. Since the absorber stack is uniform throughout the full depth and both transverse directions are sampled with equal density in every layer, a shower is measured with the same sampling regardless of where it starts. Consequently, the pointing performance for photons is expected to be identical to that of electrons at the same energy and incident angle, provided the shower is fully contained. The only mechanism by which photon showers can underperform electrons is longitudinal leakage at the rear of the calorimeter, which depends on energy and the depth at which the photon converts, but is not a function of the lateral granularity and cannot be mitigated by a change in detector design within the given boundary conditions. As the geometry is symmetric in x and y , angle scans are performed only in the x direction.

Tile sizes between $0.5 \times 0.5 \text{ cm}^2$ and $5 \times 5 \text{ cm}^2$ were tested across a range of energies and incident angles. At low incident angles and particle energies, a characteristic failure mode was identified. If the granularity is too coarse, the shower deposits energy in only a single column of tiles along its trajectory, providing no transverse lever arm from which the incident angle can be extracted. The reconstruction treats such showers as perfectly orthogonal to the detector layers regardless of their true direction, producing large systematic residuals in the angle reconstruction. This effect is illustrated schematically in Figure 32.

A benchmark energy of 2 GeV is chosen for this study, as it is the lowest energy still providing sufficient shower penetration to enable meaningful directional reconstruction given the sampling fraction. In Figure 33, it becomes visible that this failure mode becomes negligible at a granularity of 1 cm. Below this threshold, further reducing the tile size yields no measurable improvement in pointing resolution, as the limitation shifts from geometry to intrinsic shower fluctuations. A tile size of 1 cm^2 is therefore identified as the optimal granularity. This readout pitch simultaneously satisfies the requirement to resolve two-photon showers at the calorimeter front face (see Section 4.5.1).

As discussed in Section 4.5.3, the granularity requirements for shower direction reconstruction are directly linked to the Molière radius of the absorber material. The conclusions from this optimisation study are therefore applicable to any sampling




Figure 32: Visualisation of the effect the detector segmentation has on the directional reconstruction of showers at low energies and angles. The yellow crosses indicate the reconstructed hit positions showing the reconstruction bias at coarser segmentation. [1–3]

calorimeter design which uses an absorber material with a similar Molière radius, such as lead ($R_m \approx 1.6$ cm [8]).




Figure 33: Angular residual for a 2 GeV electron shower directional reconstruction at different angles and tile dimensions. [1–3]

5.3.3 Channel Count Optimisation

A tile calorimeter with 1 cm^2 tiles across the full $2.5 \times 2.5 \text{ m}^2$ active area and 40 layers results in

$$N_{\text{channels, tile}} = \left(\frac{250 \text{ cm}}{1 \text{ cm}} \right)^2 \times 40 = 2.5 \times 10^6 \quad (39)$$

readout channels, placing the design well outside the available instrumentation budget. To preserve the 1 cm spatial resolution while dramatically reducing channel count, an alternative concept with 1 cm wide scintillator strips oriented alternately in the x and y directions in successive layers was studied, reducing the total channel count by two orders of magnitude to:

$$N_{\text{channels, strip}} = \left(\frac{250 \text{ strips}}{\text{layer}} \times 40 \text{ layers} \right) \times 2 = 20,000 \text{ channels.} \quad (40)$$

By introducing the strip geometry, the spatial sampling in each transverse direction is reduced to every second layer. Studies showed that this only marginally affects the pointing reconstruction of high-energy showers, which develop over many layers, but introduces a performance degradation for low-energy showers penetrating only a few layers. Configurations combining tile layers with strip layers were also evaluated, but the additional cost and mechanical complexity of a hybrid design were found to be unjustified given the small performance gain, and a full-strip geometry was adopted as the baseline.

5.4 Baseline Design

For the proposal, the baseline ECAL design is a $20 X_0$ sampling calorimeter with an active area of $2.5 \times 2.5 \text{ m}^2$, consisting of 40 identical sampling layers. Each layer comprises a 9 mm iron absorber plate followed by an active layer of $1 \times 1 \times 250 \text{ cm}^3$ plastic scintillator strips, with strips in adjacent layers oriented alternately in the x and y directions. The first active layer is in x orientation.

5.5 Performance

The energy and angular resolution of the baseline design were evaluated using the described simulation framework. To account for the change from tiles to strips, the MIP conversion factor was adjusted to 50 p.e./MIP, summed over both SiPMs based on experimental results with similar strips. In Figure 34, the measurement of the energy resolution over the relevant energy range for SHADOWS signal decays and at angles of 0 and 30° can be found, yielding $\sigma_E/E \approx 11\%/\sqrt{E/\text{GeV}}$ for any angle, satisfying the design target of $10\text{--}15\%/\sqrt{E/\text{GeV}}$. The noise term ranges from $8\text{--}14\%/\sqrt{E/\text{GeV}}$, increasing with the incident angle, consistent with the reduced sampling fraction experienced by a shower traversing the layers at an angle.

The angular resolution is also characterised as a function of both energy and incident angle, with the pointing residual defined as the difference between the reconstructed and true electron direction. The results for both the x and y directions can be found in Figure 35.

5.6 Benchmark Scenarios

To evaluate the physics performance of the baseline ECAL design, two signal scenarios for the $a \rightarrow \gamma\gamma$ decay were investigated in the scope of the proposal by Torben Ferber. These represent distinct kinematic regimes for ALPs:

Scenario A (Low Mass): $m_a = 0.35 \text{ GeV}$ and $c\tau = 0.5 \text{ m}$.




Figure 34: Energy resolution of the SHADOWS ECAL baseline design for energies between 2 and 20 GeV at incident angles between 0° and 30°. [1–3]




Figure 35: Angular resolution of the SHADOWS ECAL baseline design for energies between 2 and 20 GeV at incident angles between 0° and 30°. In the x direction on the left (a) and the y direction on the right (b). [1–3]

Scenario B (High Mass): $m_a = 2.0$ GeV and $c\tau = 0.85$ m.

For each scenario, 10^4 events were generated using the ALPINIST software [69]. The generation was constrained such that the decay vertex is located within the fiducial decay volume, and both photons are required to have a minimum energy of $E_\gamma > 1$ GeV and fall within the ECAL acceptance. A full set of plots for each benchmark scenario is provided in Appendix A.3.

5.6.1 Invariant Mass Resolution

The resulting invariant mass resolutions for both scenarios are presented in Figure 36. The resolution for the $m_a = 2.0$ GeV case shows higher precision compared to the lighter candidate. This improvement is driven by the higher secondary photon energies, which result in superior angular reconstruction, and the larger opening

angles associated with higher masses, which reduce the relative error in the triangulation of the decay vertex.




Figure 36: Invariant mass resolution of the SHADOWS ECAL baseline design for the two benchmark scenarios. The left panel corresponds to Scenario A, and the right panel corresponds to Scenario B. [3]

5.6.2 Vertex Reconstruction and Background Rejection

The vertex resolution for the benchmark scenarios is shown in Figure 37. While not strictly required for the reconstruction of the event kinematics, the localisation of the decay vertex is a powerful tool for background suppression. Precise vertexing allows for the spatial vetoing of events originating from interactions in the detector material or beam-induced neutral backgrounds.




Figure 37: Vertex resolution of the SHADOWS ECAL baseline design for the two benchmark scenarios. The left panel corresponds to Scenario A, and the right panel corresponds to Scenario B. [3]

5.7 Conclusion

The simulation study presented in this chapter demonstrates that a longitudinally segmented sampling calorimeter based on $1 \times 1 \text{ cm}^2$ plastic scintillator strips, arranged in a $20 X_0$ iron–scintillator stack with 40 layers, achieves the energy and angular resolution required for the reconstruction of $a \rightarrow \gamma\gamma$ decays in the SHADOWS environment. The baseline design yields an energy resolution of $\sigma_E/E \approx 11\%/\sqrt{E/\text{GeV}}$ and sufficient angular resolution to reconstruct the invariant mass of both low- and high-mass ALP benchmark scenarios with adequate precision. The vertex resolution can further provide a handle for suppressing backgrounds.

6 The SHiP ECAL Design

The physics requirements for the SHiP ECAL are similar to those established for the SHADOWS ECAL in Section 5.1: the calorimeter must measure final state energies and reconstruct the directions of electromagnetic showers with sufficient precision to resolve $X \rightarrow \gamma\gamma$ decays and reconstruct the invariant mass. The pointing resolution remains the primary design driver, and the same $20 X_0$ total absorber depth with $40 \times 0.5 X_0$ sampling layers is used. However, the boundary conditions imposed by the SHiP experimental geometry are substantially different from the off-axis SHADOWS case and lead to several key differences in the detector concept.

6.1 Boundary Conditions and Design Drivers

Through its on-axis position and the 50 m long decay volume, the SHiP ECAL design is dominated by the following three factors:

1. **Large aperture:** The transverse area cleared by the muon shield combined with the long fiducial decay volume results in a required spectrometer aperture of $4 \times 6 \text{ m}^2$.
2. **Higher particle energies:** The on-axis position results in high boosts, producing final states with energies extending up to several hundred GeV and accordingly small opening angles.
3. **Higher muon and hadron flux:** The on-axis position exposes the SHiP ECAL to a significantly higher muon and hadron flux.

The size of the ECAL motivates the choice of lead as absorber material, as it is, despite the higher density, significantly lighter per X_0 than iron. This weight management is necessary due to the small footprint of the ECAL, which is close to the structural limits of the ECN3 hall in terms of area load. Secondly, the $4 \times 6 \text{ m}^2$ face needs to be subdivided into segments, which reduces the per-channel event rate, reconstruction ambiguities (see Figure 38), and makes the construction and installation at all feasible.

6.1.1 Longitudinal Gap

As discussed in Section 4.5.3, the pointing resolution improves with the longitudinal lever arm between the shower start and the downstream active layers. Introducing an air gap within the calorimeter stack increases this lever arm without adding absorber material or active channels.

The on-axis position of SHiP and the high boosts environment result in photon incident angles that can be very small. At these angles, the transverse shower axis shift between successive layers becomes comparable to or smaller than the strip pitch, making the channel granularity alone insufficient to reliably reconstruct the shower direction. An increased lever arm is therefore essential to resolve small incident angles, independently of any further reduction in strip pitch.

Unlike SHADOWS, where the available cross-section in the ECN3 hall constrains the active area to $2.5 \times 2.5 \text{ m}^2$ (see Section 4), the on-axis SHiP geometry imposes




Figure 38: Reconstruction ambiguities of the SHiP ECAL baseline design. The red crosses represent real hits, while the blue circles indicate reconstructed ambiguities introduced by the strip architecture. The segmentation into six independent modules helps to reduce these ambiguities. [57]

no equivalent hard constraint on the transverse detector size. The gap is placed after the first $6 X_0$, where photon conversion efficiency before the gap exceeds 99 %⁶. The $14 X_0$ downstream then sample the shower at increased distance from its start, directly enhancing the angular resolution for all showers traversing the gap, as illustrated in Figure 39. The effect of the gap on the angular resolution and detector acceptance is studied in detail in Chapter 10.




Figure 39: Illustration of the longitudinal gap in the SHiP ECAL, showing the increased lever arm for improved angular resolution.

6.1.2 Strip Width and Channel Count

A $4 \times 6 \text{ m}^2$ ECAL cannot be built as a monolithic detector. At strip lengths of 4–6 m, the attenuation of WLS fibers becomes significant, reducing the signal uniformity

⁶The systematic optimisation study in Chapter 10 will show that this initial gap placement is not optimal. Even though photon conversion efficiency is lower at smaller depths, an earlier gap placement yields better overall performance.

across the detector. Additionally, the resulting mechanical structure is impractical to handle and install, motivating a segmentation into smaller modules. A module size of $2 \times 2 \text{ m}^2$ is a good compromise between the need to limit the strip length and the desire to minimise the number of segments and associated dead space in between. Considering this module size, instrumenting all 40 layers with 1 cm strips would require

$$N_{\text{channels, 1 cm}} = 216 \times 2 \times 6 \times 40 \approx 100,000 \quad (41)$$

readout channels, where the factor of two accounts for double-sided readout. This channel count exceeds the foreseen instrumentation budget. To reduce the number of channels while maintaining the required performance, layers with wider scintillator strips are introduced between the 1 cm strip layers. The wider strips do not have to provide independent pointing information, but must allow separation of overlapping showers and correct assignment of energy deposits, especially for particle identification and shower pairs with small opening angles. At 6 cm width, the strips are approximately 4 Molière radii, providing sufficient granularity to resolve individual showers and localise energy deposits [57]. Going forward, the 1 cm strips will be referred to as *thin strips* and the 6 cm strips as *wide strips*.

The ASICs considered for the SHiP ECAL have 36 channels [70, 71], which leads to a $2.16 \times 2.16 \text{ m}^2$ module active area using the 6 cm strip pitch. Considering a design, with alternating 1 cm and 6 cm strip layers, requires

$$N_{\text{channels, mixed}} = (36 + 216) \times 2 \times 6 \times 20 \approx 60,000 \quad (42)$$

readout channels, reducing the channel count by 40 %. The alternating pattern of thin and wide strip layers ensures that every shower is sampled consistently. The reduction in sampling frequency from the thin strip layers is compensated for by the improved pointing resolution introduced by the gap, resulting in an overall improved performance. This will be discussed in detail in Chapter 10.

6.2 Baseline Design

The SHiP ECAL baseline design is a $20 X_0$ lead-scintillator sampling calorimeter consisting of 40 $0.5 X_0$ sampling layers. Each layer comprises a 3 mm lead absorber plate followed by a 1 cm thick active scintillator layer. To measure shower directions in both transverse coordinates, adjacent layers alternate between x - and y -orientation. As discussed above, the layer structure follows an alternating pattern of thin and wide strip pairs as shown in Figure 40. A longitudinal air gap of 50 cm is introduced after the first $6 X_0$.

The calorimeter face is divided into six independent segments with dimensions of $2.16 \times 2.16 \text{ m}^2$, bringing the full ECAL aperture up to $4.32 \times 6.48 \text{ m}^2$. Dead areas between the segments, which will accommodate readout electronics and mechanics in the final design, are indicated here but not accounted for. Their effect on the physics performance and acceptance is under study, but is outside the scope of this thesis. A to-scale rendering of the SHiP ECAL geometry is shown in Figure 41.




Figure 40: Render of the SHiP baseline layer structure showing the alternating pattern of thin strips in green with blue SiPMs and the wide strips in light brown with larger brown SiPMs. The lead absorber plates are shown in grey, with mechanical supports indicated in light grey and yellow.




Figure 41: To scale 3D-model of the SHiP ECAL baseline design, indicating the segmentation and the gaps between the modules.

7 Conceptual Prototype Testing at Test Beams

Following the approval of SHiP by CERN management, the pointing capabilities of the strip-based design require experimental validation. A conceptual prototype was therefore built and tested in a dedicated *test beam* (TB) campaign. The tested prototype represents a scaled-down but functionally complete implementation of the proposed SHiP ECAL design.

The CERN SPS TB [72] provides an ideal environment for studying the pointing performance of the ECAL and the particle identification capabilities of the whole calorimeter system, offering high-energy electron, muon, and hadron beams from 3 up to 288 GeV. Before the SPS TB, technical validation of the prototype's functionality was performed at lower energies at the DESY TB facility [73] and the CERN *Proton Synchrotron* (PS) TB [74].

This campaign provides the first experimental demonstration of calorimetric pointing in a scintillator strip-based sampling calorimeter under realistic SPS beam conditions and provides essential data for the simulation and optimisation of the SHiP ECAL.

7.1 The Electromagnetic Calorimeter Prototype

The ECAL prototype is constructed from alternating absorber and scintillator layers. The baseline design specifies 40 layers of $0.5 X_0$ each. In a balanced configuration of thin and wide strips, this corresponds to 10 layers of each type per orientation. The prototype closely mirrors the instrumentation of the SHiP ECAL baseline design in one orientation, with 9 thin strip layers and 8 wide strip layers placed between absorbers with a total depth of approximately $20 X_0$. The absorbers are 4 mm iron plates, chosen for availability and structural simplicity.

7.1.1 Thin Strip Layers

Each thin strip layer consists of 20 scintillator strips with dimensions of $1 \times 1 \times 20 \text{ cm}^3$. The strips are manufactured from coextruded *polystyrene* (PS) at Fermilab and are coated with *titanium dioxide* (TiO_2) for light containment [75]. A central hole houses a 1.2 mm diameter Kuraray YS-2 WLS fiber [76]. The coextruded strips are chosen for their cost-effectiveness and ease of handling. The WLS fibers are placed inside the center hole without glue or additional optical coupling, minimising the complexity of the assembly and allowing for easy replacement if necessary.

Readout is performed on one side using Hamamatsu S13360-1325PE SiPMs ($1.3 \times 1.3 \text{ mm}^2$ active area) [77]. Due to space constraints, the readout alternates between the two sides of adjacent strips. The opposite fiber ends are left open, so that half of the captured light is lost. This is a deliberate choice: in the final detector, double-sided readout will recover this light, and capping the open ends would artificially inflate the measured light yield (LY) on the readout side, biasing the characterisation of the achievable LY. This would bias the evaluation of the achievable LY with this combination of scintillator, WLS and SiPM and reduce the dynamic range of the SiPM in terms of MIP signals. The opposite polished fiber ends are glued into a 3D-printed mask and mounted to the SiPM PCB as a modular unit. The signals are routed from the SiPMs to SMA connectors on shielded PCBs, with ground



Figure 42: Picture of the $1 \times 1 \times 20 \text{ cm}^3$ coextruded scintillators from Fermilab under UV light with the TiO_2 coating on the outside and the fiber channel in the center.

provided via separate SMAs. Since the prototype strips are considerably shorter than those foreseen in the final detector, their open ends would introduce light loss not present in a full-length strip. To mitigate this, all strip ends are painted with TiO_2 varnish, which acts as a diffuse reflector and reduces the light loss at the boundaries, reproducing the performance of a longer strip more closely. The layout of such a layer is visualised in Figure 43.




Figure 43: CAD visualisation of a thin strip layer dimensions and readout. The 3D printed mask holding the WLS and coupling them to the SiPMs is depicted in transparent grey.

WLS to SiPM coupling A small air gap is maintained between the polished fiber end and the SiPM to maximise the illuminated area on the sensor. While round

fibers on square SiPMs leave unused corners, these corners contain pixels that increase the dynamic range of the SiPM if illuminated. Some light is lost in this setup, but this is considered a favourable compromise as long as the signal of one MIP remains sufficiently large. The coupling between WLS and SiPMs is covered in depth in the author's Master's thesis [78] and in [57].




Figure 44: Schematic of the illumination of the square sensitive area of a SiPM by a round light cone from a WLS fiber. [78]


7.1.2 Wide Strip Layers

Each of the eight wide strip layers consists of 6 scintillator strips made of Eljen EJ200 [79]. This *polyvinyltoluene* (PVT) based scintillator offers higher light yield than PS, but comes with a higher price tag and requires more careful handling. Each $6 \times 36 \times 1 \text{ cm}^3$ strip features two WLS fibers glued into parallel milled grooves, which are merged at both ends and coupled to $6 \times 6 \text{ mm}^2$ Hamamatsu S14160-6050HS SiPMs [80] using optical grease. The strips are wrapped in aluminised Mylar and mounted on $56 \times 46 \text{ cm}^2$ iron absorber plates. A more in depth characterization and optimisation of these strips can be found in [57].

The wide strip layers are adapted from an earlier prototype version with a $36 \times 36 \text{ cm}^2$ active area. Only the four central channels are read out, reducing the effective active area to $36 \times 24 \text{ cm}^2$, as the outer strips would extend beyond the active area of the thin strip layers. The layer orientation alternates between x and y . All 64 channels are read out by a single frontend module to accommodate the specific electronics configuration required by the larger SiPMs. Pictures of both a thin strip layer and a wide strip layer are shown in Figure 45.

7.1.3 The Prototype Assembly

The assembly is installed in a custom dark box designed for maximum flexibility. Internal slots with an 8 mm pitch (4.5 mm gaps and 3.5 mm spacers) ensure flexible and reproducible layer positioning. The box is mounted on a rotatable stage to be able to vary the beam incidence angles, and a lifting stage for vertical scans and adaptability to various beam heights. 3 m SMA cables [81] are routed through feed-throughs to the electronics below, allowing the electronics to remain outside the



(a) Thin strip layer.

(b) Wide strip layer.

Figure 45: Pictures of thin and wide strip scintillators mounted on their respective absorber plates.

dark box. The full setup is shown in Figure 46, while Figure 47a shows a close-up of the ECAL prototype layer stack.

To have equally spaced layers over the full $20 X_0$, the number of absorber layers varies by one between the thin and wide strip layers. The sampling fraction between each layer type stays constant. A schematic overview of the layer ordering in the prototype is shown in Figure 47b and Table 1.

During angular scans, the whole ECAL is rotated to change the beam incidence angle. To keep the shower center within the active area, the layers are shifted horizontally in increments of the strip granularity. This ensures that the alignment between layers is maintained and that the shower is fully contained within the active area for accurate reconstruction. A visualisation of the layer shifting can be seen in Figure 48.

7.2 The Hadronic Calorimeter Prototype

A HCAL prototype was planned and built for the SPS TB to validate the PID performance of the full ECAL+HCAL setup. It consists of five layers, each made of a 16 cm iron absorber, totalling 5λ . Combined with the $20 X_0$ ECAL, this provides the full depth of the SHiP calorimeter system proposed in [4]. The active elements





Figure 46: Picture of the open ECAL prototype box in lowered position at the SPS TB. In this picture, the beam would come from the right, and the HCAL is positioned to the left. The trigger scintillator (blue box) is positioned right of the ECAL centered on the beam.

are identical to the 6 cm wide ECAL strips, all oriented horizontally. The HCAL is positioned 119 cm behind the ECAL in a separate dark box on a lifting table for vertical muon beam calibration scans. The HCAL has no rotation or layer displacement capabilities. A picture of the assembled HCAL is shown in Figure 49.



(a) Picture of the open ECAL layer stack at the SPS TB showing the slot structure and layer ordering.



(b) Schematic layer ordering in the ECAL prototype. Blue layers are thin strip layers, green layers are wide strip layers, and grey blocks are absorber material.

Figure 47: (a) Picture of the open ECAL layer stack at the SPS TB. (b) Schematic overview of the layer ordering in the prototype. In both figures, the beam enters from the right.

Table 1: Layer positions in the ECAL prototype. The slot index refers to the 8 mm-pitch slot number in the dark box. The position is the physical distance from the first scintillator layer. The number of absorbers in front counts all 4 mm iron absorber plates upstream of the layer.

Layer	Slot	Position (mm)	Absorbers in front	Depth (X_0)
1	1	0	1	0.23
2	9	64	7	1.59
3	16	120	12	2.73
4	24	184	18	4.10
5	31	240	23	5.23
6	39	304	29	6.60
7	46	360	34	7.74
8	54	424	40	9.10
<i>17 cm longitudinal gap</i>				
9	80	632	45	10.24
10	88	696	51	11.60
11	95	752	56	12.74
12	103	816	62	14.11
13	110	872	67	15.25
14	118	936	73	16.61
15	125	992	78	17.75
16	133	1056	84	19.11
17	140	1112	89	20.25




Figure 48: Layer shifting in the ECAL prototype during angular scans. (a) When the ECAL is rotated to change the beam incidence angle, the layers are shifted horizontally to keep the shower center within the active area. (b) The shift in each layer is done in increments of the strip width to maintain the alignment between layers. Both figures are top views.



Figure 49: Picture of the HCAL prototype at the SPS TB. The ECAL and beam pipe are in the top right direction.

7.3 Readout

The *data acquisition* (DAQ) utilises the CAEN FERS system [82], comprising four DT5202 frontend modules [83] for the ECAL (three for the thin-strip layers and one for the wide-strip layers) and one for the HCAL. A DT5215 concentrator [84] is used to aggregate the data from the frontend modules and maintain the synchronisation between all frontend modules. Each module uses the Citiroc1A chip [85] (2 chips per module, 64 channels total) and allows parallel readout⁷. Features per channel include adjustable bias voltage, separate high/low gain paths, and channel masking. Most relevant for the readout and digitisation of the given detector system is the trigger and peak sensing logic. Upon trigger arrival, the *peak sensing* is started in all 64 channels, which records the maximum integrated signal occurring during the time of the *hold delay*. The *integration time* can be set between 12.5 and 87.5 ns according to individual resolution and dynamic range requirements. During the *hold phase*, each channel is read out on both high and low gain paths. A chronogram of this logic is shown in Figure 50. The length of this cycle defines the system dead time of about $1\ \mu\text{s}$ and is dominated by the hold time required to read out all channels. The most important operating settings are summarised in Table 2.

Each frontend module is mounted on a custom-designed adaptor board connecting the 64 pairs of anode and cathode pins on the front end module to one SMA plug each. Each frontend module is powered individually and is connected to a DT5215 concentrator board via an optical link, which facilitates the data transfer, timing synchronisation and command infrastructure. From the concentrator, a single Ethernet cable is linked to a computer running the CAEN Janus Software [87] to

⁷In *parallel* readout, all channels are sampled simultaneously upon a common external trigger and held during readout, rendering the entire detector blind for the duration of the readout cycle and resulting in high data throughput. In *streaming* readout, each channel independently sends data whenever its signal exceeds a programmable threshold. Only the hit channel and potentially its neighbours are temporarily blind, while the rest of the detector remains active, and the overall data rate is significantly lower. The final SHiP ECAL foresees streaming readout.




Figure 50: Peak detector chronogram of the Citiroc1A. [86]

Table 2: Key operating parameters of the DAQ system used during the SPS TB. The wide strips are read out with the minimal gain to account for their high light yield. The gain for the thin strips is optimised for dynamic range and resolution for MIP signals.

Parameter	Thin-strip layers	Wide-strip layers
HG preamplifier gain	55	1
LG preamplifier gain	45	1
Hold delay	150 ns	150 ns
LG integration time	50 ns	50 ns
HG integration time	50 ns	50 ns

control the system setup and data taking. A schematic of the DAQ setup is shown in Figure 51.

7.4 Trigger

The DAQ system features various self-triggering modes, which can be set to trigger on different majority levels or basic logic combinations of the channels. These modes were tested during the initial commissioning phase, but proved to be unreliable at high rates, leading to desynchronisation between frontend modules and data loss. Therefore, these modes were discarded, and a dedicated external trigger system was developed instead to ensure consistent and reliable readout across the entire detector. The DAQ architecture requires the trigger signal to arrive several nanoseconds before the analog detector signal to initiate the peak-sensing in time. The global trigger must be distributed through all modules to ensure all modules are read out simultaneously. The FERS modules accept NIM and LVTTTL standards for this purpose.

The trigger detector consists of a $6 \times 6 \times 1 \text{ cm}^3$ scintillator tile coupled to a Hamamatsu S14160-6050HS SiPM [80]. A custom comparator circuit converts the analog





Figure 51: Schematic of the DAQ setup used during the SPS TB, showing the full readout chain from SiPMs to the DAQ computer.

SiPM signal into an LVTTTL-compliant square pulse with a propagation delay of approximately 20 ns. The circuit is compact and can be mounted directly adjacent to the trigger tile, minimising cable-induced signal delay. Together with the 3 m cables connecting the detector SiPM boards to the frontend modules and delaying the analog detector signals accordingly, this ensures the trigger arrives sufficiently early to initiate the peak-sensing logic. Pictures of the trigger scintillator and the comparator PCB can be seen in Figure 52.



(a) Picture of the $6 \times 6 \times 1 \text{ cm}^3$ trigger scintillator. (b) Picture of the comparator PCB.

Figure 52: Hardware components of the fast trigger setup.

A setup where the signal is distributed via fanout through multiple comparator channels was tested, but only slight differences between the analog settings of the comparator circuits led to desynchronized data taking between the modules at high rates and was therefore discarded. To ensure identical trigger conditions and syn-

chronous signal processing across the system, a daisy chain approach is chosen, which distributes the exact same signal through all frontend modules. This is supported by the DT5202 modules by providing a bridged in/out connector, reducing delays to a minimum (see Appendix A.2 for details). The delay introduced by the daisy chain is consequently dominated by the cable length between modules. The proximity of the modules is only constrained by the adaptor boards, requiring the usage of 5 ns cables between modules, resulting in a total delay of the trigger between the first ECAL and HCAL module of 20 ns. Given the timely arrival of the trigger signal and the long hold delay, this delay does not affect the sensing of the full detector signal in the modules later in the chain.

This fast-timing architecture enables a serial trigger distribution which ensures that every module receives an identical trigger pulse, while allowing for the signal to be monitored via an oscilloscope after it has traversed the entire chain to verify synchronisation. This integrated trigger logic is essential for the high-rate environment of the SPS beam. A schematic of the most complete setup, including ECAL, HCAL and trigger, is shown in Figure 53.




Figure 53: Sketch of setup at the SPS TB. The beam comes from the left and first hits the ECAL, followed by the HCAL. The trigger scintillator is positioned left of the ECAL centered on the beam.

7.5 Test Beams

Before the final two-week TB at the SPS, the prototype was tested at energies up to 6 GeV in three TBs at DESY and the PS. These TBs were crucial for validating the functionality of the prototype for the SPS TB. An overview of the three TB facilities used in this campaign and their relevant beam parameters is given in Table 3.

Table 3: Overview of the TB facilities used in this campaign. The energy ranges given reflect the facility capabilities. [72–74]

Facility	Beam line	Particle types	Energy range
DESY	DESY TB	$e^\pm$	1–6 GeV
CERN PS	T10	$e^\pm, \pi^\pm, \mu^\pm$	1–6 GeV
CERN SPS	H8	$e^\pm, \pi^\pm, \mu^\pm$	10–400 GeV

First DESY TB: The first TB at DESY served as the first test of the active prototype components and their integration in a high-rate environment. Extensive tests have been performed with cosmics in the lab before, but the TB provided the first opportunity to test the DAQ at kHz rates. The main focus was on validating the functionality of the readout system, the trigger logic and the data acquisition software. Several issues with the commercial DAQ were identified during this TB and resolved in cooperation with the manufacturer afterwards.

PS TB: The PS TB was used to test the updated electronics firmware and take data with pions, muons and positrons to do the first PID studies with the ECAL. The mechanical setup was upgraded for this TB to allow rotation of the prototype for angular scans. During this TB synchronisation issues in the DAQ system were identified, which were resolved by introducing the external trigger.

Second DESY TB: The second TB at DESY was used to test the external trigger system and do a full test of the prototype functionality before the SPS TB. The trigger system was tested at high rates up to about 5 kHz and proved to be reliable and efficient. This TB provided the final validation of the prototype for the SPS campaign.

SPS TB: For the final SPS TB, the ECAL prototype was reworked to be able to accommodate the full $20 X_0$. For the tests at DESY and the PS, a maximum depth of $15 X_0$ was sufficient. Additionally, an HCAL prototype was added to validate the PID performance of the full calorimeter system. The setup described above is the final configuration for the SPS TB. The SPS TB was performed at the H8 beam line [72].

7.6 Physics Program at the SPS Test Beam

The H8 beam line at the CERN SPS North Area [72] delivers secondary and tertiary beams produced by 400 GeV primary protons impinging on a beryllium target. The beams can be tuned to get electrons, muons or hadrons at energies between 10 and 400 GeV. The beam is delivered in 4.8 s spills every 16.8 s, with typical particle rates of 10^4 – 10^6 per spill depending on energy and collimator settings. For energies up to 140 GeV a tertiary beam was used. For the measurements at 240 and 288 GeV, a secondary beam was used. Because of rate acceptance limitations experienced with the DAQ system, the beam was collimated to reduce the rate to a maximum of 5 kHz during spills.

The SPS TB was structured around the experimental demonstration of direction reconstruction in the ECAL prototype and the validation of the PID performance using the combined ECAL+HCAL system.

Pointing Studies. A central objective of the campaign was to characterise the angular resolution of the ECAL prototype as a function of beam energy and incident angle. Angular scans were performed at incidence angles between 0° and 7.5° , the latter being the practical limit imposed by the prototype dimensions. Measurements were taken at energies from 3 GeV to 288 GeV.

The SPS electron beams have between 6 and 55 % contamination with muons and pions [88, 89]. Based on the available information on beam composition, data-taking runs were planned to accumulate 100 000 or more effective electron events for each configuration.

Particle Identification Studies. In the second week of the TB, the HCAL prototype was added to the experimental setup downstream of the ECAL. Dedicated runs were taken with the pion, muon and electron beams at multiple energies to study the PID performance of the combined system. An overview of the run configurations for both physics programs is given in Table 4.

Table 4: Run configurations at the SPS TB. Pointing runs were taken at the listed combinations of beam energy and incidence angle. PID runs were taken at normal incidence with the combined ECAL+HCAL system.

Pointing Runs <i>(electrons)</i>		PID Runs <i>(electrons, pions, muons)</i>
Energy [GeV]	Angle [$^\circ$]	Energy [GeV]
3	0	15
5	1	20
10	2	30
20	3	50
50	4	100
100	5	140
140	7.5	180
240		240
288		

Muon Calibration. Dedicated muon beam runs were performed for calibration purposes. A full scan of both the ECAL and the HCAL was conducted with the muon beam to equalise the channel-by-channel response. The minimum-ionising signature of muons provides a well-defined reference signal in each scintillator layer and is used to derive inter-channel calibration constants, correcting, among others, for variations in SiPM gain and fiber coupling efficiency.

8 Analysis of Test Beam Data

This chapter discusses the analysis flow from raw detector data to the direction reconstruction of showers. The PID studies with the HCAL are not part of this thesis.

8.1 Data Preparation

The raw output of the DAQ is processed through event building, pedestal correction, intercalibration, and MIP calibration to convert hit-wise *analog-to-digital converter* (ADC) values into MIP units. A limited PID based on the ECAL data alone is then applied to select a pure electron sample for the pointing analysis.

8.1.1 Event Building

The DAQ outputs a binary file which contains batches of 64 entries, each containing the data of all channels in one frontend module. Within these bunches, the timestamp for all entries is identical. Even though the detector is set up to read out each board once a trigger arrives, there are incomplete events and individual data sets from which valid events have to be extracted. To do this, all data is sorted by timestamp and is then iterated through with a moving window, which is 4 (5 for the HCAL runs) module entries long (4×64 events). A set of entries is considered to be an event if the timestamps of all entries are within 10 ns and there is one set of entries from each of the modules in the setup. A schematic of the event building process is shown in Figure 54.




Figure 54: Schematic representation of the event building process.

8.1.2 Pedestal Correction

Each channel has a constant electronic pedestal offset, determined individually for both the HG and LG branches from dedicated pedestal runs or from the baseline noise in physics data. A Gaussian is fitted to the ADC distribution, with the peak position giving the pedestal value and the width serving as a measure of the electronic noise. Example fits for both branches are shown in Figure 55. The DAQ

is configured to centre the pedestal at 50 ADC counts to ensure the pedestal distribution is fully captured. The LG pedestal is naturally narrower than the HG pedestal due to the lower amplification factor. The extracted pedestal mean is then subtracted channel-wise from the raw ADC values.




Figure 55: Example pedestal fits for the HG and LG branches. Due to the lower amplification factor, the pedestal distribution of the LG branch is narrower than in the HG branch.

For some of the thin strip channels, the pedestal shows a double-peak structure. In these cases, the pedestal is determined by fitting a double Gaussian and using the amplitude-weighted mean of the two peaks as the pedestal value. The origin of the double peak structure is unknown, but the half-width of the double peak structure is about 75 ADC counts in the worst cases, which is small compared to the MIP signal in the thin strips (3150 MIPs on average) and therefore not considered to be an issue. An example fit for a channel with a double peak pedestal is shown in Figure 56.

8.1.3 Intercalibration

The Citiroc 1A ASIC reads out each channel simultaneously on two amplification branches with individually settable gains. Since the two branches are scaled differently, a per-channel intercalibration is needed to convert LG signals to HG units. Both branches have an ADC range of 0 to 8000. Figure 57 shows the correlation between the HG and LG ADC values for all wide strip channels, showing the clear linear dependence in the relevant range.

For each channel, the ratios between the HG and LG branches are plotted and fitted with a Gaussian. The range for the determination of the intercalibration factor is 7000 - 7900 HG ADC and LG ADC > 100. In this range, the LG is above the pedestal, while the HG is not yet saturated. Example plots for a thin-strip and a wide-strip channel are shown in Figure 58. The large difference in intercalibration factors between thin and wide strips results from the different gain settings for the two strip types.

The HG value is used up to 7500 ADC, above which the converted LG value is used. The resulting spectrum of a single channel can show a small artefact at the intercalibration point, where the LG branch takes over, as can be seen in Figure 59.




Figure 56: Double peak pedestal fit. The weighted mean of the two peaks is used as the pedestal value.




Figure 57: Heatmap showing the HG vs LG ADCs for all wide strip channels.

8.1.4 MIP Calibration

To convert the corrected ADC signals into physically meaningful units, a MIP calibration is performed for each channel individually using dedicated 160 GeV muon beam runs. At this energy, muons are MIPs throughout the full detector depth and deposit a well-defined signal in each scintillator layer. To cover the full active area, the beam spot was widened to a sigma of approximately 4 cm and 25 runs were taken in a 5×5 grid over the ECAL. For the HCAL, six runs were taken with the beam centred on each horizontal strip in turn. To select a pure muon sample, events are filtered to be track-like with the same criteria as discussed in the upcoming PID section (Section 8.1.6).

The resulting channel-wise ADC spectrum is fitted with a Landau–Gauss convolu-




Figure 58: Example intercalibration plots for thin and wide strip channels.




Figure 59: Intercalibrated channel spectrum. At the intercalibration point at 7500 HG ADC, an artefact can be observed. This plot shows a channel with a strong artefact for illustrative purposes.

tion to extract the *most-probable value* (MPV), which defines the MIP calibration constant. Fitted example spectra for a thin-strip and a wide-strip channel are shown in Figure 60. The different MIP levels are a result of the different strip types, light yields, SiPMs and gain settings. Figure 61 shows the resulting MIP calibration constants for all channels. The alternating pattern of slightly higher and lower MIP constants in some thin strip layers is due to the alternating readout of the channels with two different SiPM boards. Some channels show a particularly high or low MIP constant, which affects the dynamic range of the individual channels. There are four dead channels, which are not recoverable, one of which is in the thin strips relevant for this analysis.

The distribution of calibration constants showing the normal distribution for both types of strips is shown in Figure 62.




Figure 60: Example MIP calibration plots for thin and wide strip channels. The slight bump in the wide strip spectrum at exactly twice the MIP value originates from events in which two muons traverse the same wide strip simultaneously. Because the wide strips extend laterally beyond the active area of the thin strip layers, a second muon can pass through the outer region of a wide strip without producing a hit in the thin strips, and is therefore not rejected by the hit multiplicity cuts applied during muon selection.




Figure 61: MIP calibration overview showing the most probable values (MPVs) for all channels.




Figure 62: MIP calibration of thin (a) and wide strips (b) showing the distribution of the MPVs for both strip types.

8.1.5 SiPM Saturation Compensation

To compensate for SiPM saturation effects, the ADC signals need to be converted into p.e. At the very low end of the spectrum, the discrete nature of the p.e. signal can be observed. The gain of the SiPM is determined from the spacing between the individual peaks. In this analysis, the dark counts of the SiPMs are used, providing up to four, just well enough separated p.e. peaks to be fitted. This is only possible for a subset of channels which have the cleanest signals. For the remaining channels, the mean conversion factor is used. An example fit of such a spectrum can be seen in Figure 63. For the thin strips, the mean conversion factor is about 35 p.e./MIP.




Figure 63: Example SiPM calibration with the multi-Gaussian fit of the ADC spectrum (a) and the linear fit of the peak positions (b).

As discussed in Section 3.4.2, a SiPM consists of N_{tot} discrete microcells. When the number of incident photons gets closer to N_{tot} , multiple photons can hit the same microcell within a single integration window and the measured number of fired pixels N_{fired} underestimates the true number of photons N_γ . Assuming a Poisson distribution of photons over the microcells, the expected number of fired pixels is

$$N_{\text{fired}} = N_{\text{tot}} (1 - e^{-N_{\gamma}/N_{\text{tot}}}), \quad (43)$$

which can be inverted to recover the true photon count:

$$N_{\gamma} = -N_{\text{tot}} \ln \left(1 - \frac{N_{\text{fired}}}{N_{\text{tot}}} \right). \quad (44)$$

The correction factor $c_{\text{sat}} = N_{\gamma}/N_{\text{fired}} \geq 1$ grows with occupancy and diverges as N_{fired} nears N_{tot} . The SiPMs used for the readout of the thin strips have about 2400 pixels, which is equivalent to about 70 MIPs. However, nonlinearities already become relevant at much lower signals and must be compensated for [90].

8.1.6 Particle Identification

All SPS electron beams contain electrons, muons, and a hadron mixture dominated by pions. Depending on the beam energy setting, the electron purity varies between roughly 10 % and 90 % [91], making PID an essential preprocessing step prior to any physics analysis. The pointing studies have only the ECAL information available, which makes pion rejection the most challenging.

Each event is classified as electron, muon, or pion based on a set of criteria derived from the hit multiplicity per layer, the lateral shower spread, and the total deposited energy. If a classification cannot be made with high confidence, the event is marked as unclassified. The classification criteria are as follows:


Muons are identified by a track-like hit pattern: at least seven layers must be hit, with at least six of those containing exactly one hit each, and the hits must be confined to a single column of strips.

Electrons are identified by an early shower start: increased hit multiplicities and elevated energy deposits must be present already in the second and third thin strip layers. In addition, the total deposited energy enters the classification as a soft criterion, with increasing deviation from the beam-energy expectation penalising the electron hypothesis.


Pions are identified by a track-like front followed by a shower-like rear, with the lateral shower spread providing additional discrimination. However, hadrons that shower early can mimic the electron signature, and hadrons that traverse the full detector without showering can mimic muons. These two edge cases are the primary sources of contamination in the electron and muon samples, respectively.

In Figure 64, representative event displays for each particle type illustrate the characteristic differences in hit patterns in the ECAL.


The goal of this classification is to select a high-purity electron sample for the pointing analysis. Since no truth particle labels are available for the SPS data, the ECAL-only PID is validated against the combined ECAL+HCAL classification [88], which constitutes the most accurate available estimate, and the H8 beam composition provided by beam physicist Maarten Van Dijk [89]. The results are shown in



(a) Electron.



(b) Muon.



(c) Pion.

Figure 64: PID display for a sample of events showcasing the clear differences in the hit patterns for electrons, muons, and pions.

Figure 65. Each data point for the ECAL-only classification represents the mean over multiple runs taken at different incidence angles at the same beam energy, and the error bars reflect the spread across these runs. For the other two data sets only one point is available. The small error bars for the electron fraction confirm that the classification performance is stable across angles, while the muon fraction shows considerably larger spread. The conservative confidence threshold applied in the ECAL-only classification leads to a systematic underestimation of the electron fraction relative to the combined reference, while the difficulty of pion–muon separation with the ECAL alone accounts for the discrepancy in the muon fraction between the two calorimeter-based methods. The pronounced deviation of the H8 reference from both calorimeter-based results is attributed to the beam instrumentation and is understood [89]. There is a change from tertiary to secondary beam between 140 and 240 GeV to get higher energy beams. Overall, the good agreement between the ECAL-only and ECAL+HCAL results, which are considered the most accurate, validates the ECAL-only electron sample for the subsequent pointing analysis.




Figure 65: Comparison of particle identification results obtained from the ECAL-only classification, the combined ECAL+HCAL analysis [88], and the H8 reference provided by the H8 beam physicist Maarten Van Dijk [89]. The error bars on the ECAL-only points reflect the spread across runs taken at different incidence angles at the same beam energy. The systematic underestimation of the electron fraction in the ECAL-only result reflects the conservative confidence threshold applied for electron classification. The discrepancy in the muon fraction between the two calorimeter-based methods arises from the intrinsic difficulty of pion–muon separation with the ECAL alone. The pronounced deviation of the H8 reference from both calorimeter-based results is attributed to the beam instrumentation used for that measurement and is understood [89]. The combined ECAL+HCAL classification is considered the most accurate.

8.2 Shower Direction Reconstruction

As discussed in Chapter 7, the thin strip layers provide spatial granularity only in the horizontal dimension. The shower direction reconstruction is therefore restricted to the horizontal plane.

8.2.1 ADC Saturation

During the SPS TB, the signal integration window in the DAQ was accidentally set to 50 ns rather than the optimised 25 ns determined in the preparatory TBs. The longer integration window increases the collected charge per MIP, which causes the ADC to saturate at lower deposited energies per strip. The effective dynamic range is limited to approximately 60 MIPs per channel for thin-strip layers and 30 MIPs per channel for wide-strip layers. The central strips of an electromagnetic shower are therefore saturated at high beam energies. A channel is flagged as saturated when the LG ADC value exceeds 7990. The low saturation threshold in the wide strips makes them unusable for the analysis, which is therefore based exclusively on the thin strip layers. Figure 66 shows the layerwise event data for an example 288 GeV electron shower with saturation in the shower core. The pointing reconstruction method described in the following is designed to compensate for the missing core information by leveraging the shower tail hits and the longitudinal profile parametrisation.




Figure 66: Layerwise event data of the thin strip layers with saturated shower centers.

The general reconstruction strategy follows a global fit approach where a simultaneous fit to the longitudinal and transverse shower profiles across all active layers is performed, from which the shower axis slope θ and intercept x_0 are extracted. The longitudinal profile is described by the Longo–Sestili parametrisation (Eq. 19) while

the transverse profile is described by a two-component parametrisation.

8.2.2 Transverse Profile Model

The transverse shower profile is parametrised by a two-component function combining a Gaussian core with a Lorentzian tail that shares a common width parameter:

$$f(x; x_c, R, F_\ell) = \frac{1}{\mathcal{N}(R, F_\ell)} \left[(1 - F_\ell) \exp\left(-\frac{(x - x_c)^2}{2 R^2}\right) + F_\ell \frac{R^2}{(x - x_c)^2 + R^2} \right], \quad (45)$$

where x_c is the shower centre, R is the common width parameter, and F_ℓ is the per-layer Lorentzian tail fraction. The normalisation constant is

$$\mathcal{N}(R, F) = (1 - F) \sqrt{2\pi} R + F \pi R. \quad (46)$$

The Gaussian component describes the compact core, while the Lorentzian term captures the shower tails. This parametrisation is motivated by the shared width parameter of the core and tail components, which allows for an extrapolation of the shower core shape from the tail hits. Additional anchoring is provided by a per-layer energy sum constraint, which compares the predicted total energy deposit in each layer to the measured layer sum with a one-sided treatment for layers containing saturated channels.

8.2.3 Cost Function

The cost function extends the standard χ^2 sum with additional terms that handle saturated channels and stabilise the fit:

$$C = \sum_{i \in \text{unsaturated}} \left(\frac{E_{\text{meas},i} - E_{\text{pred},i}}{\sigma_i} \right)^2 + C_{\text{sat}} + C_{\text{layer}} + C_{\text{reg}}, \quad (47)$$

with the per-hit uncertainty

$$\sigma_i = \sqrt{1 + S_i (1 + (c_{\text{sat},i} - 1) \epsilon_{\text{sat}})} \text{ MIP}. \quad (48)$$

The constant term of 1 MIP represents the statistical floor, while the stochastic term scales with the signal size S_i and is further inflated by the saturation correction factor $c_{\text{sat},i} > 1$, which increases the uncertainty for high-signal hits where saturation compensation becomes a significant factor. The parameter $\epsilon_{\text{sat}} = 10\%$ controls the strength of this inflation and is equivalent to a 10% systematic uncertainty on the saturation correction.

For saturated hits, the measured ADC value is a lower bound on the true energy. Saturated channels are excluded from the standard χ^2 sum and are instead handled exclusively by C_{sat} , which introduces a one-sided penalty that is only applied when the model prediction falls below the saturated measurement. This allows the fit to predict more energy than recorded for saturated strips without penalty. The per-layer energy term C_{layer} compares the predicted total energy deposit in each layer to the measured layer sum. For layers containing saturated channels, the same one-sided treatment as for C_{sat} is applied, so the fit is not penalised for predicting

more energy than the saturated layer total. The regularisation term C_{reg} includes a quadratic penalty for unphysical longitudinal shape parameters α and β and a Laplacian smoothness penalty on adjacent-layer radii R_ℓ to suppress unphysical jumps between layers.

The minimiser uses MINUIT via `iminuit` [92]. The free parameters of the fit are the five global parameters and two per-layer parameters:

- E_{tot} : the total shower energy, which sets the overall normalisation of the longitudinal profile;
- α, β : the shape parameters of the longitudinal profile;
- x_0 : the transverse intercept of the shower axis at the front face of the detector;
- θ : the slope of the shower axis (the reconstructed pointing angle);
- R_ℓ : the per-layer common width parameter of the transverse profile;
- F_ℓ : the per-layer Lorentzian tail fraction.

8.2.4 Minimisation

The minimisation proceeds in two phases to manage the high-dimensional parameter space ($5 + 2N_\ell$ free parameters):

Preprocessing. Layers behind layer 4 with a total energy deposit below 5 MIPs are removed from the fit data, as they contain the diffuse longitudinal tail of the shower. The requirement that this is only done at later layers is to avoid the removal of relevant data from the shower start or short, low-energy showers. The initial shower axis (x_0, θ) is seeded from a linear fit to the per-layer energy-weighted centroids.

Phase 1: All per-layer F_ℓ and R_ℓ are temporarily fixed at their initial values and only the five global parameters ($E_{\text{tot}}, \alpha, \beta, x_0, \theta$) are optimised. This provides a well-converged starting point for the axis parameters before the full high-dimensional space is opened.

Phase 2: The full parameter space is explored, with F_ℓ and R_ℓ free to vary. The Phase-1 result is used as the starting point for the global parameters.

Figure 68 shows the layerwise profile fit for a saturated example 140 GeV electron shower alongside a 10 GeV unsaturated shower for comparison. Starting from layer three, the ADC saturation is clipping the shower cores of the saturated event. With the information from the tail hits and the longitudinal profile anchoring, illustrated by the side-by-side comparison in Figure 67, the fit can compensate for the missing core information and reconstruct a physically reasonable shower profile.




Figure 67: Event display comparison of the global shower fit for a 10 GeV unsaturated electron shower (right) and a 140 GeV saturated electron shower (left).

8.3 Discussion of Results

A primary goal of the test beam campaign was to experimentally demonstrate that the scintillator strip geometry of the SHiP ECAL could reconstruct the direction of electromagnetic showers. The key observables extracted from the global shower fit are the shower axis slope θ and the transverse intercept x_0 at the detector front face.

An energy resolution measurement is not a primary objective of this analysis. The combination of ADC saturation and the coarse transverse sampling fraction of the prototype renders any extracted energy resolution unrepresentative of the true calorimeter performance and is therefore not reported here.

8.3.1 Intercept Reconstruction

Throughout this and the following chapter, the 0° data sets at 3 GeV, 20 GeV, 100 GeV, and 240 GeV are used as representative examples. Figure 69 shows the distribution of the reconstructed intercept x_0 for these four energies. The intercept distributions show a characteristic quantisation pattern visible as discrete preferred values of x_0 . This arises from the discrete strip structure where the fit tends to place the shower axis near a strip centre or strip edge. The current level of optimisation of the fit does not suppress this discretisation artefact. The widths of the distributions reflect the beam profile at the SPS, which was asked to be set up to maximise the parallness of the beam rather than to provide a narrow spot. The truth information about neither beam divergence nor beam spot size is available. The profiles for the angled data look similar.

Two further features are visible in the distributions. First, a constant offset of the distribution centre relative to zero is present across all energies, indicating a small misalignment of the detector with respect to the beam axis. Second, the apparent beam width decreases with increasing energy. Both effects cannot be disentangled





Figure 68: Comparison of layerwise profile fits for a 10 GeV unsaturated electron shower (top) and a 140 GeV saturated electron shower (bottom). The layers tinted blue are before the shower maximum, the yellow ones after.




Figure 69: Reconstructed intercept x_0 of 0° SPS TB data at 3 GeV, 20 GeV, 100 GeV, and 240 GeV. The grid lines indicate the strip edges.

from the true beam properties without knowledge of the beam divergence and spot size, so no quantitative conclusions can be made.

Also notable is the switch in the quantisation of the distribution from the edge between two strips (grid lines) at 3 GeV to the center of the strip at the higher energies. The quantisation is generally caused by the fit struggling to properly handle the shower start, where the shower is significantly narrower than the strip pitch and only one or two strips have non-zero signals, causing minimisation artefacts which differ between fully developed showers and short and slim low energy showers.

8.3.2 Angular Resolution

Figure 70 shows the distribution of the reconstructed shower axis slope θ for the example data sets. The distributions have a Gaussian core, but especially at lower energies, there are non-Gaussian tails which are not well captured by a Gaussian fit. To account for this, both the Gaussian sigma and the statistical standard deviation are extracted from each distribution and discussed in the following.

Figure 71 summarises the angular resolution as a function of beam energy for all data sets taken during the SPS TB. Both the fitted and statistical resolutions are shown. It can be seen that the resolution continuously improves with higher energies




Figure 70: Reconstructed shower axis slope θ of 0° SPS TB data at 3 GeV, 20 GeV, 100 GeV, and 240 GeV.




Figure 71: Angular resolution σ_θ as a function of beam energy for all SPS TB data sets. Both the Gaussian fit width (filled markers) and the statistical standard deviation (open markers) are shown.

and has a minor dependence on the incident angle, as already indicated by the SHADOWS simulation discussed in Section 5.5. The statistical sigma is consistently larger, caused by non-Gaussian tails and outliers in the angle distributions (see Figure 72), which are likely caused by beam divergence and the presence of non-electron contamination in the sample that the fit can not handle properly.




Figure 72: Reconstructed shower axis slope θ in log-scale with non-Gaussian tails for the 100 GeV data set. The legend shows the divergence between the fitted Gaussian sigma and the statistical standard deviation caused by the non-Gaussian tails.

9 Simulation of Test Beam Setup

The TB results presented in the last chapter establish the pointing performance of the prototype under realistic conditions. A validated simulation is required to extrapolate from the prototype measurements to the full detector geometry and to serve as the basis for the design optimisation in Chapter 10. The detector simulation is built on top of GEANT4 [66], which handles all physics processes. Geometry and material definitions are managed through the DD4HEP framework [93]. The general simulation framework for the SHiP calorimeters was implemented by Matei Climescu [94]. The simulations discussed in the following sections are derived from this framework and have been adapted to reproduce the exact prototype geometry of the SPS TB described in Chapter 7. As in the SPS TB, the HCAL is simulated, but not discussed in the scope of this thesis. A render of the full simulated setup is shown in Figure 73.




Figure 73: Visualisation of the simulated setup with ECAL and HCAL recreated as close as possible to the actual TB setup. In the front (left) there is a 1 mm aluminum plate representing the dark box cover.

With the simulation, it is possible to recreate every configuration the physical prototype was tested with. Since the simulation for high electron energies results in long runtimes and large data output, the number of events per run is limited to 10 000 events for each angle, energy and particle combination tested.

9.1 Particle Gun

The particle gun is placed 10 m in front of the detector. The x-position and emission direction of the particles are adapted according to the desired incident angle, with the scintillating layers being shifted in the same way as in the real setup. Figure 74 shows the simulated ECAL setup from the top, where the two layer types are shifted into position on the angled beam. The vertex of the particles has a Gaussian spread of $\sigma = 5$ mm applied in three dimensions to avoid edge effects. No smearing is applied to the particle direction and energy.

9.2 Digitisation

No real digitisation has been added at the level of the simulation. The raw output is the true energy deposited in each scintillator strip, which is converted to the MIP scale using a simulation-internal MIP calibration performed with muons. A




Figure 74: Visualisation of the simulated ECAL setup from the top. The relative shifts between the layers to reproduce the angled beam configurations are clearly visible. In the front (left) there is a 1 mm aluminum plate representing the dark box cover.

mock digitisation is performed based on the experimental results from the physical prototype, covering the following effects:

Noise: The pedestal information from the physical prototype is used to add corresponding noise to the simulation.

SiPM Saturation As described in Section 8.1.5, SiPM pixel saturation affects the physical prototype and is corrected for using Eq. 44. This procedure is reverse-engineered in the simulation, where the MIP signal S is first converted to the number of fired pixels using

$$N_{\text{fired}} = N_{\text{tot}} \left(1 - e^{-S \cdot \gamma_{\text{PE}} / N_{\text{tot}}} \right), \quad (49)$$

where $\gamma_{\text{PE}} = 35 \text{ p.e./MIP}$ is the experimentally determined p.e. yield for thin strips and N_{tot} is the total number of SiPM pixels. To account for the uncertainty of the compensation, a symmetric Gaussian smearing

$$\sigma_{\text{comp}} = \epsilon_{\text{sat}} (c_{\text{sat}} - 1) S, \quad \epsilon_{\text{sat}} = 0.1, \quad (50)$$

is added to the signal, where $c_{\text{sat}} = N_{\gamma} / N_{\text{fired}} \geq 1$ is the correction factor. This term is zero for unsaturated channels and grows with increasing occupancy, reflecting the larger uncertainty of the compensation at high signal levels.

ADC Saturation: The ADC saturation is the most dominant effect in the TB data and is also recreated in the simulation. The experimental values from Section 8.1.4 of 3150 ADC/MIP (thin strips) and 2710 ADC/MIP (wide strips) were used to introduce the ADC overflow. Applying the intercalibration factors, this leads to a saturation level of 51 MIPs in the thin strips and 29 MIPs in the wide strips. However, in the analysis of the actual TB data, it became visible that only the center channels of the thin strips in the front layers

are actually saturating at high energies. Between those channels, a common saturation level of 60 ± 2 MIPs was observed, which is why 60 MIPs was chosen to represent the actual thin strip behaviour better than the statistical average over the whole detector. For the wide strips, the ADC saturation is a much more common phenomenon across all channels, which is why the statistical average of 29 MIPs was considered a reliable estimate and was applied directly. The wide strips are fully modelled and read out in the simulation, but as in the TB data, they are not used for shower reconstruction and only enter the PID.

This simplified architecture provides a flexible and transparent environment for studying the prototype ECAL, focusing on geometry and physics effects without simulating the full readout chain. Before running the shower fit on the simulated data, the same PID is performed as in the TB data to remove electrons which would not be identified as such in the TB. Due to the strict requirements of the PID, this removes a fraction of up to 10% at low energies.

9.3 Simulation Validation

To validate the simulation, several shower observables are compared between the simulated events and the TB data. For the comparison, the same runs are used as in the last chapter (0°, 3 GeV, 20 GeV, 100 GeV, 240 GeV). The comparisons include the per-layer energy deposits in the first (layer 1), middle (layer 9), and last (layer 17) thin strip layers, the reconstructed shower maximum from the Longo–Sestili fit, and the total energy sum of the thin strip layers. All distributions are normalised.

9.3.1 Per-Layer Energy Deposits

Figures 75, 76, and 77 show the energy deposit distributions in layers 1, 9, and 17 respectively, for all four beam energies. The simulation reproduces the measured distributions well across the full range of deposited energies. The agreement in the first layer reflects a correct modelling of the shower start, while the agreement in the middle and last layer validates the longitudinal shower development as well as the correct modelling of the digitisation.

In the first layer, the only notable difference between simulation and data is a small tail towards higher deposited energies in the TB data, which is likely caused by pre-showering in the material upstream of the detector⁸, which is not modelled in the simulation.

In layer 9, there are two effects visible. The simulation slightly underrepresents the deposited energy at higher energies, which is likely due to a small mismatch of the saturation effects in the digitisation, caused by the usage of a common saturation level for all channels. The second effect is a population of low-energy deposits in the TB data, which is not present in the simulation. This is caused by imperfect PID in the TB, which allows for a small contamination of hadrons and muons in the electron sample.

⁸In the TB there are beam interactions with the upstream beam pipe, beam line instrumentation and other material, which are not accounted for in the simulation.




Figure 75: Normalised energy deposit distributions in layer 1 for simulation and TB data at 0° incident angle for the sample energies. The TB data show a small tail towards higher deposited energies, not present in the simulation.

In the last layer, the agreement is very good across all energies, confirming that the shower tail and leakage are correctly modelled. A small residual tail of higher energy deposits similar to the first layer remains.




Figure 76: Normalised energy deposit distributions in layer 9 for simulation and TB data at 0° incident angle for the sample energies. At high energies, the simulation slightly underestimates the deposited energy. The TB data show an additional population of low-energy deposits not present in the simulation.




Figure 77: Normalised energy deposit distributions in layer 17 for simulation and TB data at 0° incident angle for the sample energies. The agreement between simulation and data is good across all energies. A small tail towards higher deposited energies, similar to the one in layer 1, remains in the TB data.

9.3.2 Shower Maximum

Figure 78 shows the distribution of the reconstructed shower maximum $t_{\max}$ from the Longo–Sestili fit. The simulation reproduces the measured shower maximum distributions well at all four energies, confirming that the longitudinal shower development is correctly modelled. The expected shift of the shower maximum to larger depths with increasing energy follows the scaling $t_{\max} \propto \ln E$ and is reproduced by the simulation.




Figure 78: Normalised distributions of the reconstructed shower maximum $t_{\max}$ for simulation and TB data at 0° incident angle for the sample energies.

9.3.3 Total Shower Energy

Figure 79 shows the total energy sum of the thin strip layers for electrons, showing that the simulation yields a systematically higher energy sum than the TB data. As all other observables are in good agreement, this is attributed to residual differences in the saturation implementation. The discrepancy is therefore considered a scaling effect that does not point towards any fundamental mismatch between simulation and data.




Figure 79: Normalised total corrected energy sum distributions for simulation and TB data at 0° incident angle for the sample energies. The simulation yields a systematically higher energy sum than the TB data, which is attributed to residual differences in the implementation of the saturation.

9.3.4 Comparison of Intercept and Angle Reconstruction

Figure 80 shows the reconstructed intercept distributions for the simulation and TB data. Notably, the width of the simulation data is consistent, supporting the argument that the variation in the TB data is caused by the beam profile. The simulation shows the same quantisation effects as the TB data, symmetrically around the detector center. The same shift in quantisation from the strip edges to the strip centers can be observed between the 3 and 20 GeV data as in the TB data.

The angle reconstruction distributions in Figure 81 show a good agreement between simulation and TB data across all energies. The simulation data is perfectly centered around the true beam angle of 0°, while the TB data shows a small bias,




Figure 80: Normalised intercept reconstruction distributions for simulation and TB data at 0° incident angle for the sample energies.

indicating a rotational misalignment of the detector in the TB setup. The observation that the bias is especially pronounced at 240 GeV is part of a larger picture shown in Figure 82. A pattern becomes visible in the TB data, which is only in part reproduced by the simulation, pointing towards a more complex interplay of beam profile and detector alignment. As already visible in Figure 81, the bias only becomes significant in relation to the resolution at the highest energies in the TB data.

Figure 83 shows the full comparison of the angular reconstruction between TB and the simulation. The simulation follows the TB data nicely with increasing accuracy at higher energies. It can be seen that, also for the simulation, the statistical sigma is consistently larger than the fitted sigma. At larger energies, they converge, which indicates that the outliers in the TB data are not caused by badly reconstructed showers but real stray particles not present in the simulation. Overall, the simulation provides an accurate description of the TB data across all energies, angles and observables relevant for the pointing resolution analysis. The close agreement also implies that the ECAL-based PID is sufficiently effective.




Figure 81: Normalised angle reconstruction distributions for simulation and TB data at 0° incident angle for the sample energies.




Figure 82: Angular residuals as a function of angle and beam energy. The error bars show the sigma of the reconstructed angle distribution. Both the Gaussian fit width (filled markers) and the statistical standard deviation (open markers) are shown.




Figure 83: Angular resolution as a function of beam energy for all SPS TB data sets. Both the Gaussian fit width (filled markers) and the statistical standard deviation (open markers) are shown.

10 Optimisation of the SHiP ECAL

The agreement between the TB results and the simulation established in Chapter 8 demonstrates that the simulation framework reproduces the pointing performance of the physical prototype at the level required for design decisions. This agreement provides the basis for extending the simulation beyond the prototype geometry to study the full-scale SHiP ECAL. The goal of this chapter is to determine the optimal detector configuration within the boundary condition set in Chapter 6 ($20 X_0$, 40 $0.5 X_0$ sampling layers of 3 mm lead and 10 mm plastic scintillator strips, 1:1 ratio of thin to wide strips) using the established simulation framework. The position and depth of the air gap, as well as the thin strip layer sampling strategy, are systematically studied with respect to detector performance. The chapter closes with a discussion of external factors possibly influencing final design choices and paths of further improvements.

10.1 From Prototype to Full Detector Section

The conceptual prototype described in Chapter 7 was designed to validate the pointing principle and provide the measurements necessary to tune the Monte Carlo simulation. Translating the validated concept into a full detector section requires some adaptations.

All prototype pointing studies were performed with electrons for their availability at TBs. However, the signal of interest for the ECAL pointing performance are photons. In contrast to the SHADOWS prototype simulation, the SHiP ECAL has a non-uniform sampling of the shower with thin and wide strips and an air gap in between. In this design, the shower start behaviour of photons has a direct impact on the pointing performance and therefore needs to be included in the simulation. The geometry of the full detector section follows the SHiP ECAL baseline design described in Section 6, with 40 sampling layers made of 3 mm lead, alternating thin and wide scintillator strip layers and a 50 cm airgap after $6 X_0$. To study the performance, the simulation of one $2.16 \times 2.16 \text{ m}^2$ segment of the ECAL is sufficient and reduces computational cost.

10.2 Simulation

The simulation is extended directly from the validated TB setup described in Chapter 9. In this simulation setup, all layers are implemented with thin strip segmentation and exclusively vertical strip orientation. This way, a single simulated dataset can be analysed with different thin-layer sampling configurations without repeating the full simulation, similarly to the approach used in Chapter 5. The distance between layers is set to 0.5 cm as a placeholder, knowing that this is subject to change in the final design depending on engineering decisions.

Single photons are generated at incident angles between 0° and 15° and at energies from 2 to 200 GeV, covering the full kinematic range of interest for SHiP. The processing chain for the obtained events is described in the following.

10.2.1 Particle Generation

Following the prototype simulation configuration, the particle gun vertex is placed 10 m upstream of the front face of the module, with a Gaussian position smearing of $\sigma = 5$ mm applied in all three dimensions to account for the strip granularity and suppress edge effects. No angular or energy smearing is applied. The incident angle is varied in the x - z plane perpendicular to the strip orientation, exploiting the independence of the angular reconstruction in the two transverse directions. For each energy-angle configuration, 10 000 events are simulated.

The mock digitisation pipeline developed and validated in Chapter 9 is applied, removing the ADC saturation and implementing a SiPM with 14 400 pixels based on the Hamamatsu S13360-3025PE model [77] to reduce SiPM saturation effects. This SiPM model is chosen for its high pixel count, which can handle the high energy deposits, and its suitability for the expected operating conditions.

10.2.2 Reconstruction

Each digitised event is reconstructed by a global shower fit that simultaneously extracts the longitudinal energy profile, the transverse shower shape in each layer, and the incident particle angle and impact position. Two sampling configurations corresponding to the alternating x - and y -measuring planes are tested, each sampling the shower at an effective longitudinal spacing of $2 X_0$ but offset by one layer from each other. A schematic illustration of the baseline configuration is shown in Figure 84:

- **x-fit:** Layers $\{1, 5, 9, 13, 17, 21, 25, 29, 33, 37\}$
- **y-fit:** Layers $\{2, 6, 10, 14, 18, 22, 26, 30, 34, 38\}$




Figure 84: Schematic illustration of the layer order in the SHiP ECAL baseline design. The layers are arranged in an alternating pattern of thin and wide strip layers and x and y measuring layers.

For each sampling set, the reconstruction is run independently. The global shower fit proceeds in four steps, which are described in the following. The change in the fit procedure is required due to the different shower start behaviour of photons compared to electrons. The first three steps are designed to provide a robust initialisation and data reduction for the subsequent global minimisation:

Pre-fit shower start detection. A one-dimensional Longo–Sestili longitudinal profile (see Section 3.2.2) with a floating shower-start depth t_0 is first fitted to the raw layer-wise energy sums. This threshold-free fit provides a robust shower start and shape, which is then held fixed throughout the subsequent global minimisation. Based on this fit the longitudinal leakage is compensated, and the total shower energy is determined.

Strip window selection. A straight-line pre-fit to all hits above 0.1 MIP is used to estimate the shower axis position x_0 and angle θ . For each layer, only the ± 40 strips centred around the linear fit are used for the subsequent global fit. This window spans more than twenty Molière radii, providing ample space for lateral shower containment while substantially reducing the number of hits entering the global minimisation.

Data trimming. Layers at depths $t < t_0$ are removed from the front. Layers in which the total energy falls below 5 MIP are removed from the back (identical to the TB reconstruction). For reference: a 2 GeV electromagnetic shower deposits about 250 MIPs in the ECAL.

Global minimization. The core of the reconstruction is a simultaneous fit of the longitudinal and transverse shower profiles to the per-strip signals of all sampled layers. The fit minimises the following χ^2 function:

$$\chi^2 = \sum_i \frac{(S_i - E_{\text{pred},i})^2}{\sigma_i^2} + w_E \sum_\ell \frac{(L_\ell E_{\text{tot}} - E_\ell)^2}{\sigma_\ell^2}, \quad (51)$$

where S_i is the measured strip signal and σ_i is the per-hit uncertainty from Section 8.2.3. The second term anchors the total energy to the layerwise energy sums, where $L_\ell = (dE/dt)|_{t_\ell}/E_{\text{tot}}$ is the normalised longitudinal shape at layer ℓ , E_ℓ is the total measured energy summed over all strips in that layer before any strip-window selection, $\sigma_\ell = \max(\sqrt{E_\ell}, 1 \text{ MIPs})$, and w_E is a weight that keeps E_{tot} anchored to the energy sum while leaving the lateral shape parameters free. Figure 85 shows an example event display for a 200 GeV photon at 0° with the pre-fit and strip window selection applied.

The predicted signal $E_{\text{pred},i}$ in strip i of layer ℓ is the product of the longitudinal energy density and the fraction of the transverse shower profile that falls within the geometric extent of the strip:

$$E_{\text{pred},i} = \left. \frac{dE}{dt} \right|_{t_\ell} \cdot \left[(1 - F(t_{\text{rel}})) \mathcal{I}(x_i; \sigma_c(t_{\text{rel}})) + F(t_{\text{rel}}) \mathcal{I}(x_i; \sigma_t(t_{\text{rel}})) \right], \quad (52)$$

where $t_{\text{rel}} = t_\ell - t_0$ is the shower-relative depth. The transverse shape is modelled as a two-component Gaussian mixture with a compact core of width σ_c and a broader halo of width σ_t , mixed with the depth-dependent fraction F . This parametrisation is well established in the literature [44, 95] and due to the absence of ADC saturation, it is not necessary to infer core information from the tails. Both widths and the tail




Figure 85: Example overview event display for a 200 GeV photon at 0° . At the top, the longitudinal fit of the shower profile is shown, indicating the true energy deposit before and after data trimming as well as the integral of the layerwise profile fit. The bottom shows the hits above a 0.05 MIPs threshold after the data trimming. The red line indicates the projected shower axis. The panel on the right shows the fit results of the Longo–Sestili profile and the derived total energy and longitudinal leakage, as well as the fit results of the layerwise profile fits.

fraction evolve linearly with t_{rel} :

$$\begin{aligned}\sigma_c(t_{\text{rel}}) &= R_{m,0} + R_{m,\text{slope}} \cdot t_{\text{rel}}, \\ \sigma_t(t_{\text{rel}}) &= R_{m,t,0} + R_{m,t,\text{slope}} \cdot t_{\text{rel}}, \\ F(t_{\text{rel}}) &= F_0 + F_{\text{slope}} \cdot t_{\text{rel}}, \quad F \in [0.01, 0.99].\end{aligned}$$

The strip integral $\mathcal{I}(x_i; \sigma)$ gives the fraction of a Gaussian of width σ , centred on the projected shower axis $\mu(z_\ell) = x_0 + \theta z_\ell$, that falls within the physical extent of strip i with edges $x_i^\pm$:

$$\mathcal{I}(x_i; \sigma) = \frac{1}{2} \left[\text{erf} \left(\frac{x_i^+ - \mu(z_\ell)}{\sqrt{2} \sigma} \right) - \text{erf} \left(\frac{x_i^- - \mu(z_\ell)}{\sqrt{2} \sigma} \right) \right]. \quad (53)$$

An example layerwise event display for a 200 GeV photon at 0° can be seen in Figure 86. The free parameters of the fit are:

- x_0 : the transverse intercept of the shower axis at the front face of the detector;
- θ : the slope of the shower axis (the reconstructed pointing angle);
- $R_{m,0}$, $R_{m,\text{slope}}$: the intercept and depth-slope of the core width σ_c ;
- $R_{m,t,0}$, $R_{m,t,\text{slope}}$: the intercept and depth-slope of the halo width σ_t ;

- F_0 , F_{slope} : the intercept and depth-slope of the tail fraction F .

With the longitudinal shape parameters E_{tot} , t_0 , α , and β fixed from the pre-fit.




Figure 86: Example event display of the layerwise profile fits for a 200 GeV photon at 0° .

Retry logic. If the initial minimization yields $\chi^2/\text{dof} > 5$, the fit is retried with alternative seed values for x_0 , θ , and $R_{m,0}$. The result with the lowest χ^2/dof is retained.

10.3 Air Gap Optimisation

The effective pointing resolution for photons for the SHiP baseline design depends on two key parameters. The first is the position of the air gap, which determines the fraction of showers which traverse the gap. The second is the depth of the air gap, which sets the lever arm and therefore directly impacts the pointing resolution of the photons which deposit energy in both sections. Similar to the TB results, the purely statistical and fitted sigma are evaluated and differ depending on the extent of non-Gaussian tails. For all configurations, the respective plots of the statistical data for the pointing resolution are shown in the Appendix A.4.2. If not stated otherwise, they show the same tendency as the fitted values discussed here. The reconstructed

angle and intercept distribution show the same structures and behaviour as in the TB results (see Sections 8.3 and 9.3.4) and are not discussed again for this setup.

10.3.1 Gap Depth Optimisation

The physical length of the air gap determines the lever arm for any shower traversing the gap. A larger gap improves the pointing resolution, but is practically limited by two factors. The first is the acceptance of the detector, which is reduced with a larger gap because the back part effectively covers a smaller solid angle. The second is that a larger gap increases the transverse spread of the shower inside the gap, which can lead to shower overlap and lateral leakage, complicating shower separation and reducing the energy resolution.

Starting from the 50 cm air gap proposed in the SHiP baseline design, the gap depth is increased to 100, 150, 300 and 500 cm and the pointing performance is compared. The different designs tested are defined by the gap position P in radiation lengths before the gap and the gap depth D in cm. In the following, the notation $P - D$ is used to refer to a specific design, e.g. 6-50 for the SHiP baseline design which has $6 X_0$ before the gap and a gap depth of 50 cm. The results of this comparison are shown in Figure 87. In the top plot, it can be seen that there is a systematic bias in the reconstructed angle at larger angles, especially for small energies, which is consistent for all setups and caused by the angled sampling of the shower (see Figure 88). The bottom plot shows the consistent improvement of the pointing resolution with the gap depth. The largest improvement is observed when increasing the gap depth from 50 cm to 100 cm, while the improvement with larger gaps is steady but less pronounced. There is also a dependence of the pointing resolution on the incident angle. This is a result of the lower effective sampling fraction the shower experiences passing the absorbers at larger angles, which consequently affects lower energies more strongly.

Figure 89 shows the dependence of the energy resolution on the gap depth. Starting at 300 cm gap, the resolution visibly deteriorates as a consequence of lateral leakage. Generally, the energy resolutions presented are a conservative estimate, as they are determined using only thin strip layers. The contribution of the wide strips, which are known from TB data to have a better energy resolution, in the final detector is therefore expected to improve it.

Figures 90 and 91 show the shower profile for 200 GeV photons in the last layer before the air gap and the first layer after the air gap for a 100 cm and a 500 cm gap depth, respectively. In both cases, the shower is expanding laterally, with the core width being largely unaffected by the gap depth, resulting in the improvement of the pointing resolution. The shower halo, however, is significantly more spread out for the 500 cm gap, extending beyond the 2 m aperture of the ECAL, which leads to the loss of energy and consequently energy resolution. In both cases, the shower profile regains its normal shape after the particles in the blown-up tail are fully absorbed. The shower profiles for all layers of these events are shown in the Appendix A.4.1. The pointing and energy performance figures presented above are evaluated exclusively on showers with the core within the geometric acceptance of the ECAL, which shrinks considerably as the gap depth increases. Figure 92 shows the acceptance of the ECAL as a function of the gap depth defined as the solid angle covered by the last active layer of the ECAL seen from the center point at the end of the decay




Figure 87: Pointing results as a function of energy for different air gap depths at a fixed gap position of $6 X_0$. The top panel shows the reconstructed angle bias, and the bottom panel shows the angular resolution σ_θ . For the 500 cm gap, the 15° data is excluded due to lateral leakage.

volume. Table 5 provides a detailed overview of the geometric acceptance and half-opening angles for the different air gap depths. This acceptance measure is heavily simplified but gives a good feeling for the acceptance loss with increasing gap depth. In reality, the acceptance is dependent on the kinematics of the secondary particles and their final states.




Figure 88: Schematic of the bias introduced by sampling the shower at an angle. The yellow line marks the true shower axis, while the blue marks show the shower center reconstruction in each layer (black lines), visualising the bias to smaller reconstructed angles. The shower (red) is idealised for illustrational purposes.




Figure 89: Energy resolution as a function of energy for different air gap depths at fixed gap position of $6 X_0$. The resolution is largely independent of the gap depth up to 150 cm. At 300 cm and above, a visible deterioration is observed.




Figure 90: Lateral shower profile before and after a 100 cm air gap for a 200 GeV shower showing the widening of the shower. The shower tails are extending several tens of centimetres but are still well contained within the 2 m aperture of the ECAL. The shower profiles for all layers of this event are shown in the Appendix A.4.1.




Figure 91: Lateral shower profile before and after a 500 cm air gap for a 200 GeV shower showing the widening of the shower. The shower core is significantly widened, with the tails extending beyond the 2 m aperture of the ECAL, leading to lateral leakage. The shower profiles for all layers of this event are shown in the Appendix A.4.1.




Figure 92: Geometric acceptance of the SHiP ECAL as a function of air gap depth, defined as the solid angle covered by the last active layer of the ECAL seen from the center point at the end of the decay volume.

Table 5: Geometric acceptance and half-opening angles for different air gap depths. $\theta_{horizontal}$ and $\theta_{vertical}$ are the horizontal and vertical half-opening angles. The acceptance ratio is the ratio of the acceptance for a given gap depth to the acceptance for no gap.

gap (cm)	acceptance ratio	$\theta_{horizontal}$ (°)	$\theta_{vertical}$ (°)
50	0.917	10.12	14.99
100	0.844	9.70	14.38
150	0.779	9.31	13.82
300	0.623	8.31	12.35
500	0.478	7.26	10.82

Taking into account the effects of the gap depth on the pointing resolution, energy resolution and acceptance, a gap depth of 100 cm is considered to be the best choice for the SHiP ECAL. The 100 cm gap shows a significant improvement in pointing resolution compared to the 50 cm gap and the energy resolution is not yet affected by lateral leakage. The loss in acceptance mainly affects photons with angles larger than 10 or 15°, depending on the direction, which are kinematically suppressed in the SHiP signal and therefore not expected to have a significant impact on the overall physics performance of the ECAL. This measure also assumes a decay of the FIP at the very end of the detector, where the acceptance is influenced the most by the air gap. For earlier decays, the relative acceptance loss is smaller.

10.3.2 Gap Position Optimisation

Equation (14) gives approximately 99% photon conversion after $6 X_0$, which motivates the gap position in the baseline design, but Figure 93 shows that the pointing resolution only benefits from showers that traverse the gap. At the same time, the front part of the calorimeter can not be too deep, as this would prevent early show-

ering of low-energy photons from depositing energy in the back section, reducing the pointing resolution in this regime significantly. The optimal gap position is therefore a trade-off between the photon conversion probability and the shower penetration at different energies.




Figure 93: Schematic illustration of the three shower categories relative to the air gap. Showers converting before the gap and depositing energy in both sections (centre) benefit from the full lever arm. Showers converting after the gap (right) or not reaching the back section (left) do not benefit from the improved pointing resolution.

Figure 94 shows the pointing resolution for gap positions between 2 and $6 X_0$. A consistent asymmetry between the x - and y -directions is visible across all configurations. This asymmetry is a direct consequence of the one-layer offset between the two sampling sets. A photon shower starting just after an x -measuring layer will be sampled before the gap in y but not in x , effectively reducing the lever arm for the x -fit in those events. At later gap positions, the asymmetry between x and y reduces, as the higher photon conversion probability ensures that more showers deposit energy before the gap for both strip orientations. Figure 95 illustrates this effect by showing the reconstructed angle distributions separately for showers that traverse the gap and those that do not, for 30 GeV photons at 0° in the 3-100 setup. At 30 GeV, showers that do not traverse the gap have typically converted late and deposit energy exclusively in the back section and produce a significantly wider distribution, contributing mainly to the non-Gaussian tails, and the ratio between the two distributions confirms that the y -direction benefits from the one-layer offset, with a larger fraction sampled before the gap. Figure 96 shows that the energy resolution is independent of the gap positions. A plot of the linearity of the energy response for the 3-100 setup can be found in Appendix A.4.4.

Moving the gap earlier than $3 X_0$ degrades the low-energy resolution because a larger fraction of low-energy photons convert after the gap and never deposit energy in the front section. Moving the gap later improves the conversion probability before the gap, but causes showers that start early to not reach the back section, equally degrading the low-energy resolution. At high energies, both effects are less pronounced, and later gap positions yield a marginal improvement because the higher photon conversion probability ensures more showers traverse the gap. Balancing these competing effects, the best overall pointing performance across the full energy range is achieved at a gap position of $3 X_0$.




Figure 94: Pointing resolution as a function of energy for gap positions between 2 and $6X_0$ with a fixed gap depth of 100 cm. The difference between the x - and y -directions visible in all data sets reflects the offset by $0.5X_0$ between the two sampling sets.




Figure 95: Reconstructed angle distributions for showers that traverse the gap vs. showers that don't traverse the gap for 30 GeV showers at 0° in the 3-100 setup. On the left (a), the x -direction is shown; on the right (b), the y -direction is shown. At 30 GeV, showers that do not traverse the gap have typically converted late and deposit energy exclusively in the back section.




Figure 96: Energy resolution as a function of energy for gap positions between 2 and $6 X_0$ at a fixed gap depth of 100 cm. The energy resolution is independent of the gap position across the full energy range.

10.4 Sampling Optimisation

Considering the difference in pointing resolution between the x and y directions, which is caused by just one layer difference in the sampling, the sampling of the shower could have a significant impact on the pointing performance. For the 3-100 setup, various alternatives to the alternating thin and wide strip layer configuration are tested.

The first set of tests is based on the idea of keeping the back of the ECAL in an alternating configuration, starting with a thin strip layer. The front section is then modified to have different sampling configurations, which can be seen in figure 97a. All configurations are compared to the baseline configuration R , which is the baseline sampling in the x-direction. It needs to be noted that these configurations generally increase the number of thin layers in the ECAL by 2 or 4 layers, which would drive the cost.

Figure 97b shows the pointing resolution for these alternative sampling configurations. At all energies, there are only minor differences within the errors, with no significant systematic improvement over the reference configuration R for any of the tested layouts, not justifying the increased complexity and cost which would be introduced by the alternative configurations. At high energies, however, a significant dependence of the pointing resolution on the incident angle can be observed with 0° incidence, yielding a noticeably worse resolution than at larger angles. This is the opposite of the low-energy behaviour discussed in Section 10.3.1, where the reduced

effective sampling fraction at larger angles dominates. At high energies, shower statistics are large, and the stochastic term is small, so the resolution is instead limited by the geometric information available to the fit. At small angles, the shower axis is nearly collinear with the absorber stack and the angular information is carried only by transverse shower fluctuations and strip granularity. At larger angles, the shower centroid shifts between successive layers, providing a direct geometric signal that suppresses the relative contribution of granularity and shower fluctuations and improves the pointing resolution. The energy resolution is also independent of the sampling configuration.

The second approach is to concentrate the thin strip layers around the air gap. Figure 98a shows the different configurations tested, and Figure 98b shows the pointing resolution for these configurations. It can be observed that halving the number of thin strip layers to five decreases the pointing resolution at all energies. Between the other configurations, no significant systematic improvement can be observed with respect to the reference either. Not justifying any change in the sampling configuration.

A third study removed thin strip layers in the back of the ECAL to evaluate potential cost savings. As expected, the pointing resolution deteriorates with the number of layers removed, affecting high energies more strongly. The respective plots to this study can be found in the Appendix A.4.3. This observation supports the conclusion that the pointing resolution is dominated by the air gap position and depth.

10.5 Proposed Setup


The optimisation is performed within the boundary conditions established in Chapter 6, namely $20 X_0$ total depth, 40 sampling layers, and a 1:1 ratio of thin to wide strip layers. Within these constraints, the study identifies the ideal detector geometry and channel structure: a 100 cm longitudinal air gap placed at $3 X_0$ combined with the baseline alternating thin and wide strip layer configuration in both the front and back sections. This 3-100 geometry achieves the best pointing resolution across a flat angular and energy distribution, and the sampling studies confirm that no alternative thin-layer distribution offers a significant improvement. The conservative energy resolution estimate of $14.4\%/\sqrt{E/\text{GeV}}$ meets the $10\text{--}15\%/\sqrt{E/\text{GeV}}$ design target established in Section 5.1.

Beyond the identification of the optimal configuration, the study establishes a full understanding of the interplay between the design parameters and their impact on detector performance. These results provide a quantitative basis for the subsequent design stages. Proposed future studies are discussed in the next section.


10.6 Future Studies and Potential Improvements

The results presented in this chapter represent a lower bound on the achievable performance for the given configuration. Several concrete improvements remain to be explored.

Cost and Complexity Considerations. Beyond pointing and energy performance, the practical viability of a design depends critically on its manufacturing complexity.




(a) Schematic illustration of the alternative front-section sampling configurations tested. The back section remains in an alternating layout starting with a thin strip layer, while the front section is modified to explore different sampling sequences. Configuration *R* is the reference.




(b) Pointing resolution as a function of energy for the alternative front-section sampling configurations shown in (a), evaluated at a gap position of $3 X_0$ and a gap depth of 100 cm. No significant change with respect to the reference configuration *R* is observed for any of the tested layouts.

Figure 97: Alternative front-section sampling configurations and their pointing resolution.



(a) Schematic illustration of the sampling configurations that concentrate thin strip layers in the vicinity of the air gap. The configurations differ in the number of thin layers placed immediately before and after the gap, with wide strip layers filling the remaining depth.



(b) Pointing resolution as a function of energy for the gap-concentrated sampling configurations shown in (a), evaluated at a gap position of $3X_0$ and a gap depth of 100 cm. Halving the number of thin strip layers to five visibly reduces the pointing resolution across all energies, while the remaining configurations show no systematic improvement over the reference.

Figure 98: Gap-concentrated thin strip layer sampling configurations and their pointing resolution.

The SHiP ECAL requires a segmentation of the active area into six modules and a further segmentation of the back section into at least two stations, as shown in Figure 99.




Figure 99: To scale 3D-model of the optimised SHiP ECAL design, indicating the necessary segmentation.

The most favourable scenario for production is one in which the full ECAL can be assembled from identical repeating modules. This requires a uniform pattern in the layer structure and a module thickness which matches the front section, and can be stacked up to match the required thickness for the back section, too. For the target depth of $20 X_0$, a module of $3 X_0$ thickness would require the depth to be adjusted to $18 X_0$ or $21 X_0$. A module thickness of $4 X_0$ or $5 X_0$ divides $20 X_0$ exactly and is directly compatible with the baseline depth. With a single module type, tooling, assembly jigs, and quality control procedures need to be designed and validated only once, simplifying production. Every extra module type adds to the cost and complexity, with a fully customised layer stack being the most complicated scenario.

Addition of High Precision Layers. The addition of two or three high-precision layers with a channel pitch at the millimetre scale or below could further boost the pointing resolution. Several detector technologies have been considered over time, including Micromegas, Resistive Plate Chambers, and scintillating fibre trackers. The benefit of such layers depends critically on their longitudinal placement, and a dedicated study of their optimal number and positioning would be needed. The

addition of such layers would also add to the cost and complexity of the detector.

Multiple air gaps. The current design uses a single air gap to provide one lever arm between the front and back sections. A second air gap placed deeper in the absorber stack would introduce a second lever arm, potentially improving the pointing resolution for showers that start late and are not well-sampled before the first gap. A double gap setup would need to undergo a similar evaluation as presented in this thesis.

Wider strips in the back section. As discussed in Section 10.3.1, the electromagnetic shower halos broaden substantially while traversing the air gap. The back section of the calorimeter, samples a considerably wider transverse profile than the front, and the pointing resolution in this regime is presumably dominated by the lever arm rather than the strip granularity. This suggests that the fine 1 cm segmentation of the thin strips may be unnecessary in the rear section, and that wider strips could be substituted with little impact on performance. Such a replacement would directly reduce the channel count and associated readout cost, and motivate a dedicated quantitative study weighing the performance trade-off against cost savings and complexity considerations.

Module boundaries, mechanical support, alignment, and upstream material. The results presented here assume a uniform, uninterrupted active volume throughout the whole detector with no material between the decay vertex and the first detector layer. In practice, the $4 \times 6 \text{ m}^2$ ECAL face needs to be segmented into six modules with dead regions at the module boundaries and mechanical support structures and is placed behind other instrumentation upstream. The impact of dead regions and additional material on angular resolution, acceptance, and reconstruction efficiency has not been studied here and represents an important next step towards a realistic performance estimate of the final detector. Furthermore, the pointing reconstruction relies on the precise knowledge of the relative layer positions. The sensitivity of the angular resolution to realistic inter-layer misalignment has not been evaluated.

Reconstruction algorithm. The global shower fit described in Section 10.2.2 is a shower physics motivated implementation which has not been optimised to its full potential, and a dedicated optimisation of the fit could yield measurable improvements. The systematic bias in the reconstructed angle at large incident angles, visible in Figure 87, is only one indicator for potential improvements in the reconstruction, and a more detailed study of the fit performance and failure modes could identify further areas for improvement.

Realistic signal kinematics and multi-photon final states. All simulations in this chapter use flat distributions in energy (2–200 GeV) and angle (0° – 15°) for isolated photons. The physics goal, however, is the reconstruction of $a \rightarrow \gamma\gamma$ decays or final states with higher multiplicity, which requires the simultaneous reconstruction of two photons or more and the determination of a decay vertex from their trajectories. Folding in realistic kinematic distributions from different BSM models would provide

a more sophisticated figure of merit for the pointing performance. The performance under these conditions, including vertex reconstruction quality and invariant mass resolution, has not yet been assessed and is the most important missing piece for a realistic physics performance estimate.

Conclusion

The Standard Model leaves several fundamental questions unanswered, including the origin of dark matter and the mechanism behind neutrino masses. *Feebly Interacting Particles* (FIPs) are a broad class of well-motivated beyond-Standard-Model candidates that address several of these shortcomings. Their defining feature is a very weak coupling to Standard Model particles, which suppresses their production cross-section and results in macroscopic decay lengths, making them experimentally accessible only at high-luminosity beam-dump experiments. *Axion-like particles* (ALPs) are a minimal, model-independent subclass of FIPs which have a clean experimental signature: $ALP \rightarrow \gamma\gamma$. Reconstructing a displaced neutral decay vertex inside the fiducial volume is the common experimental challenge for all such searches. The *electromagnetic calorimeter* (ECAL) must therefore measure the energy and the direction of the incident electromagnetic showers. Both experiments studied here use a plastic scintillator strip sampling calorimeter. A derivation of the ALP decay kinematics in the boosted regime relevant at the SPS shows that the propagated invariant-mass uncertainty is dominated by the angular resolution. Pointing resolution is therefore the primary figure of merit for all studies performed. This thesis studied the ECAL design for two experiments at the SPS that differ substantially in geometry and boundary conditions: SHADOWS, a proposed off-axis detector, and SHiP, the approved on-axis experiment. For SHADOWS, the simulation study showed that the relevant limitation of the ECAL is set by shower-direction reconstruction at low energies and low angles, limited by the channel-pitch. For a $20 X_0$ iron–scintillator sampling calorimeter with $0.5 X_0$ layers and a lower energy bound of 2 GeV, a channel pitch of 1 cm is sufficient to lift this granularity limitation. The resulting energy resolution of about $11\%/\sqrt{E/\text{GeV}}$ is within the requirements, and the reconstructed vertex and invariant-mass performance for the two ALP benchmark scenarios confirms sufficient pointing resolution for an $X \rightarrow \gamma\gamma$ search in the SHADOWS environment.

For SHiP, the same reconstruction task is constrained by substantially different boundary conditions. The on-axis position of SHiP produces final states with higher energies and correspondingly smaller opening angles than the off-axis SHADOWS case, placing significantly stricter requirements on the angular resolution. At these energies and angles, the centroid displacement between successive layers is small and strip granularity alone is insufficient to resolve it. A longitudinal air gap is therefore introduced as the central design element, increasing the lever arm and consequently the pointing resolution without changing the total absorber depth, the number of active layers, or the strip segmentation. The larger dimensions require alternating 1 cm (thin) and 6 cm (wide) scintillator strip layers, reducing the channel count while preserving the pointing performance. The design challenge becomes a trade-off between the lever arm, shower containment, geometric acceptance, and implementation complexity.

A conceptual prototype was built to experimentally validate the pointing capability and to provide the data needed to tune and validate the simulation framework used for the detector optimisation. The scaled-down but functionally complete prototype was progressively commissioned at DESY and the CERN PS before the final test beam at the CERN SPS with electron beams from 3 to 288 GeV experimentally

established that the plastic scintillator strip-based concept can perform pointing reconstruction under realistic SPS conditions.

A GEANT4-based simulation of the prototype was developed and tuned to the measured detector response using a mock digitisation that considers noise, SiPM saturation, and ADC overflow. The simulation reproduces the test-beam observables closely, establishing that the underlying physics and detector model are correct at the level required for design decisions. This validated simulation can be scaled to full detector modules or the complete SHiP ECAL geometry to be used for performance and optimisation studies.

Within this framework, the SHiP ECAL geometry is optimised starting from the baseline design, with the air gap depth and position varied systematically to achieve the best pointing resolution over the full angular and energy range of interest. The study shows that increasing the gap depth gives a clear gain in pointing resolution, which is limited at the upper end by the widening of the shower halo in the gap and the acceptance loss from the reduced solid angle coverage. Shifting the gap position has a more complex effect on the pointing resolution at different energies and angles, where generally an earlier position improves the low-energy performance and a later position improves the high-energy performance, limited on both ends by the fraction of showers traversing the gap and benefiting from the lever arm.

Within the boundary conditions of $20 X_0$ total depth, 40 sampling layers, and a 1:1 ratio of thin to wide strip layers, the study identifies the optimal detector geometry and channel structure: a 100 cm gap at $3 X_0$ combined with the baseline alternating thin and wide strip layer configuration. The conservative estimate for the energy resolution of $14.4\%/\sqrt{E/\text{GeV}}$ meets the $10\text{--}15\%/\sqrt{E/\text{GeV}}$ design target established in Section 5.1, and the pointing resolution is significantly improved over the baseline design. Alongside the identification of the optimal configuration, the study establishes a systematic understanding of how each design parameter influences each performance observable. These results are direct input for future optimisation with theory-motivated kinematic spectra or under modified boundary conditions and design decisions.

The natural next step is to move from the flat single-photon performance metric to theory-motivated final state kinematics, and to evaluate vertex and invariant-mass reconstruction starting from the optimised design presented here. With an integration of the ECAL into the full experimental simulation of the SHiP experiment, the impact of module boundaries and dead space on acceptance and shower fitting must be quantified, as these are now among the leading detector-specific uncertainties. On the hardware side, a technological prototype of a $2 \times 2 \text{ m}^2$ module with dedicated electronics and readout is the next step towards the final detector.

A Appendix

A.1 Comparator schematic

The trigger detector used in the beam test setup is based on a $6 \times 6 \times 1 \text{ cm}^3$ scintillator tile read out by a Hamamatsu SiPM. To convert the analog SiPM signal into a digital LVTTTL-compliant square signal required by the FERS frontend modules, a custom fast comparator circuit was designed by Antonin Zeman. The circuit introduces a propagation delay of approximately 20 ns and is compact enough to be mounted directly adjacent to the trigger tile, minimising cable-induced delays. The full schematic of the comparator board is shown in Figure 100.




Figure 100: Schematic of the fast comparator circuit designed by Antonin Zeman.

A.2 FERS Daisy Chain

To ensure synchronous data acquisition across all ECAL and HCAL frontend modules, a daisy chain trigger distribution scheme is used. Figure 101 shows a schematic of this trigger distribution chain taken from the DT5202 manual [86].




Figure 101: Schematic diagram of a daisy-chained trigger distribution with multiple DT5202. [86]

A.3 SHADOWS ALP Benchmark Scenarios

The performance of the SHADOWS ECAL baseline design is evaluated on two ALP benchmark scenarios with representative kinematic properties provided by Torben Ferber [96]. The full overview plots for both scenarios are shown in Figures 102 and 103, including the decay vertex distributions, ECAL impact points, photon energy spectra, and opening angle distributions.





Figure 102: Overview over ALP benchmark scenario with $m_a = 0.35$ GeV and $c\tau = 0.5$ m. In the top row, the decay vertex distribution is shown in all three dimensions; in the bottom left, the impact points of the two gammas on the ECAL, as well as the energy distributions of the gammas and the opening angles. This figure is high resolution and can be zoomed into in the digital version. [96]

ALP kinematics: $m=2$ GeV, $c\tau=0.85$ m



6 SHADOWS - Torben Ferber

Institute of Experimental Particle Physics (ETP)

Figure 103: Overview over ALP benchmark scenario with $m_a = 2.0$ GeV and $c\tau = 0.85$ m. In the top row, the decay vertex distribution is shown in all three dimensions; in the bottom left, the impact points of the two gammas on the ECAL, as well as the energy distributions of the gammas and the opening angles. This figure is high resolution and can be zoomed into in the digital version. [96]

A.4 Additional Plots for Chapter 10

A.4.1 Shower Profiles

The lateral shower profiles for all detector layers are shown here for two representative single-photon events at 300 GeV, one for a 500 cm air gap configuration and one for a 100 cm gap. These profiles complement the single-layer comparisons shown in the main text (Figures 90 and 91). For the 500 cm gap, the shower halo expands significantly beyond the 2 m aperture of the ECAL after the gap, leading to lateral leakage and a degradation of energy resolution. For the 100 cm gap, the shower remains well contained throughout all layers. The layer-by-layer profiles illustrate the qualitative difference in shower containment between the two configurations.




Figure 104: Lateral shower profiles for a 300 GeV photon shower in a detector configuration with a 500 cm air gap. The profiles are shown for each layer, illustrating the lateral spread of the shower as it progresses through the detector.




Figure 105: Lateral shower profiles for a 300 GeV photon shower in a detector configuration with a 100 cm air gap. The profiles are shown for each layer, illustrating the lateral spread of the shower as it progresses through the detector.

A.4.2 Optimisation plots with statistical sigma

This section collects additional plots related to the optimisation studies of the SHiP ECAL design in Chapter 10. Figures 106 and 107 show the pointing resolution (statistical sigma) as a function of energy for different gap positions and depths, complementing the fitted pointing resolution plots in the main text. Figures 108 and 109 show the pointing resolution for alternative sampling configurations.




Figure 106: Pointing resolution (statistical sigma) as a function of energy for gap positions between 2 and 6 X_0 at a fixed gap depth of 100 cm. The difference between the x - and y -directions visible at all positions reflects the offset by $0.5 X_0$ between the two sampling sets.




Figure 107: Pointing resolution (statistical sigma) as a function of energy for gap depths between 50 and 500 cm at a fixed gap position of $6 X_0$.




Figure 108: Pointing resolution (statistical sigma) as a function of energy for the alternative front-section sampling configurations shown in Figure 97a, evaluated at a gap position of $3 X_0$ and a gap depth of 100 cm. No improvement with respect to the reference configuration R is observed for any of the tested layouts.




Figure 109: Pointing resolution (statistical sigma) as a function of energy for the gap-concentrated sampling configurations shown in Figure 98a, evaluated at a gap position of $3X_0$ and a gap depth of 100 cm. No configuration shows any systematic improvement over the reference.




Figure 110: Pointing resolution as a function of energy for the gap-concentrated sampling configurations shown in Figure 98a, evaluated at a gap position of $3 X_0$ and a gap depth of 100 cm. Halving the number of thin strip layers to five visibly reduces the pointing resolution across all energies, while the remaining configurations show no systematic improvement over the reference.




Figure 111: Pointing resolution (statistical sigma) as a function of energy for the configuration with layers removed in the back, evaluated at a gap position of $3 X_0$ and a gap depth of 100 cm.

A.4.3 Sampling with layers removed in the back

Figures 112, 113, and 111 show the sampling configuration with thin strip layers removed in the back and its performance.




Figure 112: Schematic illustration of the sampling configurations remove thin strip layers in the back. The configurations differ in the number of thin layers, with wide strip layers filling the remaining depth.




Figure 113: Pointing resolution as a function of energy for the configuration with layers removed in the back, evaluated at a gap position of $3X_0$ and a gap depth of 100 cm.

A.4.4 Energy reconstruction linearity




Figure 114: Linearity in energy reconstruction for setup 6-100.

List of Figures

1	The particles of the Standard Model.	5
2	The Mexican-hat potential and its ground states.	7
3	Exclusion limits on ALPs with photon coupling.	9
4	ALP parameter space accessible to accelerator-based experiments. . .	10
5	Primakoff production of ALPs.	13
6	Kinematics of FIPs produced at the SPS.	15
7	The average energy loss per unit path length for positive muons in copper.	17
8	The energy domains in which three possible photon interactions happen, as a function of the Z value of the absorber material.	17
9	Energy deposit as a function of depth, for 10 GeV electron showers developing in aluminium, iron and lead.	22
10	The energy deposit for 1, 10, 100 and 1000 GeV electron showers in a block of copper as a function of depth.	22
11	Longitudinal and lateral development of an electron shower in lead. .	24
12	Simulation of showers electromagnetic and hadronic in an iron block. .	24
13	Schematic of a homogenous and a sampling calorimeter.	25
14	Schematic of calorimeters with two different sampling fractions. . . .	26
15	Schematic illustration of particle identification in a detector at a high energy accelerator.	27
16	Conceptual sketch of a readout chain of plastic scintillator, wavelength-shifting fibre and photosensors.	28
17	Absorption and emission spectrum of scintillators (schematic).	28
18	Schematic of the array structure of a SiPM and the electric circuit representation.	29
19	Single photon spectrum of a SiPM.	30
20	Illustration of the captured light fraction for a typical fibre with two cladding layers.	31
21	Illustration of the different contributions to the geometric acceptance of beam dump experiments.	33
22	The SHADOWS off-axis detector setup.	34
23	Illumination at the first SHADOWS tracking layer from $\pi^+\mu^-$ pairs of HNL decays.	34
24	The SHiP experimental setup at ECN3.	36
25	Simulated muon flux distribution at the first tracking station of the SHADOWS detector system.	37
26	Muon flux distribution after the SHiP active muon shield.	37
27	Effect of lever-arm length on diphoton separation in $X \rightarrow \gamma\gamma$ decays. .	40
28	Photon separation at the SHiP ECAL face for ALP decays with different masses and lifetimes.	41
29	Reconstructed invariant mass distributions for $X \rightarrow \gamma\gamma$ under different detector configurations.	42
30	Picture of the SiPM-on-tile concept.	45
31	Render of the SHADOWS ECAL baseline design simulation.	46
32	Visualisation of segmentation effect at low angles and energies.	48

33	Angular residual for a 2 GeV electron shower directional reconstruction at different angles and tile dimensions.	48
34	Energy resolution of the SHADOWS ECAL baseline design.	50
35	Angular resolution of the SHADOWS ECAL baseline design.	50
36	Invariant mass resolution of the SHADOWS ECAL baseline design. .	51
37	Vertex resolution of the SHADOWS ECAL baseline design.	51
38	Reconstruction ambiguities in the SHiP ECAL baseline design.	54
39	Illustration of the longitudinal gap in the SHiP ECAL.	54
40	Render of the SHiP baseline layer structure.	56
41	To scale 3D-model of the SHiP ECAL baseline design.	56
42	Picture of the coextruded scintillators from Fermilab.	58
43	CAD visualisation of a thin strip layer.	58
44	Schematic of the illumination of the square sensitive area of a SiPM by a round light cone from a WLS fiber.	59
45	Pictures of thin and wide strip layers.	60
46	Picture of the ECAL prototype at the SPS test beam.	61
47	ECAL prototype layer stack and ordering at the SPS test beam. . . .	62
48	Layer shifting in the ECAL prototype during angular scans.	63
49	Picture of the HCAL prototype at the SPS test beam.	64
50	Peak detector chronogram of the Citiroc1A.	65
51	Schematic of the DAQ setup.	66
52	Hardware components of the fast trigger setup.	66
53	Sketch of setup at the SPS test beam.	67
54	Event building schematic.	70
55	Example pedestal fits for the HG and LG branches.	71
56	Double peak pedestal fit.	72
57	Intercalibration heatmap for wide strips.	72
58	Example intercalibration plots for thin and wide strip channels. . . .	73
59	Intercalibrated channel spectrum.	73
60	Example MIP calibration plots for thin and wide strip channels. . . .	74
61	MIP calibration overview.	74
62	MIP calibration of thin and wide strips.	75
63	Example SiPM calibration.	75
64	PID displays for a sample of events.	77
65	Comparison of particle identification results from ECAL-only, combined ECAL+HCAL, and H8 reference.	78
66	Layerwise event data with saturated shower centers.	79
67	Event display comparison of global shower fit.	82
68	Comparison of layerwise profile fits.	83
69	Reconstructed intercept x_0 of SPS TB data for example energies at 0° incidence angle.	84
70	Reconstructed shower axis slope θ of SPS TB data for example energies at 0° incidence angle.	85
71	Angular resolution as a function of beam energy.	85
72	Reconstructed shower axis slope θ in log-scale with non-Gaussian tails. .	86
73	Visualisation of the simulated setup with ECAL and HCAL.	87
74	Visualisation of the simulated ECAL setup from the top.	88

75	Simulation vs. test beam energy deposit in layer 1.	90
76	Simulation vs. test beam energy deposit in layer 9.	91
77	Simulation vs. test beam energy deposit in layer 17.	92
78	Simulation vs. test beam reconstructed shower maximum.	93
79	Simulation vs. test beam: total corrected energy sum.	94
80	Simulation vs. test beam: intercept reconstruction.	95
81	Simulation vs. test beam: angle reconstruction.	96
82	Angular residuals as a function of angle and beam energy.	96
83	Angular resolution as a function of beam energy.	97
84	Schematic illustration of the layer order in the SHiP ECAL baseline design.	99
85	Example overview event display.	101
86	Example event display of the layerwise profile fits.	102
87	Pointing results as a function of energy for different air gap depths. .	104
88	Schematic of the bias introduced by sampling the shower at an angle.	105
89	Energy resolution as a function of energy for different air gap depths.	105
90	Lateral shower profile before and after a 100 cm air gap for a 200 GeV shower.	106
91	Lateral shower profile before and after a 500 cm air gap for a 200 GeV shower.	106
92	Geometric acceptance of the SHiP ECAL as a function of air gap depth.	107
93	Schematic illustration of the three shower categories relative to the air gap.	108
94	Pointing resolution as a function of energy for different gap positions.	109
95	Reconstructed angle distributions for showers that traverse the gap vs. showers that don't traverse the gap.	109
96	Energy resolution as a function of energy for different gap positions. .	110
97	Alternative front-section sampling configurations and pointing resolution.	112
98	Gap-concentrated thin strip layer configurations and pointing resolution.	113
99	To scale 3D-model of the optimised SHiP ECAL design, including segmentation.	114
100	Schematic of the fast comparator circuit.	119
101	Schematic diagram of a daisy-chained trigger distribution with multiple DT5202.	120
102	Overview over ALP benchmark scenario with $m_a = 0.35$ GeV and $c\tau = 0.5$ m.	121
103	Overview over ALP benchmark scenario with $m_a = 2.0$ GeV and $c\tau = 0.85$ m.	122
104	Lateral shower profiles for a 300 GeV photon shower in a detector configuration with a 500 cm air gap.	123
105	Lateral shower profiles for a 300 GeV photon shower in a detector configuration with a 100 cm air gap.	124
106	Pointing resolution (statistical sigma) as a function of energy for different gap positions.	125

107	Pointing resolution (statistical sigma) as a function of energy for different gap depths.	126
108	Pointing resolution (statistical sigma) for alternative front-section sampling configurations.	126
109	Pointing resolution (statistical sigma) for gap-concentrated thin strip layer configurations.	127
110	Pointing resolution for gap-concentrated thin strip layer configurations in log scale.	128
111	Pointing resolution (statistical sigma) for configuration with layers removed in the back.	129
112	Sampling configurations with thin strip layers removed in the back. .	130
113	Pointing resolution for configuration with layers removed in the back.	130
114	Linearity in energy reconstruction for setup 6-100.	131

List of Tables

1	Layer positions in the ECAL prototype.	63
2	Key operating settings for the test beam readout.	65
3	Overview of the test beam facilities used in this campaign.	68
4	Run configurations at the SPS test beam.	69
5	Geometric acceptance and half-opening angles for different air gap depths.	107

References

- [1] C. C. Delogu et al. “The SHADOWS calorimeter — A pointing calorimeter for future beam dump experiments”. In: *Nucl. Instrum. Meth. A* 1069 (2024), p. 169873.
- [2] Sebastian Ritter et al. “A Pointing Electromagnetic Calorimeter for FIP Experiments leveraging $X \rightarrow \gamma\gamma$ Decays”. In: *EPJ Web Conf.* 320 (2025), p. 00011.
- [3] M Alviggi et al. *SHADOWS Technical Proposal*. Tech. rep. Geneva: CERN, 2023. URL: <https://cds.cern.ch/record/2878470>.
- [4] R Albanese et al. *BDF/SHiP at the ECN3 high-intensity beam facility*. Tech. rep. Geneva: CERN, 2023. URL: <https://cds.cern.ch/record/2878604>.
- [5] Fabiola Gianotti. *News from the March CERN Council Session*. Accessed: 2026-04-20. Mar. 2024. URL: <https://home.cern/news/opinion/cern/news-march-cern-council-session>.
- [6] David J Griffiths. *Introduction to elementary particles*. Physics textbook. New York, NY: Wiley, 2008. URL: <https://cds.cern.ch/record/111880>.
- [7] Mark Thomson. *Modern Particle Physics*. Cambridge University Press, 2013.
- [8] P. A. Zyla et al. “Review of Particle Physics”. In: *PTEP* 2020.8 (2020), p. 083C01.
- [9] ATLAS Collaboration. *The Standard Model and Beyond*. CERN Open Data Portal. https://opendata.atlas.cern/docs/documentation/introduction/SM_and_beyond, Accessed: 2026-02-10. 2023.
- [10] X. Fan et al. “Measurement of the Electron Magnetic Moment”. In: *Phys. Rev. Lett.* 130.7 (2023), p. 071801. arXiv: 2209.13084 [physics.atom-ph].
- [11] G. Aad et al. “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”. In: *Physics Letters B* 716.1 (2012), pp. 1–29. URL: <https://www.sciencedirect.com/science/article/pii/S037026931200857X>.
- [12] Vincenzo Cirigliano and Michael J. Ramsey-Musolf. “Low energy probes of physics beyond the standard model”. In: *Progress in Particle and Nuclear Physics* 71 (2013). Fundamental Symmetries in the Era of the LHC, pp. 2–20. URL: <https://www.sciencedirect.com/science/article/pii/S0146641013000215>.
- [13] D. P. Aguillard et al. “Measurement of the Positive Muon Anomalous Magnetic Moment to 0.20 ppm”. In: *Phys. Rev. Lett.* 131.16 (2023), p. 161802. arXiv: 2308.06230 [hep-ex].
- [14] T. Aoyama et al. “The anomalous magnetic moment of the muon in the Standard Model”. In: *Phys. Rept.* 887 (2020), pp. 1–166. arXiv: 2006.04822 [hep-ph].
- [15] Gianfranco Bertone, Dan Hooper, and Joseph Silk. “Particle dark matter: Evidence, candidates and constraints”. In: *Phys. Rept.* 405 (2005), pp. 279–390. arXiv: hep-ph/0404175.

- [16] Y. Fukuda et al. “Evidence for oscillation of atmospheric neutrinos”. In: *Phys. Rev. Lett.* 81 (1998), pp. 1562–1567. arXiv: [hep-ex/9807003](#).
- [17] N. Aghanim et al. “Planck 2018 results. VI. Cosmological parameters”. In: *Astron. Astrophys.* 641 (2020). [Erratum: *Astron. Astrophys.* 652, C4 (2021)], A6. arXiv: [1807.06209](#) [[astro-ph.CO](#)].
- [18] Jihn E. Kim and Gianpaolo Carosi. “Axions and the Strong CP Problem”. In: *Rev. Mod. Phys.* 82 (2010). [Erratum: *Rev. Mod. Phys.* 91, 049902 (2019)], pp. 557–602. arXiv: [0807.3125](#) [[hep-ph](#)].
- [19] J. Goldstone. “Field Theories with Superconductor Solutions”. In: *Nuovo Cim.* 19 (1961), pp. 154–164.
- [20] Manuel Endres et al. “The ‘Higgs’ Amplitude Mode at the Two-Dimensional Superfluid-Mott Insulator Transition”. In: *Nature* 487 (2012), pp. 454–458. arXiv: [1204.5183](#) [[cond-mat.quant-gas](#)].
- [21] R. D. Peccei and Helen R. Quinn. “CP Conservation in the Presence of Instantons”. In: *Phys. Rev. Lett.* 38 (1977), pp. 1440–1443.
- [22] Jihn E. Kim and Gianpaolo Carosi. “Axions and the Strong CP Problem”. In: *Rev. Mod. Phys.* 82 (2010). [Erratum: *Rev. Mod. Phys.* 91, 049902 (2019)], pp. 557–602. arXiv: [0807.3125](#) [[hep-ph](#)].
- [23] Giovanni Grilli di Cortona et al. “The QCD axion, precisely”. In: *JHEP* 01 (2016), p. 034. arXiv: [1511.02867](#) [[hep-ph](#)].
- [24] J. Beacham et al. “Physics Beyond Colliders at CERN: Beyond the Standard Model Working Group Report”. In: *J. Phys. G* 47.1 (2020), p. 010501. arXiv: [1901.09966](#) [[hep-ex](#)].
- [25] Georg G. Raffelt. “Astrophysical axion bounds”. In: *Lect. Notes Phys.* 741 (2008). Ed. by Markus Kuster, Georg Raffelt, and Berta Beltran, pp. 51–71. arXiv: [hep-ph/0611350](#).
- [26] Prateek Agrawal et al. “Feebly-interacting particles: FIPs 2020 workshop report”. In: *Eur. Phys. J. C* 81.11 (2021), p. 1015. arXiv: [2102.12143](#) [[hep-ph](#)].
- [27] F. Englert and R. Brout. “Broken Symmetry and the Mass of Gauge Vector Mesons”. In: *Phys. Rev. Lett.* 13 (1964). Ed. by J. C. Taylor, pp. 321–323.
- [28] Peter W. Higgs. “Broken Symmetries and the Masses of Gauge Bosons”. In: *Phys. Rev. Lett.* 13 (1964). Ed. by J. C. Taylor, pp. 508–509.
- [29] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble. “Global Conservation Laws and Massless Particles”. In: *Phys. Rev. Lett.* 13 (1964). Ed. by J. C. Taylor, pp. 585–587.
- [30] Oliver Fischer et al. “Unveiling hidden physics at the LHC”. In: *Eur. Phys. J. C* 82.8 (2022), p. 665. arXiv: [2109.06065](#) [[hep-ph](#)].
- [31] C.C. Ahdida et al. *SPS Beam Dump Facility - Comprehensive Design Study*. CERN Yellow Reports: Monographs. Geneva: CERN, 2020. URL: <https://cds.cern.ch/record/2703984>.
- [32] Manuel Meucci. “MEG II experiment status and prospect”. In: *PoS NuFact2021* (2022), p. 120. arXiv: [2201.08200](#) [[hep-ex](#)].

- [33] Letizia Peruzzo. “Physics beyond the Standard Model with the NA62 experiment at CERN”. In: Jan. 2024, p. 071.
- [34] L. Bartoszek et al. “Mu2e Conceptual Design Report”. In: (Nov. 2012). arXiv: 1211.7019 [physics.ins-det].
- [35] Babette Döbrich et al. “ALPtraum: ALP production in proton beam dump experiments”. In: *JHEP* 02 (2016), p. 018. arXiv: 1512.03069 [hep-ph].
- [36] R Albanese et al. *BDF/SHiP Annual Report 2025*. Tech. rep. Geneva: CERN, 2025. URL: <https://cds.cern.ch/record/2948477>.
- [37] Babette Döbrich, Joerg Jaeckel, and Tommaso Spadaro. “Light in the beam dump - ALP production from decay photons in proton beam-dumps”. In: *JHEP* 05 (2019). [Erratum: *JHEP* 10, 046 (2020)], p. 213. arXiv: 1904.02091 [hep-ph].
- [38] Gaia Lanfranchi, Maxim Pospelov, and Philip Schuster. “The Search for Feebly Interacting Particles”. In: *Ann. Rev. Nucl. Part. Sci.* 71 (2021), pp. 279–313. arXiv: 2011.02157 [hep-ph].
- [39] Martin Bauer, Matthias Neubert, and Andrea Thamm. “Collider Probes of Axion-Like Particles”. In: *JHEP* 12 (2017), p. 044. arXiv: 1708.00443 [hep-ph].
- [40] Claudia Ahdida et al. “Sensitivity of the SHiP experiment to light dark matter”. In: *Journal of High Energy Physics* 2021 (Apr. 2021), p. 199.
- [41] Kyrylo Bondarenko et al. “Towards the optimal beam dump experiment to search for feebly interacting particles”. In: *Eur. Phys. J. C* 83.12 (2023), p. 1126. arXiv: 2304.02511 [hep-ph].
- [42] Claus Grupen and Boris Schwartz. *Particle Detectors*. ISBN: 9781009401531. 2008.
- [43] Hermann Kolanoski and Norbert Wermes. *Particle Detectors: Fundamentals and Applications*. Oxford University Press, June 2020. URL: <https://doi.org/10.1093/oso/9780198858362.001.0001>.
- [44] Richard Wigmans. *Calorimetry: Energy Measurement in Particle Physics*. Oxford University Press, Sept. 2017. URL: <https://doi.org/10.1093/oso/9780198786351.001.0001>.
- [45] S. Navas et al. “Review of particle physics”. In: *Phys. Rev. D* 110.3 (2024), p. 030001.
- [46] Bruno Rossi and R Ronald Rau. *High-Energy Particles*. 1953.
- [47] E. Longo and I. Sestili. “Monte Carlo Calculation of Photon Initiated Electromagnetic Showers in Lead Glass”. In: *Nucl. Instrum. Meth.* 128 (1975). [Erratum: *Nucl. Instrum. Meth.* 135, 587 (1976)], p. 283.
- [48] W. Ralph Nelson, H. Hirayama, and David W. O. Rogers. “The Egs4 Code System”. In: (Dec. 1985).
- [49] HAMAMATSU. *What is a silicon photomultiplier?* Website. accessed on 3.4.26. URL: <https://hub.hamamatsu.com/us/en/technical-notes/mppc-sipms/what-is-an-SiPM-and-how-does-it-work.html>.

- [50] Broadcom. *Single-Photon Measurements*. Application Note. accessed on 3.4.26. URL: <https://docs.broadcom.com/doc/AFBR-S4NxxPyy4M-Single-Photon-Measurements>.
- [51] W Baldini et al. *SHADOWS (Search for Hidden And Dark Objects With the SPS): Expression of Interest*. Tech. rep. Geneva: CERN, 2022. URL: <https://cds.cern.ch/record/2799412>.
- [52] Eric van Herwijnen. “The SHiP (Search for Hidden Particles) proposal”. In: (2018), pp. 236–240. URL: <https://cds.cern.ch/record/2677499>.
- [53] C. Ahdida et al. “The SHiP experiment at the proposed CERN SPS Beam Dump Facility”. In: *Eur. Phys. J. C* 82.5 (2022), p. 486. arXiv: 2112.01487 [physics.ins-det].
- [54] M. U. Ashraf et al. “High Intensity Kaon Experiments (HIKE) at the CERN SPS Proposal for Phases 1 and 2”. In: (Nov. 2023). arXiv: 2311.08231 [hep-ex].
- [55] Guglielmo Frisella and Massimiliano Ferro-Luzzi. “Study of combinatorial background tagging for different muon shield variants”. In: (2025). URL: <https://cds.cern.ch/record/2938479>.
- [56] M. Anelli et al. *A Facility to Search for Hidden Particles (SHiP) at the CERN SPS*. Tech. rep. Technical Proposal. Geneva: CERN, 2015. arXiv: 1504.04956. URL: <https://cds.cern.ch/record/2007512>.
- [57] Matei Climescu. *Design and development of calorimetry at SHiP and SND@LHC*. PhD Thesis. 2025.
- [58] Josue Daniel Jaramillo Arroyave. *EventCalc: An Event Generator for the SHiP Experiment at CERN*. August 2024.
- [59] Eduardo Cortina Gil et al. “The Beam and detector of the NA62 experiment at CERN”. In: *JINST* 12.05 (2017), P05025. arXiv: 1703.08501 [physics.ins-det].
- [60] “ATLAS liquid argon calorimeter: Technical design report”. In: (Dec. 1996).
- [61] “The CMS electromagnetic calorimeter project: Technical Design Report”. In: (1997).
- [62] G. Dellacasa et al. “ALICE technical design report of the photon spectrometer (PHOS)”. In: (Mar. 1999).
- [63] C. Adloff et al. “Construction and Commissioning of the CALICE Analog Hadron Calorimeter Prototype”. In: *JINST* 5 (2010), P05004. arXiv: 1003.2662 [physics.ins-det].
- [64] Felix Sefkow and Frank Simon. “A highly granular SiPM-on-tile calorimeter prototype”. In: *J. Phys.: Conf. Ser.* 1162.1 (2019). 5 pages, 5 figures, proceedings of CALOR 2018, Eugene, OR, USA, May 2018, p. 012012. arXiv: 1808.09281. URL: <https://cds.cern.ch/record/2636699>.
- [65] Daria Selivanova. “High Granularity Scintillator Tiles for the High Luminosity Upgrade of the CMS Endcap Calorimeter”. In: *Proceedings of the International Conference on Technology & Instrumentation in Particle Physics (TIPP 2026)*. Mumbai, India, Feb. 2026. URL: <https://bib-pubdb1.desy.de/record/646262>.

- [66] S. Agostinelli et al. “Geant4—a simulation toolkit”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 506.3 (2003), pp. 250–303. URL: <https://www.sciencedirect.com/science/article/pii/S0168900203013688>.
- [67] The CALICE collaboration et al. “Electromagnetic response of a highly granular hadronic calorimeter”. In: *Journal of Instrumentation* 6.04 (Apr. 2011), P04003–P04003. URL: <http://dx.doi.org/10.1088/1748-0221/6/04/P04003>.
- [68] Oskar Hartbrich. “Scintillator Calorimeters for a Future Linear Collider Experiment”. PhD thesis. Hamburg: Bergische U., Wuppertal (main), 2016.
- [69] Jan Jerhot et al. “ALPINIST: Axion-Like Particles In Numerous Interactions Simulated and Tabulated”. In: *JHEP* 07 (2022), p. 094. arXiv: 2201.05170 [hep-ph].
- [70] M. Dorn et al. “KLauS: A charge readout and fast discrimination chip for silicon photomultipliers”. In: *JINST* 7 (2012), p. C01008.
- [71] F. Dulucq et al. “CALOROC-B: A Fully Integrated 36-Channel ASIC for SiPM Readout at EIC”. In: *2025 IEEE Nuclear Science Symposium (NSS), Medical Imaging Conference (MIC) and Room Temperature Semiconductor Detector Conference (RTSD)*. 2025, pp. 1–1.
- [72] CERN SPS. *H8*. Website. accessed on 9.5.25. URL: <https://be-dep-ea.web.cern.ch/experimental-areas/north-area/h8>.
- [73] R. Diener et al. “The DESY II Test Beam Facility”. In: *Nucl. Instrum. Meth. A* 922 (2019), pp. 265–286. arXiv: 1807.09328 [physics.ins-det].
- [74] Eva Montbarbon et al. “The new CERN East Area primary and secondary beams”. In: *10th International Particle Accelerator Conference*. 2019, TH-PGW062.
- [75] D. Beznosko et al. “FNAL-NICADD extruded scintillator”. In: (Sept. 2005).
- [76] kuraray. *Plastic Scintillating Fibers*. Datasheet. accessed on 23.7.25. URL: <http://kuraraypsf.jp/pdf/all.pdf>.
- [77] HAMAMATSU. *MPPC S13360 series*. Datasheet. accessed on 23.7.25. URL: https://www.hamamatsu.com/content/dam/hamamatsu-photonics/sites/documents/99_SALES_LIBRARY/ssd/s13360_series_kapd1052e.pdf.
- [78] Sebastian Ritter. *Studies to Optimize the Electromagnetic Calorimeter of the DUNE ND-GAr Detector*. Master’s Thesis. 2022.
- [79] Eljen Technology. *EJ200*. Website. accessed on 9.5.25. URL: <https://eljentechnology.com/products/plastic-scintillators/ej-200-ej-204-ej-208-ej-212>.
- [80] HAMAMATSU. *MPPC S14160 series*. Datasheet. accessed on 23.7.25. URL: https://www.hamamatsu.com/content/dam/hamamatsu-photonics/sites/documents/99_SALES_LIBRARY/ssd/s14160-1310ps_etc_kapd1070e.pdf.
- [81] Fairview Microwave. *FMC0202316*. Datasheet. accessed on 23.5.26. URL: <https://www.mouser.at/datasheet/3/6148/1/FMC0202316.pdf>.

- [82] CAEN S.p.A. *FERS-5200 Boards*. Website. accessed on 23.3.2026. URL: <https://www.caen.it/subfamilies/fers-5200/>.
- [83] CAEN S.p.A. *A5202*. Website. accessed on 23.3.2026. URL: <https://www.caen.it/products/a5202/>.
- [84] CAEN S.p.A. *DT5215*. Website. accessed on 23.3.2026. URL: <https://www.caen.it/products/dt5215/>.
- [85] weeroc. *Citiroc1A*. Datasheet. accessed on 23.3.2026. URL: <https://www.weeroc.com/~documents/products/citiroc-1a/citiroc1a-datasheet-v2-53/?layout=file>.
- [86] CAEN S.p.A. *A5202/DT5202*. Manual. Rev. 4 - Jan 25th, 2023.
- [87] CAEN S.p.A. *Janus*. Website. accessed on 23.3.2026. URL: <https://www.caen.it/products/janus/>.
- [88] Matei Climescu. *Developments of the PID detector*. 34th SHiP Collaboration Meeting. Sept. 18, 2025. URL: https://indico.cern.ch/event/1578670/contributions/6689130/attachments/3138309/5569254/SHiPCM_PID_34.pdf.
- [89] Maarten Van Dijk. Private communication with Maarten Van Dijk (H8 beam physicist at CERN) on 13.5.2026 via Mattermost.
- [90] Sascha Krause. “Studies of the Response of Silicon Photomultipliers and Test-beam Data Analysis of a Highly Granular Analog Hadron Calorimeter Prototype”. eng. PhD thesis. Mainz: Johannes Gutenberg-Universität Mainz, 2020.
- [91] Maarten Van Dijk. *Electron Intensity in the H8 Beamline*. CERN EDMS. URL: <https://edms.cern.ch/ui/#!/master/navigator/document?D:101585345:101585345:subDocs>.
- [92] Python Software Foundation. *iminuit 2.32.0*. URL: <https://pypi.org/project/iminuit/>.
- [93] Markus Frank et al. *AIDAsoft/DD4hep*. webpage: <http://dd4hep.cern.ch/>. Oct. 2018. URL: <https://doi.org/10.5281/zenodo.592244>.
- [94] Matei Climescu. *DD4SHiP*. URL: <https://github.com/matclim/DD4SHiP>.
- [95] Guenter Grindhammer and S. Peters. “The Parameterized simulation of electromagnetic showers in homogeneous and sampling calorimeters”. In: *International Conference on Monte Carlo Simulation in High-Energy and Nuclear Physics - MC 93*. Feb. 1993. arXiv: [hep-ex/0001020](https://arxiv.org/abs/hep-ex/0001020).
- [96] Torben Ferber. Private communication with Torben Ferber on 26.7.2023 via e-mail.