



CERN-SPSC-2025-039 / SPSC-SR-370

11 November 2025

BDF/SHiP Annual Report 2025

SHiP Collaboration, HI-ECN3 Project Team ¹

Abstract

The report covers the progress of the SHiP collaboration and the HI-ECN3 Project Team since the start of the TDR phase mid-2024.

Keywords: annual report, SHiP, BDF, HI-ECN3, SPS

Contacts: Andrey.Golutvin@cern.ch, Richard.Jacobsson@cern.ch, Matthew.Fraser@cern.ch

¹Complete author lists at the end.

Contents

1	Introduction	1
2	Progress with the BDF	1
2.1	High Intensity ECN3 Project	1
2.2	Preparations for LS3	2
2.3	CERN-wide proton sharing	3
2.4	SPS slow extraction R&D	3
2.5	SHiP beam-induced background studies	4
2.6	BDF target R&D and XTAX	5
2.7	BDF target station R&D	6
2.8	Infrastructure	6
2.8.1	BDF target service building	6
2.8.2	Service cell and synergy with radioactive waste management at CERN	7
2.8.3	n-ACT	7
2.9	RP and safety	7
2.10	Evaluation of radiation to electronics	7
3	Progress with the SHiP detector and major changes since the approval	10
3.1	Muon shield	11
3.2	Neutrino and LDM Scattering Detector	17
3.3	FIP Decay Detector	21
3.3.1	He decay volume	21
3.3.2	Upstream Background Tagger	23
3.3.3	Surrounding Background Tagger	27
3.3.4	Spectrometer Straw Tracker	32
3.3.5	Spectrometer Magnet	34
3.3.6	Spectrometer Timing Detector	40
3.3.7	PID system	41
3.4	Common electronics, online system and computing	44
3.4.1	Electronics and online system	44
3.4.2	Computing	45
3.5	Physics software	47
4	Reassessment of physics performance with SHiP-2026	48
4.1	Large-scale production of background samples	49
4.2	ML-based simulation of the large background data samples	49
5	Test beams at CERN in 2026	50
5.1	DUBT and SMTC scintillating tiles	50
5.2	High-granularity calorimeter (SHGC) prototype	50
5.3	Surrounding background tagger (DSBT)	50
5.4	Spectrometer straw tracker (DSST)	50
5.5	Spectrometer timing detector (DSTD)	51

6	Plan and objectives for Run 4	51
7	Status of the SHiP collaboration	55
	References	56
	Author list	60
	Acknowledgements	63

1 Introduction

This document outlines developments in the High-Intensity ECN3 (HI-ECN3) project and by the Search for Hidden Particles (SHiP) collaboration [1, 2, 3, 4, 5, 6, 7, 8, 9, 10] (and references therein) since the approval of the experiment in March 2024. The SHiP collaboration plan to install the initial configuration of the detector in the years 2029 - 2032 in order to start the detector commissioning in the second half of 2032 and data taking in 2033, ahead of Long Shutdown 4 (LS4). This time scale is primarily motivated by the hard goal to commission the Beam Dump Facility (BDF) and the experiment before LS4, and thus have the shutdown to complete the detector and take any remedial action that is required.

The CERN management have fully endorsed the physics case and approved construction of the High-Intensity ECN3 facility and BDF. Based on their assessment in 2024 of the external funding prospects for the full 50 MCHF SHiP experiment, they have requested a staged approach to the construction and installation of the SHiP detector. The collaboration is therefore preparing the initial configuration with a simplification of several key subsystems, in particular the muon shield and the decay volume. The collaboration is prioritising a detector in Run 4 that is capable of primarily exploring fully reconstructed final states in the initial phase without neutral particles and neutrinos. This allows staging several detector components. It remains the collaboration's aim to have the muon shield, decay volume, background taggers and the main spectrometer installed such that commissioning with the first data can take place and the background suppression be verified. Although the reduced initial configuration would not be capable of guaranteeing zero-background with several years of operation, it will still outperform previous and current experiments with one year of data (4×10^{19} PoT) in all benchmark models by 1-2 orders of magnitude. The complete detector is necessary to fully exploit the physics potential of the SPS and the facility, and to carry out the full physics programme with the complete set of final states at the required zero-level background. The collaboration believes the staging strategy to be an asset, given that the ultimate sensitivity will be driven by the understanding of the backgrounds. Long shutdown 4 is an essential step with the opportunity to analyse the information from Run 4, to make adjustments or modifications to the facility, detector and software tools, and launch upgrades.

Beyond progressing towards the initial configuration with the technical and design challenges, which are discussed in Chapter 3, a key obstacle for the collaboration remains getting firm commitments from funding agencies on the compressed timescale.

2 Progress with the BDF

2.1 High Intensity ECN3 Project

The HI-ECN3 project was officially approved and mandated in June 2024 to implement the BDF for SHiP [11] following ratification by the CERN Council of an additional 64 MCHF for accelerator and experimental area infrastructure. This commitment complements the investment already made in CERN's MTP for the consolidation of the North Area to secure CERN's long-term future as the world leader in high-energy fixed-target physics. To make BDF/SHiP operational in Run4, the HI-ECN3 project is currently undergoing a rapid Tech-

nical Design Report (TDR) phase building on the CDS [4]. The publication of the TDR is expected in 2026 before execution starts in Long Shutdown 3. The timeline is short, but conservative margins and clear constraints have been made (scope, volumes & loads) with tendering for the civil engineering works expected before the end of 2026. Authorisation to go ahead with the civil engineering procurement process was recently granted by the CERN management via the Infrastructure Request Process (IRP) and completion of the IRP Design Study [12].

2.2 Preparations for LS3

Ground was already broken in 2025 and the site around ECN3 prepared for the construction works scheduled to start in mid-2027. Included in the scope of the IRP Design Study are a new target service building (B754), underground excavations to integrate the SHiP spectrometer and consolidation of the existing surface buildings, shown in Fig. 1. The previously foreseen iron inserts in the TCC8 floor under the BDF target complex have been de-scoped following further radiological assessment and optimization, including soil activation studies, comparisons with today's operational situation for NA62 and hydrogeological analyses of the site. The additional access shaft in ECN3 has also been de-scoped after studies concluded that SHiP's infrastructure can be installed via the existing B911 shaft together with the routing of the services between B754 and the new target complex. To ensure uninterrupted site access and continuity of power and service networks during the works, a service trench was rerouted to allow for the future enlargement of the road for vehicle access and manoeuvring, as well as creating space for parking and an assembly area for SHiP. The position of the personnel and material doors of B911 were moved to facilitate a clear separation of underground dismantling activities of the NA62 experiment in parallel to the surface construction works.




Figure 1: BDF/SHiP integrated at ECN3 with the new surface building containing ancillaries for the new target complex.

NA-CONS Phase 1 is on schedule to start execution at the beginning of LS3. The systems around the T4 target in TCC2 will also be upgraded for BDF/SHiP in LS3 to allow for the fast-switching of the beam's destination and operation of EHN1 and EHN2 in parallel to ECN3. The beam will be bumped vertically onto the T4 production target to generate secondary beams for EHN1, whilst bypassing it when being delivered to SHiP. The T4 XTAX

on the P4 beamline towards ECN3 will no longer be moveable so that the beam line can be placed under vacuum to minimise the beam’s interaction with matter upstream of ECN3. This is an essential step to reducing beam loss in the transfer line and minimizing beam induced background at SHiP. A collimation system is also under design and will be installed on the beamline downstream of TCC2 to make sure that the proton beam is delivered cleanly onto the BDF target. A Beam Interlock System is being deployed in the North Area and its design has been updated to guarantee the level of machine protection required to safely transfer the high intensity beam destined for SHiP through TCC2 to ECN3 [13]. New Elements Important for Safety (EIS) are to be installed in TCC2 as part of the Personnel Protection System (PPS) to guarantee safe access to TCC8 and ECN3 throughout the various project phases — post-LS3, BDF beam commissioning, SHiP installation, and BDF/SHIP operation. Significant progress has been made to ensure the safe transfer of the high-intensity beam to the BDF facility in TCC8, as well as a comprehensive analysis of the temporary beam dump required for the TCC2 commissioning phase after LS3 and before BDF/SHiP is ready to start beam commissioning.

Under the auspices of the HI-ECN3 project during 2025, the Hadron Stopper of the de-commissioned PS Neutrino Facility located between the AD and Restaurant 2 was excavated to recover 618 tonnes of cast-iron shielding buried at the end of the TT7 tunnel in the early 1980s. The cast-iron blocks have been removed, cleaned and are now in storage and ready for reuse in the BDF Target Complex.

2.3 CERN-wide proton sharing

A complex-wide beam sharing study was carried out this year to expand on the previous SPS and ECN3-specific proton sharing studies [14]. The study is part of a process to develop a coherent strategy to deliver the desired flux and quality to all the different CERN facilities, not only at SPS, and including both protons and ions, during the SHiP operational era in Run4 and beyond. The study [15] identified potential bottlenecks for specific facilities and has outlined a roadmap to accommodate SHiP’s request without impacting these facilities. As input to this study the specific requirements and capabilities of current and future experiments and test-beam users in the North Area were investigated and recommendations made to optimise beam delivery in the future [16].

2.4 SPS slow extraction R&D

The DECRYCE (Development of CRYstals for Collimation and beam Extraction) project has been established to develop in-house capabilities at CERN for designing, assembling, and characterising bent crystal devices. These devices will support present and future accelerator applications. In particular, one of DECRYCE’s main objectives is to turn the SPS prototype crystal shadowing system into an operational system that will open the door to SHiP’s high intensity physics programme in the North Area. Last year, optimised crystal devices were manufactured and validated for installation in the TECA goniometer located in LSS4 of SPS to reduce the beam loss in the slow extraction straight of LSS2. Today, the prototype is operational and can achieve a beam loss reduction factor of up to 2. The use of Single Crystals with an optimised thickness offer a potential beam loss at reduction factor of 3,

whilst a Multi Volume Reflection Array (MVRA) with four strips aligned to deflect the beam using the Volume Reflection (VR) process is used to reduce the beam loss by up to a factor of 10. Figure 2 shows the assembly of a first CERN-produced MVRA in 2025. Preliminary results taken on the H8 beamline this year using a new CERN-developed telescope showed that three crystal strips were nearly fully aligned, achieving a deflection of close to the expected $3\times VR$ deflection. Further optimisation of the benders and clamp designs for both the MVRA and Optimised Single Crystal is ongoing and new crystal assemblies are planned for installation during the YETS25/26 and for beam extraction tests in the SPS in 2026.




Figure 2: Left: Multi-volume Reflection Array assembly, Right: simulated beam loss as the electrostatic extraction septum in LSS2 as a function of crystal number.

Low-density electrostatic septum tanks made from aluminium rather than stainless steel have finished prototyping and will be installed during LS3 as part of the planned consolidation of the septa. The residual dose next to these tanks will be reduced by a factor of two from the low-Z material choice alone. R&D has identified a key source of the straightness imperfections observed on the anodes of the electrostatic anodes, which will be improved where possible to further reduce the beam lost during slow extraction.

2.5 SHiP beam-induced background studies

The background of concern for the experiment from the beamline is muons. Any upstream loss of primary protons will lead to muons being produced, predominantly from the decay of secondary hadrons. The BDF Target Complex and the downstream muon shield consist of significant material that would degrade incoming muons by at least 30 GeV/c. The muon shield will use magnetic fields to deflect muons created in the target. A leakage rate of approximately 10^4 muons is expected per spill of 4×10^{13} protons on target. The beam-induced muon background from the upstream beamline therefore needs to be of a similar magnitude or less than this leakage rate, assuming that every muon is problematic for the experiment. Muons created upstream and within the radius of the BDF target (and with sufficient energy) will pass through it and the muon shield will reject them. Muons at larger radii from the beam axis may pass through the shielding of the BDF Target Complex and

potentially enter into the detector’s acceptance if they are captured by the return fields of the muon shield.

Empirical experience in various North Area beamlines has shown that one can expect 10^{-3} muons downstream per proton lost. Assuming a tolerable 10^4 muons at BDF from the beamline, an integral of 10^7 protons could therefore be lost at most, corresponding to a fractional beam loss of $\sim 10^{-7}$. Today, typical fractional losses in the later part of P42 in line-of-sight of TCC8 and ECN3 are approximately $10^{-4} - 10^{-5}$, caused by a halo of large amplitude particles traveling with the beam. The halo is generated by a significant amount of material seen by the beam in TCC2 (the T4 production target itself, necessary diagnostics, surrounding air and vacuum windows). The material seen by the BDF/SHiP beam in TCC2 will be reduced significantly after the upgrades in LS3.

In addition, a collimation system is envisaged in P42 after TCC2, but well before the line-of-sight of the experiment, to define the maximum momentum acceptance of protons as well as to clean scattered protons from the upstream material interactions. Once the beamline turns away from the M2 beamline and towards ECN3, the beam optics and magnet layout will transport the beam as cleanly as possible. Additionally, two magnetised iron blocks (MIBs) with a crude toroidal magnetic field are being considered for installation around the beam pipe 100 – 300 m upstream of BDF to deflect muons to large amplitudes before BDF/SHiP.

A simulation model of the full P42 beamline from TCC2 through to ECN3, including a detailed tunnel geometry, has been created in the Geant4-based BDSIM code. The model is being used to optimise the beamline layout for collimation and interceptive diagnostics by generating the predicted muon flux and distribution at the front of BDF. The proton beam distribution from this simulation, i.e. with tails from material interaction and correct time-momentum correlations, will be used to understand the background from protons impacting on the BDF target at large amplitude or the upstream aperture in the shielding of the BDF Target Complex.

2.6 BDF target R&D and XTAX

Multiple target concepts have been studied resulting in an optimized full-W, He-cooled target for SHiP. A comprehensive simulation package is being developed to optimize the core geometry and address operational conditions such as the Loss-of-Coolant-Accident scenario. A comprehensive material R&D campaign is following the design studies, particularly on the characterization of the pure W material and the manufacturing and testing of the HIP-produced core blocks. In parallel, a two-stage prototyping strategy with beam tests has been put in place to (i) test multiple core manufacturing technologies in 2025 and (ii) assess the target operation with a He-cooling skid in 2026. The prototype targets tested in 2025 are shown in Figure 3.

Design and engineering works is also underway on other beam intercepting devices required in the transfer line to ECN3: three new beam stoppers are under production; studies to integrate the TIDVG4 as a P42 commissioning Dump have been concluded; design work is ongoing for three new primary collimators. Finally, the new P4 XTAX is being designed, with procurement activities starting soon.



Figure 3: The prototype targets for test in 2025 [17]. Left: one target shown before assembly in its helium jacket. Right: both targets installed in the slow extraction test facility on T6 in TCC2.

2.7 BDF target station R&D

The engineering of the target station has progressed on the design of the interfaces between the target and the surrounding systems, in particular on the cooling system, proximity water cooled shielding and retraction/extraction. The different layers of shielding surrounding the target station have been made to comply with radiation protection requirements and background attenuation for the experiment. The length of the Magnetized Hadron Stopper (MHS) and its surrounding shielding could be reduced while maintaining optimized residual dose rates, radiation levels to equipment, and experimental background conditions. The target system vacuum vessel is currently under a detailed design study through a design contract and the foreseen size of the structure of the vacuum vessel is being confirmed. In addition, there are ongoing qualifications of different types of sealing solutions under irradiation. The maintenance scenarios for the target station are being developed, incorporating the handling systems. The new overhead traveling crane in TCC8 has been fully specified and the procurement of the fully remotely controlled crane is underway. The overall integration of the underground facility has been developed together with the SHiP experiment and the upstream beam delivery.

2.8 Infrastructure

2.8.1 BDF target service building

The design of the new building B754 has reached an advanced stage, defining a concrete structure with a footprint of 960 m² directly connected to B911 to house the ancillary systems for the BDF's target complex. The design contains airtight and under-pressurised zones to safely house a Class A Service Cell for the handling of radioactive components, as well as dedicated areas for cooling and ventilation systems for the BDF target, electrical installations, control racks, and safety systems. The building's conceptual and schematic layouts, load envelopes, and foundation design are complete, providing a robust basis for the

detailed design phase ahead of construction in 2027. The power converters for BDF/SHiP are no longer planned to be housed in B754, but rather in BA82, once the the K12 beamline power converters have been dismantled in LS3.

2.8.2 Service cell and synergy with radioactive waste management at CERN

A service cell will be required to handle, inspect and maintain the radioactive components of the BDF target station and to package radioactive waste before disposal upon decommissioning. This is today a stringent design requirement. The service cell design is also being developed to treat and eliminate legacy radioactive waste from other facilities as part of the Radioactive Waste Management Programme at CERN. The RWM Steering Board has endorsed the extended use of the BDF service cell for the processing and packaging of radioactive waste, which guarantees CERN’s long-term commitment and strategy to dispose of radioactive waste. The design of the service cell is currently being developed through a design contract with a specialized company in this specific field.

2.8.3 n_ACT

A parasitic, high-flux neutron activation station embedded within the BDF target station has been studied as part of the n_ACT@BDF proposal by the n_TOF collaboration and its feasibility has been demonstrated. An Expression of Interest and a Technical proposal by the n_TOF collaboration has been submitted to the SPSC [18, 19].

2.9 RP and safety

Comprehensive radiation protection assessments and optimisation studies have been performed across the full facility. For the experimental area, residual dose rates were optimised to comply with a Supervised Radiation Area [20]. Radioactive waste zoning studies were conducted and demonstrated negligible activation risks in the SHiP detector region. Studies of muon deflection from the new muon shield design towards the above-ground fenced and public areas demonstrated that the resulting effective doses remain well below CERN’s optimisation objective.

The FIRIA project has delivered the initial draft proposal of the general fire safety requirements for the underground areas — BDF (TCC8) and SHiP (ECN3). The proposed fire-safety concept addresses the fire resistance and reaction of construction materials, compartmentalisation requirements, a preliminary ventilation and smoke-extraction strategy, the performance of safety systems, life-safety provisions, and fire-brigade intervention measures. Furthermore, the FIRIA studies investigated the minimum fire-resistance rating for the confinement walls through numerical modelling of a fire scenario on the BDF side. A radiological dose assessment of potential environmental smoke release from an MHS coil epoxy fire is also currently underway.

2.10 Evaluation of radiation to electronics

Radiation interacting with electronic components can induce cumulative effects and destructive and non-destructive single event effects (SEE). Protons with a momentum of 400 GeV/c




Figure 4: Implemented modifications in the FLUKA model of BDF in the TCC8/ECN3 beam facility.

impinging on a 150 cm-long tungsten (W) target generates a strong radiation field leading to both effects. Radiation simulations for BDF/SHiP have been performed using the FLUKA Monte Carlo code [21, 22] to evaluate the radiation environment in TCC8 and ECN3 and its potential impact on electronics (R2E). The objective was to assess the risk that BDF/SHiP in TCC8 and ECN3 impose restrictions on the use of non-radiation tolerant commercial electronics or impose the need for radiation-tolerant equipment.

The FLUKA [21, 22] model has been updated with recent modifications, shown in Figure 4. The main modifications to the geometry include:

- Full 1.5 m tungsten target with a diameter of 25 cm and the shorter hadron stopper of 2.3 m.
- New muon shield design, featuring a geometry with twice the volume, copper coils in the magnets, and an updated magnetic field map;
- Confinement wall composed of 5 cm of 30% borated polyethylene and 2 cm of steel;
- Additional concrete shielding around the upstream part of the muon shield;

In order to have a quantitative measure of the radiation effects on electronics, the quantities listed in Table 1 were assessed. The quantities are listed with the limits that characterise a zone as a safe R2E area. The results are reported in Figure 5, considering a proton beam with a transverse size of $\sigma = 16$ mm and a sweep radius of 50 mm, under standard BDF operating conditions corresponding to 4×10^{19} protons on target (PoT) per year.

The radiation levels in ECN3 are expected to remain below the R2E limits, with the exception of the areas along the muon flames, the muon shield region itself and, in the case of ThNeq, its surroundings. The optimised shielding configuration composed of concrete blocks around the first magnet of the muon shield and the confinement wall made with 30% borated polyethylene, effectively reduces radiation effects. In particular, the ThNeq decreases by approximately two orders of magnitude beyond the wall, and the enclosed geometry helps




Figure 5: R2E estimation in ECN3 considering nominal BDF operation of 4×10^{19} PoT/year. Values represented in xz -plane averaging y between -75 cm and 75 cm with beam line at $y = 0$ cm.

Effects	R2E quantity	Limit value	Unit
Cumulative	TID	1	Gy/(10 years)
	Si1MeVNE	1e10	$\text{cm}^{-2}/(10 \text{ years})$
SEE	HEHeq	3e6	$\text{cm}^{-2}/\text{year}$
	ThNeq	3e7	$\text{cm}^{-2}/\text{year}$

Table 1: General R2E limits derived from the LHC experience to characterise an area as radiation-safe for electronics [23].

to prevent HEH leakage. The new muon shield design confines the muon fluence within its structure, resulting in significantly reduced fluence outside the shield, see Figure 6.




Figure 6: Muon fluence in ECN3 considering nominal BDF operation of 4×10^{19} PoT/year. Values represented in xz -plane averaging y between -75 cm and 75 cm with beam line at $y = 0$ cm.

The radiation exposure to the SND detector and electronics services in ECN3 remains below the thresholds requiring special mitigation measures, indicating that radiation-tolerant electronics are not necessary. Further optimization could be considered in the areas housing the compressor and superconducting magnet cryocoolers, where R2E limits are still exceeded. All radiation effects in the potential R2E-safe location inside the GHN300 tunnel are within acceptable limits.

3 Progress with the SHiP detector and major changes since the approval

Fig. 7 shows the overview of the current baseline detector with all of the sub-components indicated. Acronyms for the subsystems associated with the scattering detector are preceded by "S" and those associated with the decay detector by "D" in what follows.




Figure 7: Overview of the SHiP detector. Note that only a partial coverage of the DSBT is shown for illustration.

Since the approval, the SHiP collaboration have made significant progress in all areas. The major developments are described throughout Section 3 of this document. The developments involve several substantial changes:

- Development of the hybrid (superconducting/warm) and fully warm options of the muon shield to prepare the choice between the two options by the end of 2025. The reduction of the muon flux has been found satisfactory in both options.
- New layout of the neutrino and light dark matter scattering detector (SND) integrated within the last section of the muon shield. SND consists of two components: high granular Si/W calorimeter (SHGC) and magnetised tracking calorimeter (SMTC) with iron absorbers and coarser segmented active layers made of scintillating tiles and scintillating fibres. MTC is using the magnetic field of the muon shield to reconstruct the charge of muons in muon neutrino interactions and in $\tau \rightarrow \mu\nu\nu$ decays to discriminate between ν and $\bar{\nu}$ neutrino interactions.
- Replacement of the SHiP vacuum vessel with a helium-filled gas balloon. Despite of the increased number of neutrino and muon interactions with He, their resulting background to signal has been found to be kept at acceptable level. At the same time, the background caused by the interactions with the walls of the decay vessel has been dramatically reduced. The 300 tonnes of steel needed to form the vacuum vessel was a significant source of secondary particles. Tagging this background was one of the principal goals of the Surrounding Background Tagger (DSBT). The requirements on the rejection power of the DSBT surrounding He-filled vessel are by far less challenging.
- Technologies and layout of the Upstream Background Tagger (DUBT). The baseline DUBT configuration is now based on two technologies. A layer of scintillating tiles, readout by SiPMs with the time resolution $O(200\text{ ps})$, is followed by few layers of low-material straw tubes to reconstruct tracks of penetrating muons and secondaries from muon and neutrino inelastic interactions in the last part of the muon shield. The transverse dimension of DUBT as well as its rejection power are currently under study.

In addition, a memorandum [24] has been established with the CERN Technology department outlining the contributions to the SHiP detector and collaboration on the magnetised hadrons stopper, the spectrometer magnet, and the consultancy on the muon shield and the helium decay volume and purification system. A technical note [25] has also been finalised, covering the requirements of the SHiP detector on the TCC8/ECN3 facility as required to prepare the HI-ECN3 project TDR.

3.1 Muon shield

The active muon shield is a crucial component of the SHiP experiment. Muons are the only particles, among those that can interact strongly or electromagnetically, which emerge from the BDF target and hadron absorber in large numbers. From simulation, about 2×10^{10} muons per spill are expected to come out of the hadron absorber with an energy larger than 10 GeV. Muons below that energy can be stopped with adequate shielding. The muon

flux at the exit of the BDF/SHiP hadron absorber as a function of the muon momentum p and the transverse momentum p_T , as expected from Monte Carlo simulation, is shown in Fig. 8. These muons can interact with matter and create fake signatures of a FIP decaying in the SHiP experimental setup. Random combinations of two muons in the detector can also mimic such signatures. SHiP has invested considerable effort to predict as accurately as reasonably possible the flux of muons emerging from the BDF target. A large muon background simulation sample, containing about 412 million events with at least one muon of $E > 10$ GeV having traversed the hadron absorber, is heavily used to study and optimise the design of the muon shield and establish its performance with adequate reliability.




Figure 8: Left: Expected muon flux from Monte Carlo simulation at the exit of the SHiP/BDF hadron absorber as a function of momentum p and transverse momentum p_T . Fig. from [26]. Right: Measured muon momentum distributions compared to Monte Carlo simulation. Fig. from [27].

The role of the active muon shield is to sweep out or stop the muons. For the SHiP experiment, an alternate-polarity scheme was proposed and developed [26]. The working principle of this scheme is displayed in the left graph of Fig. 9 which shows schematically, in green and blue, the regions of positive and negative vertical magnetic field (here, $|B_y| \approx 1.8$ T) in the horizontal mid plane ($y = 0$) of the muon shield. The high-energy muons are deflected by the core of the first section (with a given polarity) and continue to be deflected by the return field in the yoke of the second section (opposite polarity). The mid-momentum muons are “refocused” by the return field of the first section onto the yoke of the second section, where they are then fully “re-deflected” towards the outside, out of the detector acceptance. Low-momentum muons are stopped. The muon rate in the SHiP Decay Detector is suppressed by about six orders of magnitude. As a result, the muon shield splits the positive and negative muons emerging from the target into two prominent flames to the left and to the right, thus creating a “sanctuary” in the middle where the low background experiment can be implemented. This is illustrated in the right graph of Fig. 9 which shows the yearly muon fluence in the horizontal mid plane as expected from FLUKA simulations with a recent version of the muon shield adapted to TCC8/ECN3.

Since the original SHiP experiment proposal [2] a number of muon shield variants were developed, mostly in the context of the “ECN4 proposal” which involved the construction of a new SPS underground facility [26, 28, 29, 30, 31, 32, 33, 34, 35]. First explorations for an implementation in TCC8/ECN3 showed that a similar performance was within reach,




Figure 9: Principle of the alternate polarity scheme of the SHiP/BDF muon shield (see text for explanation). Left: the green/blue colors and $\odot/\otimes$ symbols indicate the polarity of the field in the iron. Fig. taken from [26] and adapted. Right: expected muon fluence in the TCC8/ECN3 facility (courtesy of Dominika Wasik, SY-STI-BMI).

despite the reduced space available and the unfavorable vicinity of the ECN3 Salève-side wall [9, 10]. A “hybrid” variant that includes a superconducting (SC) magnet in the first section of the muon shield was also explored and simulation results indicated that significant advantages could be expected with such a configuration, mostly in terms of longitudinal space and integration of emulsions² within the Scattering Detector.




Figure 10: Residual dose rate in the first few meters downstream of the BDF target ($X \in [-12, 12]$ cm, $Y \in [-13, 15]$ cm). Figure taken from reference [36].

The muon shield consists of a series of electromagnets. The first magnet, the magnetized hadron stopper (MHS), being located just downstream of the BDF target, will be exposed to severe radiation and is therefore part of the target complex. It will be designed and installed

²Emulsion detectors have to be exchanged after reaching a charged particle track density of $\mathcal{O}(10^6 \text{ cm}^{-2})$ which, assuming a yearly exchange, imposes a rate limit of the order of 1 Hz/cm^2 .

by the HI-ECN3 project, with functional requirements defined by the SHiP Collaboration. The MHS guarantees that the vast majority of hadrons and electromagnetic debris are stopped in the target complex and that muons experience a sweeping magnet field as early as possible. The last part of the muon shield will also require a dedicated design, as it must host the scattering detector throughout its middle (see Section 3.2).

In 2025, the muon shield studies led to several important results and a number of actions were taken. The design of the muon shield, following the original idea of the alternate-polarity scheme, was further developed in several steps, . Configurations based solely on normal-conducting (NC) magnets and a hybrid (SC/NC) configuration were optimised with constraints of the TCC8/ECN3 facility. For this, a new machine-learning Bayesian optimisation framework (dubbed MuonsAndMatter, abbreviated MaM) was developed by the Collaboration and deployed on a CPU/GPU cluster of the University of Zürich for optimising the muon shield layout. MaM implements a machine learning-based global optimisation method that uses probabilistic surrogate models to efficiently explore and optimise complex, computationally expensive functions. It includes a complete GEANT4 implementation of the target and muon shield magnets. Thousands of muon shield layouts are tested within each optimisation run. Additionally, a fast magnetostatic solver (Sn ∞ .py) was developed [37] to facilitate field map calculations “on the fly” within the MaM optimisation framework, an important improvement compared to the previous use of blocks of constant fields.




Figure 11: Total ionizing dose in the region of the MHS. Left: XY plane at $Z \in [207.7, 217.7]$ cm. Middle: YZ plane at $X \in [-4.9, 4.9]$ cm Right: XZ plane at $Y \in [-5.1, 4.8]$ cm. Courtesy of Dominika Wasik, SY-STI-BMI.

MaM and Sn ∞ .py were extensively used to optimise the fully warm and hybrid muon shield variants in the ECN3/TCC8 context. Satisfactory solutions were found for all variants. In a second step, the obtained solutions were re-optimised with constraints to include sufficient space in the core of the last 5 m of the muon shield for integrating the Scattering Detector. Finally, further studies were conducted, both with manual scans of parameters and with an advanced random sampling algorithm, to explore the robustness of the optimised muon shield variants. In the latter approach, a neural-network-based local surrogate model was trained to predict the expected number of muons within the acceptance of the spectrometer tracker for a given parameter configuration. This differentiable model enables efficient robustness analyses by capturing the stochastic behavior of the system and allowing system-

atic exploration of parameter variations. A study of the combinatorial muon background for different muon shield configurations was also performed [38].

The length of the MHS was revisited in the light of recent radiation simulations by our colleagues from HSE-RP [36]. Fig. 10 shows the residual dose rate in the first few meters downstream of the BDF target indicating that the downstream end of the MHS is compatible with “Simple Controlled” area classification limits. Based on these studies, it was decided to shorten the MHS from 4.5 m to 2.3 m. Furthermore, one of the muon shield robustness studies indicated that the core of the MHS could be reduced from 1 m to 0.8 m width and still act efficiently. A preliminary functional specification for the MHS was produced [39] and a first conceptual design study was carried out [37]. Fig. 11 shows the most recent total ionising dose calculation in the region of the MHS, which is used to determine the allowed location of the MHS coil.




Figure 12: Top left: dummy coil shown during insertion into impregnation mold. Bottom left: dummy coil after vacuum-pressure impregnation. Top right: microscope image of windings and epoxy quality. Bottom right: zoom of microscope image indicating complete epoxy filling.

In parallel, an R&D program to assess the feasibility of a 5 T magnet based on non-insulated (NI) gen-2 high temperature superconducting (HTS) coils was launched. With the help of TE-MSR, CERN procured and delivered to KIT a total of 4.7 km of REBCO tape (4.1 km in width of 4 mm and 0.6 km in width of 12 mm, in unit lengths of 100, 150 and 300 m). The R&D program at KIT progressed well, with first dummy coils used to select the correct winding tension giving optimal epoxy impregnation, and in parallel the design and fabrication of tools necessary for producing the HTS test coils. Fig. 12 shows some images related to the studies of epoxy impregnation quality. A 4 mm wide,

87 μm thick dummy tape was used to produce coil series under several controlled winding tensions (10 N, 20 N, 30 N, and 40 N). The coils were vacuum-pressure impregnated using an epoxy type selected for its low viscosity³. In addition, progress was made on developing a model devoted to the understanding of stresses in epoxy-impregnated NI HTS coils based on bespoke analytical/numerical tools. The first series of 4 mm width HTS test coils is currently under production, which involves dedicated programming of the KIT robotic stations for coil winding. A realistic R&D schedule was worked out which, together with the estimated cost of a full SC magnet, indicates that it will be extremely challenging to align the “hybrid” muon shield variant with both the desired start date of the SHiP experiment in 2032 and the injunction to reduce the cost of the initial configuration of the experiment.




Figure 13: Preliminary design of a template magnet with a tunable return yoke consisting of interchangeable blocks of magnetic and non-magnetic material (left) located in the magnet skeleton with the magnet core of high-quality electric steel (right). The base support consists of profiles (I-beams) that are approximately 4.8 m in length.

Currently the most promising muon shield variant is a fully NC option with a tuned return yoke. This approach is described in more details in Section 3.2. The idea is to segment the return legs of the yoke by alternating non-magnetic dense slabs with magnetic iron slabs, and extending those slabs sideways such as to maintain a full flux return in the magnetic iron. This means that individual mid-energy muons (the most difficult ones to treat) see a reduced integral field on their trajectory, while the energy loss is maintained. This method offers the advantage that it will grant some flexibility in the shaping of the magnetic return field and allow inserting muon tracking detector modules in the most interesting regions of the muon shield in order to characterise and understand the muon flux. At CERN, a first preliminary design of a template magnet with a tuned return yoke has been developed. It has allowed starting the discussions with several manufacturing companies. These discussions will be continued to revisit (and optimise) the production cost of such electromagnets. A preliminary view of this concept is shown in Fig. 13.

³Stycast 1266.

A preliminary procurement planning was worked out, indicating that a fully NC variant of the muon shield can be installed before mid-2031, in time for a full participation in the commissioning of BDF with beam in the second half of 2031 (see Section 6). This planning assumes the completion of the full engineering design in early 2027 and the start of the tendering procedure in 2027. A conceptual design review (CDR) is to be held by the end of 2025 in order to compare the studied muon shield variants, evaluate the risks, weaknesses and strengths of each variant, and decide on a single variant to be further elaborated in the context of an engineering design.




Figure 14: Annual neutrino flux on the detector surface at the SND location, as a function of the slope Tx in the horizontal plane.

3.2 Neutrino and LDM Scattering Detector

The neutrino and LDM scattering detector (SND) is designed to profit from the huge neutrino flux of all three flavours available at BDF as reported in Figure 14. The planned physics programme covers a wide range of interests including the cross-section of tau neutrinos and anti-neutrinos, investigation of the strange quark content of the nucleons, and probing lepton flavour violation in the neutrino sector. The detector is also sensitive to the scattering of dark matter particles off the electrons and the nucleons of its constituent detector-target planes.

The timeline set for the start of the SHiP data taking in 2032 makes the fully warm option for the muon shield the preferred solution. This fact, together with the need for a magnetised region in the SND for the measurement of the muon charge and momentum, has driven the idea to embed the detector inside the most downstream magnet of the muon shield. After carrying out optimisation studies for the muon shield, it turned out that the muon rates expected in this region, shown in Fig. 15, exceed by 1 to 2 orders of magnitude the value acceptable by an emulsion detector. On the other side, the emulsion detector can be made very compact and therefore the Collaboration is confident that the allocation of a dedicated space in a strategic position is possible. It will be investigated in the near future. While the identification of a region for the emulsion detector is ongoing, this document




Figure 15: Muon rate at the SND location.

focuses on the fully electronically readout SND detector that will be integrated in the muon shield.

The schematic drawing of the detector is shown in Fig. 16. The detector consists of two parts, the most upstream is a high-granularity calorimeter (SHGC) based on a sequence of Si/W modules in a region which is not magnetised, while the second part is a magnetised tracking calorimeter (SMTC), based on a sequence of magnetised iron slabs interleaved with scintillating detectors providing both the muon tracking and the energy measurement. The detector is very similar and highly synergic with the one currently being built for the high-luminosity upgrade of the SND@LHC experiment [40].




Figure 16: Layout of the SND detector.

Neutrino interactions occurring in both the SHGC and SMTC will be analysed. Those occurring in the high-granularity calorimeter will use the MTC as muon spectrometer and hadronic calorimeter while the MTC can identify and measure neutrino interactions by itself.

The baseline option for the SHGC is the use of tungsten plates, a few mm thick, interleaved with silicon microstrip sensors with a pitch around $100\text{ }\mu\text{m}$, ensuring tracking resolution around $30\text{ }\mu\text{m}$. The tungsten thickness is being optimised, and it is a trade-off between two conflicting requirements, i.e. maximising the detector mass and the wish to increase the segmentation, thus improving tracking/vertexing performance and electromagnetic energy resolution. The energy deposited in the SHGC, essentially all the electromagnetic component with a part of the hadronic one, must be measured in this detector. Given its large density, tungsten is the best choice for the passive material. Its short radiation and interaction lengths also make the shower development more confined, thus also improving the containment of hadronic showers. The high-granularity and high spatial resolution together with the large density is a key ingredient in the identification of light dark matter scattering off electrons, where a single electron appears in the final state, leading to an electromagnetic shower. The compactness of the shower development in tungsten helps avoiding overlaps with other particle tracks at the primary vertex, thus suppressing the background to light dark matter searches from inelastic interactions. The configuration foresees 120 tungsten plates with a cross-section of $40 \times 40\text{ cm}^2$. For 3.5 mm thick W plates, a mass of 1.3 tonnes is obtained. Since the detector develops over more than four interaction lengths, neutrino interactions without muons in the final state, including tau neutrinos, occurring in the upstream part can be reconstructed almost entirely within the SHGC.




Figure 17: Simple geometrical configuration for the first solution of the most downstream magnet of the muon shield to host the SND detector. The field map in the iron is shown in the two insets.

The baseline option for the SMTC is the use of 5 cm thick iron slabs, magnetised, and interleaved with scintillating fibres (SciFi), with $250\text{ }\mu\text{m}$ diameter, and 1 cm thick plastic

scintillator pads. A total of 42 planes are envisaged with transverse size progressively increasing from $40 \times 40 \text{ cm}^2$ to $60 \times 60 \text{ cm}^2$. The 42 iron slabs amount to more than 10 interaction lengths thus providing a good hadronic calorimeter. The transverse size of the pads is $5 \times 5 \text{ cm}^2$, providing a coarse granularity and avoiding saturation effects in the energy measurement. The main function of the pads is the energy measurement while SciFi provides good tracking capabilities for isolated tracks, i.e. muons and some hadrons. The muon tracking is essential to provide its charge and momentum measurement. The total mass of the iron slabs is 4.2 tonnes.




Figure 18: (Left) Schematic drawing of the second solution, with tuned return yoke, identified for the most downstream magnet of the muon shield to host the SMTC. For symmetry reasons only one quarter of the system is shown. (Right) Field map of the configuration shown in the left figure. Given the symmetry, only one quarter of the magnet is shown.

The first challenge faced during the first year of the TDR phase was to identify engineering solutions to embed the SND detector inside the muon shield while keeping its full muon sweeping functionality. The challenge was to make sure that all the detector modules and their ancillary services would be accessible for installation and maintenance. Two solutions have been identified, both with a common element: the upstream part hosting the SHGC foresees the magnetised region only outside the detector volume, with the shape illustrated in Fig. 17. In this compact region (about 1.5 m in length), the outer legs of the muon shield continue playing their function to sweep out high-energy muons while avoiding the return towards the centre of softer momentum muons. For the SMTC region, which is about 3.5 m long, there are two different solutions developed. One consists of connecting the outer magnetic circuit, where the main functionality of the muon shield is in place, with the inner part hosting the SMTC, without perturbing the magnetic field in the outer part. The field map of this solution, technically viable, is shown in Fig. 17 in the top-right inset. The field is rather uniform and the integrated field is above the design value of 3 Tm.

Although this solution is still under study and the field uniformity and strength could be improved, another option seems particularly suitable. Since high-energy muons have already integrated a $P_{\perp}$ kick sufficient to be swept out of the decay volume, in this magnetic section the main function of the outer legs is to deflect outwards the mid-momentum muons and continue degrading their total momentum by energy loss. This is better achieved with a sideways extended region of dense material with moderate magnetic field. This concept

seems particularly suited to host a detector in the central region where a saturation field is needed while diluting the field in the outer legs: the space left for the detector can be accommodated and the field in the outer legs diluted accordingly. An illustration is given in Fig. 18(left). The nature of the non-magnetised material and the ratio of its thickness to the iron slabs is subject to the ongoing optimisation. The field map reported in Fig. 18(right) shows that the value in the core, where the SMTC is located, is sufficient and that the desired field in the outer legs is achieved.

Table 2 reports the expected yield of neutrino and anti-neutrino interactions for the three different neutrino flavours in the SHGC and in the SMTC.

	CC DIS (W)	Charm CC DIS (W)	CC DIS (Fe)	Charm CC DIS (Fe)
N_{ν_e}	6.8×10^4	4.1×10^3	1.6×10^5	9.8×10^3
N_{ν_μ}	2.0×10^5	8.7×10^3	4.6×10^5	2.0×10^4
N_{ν_τ}	2.1×10^3	1.2×10^2	5.1×10^3	2.8×10^2
$N_{\bar{\nu}_e}$	1.4×10^4	7.4×10^2	3.7×10^4	1.9×10^3
$N_{\bar{\nu}_\mu}$	4.4×10^4	1.7×10^3	1.1×10^5	4.1×10^3
$N_{\bar{\nu}_\tau}$	1.5×10^3	8.6×10^1	3.8×10^3	2.1×10^2

Table 2: Neutrino interaction yields per year in the SND High Granularity Calorimeter, and in the Magnetised Tracking Calorimeter.

3.3 FIP Decay Detector

3.3.1 He decay volume

The decay volume for the FIP decay search is located immediately downstream of the muon shield. The geometry and dimensions of the detector volume that define the decay acceptance have been obtained by an optimisation based on a wide range of FIP physics models and particle masses, the performance of the muon shield, and the aperture of SHiP’s decay spectrometer. The resulting fiducial decay volume delineates a pyramidal frustum with a length of 50 m and upstream dimensions of $1.0 \times 2.7 \text{ m}^2$ and downstream $4 \times 6 \text{ m}^2$.

Originally, the decay volume consisted of a vacuum vessel with a pressure of $< 10^{-2}$ bar in order to suppress the background from neutrinos interacting in the fiducial volume. Studies undertaken since the start of the TDR phase have shown that this was overly conservative and that it is possible to replace the heavy-walled vacuum vessel with a helium-gas balloon of the same geometry and still guarantee the background requirements. This reduces enormously the amount of material around the decay volume, and hence interactions of muons and neutrinos with the walls, and removes entirely the need for a vacuum vessel throughout the spectrometer. Instead, the spectrometer tracker can be located in air downstream of a thin helium membrane at the end of the decay volume. The design change leads to a major reduction in complexity, development time and cost, making the implementation realistic for operation in Run 4.

As previously, the decay volume must be instrumented with a system of background taggers upstream and on the sides described in Section 3.3.2 and 3.3.3, respectively.

The decay volume geometry results in a balloon with a total area and volume of about 685 m^2 and 543 m^3 , respectively. The expected helium pressure is 101325 Pa plus 20 to 50 mbar to maintain the required geometry. The helium purity shall be kept above 99% with the help of a helium purification system.

The liner material, more specifically its gas diffusion rate, has been identified as the key parameter to, firstly, estimate the helium losses and, secondly, to size the purification system. Particularly critical is the thickness of the liner at the front of the decay volume and at the end towards the spectrometer tracker to control multiple scattering between the background taggers and the spectrometer tracker. Several manufacturers were contacted to list potential candidate materials and identify the best fit. While the He and air diffusion rates together with the wall thickness shall be kept as low as possible, the selection criteria of the liner material will also take into account its manufacturability (welding, balloon handling, safety) for the given geometry, cost and fire load. Several sample multilayer samples based on fabric matrices with polyurethane, PVC and nylon layers (among others), were delivered to CERN to measure their helium, nitrogen and dry air diffusion rates, including CO_2 . A measurement campaign is ongoing in the permeability test set up of the TE-VSC group (see Fig. 19). The samples are glued by the CERN Polymer Laboratory on a copper gasket. The set up is now running and diffusion rate measurements of at least 14 samples will be completed in the next months. The samples thickness ranges from 1.44 mm to 0.12 mm.




Figure 19: Diffusion rate measurement set up.

Several technologies that are already available on the market have been identified for the helium purification system, such as membrane, Pressure Swing Absorption (PSA) and cryogenic purifiers. The required 99% helium purity excludes, a priori, the need for the cryogenic purifier as it is commonly used to obtain purities at ppm level and comes at a higher cost and maintenance. Therefore, membranes and PSA technologies are under

investigation, the last of which is the preferred based on its higher purification efficiency and following the discussion with several manufacturers. The visit to the Rutherford Appleton Laboratory Helium recovery system in United Kingdom confirmed that the PSA can reach a helium purity above 99.7% at large scale with two 1 m long molecular sieves and a charcoal 1 m absorber.

Fluid dynamics processes and helium versus impurity stratification inside the balloon are to be understood. To do so, first analytical computations with dimensionless numbers were performed. As a small temperature difference between the top and bottom of the balloon is expected due to the temperature gradient in the heat radiated from the detector electronics and the ECN3 ventilation system, a natural convection flow inside the balloon is expected. For a more detailed study, first steps in building a 3D Computation Fluids Dynamics model in ANSYS Fluent were carried out. A first simple model to analyse the natural convection inside the balloon with a closed volume and a ΔT along the height of (first guess) $0.5^\circ\text{C}/2\text{m}$ is being worked out. Helium, oxygen and nitrogen losses through the balloon wall will be included to see the helium concentration evolution and pressure decrease over a month. Later, a circulation system with the helium purification system will be included, that is, a flow rate injection of fresh helium and extraction of a contaminated He-O₂-N₂ mixture. This model will help to optimise the location and number of the purification connection flanges along the balloon and the gas analyser output flanges along the geometry.

The instrumentation necessary to monitor the balloon has been identified, including a gas mass spectrometer analyser with a multi-port valve to monitor the concentration and balloon contaminants, pressure sensors and Oxygen Deficiency Hazard (ODH) sensors.

With regards to the supporting structure, it shall withstand more than 500 kg buoyancy force due to the helium balloon together with the expected mass loads from the SBT and UBT self-weights. A carbon fibre versus aluminum trade-off study is underway to reduce as much as possible the mass of the structure along the sides of the decay volume.

A prototype that includes a supporting structure, balloon, purification system, and instrumentation is planned in 2026 to validate the Computational Fluid Dynamics model, ensure the understanding of pressure and helium losses, and to confirm the balloon material selection and purification technology.

3.3.2 Upstream Background Tagger

The upstream background tagger (DUBT) is placed between the muon shield and the decay volume, as shown in Fig. 7. It should provide position and time information on the residual charged particles that enter the decay volume. They may come from two sources. Firstly charged particles, mainly residual muons, may lead to background in the form of fake vertices formed from accidental combinations of charged tracks in the decay volume. Secondly, charged particles may be used to flag neutrino or muon inelastic scattering in the last interaction length of the muon shield, potentially leading to physical V^0 s in the decay volume.

A study of the combinatorial muon background for different muon shield configurations has been performed [38]. Fig. 20 (left) shows that the UBT detector should have lateral coverage of up to $\pm 2\text{m}$ to be able to detect all the background muons reaching the decay spectrometer. The studies show that excellent time and spatial resolution are essential to

guaranty the required tagging efficiency. As demonstrated in previous experiments, a time resolution of 200 ps can be achieved with scintillating tiles of $5 \times 5 \text{ cm}^2$. With this granularity, the detector must be able to cope with a rate of charged particles of up to 10 MHz per tile in the region of the “muon flames”, located at $x \sim \pm 2 \text{ m}$ at the level of the entrance of the decay volume (see Fig. 20 right). Given the need for tracking with a spatial resolution of $\sim 100 \text{ } \mu\text{m}$, and $>99\%$ efficiency and low-mass, the choice for the DUBT consists of a combination of the scintillating tile timing detector and a dedicated tracker based on straw tubes.




Figure 20: (left) Muon rate per 16 cm in x -coordinate, integrated over y -coordinate. The green histogram corresponds to the full rate; the salmon – to the rate of muons, reaching HSDS and forming a good common vertex; the red – to the rate of muons, bouncing off the walls and reaching the decay spectrometer. (right) Muon hit map at the DUBT position, rate per $5 \times 5 \text{ cm}^2$ area. Figures from ref. [38].

DUBT Timing - Scintillating tiles

To achieve the necessary time resolution and cope with the occupancy and rate, a timing detector formed out of scintillating tiles is considered. The typical cell size of the scintillating tile is $5 \times 5 \times 1 \text{ cm}^3$. Transverse dimensions can be reduced for the tiles close to the “muon flames” to correspondingly decrease the muon rate per channel below $\sim 1 \text{ MHz}$. The thickness of the scintillating tiles will be optimised to minimise the amount of material in the acceptance. The thickness is driven by the necessity to have a large enough light yield to achieve the required time resolution.

At the moment, two solutions are considered for the scintillating light readout, as shown in Fig. 21. In the first, the scintillating tiles are coupled with a WLS fibre passing through and read out on each side with silicon photomultipliers (SiPMs). The advantage of this approach is in the lower amount of material in the detector acceptance and the lower number of SiPMs necessary to equip the detector. On the other hand, in this way, the benefit of the tile granularity is lost, and the rate seen by the SiPM has to be integrated along the fibre. In the second solution, each tile is read out by a dedicated SiPM, pressed to it. In this case, the pixelated granularity of the detector is maintained, and the individual SiPM rates are

controlled by the size of a tiles. However, this comes at the expense of having to equip the detector with long flex cables, transmitting the SiPM signal to the readout electronics, and with the necessity to have a larger number of SiPMs and readout channels.




Figure 21: DUBT timing layer design with scintillating tiles, either (left) coupled with a WLS fiber, read out at the two ends, or (right) read out individually with a SiPM.

The charged particle rate to which the detector is exposed poses a stringent requirement on the readout electronics. These rates are compatible with the demands of the proposed LHCb Upgrade II, and therefore some of the solutions developed there can be adopted for the DUBT, with the only difference being in the asynchronous readout in the SHiP experiment. The PACIFIC++ chip, designed to read out the SciFi detector in the LHCb Upgrade II can be adapted for these purposes. Other solutions include TOFPET2 ASIC, already used in the SND@LHC for the SciFi detector readout, or FastIC+ ASIC, developed for a Ring-imaging Cherenkov detector readout in LHCb, and also considered for SHiP's SBT and spectrometer timing detector. To account for the high count rate in the muon flames, an FPGA-based TDC with a 100-200 ps time bin may also be an option, as it provides fast timing acquisition and a relatively simple implementation.

DUBT Tracking - Light straw tubes

The main requirements for the tracking system of the DUBT are very low mass and high efficiency. Ideally, no additional neutrino or muon interactions (< 1) should be generated in the acceptance from 6×10^{20} PoT and the tracking efficiency should be close to unity (more than 99.9%). This substantially restricts the choice of technology. The proposed solution is a low-mass straw detector, operating approximately 20 mbar above atmospheric pressure.

Straws are limited in their maximum rate of charged particles and total flux. For example, ATLAS TRT with straws of 4 mm diameter can operate up to 0.5 MHz/cm^2 [41] without loss of efficiency. Given the predicted muon flux (Fig. 20), the straws outside of the central region will rapidly become overwhelmed. Thus, the active region of the UBT straws is restricted to the entrance window of the decay volume ($1.3 \times 2.7 \text{ m}$).

The efficiency, $\epsilon = 1 - \delta$, of a single straw to a muon traversing its active area can exceed

99%, [42]. There will also be, in general, a dead space between straws, d_{dead} , of diameter d_{straw} . Hence, a minimum of four layers is assumed to construct a track. For $d_{\text{straw}} = 5$ mm, $d_{\text{dead}} = 1$ mm, and $\delta = 0.5$ %. In this case, the expected efficiency for tracking a muon traversing a six-layer detector will, with suitable staggering, exceed $\epsilon_{\text{track}} = 99.99\%$. Additional layers would increase this efficiency.

Straw material, e.g. for mu2e/g-2m, is dominated by the thickness of the wall material. A particle, at normal incidence, will traverse approximately 0.02% of a radiation length, Table 3. In a six-layer detector this amounts to a total of 1.2% of a radiation length. For reference a muon rate of 100 Hz/cm², integrated over the active area of a six-layer detector, would generate an additional 3 MHz of particles into the decay volume. Additionally, in a six-layer detector, exposed to a total of 10²⁰ PoT, can be estimated to produce approximately 100 neutrino interactions in the DUBT, some of which will occur in the latter half of the detector, making detection (from tracking) harder. This further emphasizes the importance of minimising material. Simulation is being performed to quantify the impact of the DUBT material on the muon and neutrino interactions.

Layer	Material	Thickness [μm]	Fraction of X_0 [%]
Outer coating	Aluminum	0.03	~ 0.000034
Outer substrate	Mylar (PET)	6.0	~ 0.0029
Adhesive (bonding)	Epoxy (laminate)	2.5 (nom)	~ 0.0007
Inner substrate	Mylar (PET)	6.0	~ 0.0029
Inner coating (Al)	Aluminum	0.02	~ 0.000022
Inner coating (Au)	Gold flash (over Al)	0.02	~ 0.00060
Total wall	—	≈ 14.6	~ 0.007
Sense wire(Au plate)	W-3% Re,	25 μm diameter	~ 0.001
Gas volume	Ar: Ethane (50:50)	—	~ 0.003

Table 3: Muon $g-2$ straw laminate using Liverpool numbers (epoxy $\approx 2-3 \mu\text{m}$; $2.5 \mu\text{m}$ nominal). Radiation-length fractions use $X_0(\text{Mylar})= 20.6$ cm, $X_0(\text{Al})= 8.9$ cm, $X_0(\text{Au})= 0.334$ cm, epoxy $X_0 \approx 37$ cm. Note each muon traversing a straw will see two layers of the straw, on entrance and exit.

A design [UBT0] for the straw DUBT is being optimised. In addition to minimising mass other considerations are being taken into account.

- Conservative design that reduces risk
- Ease of construction
- Lowering cost and enabling early prototypes

The UBT0 concept is to build straw planes approximately 1×1 m in size, from which the larger 3×1 m detector can be constructed. These planes would be made of tensioned, rather than freely supported, straws. This builds on our experience of building straw detectors. It also ameliorates the problem of dealing with creep in long straws [43]. We would aim to tension the wires at approximately 50 g and the straws at about 500 g. The choice of shorter, lower-tension, straws also permits us to reduce straw occupancy. However, this then dictates

the use of horizontal rather than vertically oriented straws. This is because the end of the straws are structural elements to support tension. Vertically oriented short straws would require undesirable material in the sensitive volume.




Figure 22: A view of the straw DUBT superimposed on the muon background. The 3×1 m area is composed of six layers of approximately 1×1 m modules.

As shown in Fig. 22, the central detector area (1.3×3.0 m²) is covered by modules of approximately 1×1 m². The layout consists of two horizontal layers and two additional layers oriented at $\pm 15^\circ$. Each module contains two staggered straw layers with identical orientation. With 5 mm diameter straws and 1 mm spacing, each module is 1152 mm wide and includes 384 straws. During installation, the total module tension of approximately 2500 N is transferred from the module tensioning frame to a larger detector tension frame, which can be mounted on the muon shield. This design provides a clear active area without the need for excessively long straws.

Currently, we are investigating the influence of the tensioning frame material and the straw manifolds—which define the straw positions—on detector performance. Further optimisation may be required. An open question is whether the additional material could significantly affect the physics measurements. If this proves to be a serious issue, an alternative design using longer (3 m) vertically oriented straws may need to be considered; however, this would substantially increase the expected construction cost.

3.3.3 Surrounding Background Tagger

The Surrounding Background Tagger (DSBT) is an essential part of SHiP’s background tagger system. Its main purposes are: 1. Tagging inelastic interactions of muons and

in particular of neutrinos on helium inside the decay volume by detecting the outgoing particles from the inelastic interaction. 2. Tagging of muons entering the decay volume, complementing the DUBT background rejection power.

The SHiP DSBT will be built from cells consisting of an AlMg4.5 alloy (EN AW-5083) [44] filled with liquid scintillator (LS). Each cell is equipped with two Wavelength-Shifting Optical Modules (WOMs), which absorb the UV scintillation photons and guide the secondary photons, emitted in the blue range, to an array of 40 SiPMs, grouped in eight groups of five SiPMs. The LS consists of linear alkylbenzene (LAB) enriched with 2 g/l 2,5-diphenyloxazole (PPO). Self-absorption is mitigated by optical purification of the LAB using column filtration with alumina (Al_2O_3) [45, 46]. The inner cell walls are mirror polished in order to increase the integrated light yield collected by the WOMs [44]. To guarantee large detection efficiencies also for neutral particles, produced in inelastic interactions on helium, the thickness of the liquid scintillator is chosen to be 20 cm. The dimension of all cells along the detector axis is 80 cm throughout the length of the decay volume while their size perpendicular to the axis depends on the position of the cell and goes up to 150 cm.

Test beam measurements with a polished AlMg4.5 cell of $110\text{ cm} \times 80\text{ cm} \times 20\text{ cm}$ shows excellent performance in terms of light-yield collection [44]. The time resolution of about 1 ns achieved in previous test beams with different prototypes have also been confirmed with this prototype [46, 44]. Reconstruction of a particle crossing point in a cell uses three informations: 1) the light-yield ratio of both WOMs, 2) the difference in signal arrival time between both WOMs, 3) the light yield distribution across the eight SiPM groups [46, 44]. This correction method allows to obtain good homogeneity in the light yield and time response across the detector cell (A): Fig. 23 shows the mean integrated signal yield and Fig. 24 the average signal arrival time ($\bar{T}_{12}$) of the two WOMs for different distances between the muon crossing point and the cell centre.

To prepare for mass production of the WOMs – PMMA tubes that are dip-coated twice with the wavelength-shifting dye, a quality control (QC) set-up has been designed and constructed, which is shown in Fig. 25. The QC set-up allows to position via a computer-controlled motor an LED and diode (in the future a SiPM) along the WOM tube to measure the transmitted light as well as the emitted light at the end of the WOM tube using a large-area PMT.

The LS cells are parts of larger containers built from 6 m (preliminary) long polished AlMg4.5 sheets that are then bent to shape the container sidewalls. The individual cells are obtained by adding thin AlMg4.5 separation walls inside the container, which is then closed by welding a polished AlMg4.5 sheet. The containers are then used to cover the surroundings of the decay volume/helium balloon, with I-beams as their support structure. This modularity provides flexibility in case not all parts of the DSBT are instrumented at the start of the data taking in 2032.

The LS transparency as a function of time continue to be monitored including ageing tests.

During the R&D phase of the DSBT, the development of the photon detector modules focused on optimising the detector geometry, reflectivity and the light collection. The electronic signal processing of the WOMs was mostly kept unchanged for better comparability.

For the detector prototype, impedance-matched printed circuit boards (PCBs) were designed to host 40 Silicon Photomultipliers (SiPMs), arranged in 8 groups of 5 SiPMs




Figure 23: Mean integrated yield as a function of distance R between the particle crossing point and the centre of the DSBT cell for the respective prototype: Cell A (*top*), Cell B (*middle*), Cell C (*bottom*). The error bars are given by the standard deviation of the integrated yield distribution. Dark squares represent the measured values, light circles the same measurements after the likelihood-based correction, with the horizontal line indicating the respective reference value. For visual purposes, corrected values and points with the same R value are shown at a slight horizontal offset.

each. Different types of SiPM like Hamamatsu S14160-3050HS SiPMs, Broadcom AFBR-S4N44P014M or SensL J-30035-TSV have been tested. While each type has shown certain advantages in specific features, the Hamamatsu SiPM has been selected for future studies for its best overall performance. According to the manufacturer’s datasheet, overvoltages of 3.7 V correspond to a photon detection efficiency of about 55 % for the Hamamatsu S14160-3050HS. The circular SiPM-PCB was designed to match the diameter of the WOM tube to ensure optimal optical coupling. During assembly, particular care was taken to achieve uniform SiPM height above the PCB to ensure equal optical contact conditions. For the detector characterisation, the SiPMs were optically coupled to the WOM using optical grease, while in the final detector configuration they will be permanently glued using Evo-Tec 301-2 glue. A dedicated gluing station has been developed for this purpose to ensure reproducibility and precise alignment.

The SiPM-PCBs are impedance-matched and connected to the front-end amplifiers via high-density connectors to preserve signal integrity. For the prototype characterisation, all detectors were read out using the eMUSIC8 amplifier-shaper ASIC distributed by Scientifica. The analogue signals were subsequently digitised with a WaveCatcher system at a sampling rate of 3.2 GSa/s. This setup provided high-resolution waveform information, allowing detailed studies of signal shapes, timing characteristics, and photon collection efficiency for different detector geometries and reflective coatings.

Figure 26 shows representative waveforms recorded with the WaveCatcher and normalised




Figure 24: Estimated particle arrival time $\bar{T}_{12}$ at the detector, as a function of distance between the particle crossing point and the centre of the cell for three tested DSBT prototypes. Before (*left*) and after (*right*) application of the likelihood-based correction. The error bars represent the standard deviations of the distributions, the shaded bands depict the time variation over the entire cell. The offset in average time between the three cells is arbitrary.

in amplitude. Despite the varying detector geometries, reflective properties, and measurement positions (central versus detector corner regions), the recorded signal shapes were found to be remarkably similar. This observation is crucial for the definition of the future readout strategy.

Based on the uniformity of the signal shapes, the future readout concept for the DSBT can be simplified significantly. Instead of performing high-frequency waveform digitisation and subsequent charge integration, the integrated signal charge can be inferred from a Time-over-Threshold (ToT) measurement. This approach substantially reduces data volume and power consumption while maintaining the required precision in time and amplitude determination.

To implement this concept, a new readout board based on the FastIC+ ASIC [47] has been developed. The FastIC+ is a highly versatile, low-noise front-end chip designed for precise measurement of the Time-of-Arrival (ToA) and energy deposition of photon signals from various types of photodetectors, including SiPMs, PMTs, and MCPs. It features a low-input impedance analogue front-end ($30\ \Omega$) that allows direct connection to the SiPM anodes or cathodes with programmable polarity. Each ASIC supports up to eight single-ended channels, which can be configured individually or summed in groups of four (SUM4 mode) to optimise signal-to-noise ratio and timing performance.

The FastIC+ generates three parallel signal branches per channel:

- Timing branch: Processes the fast discriminator output to measure the photon time-of-arrival with a precision of 25 ps using an internal Time-to-Digital Converter (TDC).
- Energy branch: Determines the deposited energy through a Time-over-Threshold (ToT) or peak-and-hold method, providing either linear or non-linear ToT encoding depending on the configuration.




Figure 25: (Left) CAD design of the Quality-Control (QC) setup to prepare for the DSBT WOM mass production. (Right) Photo of the QC setup.

- Trigger branch: Produces programmable low-level and cluster-level triggers, allowing suppression of dark counts and enabling external or self-triggered acquisition modes.

The output stage of the FastIC+ can be configured to provide either binary signals (encoding ToA and/or ToT information) or analogue outputs for diagnostic purposes. Binary outputs can be transmitted via Scalable Low-Voltage Signalling (SLVS) or CMOS, with SLVS preferred for low-jitter, high-precision timing. The ASIC can be fully configured via an I²C interface, allowing per-channel adjustment of bias voltages, thresholds, and shaping parameters.

A key feature of the FastIC+ is its ability to combine the time and energy information within a single digital pulse, where the rising edge encodes the time-of-arrival and the pulse width corresponds to the energy (ToT). This enables compact, low-bandwidth data encoding suitable for large-scale detector systems such as the DSBT. The first prototype boards equipped with the FastIC+ ASIC will undergo beam tests at the CERN PS T9 beamline in November, where the full timing and ToT performance will be characterised under realistic operating conditions.

In summary, the transition from waveform sampling with the eMUSIC8 and WaveCatcher system to a ToT-based readout with the FastIC+ represents a crucial step towards a scalable, radiation-tolerant, and power-efficient readout concept for the DSBT. The ongoing development aims to establish a compact front-end architecture capable of precise photon time-stamping and amplitude reconstruction, fulfilling the stringent requirements of the SHiP experiment at the CERN Beam Dump Facility.

In the near-future plans, there will be the next testbeam campaign in november 2025 at CERN's PS T9 to study:

- long-term performance with the same prototype detector cell by comparison with former measurements




Figure 26: Averaged signal response for energy loss measurements of minimum ionizing muons for different signal generation positions, detector geometries, and inner surface treatments. Due to signal shaping in the readout electronics, the signals differ only in amplitude but not in shape.

- the light yield and the light yield share between WOMs when enlarging detector cells by 50%
- for the first time SiPM-arrays glued to the WOMs for optimal optical coupling
- further improvements concerning detector mechanics
- new front-end electronics utilizing the new FASTIC+ ASIC

3.3.4 Spectrometer Straw Tracker

The spectrometer straw tracker (DSST) will be responsible for measuring the momentum of particles with high accuracy. It is located behind the decay vessel and will consist of four identical tracking stations, two stations on each side up- and downstream of the spectrometer magnet. The stations will be equipped with straw tubes. Each station will have four views, arranged in a y - u - v - y setup, with two horizontal-orientated planes y and two stereo planes u and v , tilted at a small stereo angle of about 5° for 3D reconstruction.

The spectrometer tracker has undergone several changes since the start of the TDR phase. Most importantly, the spectrometer will now be placed in air. This simplifies the design and construction significantly. For the placement in air, the exact positioning of the tracker has been reevaluated and are depicted in Fig. 27. The stations (centre) will be placed at -5.5 m, -3.5 m, 3.5 m, and 5.5 m with respect to the magnet centre. The updated design has been fully implemented in the FairShip software. Background simulations performed with a preliminary set of coordinates and frame design are foreseen to be updated as soon as possible. The previous simulations showed moderate background rates of up to

~ 10 kHz per straw, however they need to be re-run due to the multiple layout changes in the upstream detectors, as well as a possible different choice for the frame material. For this, the straw tracker geometry definitions have already been moved to yaml, which makes the implementation of different parameters such as the frame layout or stereo angle easier. Updated simulations will be run for the baseline option of a stainless-steel frame, with alternative options such as aluminum or carbon fibre to be evaluated.




Figure 27: Layout of the spectrometer straw tracker. Four identical stations are placed up- and downstream of the spectrometer magnet. The two inner stations will be placed inside the magnet pit.

The frames themselves consist of two parallel I-beam structures that extend over 1 m with a 40 cm gap in the middle. The straw tubes will be assembled in a modular concept outside CERN. They are then inserted into the frame from the top by crane and accessed from the side. A module will comprise two layers of sixteen straws, and will contain a support element for mechanical robustness. For this support, several options are under study, such as carbon fibres or thin carbon plates.

Straw-tube module prototyping is about to start, with two types of straw tubes under consideration. The first option is based on the NA62 design: longitudinally ultrasonically welded mylar tubes, with a wall thickness of $36\text{ }\mu\text{m}$ and a coating of 50 nm copper and 20 nm gold. Tubes with 2 cm diameter and 4 m length have been produced and are currently awaiting shipment from INP, Almaty, Kazakhstan to the University of Hamburg, Germany. In Almaty, a site for mass production of straws is currently being setup, expected to be operational by the end of the year. As an alternative option, commercial straws produced by Lamina in the UK are under consideration. These straws are produced through winding and come with a wall thickness of $30\text{ }\mu\text{m}$ and an aluminum coating of 100 nm on both the

inner and outer surface. Straws of this kind have been ordered by the University of Hamburg and are expected to be delivered in early November 2025.




Figure 28: Ongoing test beam efforts for the DSST readout with the Sampic chip. (left) Straw tracker in the test beam telescope. (right) Dedicated pre-amp developed for the Sampic readout.

Progress has also been made on the evaluation of different readout options. Two options are under consideration: the use of the Sampic chip developed at CEA Saclay and LAL Orsay / IJCLab in France, or the ASD chip developed for the ATLAS muon system. Test beam measurements with the Sampic chip on a smaller prototype, using a custom-developed discrete pre-amplifier look promising. The chip features a self-triggering readout and allows for a digitization of the signal leading edge with high precision timing. Thus, a higher threshold can be set, and the signal arrival time can then be determined offline with sub-nanosecond precision. The chip also features a time-over-threshold measurement, which can be used to reduce noise. Integration of the Sampic chip in the overall DAQ is expected to be relatively straight-forward, making use of possible synergies across sub-systems that are also considering using this chip. The developers have been active in SHiP for several years already. Fig. 28 shows the small prototype during the test-beam measurements together with the developed pre-amplifier. The second readout option, the ASD chip, is available at MPI Munich [48]. Recent developments of an upgraded chip include a shorter peaking time, leading to an enhancement in the spatial resolution and the capability to handle high rates [49]. Test beam measurements using the original chip have been performed in conjunction with the straw tube R&D in the scope of DRD1 [50], including tracking tests with stereo layers. The analysis of these measurements is ongoing.

3.3.5 Spectrometer Magnet

The development and construction of the SHiP spectrometer magnet follow a two-stage project plan. Stage 1 covers a proof-of-principle program called Energy-Efficient Superferriic Dipole (EESD)[51] and Stage 2 covers the design, prototyping, production and installation of the full-scale magnet in the experimental hall of ECN3 with all its ancillaries of which the main are the cryogenic satellite, the cryogenic transfer lines, the vacuum pumps, the power converters, and the magnet protection system.

The proof-of-principle program comprises the design, construction and testing of a demonstrator, and a reduced-scale magnet, called respectively EESD1 and EESD2. The coils of

these magnet models are made of round cable of 9.4 mm diameter comprising eighteen MgB_2 wires of 1 mm diameter twisted around a central stranded copper core of 4.9 mm diameter. This cable was developed, manufactured and extensively tested for the Superconducting Link of the High Luminosity LHC Project (LH-LHC) in the framework of WP6a [51]. The program is split into five phases as follows.

1. Phase I: Design and manufacture of a superferic dipole magnet model, EESD1. This demonstrator comprises a single coil pack assembled with an external iron yoke, without cryostat. The whole structure is cooled down and tested in liquid helium (LHe) at 4.2 K.
2. Phase II: Adaptation of EESD1 to retest it in gaseous helium (GHe) at 20 K.
3. Phase III: Design and manufacture of another superferic dipole magnet model, EESD2. This model comprises two coil packs housed in a common cryostat surrounded by an external iron yoke operating at room temperature (RT). The existing infrastructure of Cluster G in SM18 will be adapted for the cold tests to provide a forced flow of GHe at 20 K for the magnet cooling, at 75 K for the thermal shield cooling, and current via Nb-Ti current leads.
4. Phase IV: Adaptation of the coil cooling circuit of EESD2 to demonstrate the use of a cryocooler and a cryofan to provide forced GHe flow at 20 K for magnet cooling.
5. Phase V: Implementation of dedicated current leads and ancillaries to enable a stand-alone operation of EESD2 as foreseen for the SHiP spectrometer magnet.

The SHiP spectrometer magnet will feature two superconducting coils made from the aforementioned MgB_2 cable, or possibly an advanced version of it. The coils will be indirectly cooled by GHe at 20 K circulating in cooling pipes embedded in the mechanical structure of the coils. GHe will be provided by means of cryocoolers and cryofans, which will be part of the cryogenic satellite. There will be four cryogenic circuits, one for the current leads at 55 K, two for the magnet at 20 K (one for each coil), and one for the thermal shield at 75 K. A simplified version of the cryogenic satellite will be needed for Phase IV and Phase V.

Phase I has been completed in the third quarter of 2023. The design of EESD1 and the results of the cold tests are extensively described in [51]. Phase II has been completed in the second quarter of 2024. The results of the cold tests are extensively described in [52].

Work on Phase III is well advanced, with the design and procurement of components for EESD2 now complete. Assembly work has begun in the Large Magnet Facility (B180), starting with a blank assembly of the coil structure, thermal shield and vacuum vessel, as shown in Figure 29. The winding of the coils will be coming next. It is planned to complete the construction of EESD2 and its installation on the test bench in SM18 by end 2025. The cold tests will be carried out in Q1-2026, when the cryogenic shutdown is finished.

The main parameters of the demonstrator EESD1, the reduced-scale magnet EESD2, and those of SHiP, are given in Table 4.

Figure 30 shows the electromagnetic design of the SHiP spectrometer magnet, which features two coils in a vertical position (shown in red) that generate a horizontal dipole field




Figure 29: TOP: CAD views of EESD2 magnet assembly (LEFT) and EESD2 cold mass assembly (RIGHT) - BOTTOM: Pictures of the blank assembly of the coil structure, thermal shield and vacuum vessel in the Large Magnet Facility (B180).

Magnet	EESD1	EESD2	SHiP
Status	Tested @ 4 K (Phase I), and @ 20 K (Phase II)	In fabrication	Conceptual design started
Vertical gap (mm)	61	60	6000
Horizontal gap (mm)	180	180	4000
Iron mass (t)	4.5	5.2	660
Operating current (kA)	5	5	3
Total Nr of turns	24	24	192
Average turn length (m)	3.4	4.0	21.2
Cable unit length (m)	81	48	680
Cable total length (m)	81	96	4080
Peak field (T)	0.941	0.544	1.004
B0 @ (0,0,0) (T)	1.721	1.729	0.154
Stored energy (MJ)	34.5	28.3	2750

Table 4: Comparison of the magnet parameters for EESD1[53], EESD2[54], and SHiP.




Figure 30: Baseline parameters of the electromagnetic design of the spectrometer magnet.

across a $4 \times 6 \text{ m}^2$ aperture. The nominal magnetic field is 0.15 T at the centre and 0.65 Tm integrated along the centre line. A peak field of 1 T occurs in the horizontal legs of the coils.

The SHiP experiment requires the magnet's bending strength to be sufficient to ensure the required spectrometer resolution. To achieve this, at least 95% of the magnet aperture must experience an integrated field of at least 0.65 Tm in the longitudinal direction ($\pm 3.5 \text{ m}$). However, the integrated field should not exceed 1.5 Tm, since this would cause particles to move away from the experimental aperture. A parametric study focusing on the operating current, the overall height and length of the coils and the geometry of the iron yoke is being conducted in order to optimise the electromagnetic design and achieve the specified requirements. It is also necessary to reduce the total weight of the iron yoke in order to not exceed the load capacity of the concrete slab in the ECN3 experimental hall. Iron plates from the OPERA experience, made from grade S235JR, may be used for the construction of the yoke. Relative magnetic permeability measurements and tensile tests will be carried out to support the electromagnetic design and structural analysis. The magnet's quench protection could be ensured by a $30 \text{ m}\Omega$ dump resistor. This will need to be verified once the electromagnetic design is complete.

Integration studies of the full-scale spectrometer magnet for SHiP are well underway, with some decisions having been made. For example, it has been decided to install the power converters inside the ECN3 experimental hall to avoid losses through long RT power lines. A cold electrical transfer line will be necessary between the cryogenic satellite and the magnet, as well as between the two coils of the magnet. Work is ongoing to determine the optimal location for the cryogenic satellite: close to the magnet in the corridor as shown in Figure 31, or further downstream, behind the particle identification detectors and the power

converters. The goal is to maximize the performance of crucial systems such as cryogenics, warm and cold powering circuits, and magnet protection. This must be done while bearing in mind integration constraints and maintainability wherever necessary.



Figure 31: The integration layout shows the cryogenic satellite in the corridor next to the spectrometer magnet. Excerpt from the SHiP "Input layout for R2E and background studies", DWG SPSXXSHIP0002 EDMS 3301751 v1.2 dated 25.08.2025.

Given that the experimental hall of ECN3 is not supplied with liquid helium distribution, and that heat loads are relatively low, the cooling of the SHiP spectrometer magnet relies on cryocoolers and cryofans, which now offer sufficient level of performance combined with enhanced energy efficiency of cryocoolers. The proposed cooling scheme is shown in Fig. 32. It consists of four independent cryogenic loops, all using pressurized GHe at 10 bars: one for cooling the current leads, using high-efficiency heat exchangers, two for cooling the magnet cold mass, i.e. one for each coil, and another one for cooling the thermal shield. The operating conditions of these cryogenic loops are as follows.

1. For the current leads: 50-55 K and ~ 2 g/s.
2. For the magnet cold mass: 18-20 K and ~ 5 g/s.
3. For the thermal shield: 70-80 K and ~ 6 g/s.

Because the MgB_2 cable will be indirectly cooled by a forced flow of helium circulating in a closed loop embedded in the aluminium former, the thermal contact between the cable




Figure 32: SHiP spectrometer magnet cooling scheme.

and the former is of utmost importance. Measurements of the thermal contact resistance between the MgB_2 cable insulated with polyimide tape and an anodized, or non-anodized, aluminum former are being conducted at the CryoLab (B165) under different conditions as follows:

1. Diameter of the groove machined in the aluminum former to house the cable: 9.4 and 9.7 mm.
2. Cable insulation scheme based on non-adhesive or adhesive polyimide tape.
3. Filler material aiming at enhancing thermal conductance between the aluminum former and the cable, e.g. GE varnish, Apiezon[®] N, and Al foil.
4. Surface of former (anodized/non-anodized).

The first measurements have shown promising results. The test programme needs to be completed before a decision can be made.

3.3.6 Spectrometer Timing Detector

During the past year, substantial progress has been achieved in the development of the spectrometer timing detector (DSTD), both in terms of mechanical design and the validation of the readout electronics. The work has focused on finalizing the mechanical concept, consolidating the electronics chain based on the FastIC+ ASIC [47], and validating the system performance in laboratory and beam-test environments.

A full mechanical layout of the DSTD has been implemented in CATIA, defining the complete structure for the detector module assembly and its integration within the SHiP experimental area. The design accommodates the modular structure of the bars, the front-end electronics, and the rails for the parking state. The current version includes the three columns with c-shape supports for the 1.4 m long scintillating bars. This new c-shape support takes into account the possibility of easy extraction in case of replacement is needed. The design also consider the horizontal and vertical overlap between rows and columns as variable. This leaves room for the development to keep searching for the optimal overlapping in both directions. The design also includes the possibility to set the data concentrator nearby the bars in order to gather the cabling from neighbouring bars.




Figure 33: Spectrometer timing detector with the mechanics allowing the detector to be extracted to the side for access (left) and the scintillating bar arrangement and overlap between the three columns.

The FastIC+ front-end ASIC has proven to be a robust and well-suited solution for the timing detector, combining excellent time resolution, compact form factor, and low-power operation. A series of laboratory and test-beam measurements have demonstrated its full functionality, including the charge and time digitization features, internal trigger logic, and synchronization capabilities. The ASIC provides a clean leading-edge timing, ensuring compatibility with high-rate environments. Tests using the dedicated DAQ chain confirmed stable performance and scalability toward a multi-board system.

A major milestone has been reached with the recent laser tests, where a time resolution of approximately 30 ps was achieved for single-bar configurations. Beam tests with pions further validated the detector concept, confirming the expected spatial dependence of the timing along the bar and allowing position reconstruction using neighbouring SiPM channels. Comparative studies of the two SiPM technologies — Hamamatsu and FBK — provided conclusive results on their performance and compatibility of both with the FastIC+ front-end. The tests demonstrated stable operation under beam conditions and high reproducibility across channels.

The focus now shifts toward completing the full readout chain of the DSTD, including synchronization across multiple boards and integration with the common readout and timing distribution system. These tests will quantify the total jitter contribution and define the required precision of the global timing reference. Parallel efforts are ongoing to prepare the final mechanical frame and service integration for system-level validation during 2026.

3.3.7 PID system

Particle identification (PID) is not used in the suppression of background but it adds redundancy and may become important to reduce neutrino-induced backgrounds in searches for partially reconstructed modes. Some FIP final states also include electrons and photons, such as e.g. $\text{HNL} \rightarrow e\pi$, $\text{HNL} \rightarrow X^\mp \pi^\pm \pi^0$ or $D^0 \rightarrow \pi^+ \pi^- \pi^0$, while several other final states are completely made of neutral particles in the decay chain, such as e.g. $\text{HNL} \rightarrow \nu \pi^0$ and $\text{ALP} \rightarrow \gamma\gamma$. With only photons in the final state, mass reconstruction with the calorimeter information alone becomes essential to characterise the new physics. In this case, mass reconstruction may also help to suppress background from neutral pions, which may arise, e.g. from neutrino coherent scattering.

In the event of the observation of a significant number of FIP candidates, with the decay into one or more final state, PID with discrimination between electrons, photons, muons and pions becomes essential for full characterisation of the BSM physics explored by SHiP. Assuming a handful of FIP events are observed, in the range between a few and a thousand, a reasonable target for the particle identification of electrons, muons, and pions would be to achieve $\sim 95\%$ efficiency with $> 99\%$ purity, in a momentum range between 2 and 150 GeV/c.

The PID detector system consists of a sampling calorimeter with both ECAL (DECA) and HCAL (DHCA) sections based on scintillator bars read out by SiPMs through wavelength-shifting fibres. The ECAL follows the SplitCal concept, which is optimised for the angular reconstruction of photons. It uses both wide and thin scintillator bars arranged in layers with alternating orientation, as well as High Precision Layers (DHPL). The DHPL is based on GEMs in the baseline design, with a fibre tracker or RPCs being considered as possible alternative options. The detector will be built in six sextants to ensure modularity.

The dynamic range requirements of the calorimeters are quite demanding since they need to measure energies ranging from minimum ionising particles (MIPs) to large electromagnetic or hadronic showers. For the light readout, the SiPMs under study are the S14160-6010PS and S14160-6015PS from Hamamatsu, as well as the AFBR-S4N66C013 from FBK for the wide layers. The thin scintillator bars have until now utilised the S13360-1325PE SiPMs. The front-end electronics is expected to rely on either the CALOROC 1b chip under design from OMEGA labs or the KLauS v2 chip from the University of Heidelberg for the SiPMs.

The former allows for a fully parallel readout, whereas the latter achieves the necessary performance through a double FIFO layout. Both yield ADCs and TDCs, and would be placed directly onto the SiPM PCBs to allow for digitization as close as possible to the sensor. Both chips have sufficient dynamic range. The GEM-based HPLs would rely on the VMM3a ASIC within the scalable readout system (SRS). The backend electronics would rely on reading out each of the six ECAL sextants independently and the HCAL all together, for a total of seven sub-partitions. Each would see all detector channels assembled within an FPGA board, which would concentrate, synchronize, repackage, and transfer data to the common concentrator while conveying clock, slow, and fast control commands to the front ends. The Xilinx Zynq ZU2CG is considered for this operation.

A test beam campaign, which followed other campaigns in 2023 and 2024, was conducted in May 2025 at the SPS. The ECAL prototype, see Fig. 34, used for this test had a lower sampling fraction than the SHiP design, and comprised 17 active layers, corresponding to 20 radiation lengths (X_0), with a total depth of 112 cm and a central split of approximately 20 cm. The prototype included both thin and wide scintillator bars.




Figure 34: Test beam setup for ECAL (top) and HCAL (bottom).

An HCAL prototype was also deployed for particle identification studies, consisting of five interaction lengths (λ_a) with sampling at each λ_a . The passive absorber was iron. For this particular test, only layers with wide scintillator layers were used. Data readout was carried out with the CAEN DT5202 system and an external trigger comprising a SiPM and a scintillator, as well as a discrete discriminator board. The Hamamatsu S14160-6050HS

SiPMs were used to read out the eight wide scintillator layers on both sides, whereas the nine layers of thin scintillating bars were read out from one side using S13360-1325PE SiPMs, also from Hamamatsu.

The test beam successfully demonstrated the feasibility of shower directionality reconstruction up to high energies within the ECAL. Preliminary studies indicate that the decays of axion-like particles (ALPs) with coupling to photons can be reconstructed at high efficiency. The inclusion of HPLs would further enhance this capability, allowing nearly all ALP decays to be detected. Further simulation studies are planned to confirm these results within a complete detector configuration.

The energy resolution of the ECAL prototype was measured to be about $28\%/\sqrt{E}$, which is worse than what is expected for the SHiP design due to the about three times poorer sampling fraction. The results were consistent across the different angles, with only a slight worsening at higher incident angles. These measurements will serve to cross-validate expectations from simulation studies.




Figure 35: Average longitudinal shower profiles for particles identified by the PID system in the four categories: muons, electrons, early hadrons, and late hadrons.

The energy response of the HCAL prototype was also studied. The sampling fraction in the HCAL is similar to that of the SHiP design, although the limited transverse size of $36 \times 36 \text{ cm}^2$ caused significant transverse leakage.

Particle identification analysis is still ongoing. We describe here preliminary results obtained with simple selection cuts. A particle is identified as a muon if the total deposited energy in the entire detector is smaller than 30 MIPs; as an electrons if the ratio of energy in the HCAL to the ECAL is smaller than 0.15, and the energy in the first 3 X_0 is larger than 20 MIPs. Particles are identified as early hadrons (those starting to shower in the ECAL) if the energy in the first four layers is smaller than 6 MIPs, the energy in the ECAL is larger than 50 MIPs, and the ratio of HCAL to ECAL energy is larger than 0.15; and

as late hadrons if the energy in the ECAL is smaller than 100 MIPs and the particle is not identified as a muon in the HCAL. Fig. 35 shows the average longitudinal profiles for particles identified in the four categories.

The H8 SPS beams are not pure in particle content, and for this test it was not possible to perform a particle selection using independent detectors. Therefore, to validate the test beam results, they were compared to the particle composition known (to a few %) by the accelerator experts, and we found agreement within 10%.

The future plan is to have a beam test of the ECAL, probably at DESY in 2026, with the inclusion of one or more high-precision layers with triple-GEM, a fibre tracker, and possibly RPC chambers as well.

3.4 Common electronics, online system and computing

3.4.1 Electronics and online system

During the course of the past year, significant progress has been made in a few areas regarding global electronics and the Online system. First, multiple SHiP electronics workshops were held and a consensus on a baseline for the electronics architecture was reached. The schematic view of the SHiP global electronics architecture can be seen in Fig. 36. The main system-level concepts of such global architecture can be summarized as follows:

- usage of common tools already planned to be utilized around CERN. This is in particular in reference to the VTRx+ and lpGBT for the Front-End electronics and a common Back-End readout board (FLX-155 or its evolution).
- trigger-less Data Acquisition system, with Front-End electronics streaming data to the Back-End.
- centralised supervision and control of the readout system (*TFC* system)
- centralised timing distribution with a minimal requirement of up to 50ps stability.
- interface to machine synchronization (clock) and parameters (e.g. start/end of spill) via White Rabbit at the SPS.
- data divided in time slices, with time and synchronization information included in the data packet.
- time-slice building (aka Event Building) performed in FPGA or CPU, or generally speaking in the DAQ part of the readout system.
- development of a common *concentrator card* as a hardware component that performs optical-to-electrical conversion (and vice versa) between the optical signal received and decoded by the coupled VTRX+/lpGBT and the Front-End electronics on differential electrical signals (so-called e-links).
- an alternative to the common *concentrator card* is to implement VTRx+ and lpGBT without an external hardware component directly in the Front-End boards. This option

is particularly important and in fact necessary for those sub-detectors which require high time resolution (≤ 50 ps) in order to profit from signal integrity on PCBs or for those needing higher chip/hardware integration requirements.




Figure 36: SHiP global electronics architecture with its hardware sub-component and functional view.

In parallel, the development of a minimal test bench has gained momentum and been given high priority. The main aim of this test bench is to replicate on a smaller scale (*slice*) the global readout architecture with the same system-level functionalities and most of the required hardware components. The test bench will profit from the usage of the VLDB+ development card – hosting the VTRX+ and the lpGBT, the FLX-712 Back-End and timing distribution card, which is the current ATLAS FELIX card in production and predecessor of the future FLX-155 for data readout and Front-End fast and slow control, and a series of development kits based on Xilinx Kintex, Virtex and Ultrascale FPGAs. The importance of the development kits is to have the possibility to test future FPGA related features as well as being able to interface electrically to the e-link signal swing and to reproduce alternative options to the baseline architecture, while keeping the system-level functionalities unchanged. In addition, Front-End emulators are being developed for the development kits in order to emulate real data taking conditions. This phase can be seen as a minimal R&D phase for SHiP, which will culminate in the final hardware decision for the architecture. Finally, these test benches can be used in test beams in 2026, with the goal to have them functional and ready to use by Q2-2026. Schematic examples of the foreseen test benches can be seen in Fig. 37, 38 and 39.

3.4.2 Computing

SHiP Computing includes hardware and software items ranging from Online Data Acquisition to Offline Data Management and Preservation. During the last year, after identifying a potential architecture and candidate tools following the approach of reuse of existing tools




Figure 37: First type of test bench where a development kit acts as a pure Front-End emulator, communicating via optical fibers to the FLX-712.




Figure 38: Second type of test bench where a development kit acts as a pure Front-End emulator, communicating via electrical connections to the VLDB+ card, which acts as a concentrator card.

and frameworks and relying as much as possible on CERN-hosted services, the focus has been on the integration of pilot services and the implementations of minimum viable product demonstrators usable on the short-term to support Physics Software’s studies.

Discussions with the “Rucio for SMEs” team in the CERN IT-SD group have evolved into a pilot instance of Rucio for the SHiP Data Management and cataloging. The pilot instance was tested with the registration of existing simulations files from EOS and their transfer to tape storage on CTA. Moreover, remote storage elements (namely Ghent University and GridPP) were added to the SHiP pilot instance and data transfers exercised. A custom web interface has been developed in SHiP to query the Rucio Data catalog.

In parallel, the existing HSF-supported *nopayloaddb* Condition Database tool was reviewed, and its limitations identified in collaboration with the maintainer. A Rust-based server implementation was developed to address the shortcomings of the original Python version and deployed on CERN IT platforms for testing. To enable efficient use of the Conditions Database, client libraries in C++ and Python were also developed and integrated within FAIRSHIP.

Finally, an initial study of the Online Data Acquisition software has been started, based on the current approximate figures regarding data formats, detector’s hit rates and windows, and total throughput. More input is still needed from sub-detectors and electronics to emulate realistic SHiP Online data conditions.




Figure 39: Third type of test bench where the concentrator card is replaced by an FPGA-based development kit acting as a pure data transmission directly from FPGA. Note that , even in this particular architecture, the transmission of the timing and fast control is still via the concentrator card, acted by the VLDB+ in the test bench.

3.5 Physics software

The software framework of the SHiP collaboration is called FAIRSHIP and is currently based on the FAIRROOT framework. It integrates GEANT4, various event generators including PYTHIA 6 and 8 and GENFIT for track fitting and provides a detailed simulation geometry, full simulation for signal and background channels and tools for digitization, reconstruction and analysis.

The simulation has been updated to align the coordinate system with the one used for the integration to ensure consistency between the physics and engineering layouts [EDMS 3234956]. The changes reported in this document on the target complex, the muon shield and the detectors were implemented in the simulation. These also includes more details of the target station with a full tungsten target, and the different options of the optimised muon shield, both fully warm and hybrid.

After remaining largely unchanged throughout the proposal phase, the physics software is now undergoing a thorough review, to lay the foundations for the next decades. Besides routine updates of many of the packages, several core components are being reconsidered to align more closely with the software frameworks of other current and future CERN experiments and to benefit from advances in the software ecosystem over the last decade. For pattern matching and tracking, a migration to the ACTS [55] package is underway for all tracking detectors. A migration from the current geometry description using TGEO to DD4HEP[56] is in preparation and will be tested with the calorimeter first before migrating the entire geometry. Finally, FAIRROOT is being compared to alternatives to decide whether to fully commit to FAIRROOT or migrate to a different base of our framework. In addition to these modernizations, we have been improving the development practices to make the development and maintenance of the physics software more sustainable. Notably, the geometry is now tested automatically, automated tools are used to improve the code quality and robustness, and documentation has seen major improvements.

The collaboration closely follows developments by EP-SFT, such as the effort for a com-

mon build system, BITS, new technologies such as ROOT’s RNTUPLE and new packages in the `scikit-hep` python ecosystem.

A major focus of recent months has been the preparation of new high-statistics background productions, as described in Section 4.1.

4 Reassessment of physics performance with SHiP-2026

Extensive simulation campaigns have been carried out over the years to quantify the expected background for the SHiP experiment, focusing primarily on neutrino- and muon-induced interactions. In general, a very good level of agreement has been found between the ECN3 and ECN4 geometries, as discussed in Ref. [57].

For the neutrino-induced background, large samples corresponding to several years of running have been simulated with interactions forced to 100% probability in the region between the entrance of the decay volume and the second tracking station. For ECN3, a total of 1.2×10^8 neutrino interactions were generated, to be compared with 1.2×10^7 expected interactions over five years of operation. After reconstruction, 5.3×10^5 signal-like candidates were produced, corresponding to 2.1×10^4 expected in five years of operation. Applying the baseline event selection yields 0.7 events for partially reconstructed signals in 15 years, and a factor of five fewer for fully reconstructed ones. These results are obtained assuming the full detector configuration. If particle identification (PID) is not available in 2032-2033, this would lead to approximately 2 background events per year for partially reconstructed signals.

These findings are fully compatible with those obtained in the ECN4 studies, confirming that the geometry change does not affect the neutrino background level. We therefore conclude that the ongoing optimisation of the muon shield factorises well with respect to the neutrino background and can proceed independently. The main design change since the re-optimisation of the experiment is the replacement of the vacuum vessel with a helium-filled balloon. To address this, a dedicated simulation corresponding to 10^7 neutrino interactions was produced, while about 5×10^4 interactions are expected in 15 years of data taking. The residual neutrino background from interactions in the helium balloon wall or in the helium volume can be completely removed by tightening the impact parameter requirement to the target and increasing the minimum distance from the balloon wall, with negligible loss of signal efficiency. A full optimisation of the selection using multivariate techniques will be carried out in 2026, once the complete background samples have been produced with the final geometry.

For the muon-induced background, studies of inelastic scattering have been updated using a ten-spill dataset corresponding to the superconducting muon-shield configuration. After applying the fiducial, DOCA, and impact parameter cuts, and imposing SBT and UBT veto conditions, no surviving background events are found for either 1 mbar air in the decay volume or helium at one atmosphere. Assuming factorisation between the veto and impact-parameter cuts to enhance statistics, these results set an upper bound of fewer than 0.2 partially reconstructed events over the full running period. However, unlike the neutrino background, the muon background cannot be simulated with full statistics without such assumptions. Moreover, the muon-shield optimisation has an obvious and strong impact on

this background. Addressing the limited simulated statistics for the muon background and producing complete background samples are among the main physics-study goals for 2026.

4.1 Large-scale production of background samples

The simulation of the SHiP detector starts from just upstream of the target, followed by the implementation of the target, the hadron absorber, the muon shield with the scattering detector, the background taggers, the decay volume, and the FIP decay spectrometer, within the FAIRSHIP framework (see Section 3.5). The Monte Carlo sample production is broken down in two steps to allow the studies of different detector configurations downstream of the hadron absorber. The first and most critical, time and CPU consuming step is the creation of a large sample of proton interactions simulated through the target and the hadron absorber.

Since the last large production in 2018, several changes have been implemented:

- the target is now made of pure W, adding to a total W length of 1.5 m, interleaved with gaps, enclosed in a steel cylinder within a pressurized vessel containing He. It is followed downstream by a proximity shielding of 0.5 m of copper.
- the hadron absorber has been shortened to ~ 2.3 m.

Following the implementation of the new layout in the simulation software, small MC samples have been produced in order to validate the implementation and check the choice of several parameters for the final large production.

For the final production, which will run for several months, the particles leaving the hadron absorber and their history will be saved for use in the background studies, similarly to the 2018 production. The production will run at CERN, UK and University of Ghent Tier2 grid sites [58], and we expect at least statistical parity if not a significant increase in simulation statistics compared to the 2018 production.

4.2 ML-based simulation of the large background data samples

The requirement of achieving zero background events over 15 years of operation places significant demands on the available simulation resources. The most critical backgrounds originate from muons emerging from the hadron absorber that are either insufficiently deflected and consequently enter the decay volume, or are deflected toward its walls or edges where they may undergo inelastic scattering.

The optimisation of the muon shield relies on the rapid propagation of large muon samples through multiple shield configurations. Using a traditional GEANT4 approach, the propagation of the full muon sample for a single configuration typically requires around four hours on large CPU clusters. This computational cost restricts the use of full samples during optimisation and can lead to statistical instabilities in the optimisation process.

To address this limitation, a GPU-based CUDA framework has been developed to accelerate the muon propagation through the shield. The method relies on sampling from histograms of muon momentum dissipation and deflection angles, which are pre-populated using large GEANT4 reference samples. The effect of the magnetic field is implemented through a fast numerical integration of the magnetic flux predicted by the magnetostatic

solver described in Sec. 3.1. With this approach, the full muon sample can be propagated through a given shield configuration in approximately four minutes using a single NVIDIA GPU, representing a speed-up of nearly two orders of magnitude compared to the CPU-based method, without any measurable loss of simulation accuracy. Work is underway to generalise this simulation approach to simulate muon interactions with the decay volume and the decay detector.

5 Test beams at CERN in 2026

5.1 DUBT and SMTC scintillating tiles

Request: 2 weeks at PS.

Purpose: Test of prototype modules of the scintillator–tiles technology. Several configurations will be evaluated to optimise light collection and readout, addressing both calorimetric performance in the SMTC region and the the tiles as fast timing elements for the Upstream Background Tagger (DUBT).

5.2 High-granularity calorimeter (SHGC) prototype

Request: 2 weeks at SPS.

Purpose: The first test beam of the silicon–tungsten target as a high granularity calorimeter will characterise a prototype module combining refurbished CMS silicon strip detectors with 7 mm tungsten absorbers. Electrons, muons, and pions will be used to study tracking and electromagnetic–shower development, assessing the module’s suitability as a neutrino–interaction tracker. This activity follows the strong synergy with SND@HL-LHC in developing silicon–based neutrino targets.

5.3 Surrounding background tagger (DSBT)

Request: 2 weeks at PS.

Purpose: Test beam studies of two independent single–cell liquid–scintillator prototypes with different geometries, each equipped with wavelength–shifting optical modules coupled to SiPMs. The measurements will characterise detector response versus beam incidence angle and validate new analog and digital readout electronics using FASTIC+ and a first prototype of the backend readout system.

5.4 Spectrometer straw tracker (DSST)

Request: 2 weeks at DESY and additional 2 weeks at PS.

Purpose: At DESY, a full–scale prototype of the straw–tube tracker will be tested for mechanical integrity and tracking performance. The PS campaign will compare different front–end ASICs (ASD, VMM) to study charge collection, rate capability, and position resolution referenced to a TimePix telescope. A key goal is to demonstrate, for the first time, time–referenced tracking by coupling a FASTIC+ timing layer to provide a precise t_0 for the straw system.

5.5 Spectrometer timing detector (DSTD)

Request: Two test periods of 1 week each at PS.

Purpose: The first run will study the readout of long scintillator bars using SiPMs distributed along the bar to optimise sensor spacing and assess light-collection efficiency for fast scintillator materials. The second run will focus on a larger prototype read out with FASTIC+ and a data concentrator, aiming to measure coincidence-time resolution and determine the optimal horizontal and vertical staggering of bars for the final detector layout.

6 Plan and objectives for Run 4

Fig. 40 shows the preliminary master schedule for SHiP and the HI-ECN3 project. Note that the details of this plan are under discussion.

The HI-ECN3 project is aiming at submitting the TDR for all aspects of BDF by mid-2026 in order to move focus towards the final design and preparation for the procurement phase and construction. The implementation of the beamline and primary services also involves significant work together with the NA consolidation in LS3. An important fraction of the HIECN3 construction phase is taken by civil engineering at ECN3 with the principal works consisting of the creation of the service building for the target complex (B754) and modifications underground in TCC8 and ECN3. The civil engineering works will start as soon as the dismantling of TCC8 and ECN3 is completed mid-2027 and are expected to continue up to the end of 2028. During this period the access to the site and underground will be extremely limited and will not allow parallel installation activities. The installation and refurbishment of infrastructure and services will start from beginning 2029, with staging of activities between TCC8, ECN3 and the new building B754 in order to allow the installation of the target complex in TCC8 to start as soon as possible. The plan is to have the facility ready for commissioning of the beamline and the target in Q3 of 2031. This step presents an important opportunity for SHiP to measure the performance of the beamline and the target and measure rates of particles in ECN3. The HIECN3 schedule and the muon shield work plan aim to give between nine and twelve months of installation time for the muon shield in 2030-2031 such that the muon shield can already undergo first tests with beam during the commissioning of the facility. The plan requires careful coordination since the target complex and the muon shield will both be heavy users of the B911 equipment shaft and space in TCC8 is too limited to allow concurrent installation activities.

The SHiP collaboration is actively pursuing work in all subsystems towards the preparation of the TDRs with the aim to submit the majority of TDRs by end of 2027. Given the ambitious plan to operate the muon shield in 2031, the muon shield is on critical path with submission of the TDR planned by the end of 2026.

In the preliminary plan, SHiP will have access to ECN3 from mid-2029. The installation works will start in building 918 with the refurbishment of the operational space, counting room and control room, and setting up SHiP's computing infrastructure and common online components. This will provide the development and deployment platform for the central system as well as the further development and integration of the subsystems and software tools. By end of 2029, it is expected that primary services in terms of powering, cooling and ventilation, networks, and lifting and handling equipment are operational underground to




Figure 40: Preliminary BDF/SHiP master plan.

start detector installation. The detector installation will start with the construction of the spectrometer magnet at the end of 2029.

By the time of the commissioning run with beam for the facility in Q3 of 2031, SHiP expects to have finished the installation of the muon shield, with integrated detectors for muon rate measurements, decay volume and spectrometer magnet. The subdetectors will only have been partially constructed. The goal of the facility commissioning with beam will be to reach the nominal intensity of 4×10^{13} protons per spill and perform a short heat run for the target system. For SHiP, it will enable commissioning of the core part of the online system and communication systems with SPS and BDF, and to make comprehensive measurements of the rates of particles in ECN3. The run will be extremely important for studying and tuning the muon shield, as well as comparing the rates with the expectations from simulation. Depending on the state of the secondary service systems, subdetectors may also participate in stand-alone mode.

After the initial commissioning run, SHiP will have another 9-12 months to complete the detector installation, perform stand-alone commissioning and dry-runs with the detectors in global data taking mode, and finalise tools for detector and data monitoring. The commissioning of the SHiP detector with beam is planned for Q3 of 2032. The run will consist of several steps, most of which will be in a mode of as low proton intensity as possible, 5×10^{11} protons per spill or possibly less. To further reduce the instantaneous rates if needed, the spill length may also be stretched from the nominal 1.2 s to 2.4 or 4.8 s. In these conditions, the target may be operated without the beam sweep in the centre as well as at a few fixed points of the nominal beam sweep. It will allow studying the effect of the muon shield aperture and field tuning. The goal is also that these conditions allow taking data




Figure 41: Branching fractions of the fully reconstructed charged-particle final states, i.e. without photons, π^0 , K_L , neutrons, neutrinos, for a number of benchmark FIP models (here referred to as Long-Lived Particles - LLP).

with the muon shield off, making the procedure for the detector time and space alignment easier. The next step, still with beam sweep and muon shield off (or near off), and also air or nitrogen in the decay volume, will be data-taking for detailed measurements of different types of background rates. This will be a crucial step to validate the full chain of data taking and event reconstruction. With the muon shield and air/nitrogen in the decay volume the goal is to make the background worst possible in order to study detection and reconstruction efficiencies and measure the resolution functions of the quantities used for background rejection with a sample of real inelastic scattering of muons and neutrinos. A more detailed study of the ideal conditions of this run is underway. It will also be used to determine the actual worst-case rates and data sizes that the detector and the readout system must be able to cope with. In the final part of this run, the detector should be operated in the nominal mode with the muon shield off and the decay volume under helium in order to prepare for data taking for physics in 2033.

In the detector commissioning run with beam in 2032, it is crucial that the beam instrumentation delivering information about the position and size of the beam is available, as well as information from the background monitors in the beamline.

The year 2033 is intended to be a full year of data taking with a nominal beam intensity of 4×10^{13} and normal operating conditions. The goal is to collect a nominal annual data set of 4×10^{19} protons on target before LS4, which starts in 2034 in the current schedule. The data taking for physics will require that the beam instrumentation that delivers detailed information about the spill quality and beam intensity is fully available.

The priority for the physics run in 2033 is on a detector capable of exploring the FIP parameter space with a subset of FIP final states, i.e. fully reconstructed signal states with only charged particles and no neutrinos. This has mainly a large effect on the sensitivity to




Figure 42: Sensitivity projections for one year of data taking ($4 \times 10^{19} \text{PoT}$) with the SHiP detector in 2033 for various benchmark models defined by the Physics Beyond Colliders (PBC) working group. The solid contours correspond to one year of data taking assuming only fully reconstructed final states with charged particles, while the dashed contours show the expected reach after 15 years of running (6×10^{20}), including all reconstructible final states. The models shown are BC1 (dark photons), BC4 (scalars without hSS coupling), BC5 (scalars with hSS coupling fixed such that $\text{Br}(h \rightarrow SS) = 0.01$), BC7 (heavy neutral leptons mixing with ν_e), BC10 (axion-like particles with universal couplings to fermions), and BC11 (axion-like particles coupled to photons).

FIP masses above $1 \text{ GeV}/c^2$ where perturbative QCD and hadronisation play an important role and lead to higher multiplicities with neutral particles. Fig. 41 shows the branching fractions of the fully reconstructed charged-particle final states, i.e. without photons, π^0 , K_L , neutrons, neutrinos, for a set of benchmark models. A crucial goal of the run will be validating the understanding of the different backgrounds in these modes, and study ahead of LS4 the final configuration of the background taggers and the rejection techniques developed.

Despite the focus on only fully reconstructed final states, a year of data taking will have significant discovery potential in all the benchmark models and allow current limits to be improved by 1-2 orders of magnitude, as shown in Fig. 42.

7 Status of the SHiP collaboration

The collaboration is actively working to secure funding from national and international funding agencies. Given the agreed-upon timeline of the TDRs and the construction, institutions from several countries have already submitted applications or are in the process of applying. In particular, applications have been submitted or are being prepared in Germany, Italy, Switzerland, and the UK, with other countries expected to follow soon. The funding of the R&D that is necessary to complete the TDRs is obtained from the various institutes' running budgets. CERN has committed ~ 5 MCHF to SHiP in the MTP over the years 2024-2031 to fund the TDR activities of the EP SHiP team and initial investments towards construction.

The muon shield is a critical component of the initial configuration of SHiP required to ensure the low-background environment necessary for feebly interacting particle searches. The development of the muon shield involves contributions from institutes in Germany, Italy UK, and CERN. Funds for a prototype are being sought while the construction will be covered by the collaboration's Common Fund.

The SHiP collaboration is making significant efforts to expand its institutional participation. Currently, the collaboration consists of 33 institutes in 18 countries, with 9 additional associate institutes, and CERN and JINR. Last year the Institute of Experimental Physics in Almaty, Kazakhstan (strawtubes), Karlsruhe Institute of Technology (superconducting magnet expertise) and Gent University (scintillating tiles) became members. This year the University of Liverpool joined. Liverpool have extensive knowledge of straw trackers and it is expected that they will make an important contribution to the UBT subdetector. The Georgia Technical University reinforced their group with a strawtube production laboratory and IJCLab also expanded their participation (detector optimisation using experience from CALICE). Since the Upgrade of SND@LHC for HL-LHC is a prototype for the SHiP SND detector, there is a large overlap between the members of both collaborations. New institutes that recently joined SND@LHC such as the Universities of Genova, Catania and Salerno will eventually also join SHiP.

References

- [1] W. Bonivento, A. Boyarsky, H. Dijkstra, U. Egede, M. Ferro-Luzzi, B. Goddard, A. Golutvin, D. Gorbunov, R. Jacobsson, J. Panman, M. Patel, O. Ruchayskiy, T. Ruf, N. Serra, M. Shaposhnikov and D. Treille, *Proposal to Search for Heavy Neutral Leptons at the SPS*, Tech. Rep. [CERN-SPSC-2013-024](#), [SPSC-EOI-010](#), [SPSC-EOI-010 \[1310.1762\]](#), CERN, Geneva (Oct, 2013).
- [2] SHiP collaboration, *A Facility to Search for Hidden Particles (SHiP) at the CERN SPS - Technical Proposal*, Tech. Rep. [CERN-SPSC-2015-016](#), [SPSC-P-350](#), [SPSC-P-350, \[1504.04956\]](#), CERN, Geneva (Apr, 2015).
- [3] M. Anelli et al., *A Facility to Search for Hidden Particles (SHiP) at the CERN SPS - the SHiP physics case*, Tech. Rep. [CERN-SPSC-2015-017](#), [SPSC-P-350-ADD-1](#), [SPSC-P-350-ADD-1, \[1504.04855\]](#), CERN, Geneva (Apr, 2015).
- [4] SHiP collaboration, *SHiP Experiment - Comprehensive Design Study report*, Tech. Rep. [CERN-SPSC-2019-049](#), [SPSC-SR-263](#), CERN, Geneva (Dec, 2019).
- [5] SHiP collaboration, *The experimental facility for the Search for Hidden Particles at the CERN SPS*, [JINST 14 \(2019\) P03025](#).
- [6] C. Ahdida et al., *SPS Beam Dump Facility - Comprehensive Design Study*, CERN Yellow Reports: Monographs, CERN, Geneva (Dec, 2019), [10.23731/CYRM-2020-002](#), [\[1912.06356\]](#).
- [7] O. Aberle, C. Ahdida, P. Arrutia, K. Balazs, J. Bernhard, M. Brugger, M. Calviani, Y. Dutheil, R.F. Ximenes, M. Fraser, F. Galleazzi, S. Gilardoni, J.-L. Grenard, T. Griesemer, R. Jacobsson, V. Kain, D. Lafarge, S. Marsh, J.M. Martin Ruiz, R.F.M. Andrade, Y. Muttoni, A. Navascues Cornago, P. Ninin, J. Osborne, R. Ramjiawan, P. Santos Diaz, F. Sanchez Galan, H. Vincke and P. Vojtyla, *Study of alternative locations for the SPS Beam Dump Facility*, Tech. Rep. [CERN-SPSC-2022-009](#), [SPSC-SR-305](#), CERN, Geneva (2022).
- [8] SHiP collaboration, *The SHiP experiment at the proposed CERN SPS Beam Dump Facility*, [Eur. Phys. J. C 82 \(2022\) 486](#).
- [9] SHiP collaboration, *BDF/SHiP at the ECN3 high-intensity beam facility - Letter of Intent*, Tech. Rep. [CERN-SPSC-2022-032](#), [SPSC-I-258](#), CERN, Geneva (2022).
- [10] SHiP collaboration, *BDF/SHiP at the ECN3 high-intensity beam facility - Technical Proposal*, Tech. Rep. [CERN-SPSC-2023-033](#), [SPSC-P-369](#), CERN, Geneva (2023).
- [11] *Nomination as the Project Leader for High Intensity ECN3 (HI-ECN3)*, February, 2025. <https://edms.cern.ch/document/3232189>.
- [12] *Infrastructure Request Proposal (IRP) HI-ECN3: Buildings 754, 911, 912, 918, 920 & TCC8/ECN3*, Tech. Rep. [EDMS 3352794](#), [IRP-DS-0900-HI-ECN3](#) (September, 2025).

- [13] A. COLINET, *Beam Interlock System for SPS slow extraction and the North Experimental Area*, Tech. Rep. [EDMS 2810115](#), [SPSX-CIB-ES-0002](#) (October, 2025).
- [14] T. Prebibaj, G. Arduini, H. Bartosik and M.A. Fraser, *SPS Operation and Future Proton Sharing Scenarios for the ECN3 facility*, Tech. Rep. [CERN-PBC-Notes-2023-001](#) (2023).
- [15] T. Prebibaj, P. Arrutia, D. Banerjee, H. Bartosik and A. Huschauer, *Studies on the Future Beam Sharing Across the CERN Injector Complex*, Tech. Rep. [EDMS 3288709](#), [ACC-OP-EN-0001](#) (August, 2025).
- [16] P. Arrutia, D. Banerjee, R. Franqueira Ximenes, A. Gorn and T. Prebibaj, *Spill Optimisation for Experiments*, Tech. Rep. [EDMS 3348121](#), [SPSX-O-EN-0001](#) (September, 2025).
- [17] M. Cavazza, *BDF/SHiP 2025 Prototype Target for the HI-ECN3 Project*, 2025. <https://cds.cern.ch/record/2941881>.
- [18] C. Lederer-Woods, *Neutron Activation Station at the SPS Beam Dump Facility (BDF) - Expression of Interest*, Tech. Rep. [CERN-SPSC-2024-027](#), [SPSC-EOI-023](#), CERN, Geneva (2024).
- [19] N_TOF collaboration, *n_ACT@BDF: A Neutron Activation Station at the SPS Beam Dump Facility (BDF) - Technical Proposal*, Tech. Rep. [CERN-SPSC-2025-038](#), [SPSC-P-375](#), CERN, Geneva (2025).
- [20] D. Forkel-Wirth, *Zonage radiologique au CERN*, Tech. Rep. [EDMS 810149](#), CERN, Geneva (2006).
- [21] C. Ahdida et al., *New Capabilities of the FLUKA Multi-Purpose Code*, *Frontiers in Physics* **9** (2022) 788253.
- [22] G. Battistoni et al., *Overview of the fluka code*, *Annals of Nuclear Energy* **82** (2015) 10.
- [23] R. Alia, *Radiation safe level definition criteria for HL-LHC electronics*, Tech. Rep. [EDMS 2389056](#), CERN, Geneva (2022).
- [24] M. Jimenez et al., *TE Department Contributions to SHiP Experiment*, Tech. Rep. [EDMS 3331040](#) (August, 2025).
- [25] P. Santos Diaz and R. Jacobsson, *SHiP preliminary infrastructure requirements, layout and integration at ECN3*, Tech. Rep. [EDMS 3309666](#) (July, 2025).
- [26] SHiP collaboration, *The active muon shield in the SHiP experiment*, *JINST* **12** (2017) P05011 [[1703.03612](#)].
- [27] SHiP collaboration, *Measurement of the muon flux from 400 GeV/c protons interacting in a thick molybdenum/tungsten target*, *Eur. Phys. J. C* **80** (2020) 284.

- [28] SHiP collaboration, *The active muon shield*, Tech. Rep. [CERN-SHiP-NOTE-2016-005](#) (2016).
- [29] O. Lantwin, *Optimisation of the SHiP experimental design*, 2019. <https://cds.cern.ch/record/2693177>.
- [30] SHiP collaboration, *Optimising the active muon shield for the SHiP experiment at CERN*, *J. Phys. Conf. Ser.* **934** (2017) 012050.
- [31] S. Shirobokov, *Optimisation of the SHiP Beam Dump Facility with generative surrogate models*, 2021. <https://spiral.imperial.ac.uk/handle/10044/1/95975>.
- [32] P. Wertelaers, *Normal-conducting muon shield with standard yoke iron and tuned return yoke ; first performance assessment*, Tech. Rep. [CERN-SHiP-INT-2022-004](#) (2022).
- [33] P. Wertelaers, *Muon shield concept based on superconducting magnet and tuned iron return yoke*, Tech. Rep. [CERN-SHiP-INT-2022-003](#) (2022).
- [34] M. Ferro-Luzzi, *Preliminary exploration of a superconducting muon shield*, Tech. Rep. [CERN-SHiP-INT-2022-005](#) (2022).
- [35] M. Koratzinos and E. Van Herwijnen, *A possible yokeless superconducting dipole for the SHiP muon shield*, Tech. Rep. [CERN-SHiP-INT-006](#) (2022).
- [36] O. Pinto and C. Ahdida, *Presentation given at the 32nd SHiP Collaboration week: Radiation Protection Update*, Tech. Rep. [EDMS 3303812](#) (March, 2025).
- [37] SHiPCOLLABORATION collaboration, *Conceptual Design Study of a High-Field Magnetized Hadron Stopper for the SHiP Muon Shield*, Tech. Rep. [CERN-SHiP-INT-2025-003](#) (2025).
- [38] SHiP collaboration, *Study of combinatorial background tagging for different muon shield variants*, Tech. Rep. [CERN-SHiP-INT-2025-004](#) (2025).
- [39] M. Ferro-Luzzi, M. Fraser and J.-L. Grenard, *Functional specification: BDF/SHiP Magnetised Hadron Stopper*, Tech. Rep. [EDMS 3233771](#) (May, 2025).
- [40] Abbaneo, D and others, *Snd@hl-lhc, scattering and neutrino detector in run 4 of the lhc*, Tech. Rep. CERN, Geneva (2025).
- [41] ATLAS collaboration, *The ATLAS Transition Radiation Tracker (TRT) proportional drift tube: design and performance*, *JINST* **3** (2008) P02013.
- [42] MUON G-2 collaboration, *The Straw Tracking Detector for the Fermilab Muon g-2 Experiment*, *JINST* **17** (2022) P01005 [2111.02076].
- [43] J.S. Bono, J.F. Caron, S.H. Oh and C. Wang, *The stress relaxation (creep) rate of Mu2e straw tubes*, *Nucl. Instrum. Meth. A* **902** (2018) 95.

- [44] A. Brignoli et al., *Performance of prototypes with different reflector materials for the SHiP liquid scintillator surrounding background tagger*, *JINST* **20** (2025) P07023 [[2503.10250](#)].
- [45] SHiP collaboration, *The SHiP experiment at the proposed CERN SPS Beam Dump Facility*, *Eur. Phys. J. C* **82** (2022) 486 [[2112.01487](#)].
- [46] J. Alt et al., *Performance of a first full-size WOM-based liquid scintillator detector cell as prototype for the SHiP Surrounding Background Tagger*, *JINST* **19** (2024) P05024 [[2311.07340](#)].
- [47] S. Gomez, J.M. Fernandez-Tenllado, J. Alozy, M. Campbell, R. Manera, J. Mauricio, A. Mariscal, C. Pujol, D. Sanchez, A. Sanmukh, A. Sanuy, R. Ballabriga and D. Gascon, *Fastic: A highly configurable asic for fast timing applications*, in *2021 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, pp. 1–4, 2021, [DOI](#).
- [48] S. Abovyan, V. Danielyan, A. Grasser, O. Kortner, H. Kroha, R. Richter and S. Simeonov, *Certification of the amplifier-shaper-discriminator asics produced for the atlas mdt chambers at the hl-lhc*, *Journal of Instrumentation* **17** (2022) C10008.
- [49] S. Abovyan, N. Bangaru, F. Fallavollita, O. Kortner, S. Kortner, H. Kroha, E. Voevodina, R. Richter and Y. Zhao, *Performance of small-diameter muon drift tube chambers with new fast readout asic at high background rates*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **1072** (2025) 170198.
- [50] V. Bautin, S. Bulanova, T. Enik, E. Kuznetsova, A. Mukhamejanova, Y. Mukhamejanov, S. Romakhov, N. Saduyev, S. Sakhiyev, K. Salamatin and D. Sosnov, *Ultrasonic welding technology for future straw trackers and performance studies with small-size tracker prototypes*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **1081** (2026) 170767.
- [51] A. Devred et al., *Proof-of-Principle of an Energy-Efficient, Iron-Dominated Electromagnet for Physics Experiments*, Tech. Rep. IEEE TRANSACTIONS ON APPLIED SUPERCONDUCTIVITY, VOL. 34, NO. 5, AUGUST 2024 (2024).
- [52] L. Baudin et al., *Conceptual Design and Technology Demonstration Program of the SHiP Spectrometer MgB₂ Magnet*, Tech. Rep. To be published soon (was presented at MT29 in July 2025), CERN (2025).
- [53] A. Milanese, *Magnetic design of a superferric MgB₂ demonstrator*, Tech. Rep. CERN TE-MSD Internal Note, vol. EDMS, no. 2939566, 2023, CERN, Geneva (2023).
- [54] L. Baudin, *Magnetic and Mechanical Design of EESD2 magnet*, Tech. Rep. CERN TE-MSD Internal Note, vol. EDMS, no. 3212397, 2025, CERN, Geneva (2025).

- [55] X. Ai et al., *A Common Tracking Software Project*, *Comput. Softw. Big Sci.* **6** (2022) 8 [2106.13593].
- [56] M. Frank, F. Gaede, M. Petric and A. Sailer, *Aidasoft/dd4hep*, Oct., 2018. 10.5281/zenodo.592244, <https://doi.org/10.5281/zenodo.592244>.
- [57] SHiP collaboration, I. Bezshyiko, O. Lantwin, M. Ovchynnikov and A. Reghunath, *Signal and background studies for ECN3*, CERN-SHiP-INT-2024-001, 2024. <https://cds.cern.ch/record/2912101>.
- [58] J. Shiers, *The worldwide lhc computing grid (worldwide lcg)*, *Computer Physics Communications* **177** (2007) 219.

SHiP Collaboration

R. Albanese^{20,e,g}, K. Albrecht⁹, A. Alexandrov^{20,e}, F. Alessio³⁵, F. Aloschi^{20,d}, S. Aoki²³, D. Aritunov¹³, A. Anokhina²⁷, T. Arndt¹⁴, N. Azorskiy³³, A. Ballarino³⁵, L. Baudin³⁵, V. Bautin³³, A. Bay³⁶, I. Bekman¹³, C. Betancourt^j, C.A. Beteta³⁷, I. Bezshyiko³⁷, O. Bezshyyko⁴⁴, D. Bick¹², A. Blanco³², M. Bogomilov², P.B. De Sousa³⁵, I. Boiarska⁶, K. Bondarenko³¹, W.M. Bonivento¹⁹, M. Borisyak¹⁰, N. Bourcey³⁵, A. Boyarsky^{31,44}, D. Breton⁷, A. Brignoli⁹, S. Buontempo²⁰, V. Büscher¹⁵, M. Campanelli⁴³, D. Centanni²⁰, A. Cervelli¹⁸, K.-Y. Choi³⁰, A. Chukanov³³, M. Climescu¹, L. Congedo^{17,b}, M. Cristinziani¹⁶, A. Crupano¹⁸, G.M. Dallavalle¹⁸, M. Dam¹⁴, N. D'Ambrosio¹⁷, D. Davinoⁱ, R. de Asmundis²⁰, P. de Bryas³⁶, J. De Carvalho Saraiva³², G. De Lellis^{20,d}, M. de Magistris^{20,d,h}, A. De Roeck³⁵, M. De Serio^{17,b}, D. De Simone³⁷, G. Del Guidice^{20,d}, C. Delogu¹⁵, P. Dergachev^j, P. Deucher¹⁵, A. Devred³⁵, A. Di Crescenzo^{20,35,d}, C. Di Cristo^{2,3}, D. Duarte Ramos³⁵, H. Dijkstra³⁵, D. Dobur¹, C. Eckardt⁷, T. Enik³³, O. Fecarotta^{20,d}, F. Fedotovs⁴³, T. Ferber¹⁴, M. Ferrillo³⁷, M. Ferro-Luzzi³⁵, R.A. Fini¹⁷, A. Fiorillo^{20,d}, H. Fischer¹¹, P. Fonte³², R. Fresa^{20,g}, G. Frisella³⁵, T. Fukuda²⁴, G. Galati^{17,b}, E. Gamberini³⁵, L. Golinka-Bezshyyko^{37,44}, A. Golovatiuk^{20,d}, A. Golutvin⁴², D. Gorbunov²⁷, V. Gorkavenko⁴⁴, E. Graverini^{36,m}, C. Grewing¹³, S. Guaman¹¹, A. M. Guler³⁸, V. Guliaeva³⁶, G.J. Haefeli³⁶, C. Hagner¹², J.C. Helo^{3,5}, E. van Herwijnen⁴², A. Hollnagel¹⁵, C. Issever⁹, S. Izquierdo Bermudez³⁵, A. Iuliano^{20,d}, R. Jacobsson³⁵, D. Jokovic³⁴, C. Kamiscioglu³⁹, S.H. Kim²⁸, Y.G. Kim²⁹, N. Kitagawa²⁴, V. Koch¹³, T. Koettig³⁵, K. Kodama²², D.I. Kolev², L. Kolupaeva³³, M. Komatsu²⁴, V. Kostyukhin¹⁶, I. Krasilnikova³³, A. Krolla¹¹, K. Kuznetsova²⁷, S. Kuleshov^{3,4}, E. Kurbatov^j, H.M. Lacker⁹, O. Lantwin¹⁶, A. Lauria^{20,d}, K.Y. Lee²⁸, W.C. Lee¹², N. Leonardo³², V.P. Loschiavo^{20,f}, L. Lopes³², I. Lomidze⁸, F. Lyons¹¹, J. Maalmi⁷, A.-M. Magnan⁴², F.J. Mangiarotti³⁵, A. Miano^{20,p}, S. Mikado²⁵, A. Mikulenko³¹, A. Milanese³⁵, J. A. Molins I Bertram¹⁵, T. Molzberger¹¹, M.C. Montesi^{20,d}, K. Morishima²⁴, A.S. Müller¹¹, Y. Mukhamejanov²⁷, N. Naganawa²⁴, M. Nakamura²⁴, T. Nakano²⁴, S. Ogawa²⁶, O. Olshevskiy³³, M. Ovchinnikov³⁵, P. Owen³⁷, B.D. Park²⁸, A. Pastore¹⁷, M. Patel⁴², O. Pirotte³⁵, U. Parzefall¹¹, K. Petridis⁴⁰, N. Polukhina³³, L.F. Prates Cattelan³⁷, A. Prota^{20,d}, S.R. Qasim³⁷, A. Quercia^{20,d}, A. Rademakers³⁵, F. Ratnikov^j, F. Redi^l, A. Reghunath⁹, S. Ritter¹⁵, H. Rokujo²⁴, S. Romakhov³³, O. Ruchayskiy⁶, T. Ruf³⁵, K. Salamatin³³, S. Sakhiyev²⁷, P. Santos Díaz³⁵, O. Sato²⁴, F. Savary³⁵, C. Scharf⁹, W. Schmidt-Parzefall¹², M. Schumann¹¹, P. Schupp¹⁰, N. Serra³⁷, M. Shaposhnikov³⁶, A. Sharmazanashvili⁸, L. Shchutska³⁶, H. Shibuya²⁶, A. Sidoti¹⁸, S. Simone^{17,b}, K. Skovpen¹, G. Soares³², J.Y. Sohn²⁸, A. Sokolenko⁴⁴, O. Samoylov³³, O. Soto^{3,4}, O. Steinkamp³⁷, W. Sutcliffe³⁷, S. Takahashi²³, I. Timiryasov⁶, V. Tioukov²⁰, D. Treilleⁿ, R. Tsenov², P. Ulloa^{3,5}, E. Ursov⁹, A. Ustyuzhanin¹⁰, G. Vankova-Kirilova², G. Vasquez³⁵,³ S. Vilchinskii⁴⁴, C. Visone^{20,d}, S. van Waasen¹³, R. Wanke¹⁵, J.M. Webb¹¹, C. Weiser¹¹, J. Wenk¹¹, P. Wertelaers³⁵, I.M. Wösthelrich⁹, M. Wurm¹⁵, S. Yamamoto²⁴, D. Yilmaz³⁹, S.M. Yoo³⁰, C.S. Yoon²⁸, A. Zaitcev³³, J. Zamora Saa^{3,4}, T. Zholdybayev²⁷,

¹ Ghent University, Ghent, Belgium

² Faculty of Physics, Sofia University, Sofia, Bulgaria

³ Millenium Institute For Subatomic Physics At High-Energy Frontier - SAPHIR, Chile

⁴ Universidad Andrés Bello (UNAB)ⁱ, Santiago, Chile

⁵ Universidad De La Serena (ULS)ⁱ, La Serena, Chile

⁶ Niels Bohr Institute, University of Copenhagen, Copenhagen, Denmark

⁷ IJCLab, CNRS, Université Paris-Saclay, Orsay, France

- ⁸ *Georgian Technical University, Tbilisi, Georgia*
- ⁹ *Institut für Physik, Humboldt-Universität zu Berlin, Berlin, Germany*
- ¹⁰ *Constructor University, Bremen, Germany*
- ¹¹ *Physikalisches Institut, Universität Freiburg, Freiburg, Germany*
- ¹² *Universität Hamburg, Hamburg, Germany*
- ¹³ *Integrated Computing Architectures (ICA — PGI-4), Forschungszentrum Jülich GmbH (KFA), Jülich , Germany*
- ¹⁴ *Karlsruhe Institute of Technology, Karlsruhe, Germany*
- ¹⁵ *Institut für Physik and PRISMA Cluster of Excellence, Johannes Gutenberg Universität Mainz, Mainz, Germany*
- ¹⁶ *Universität Siegen, Siegen, Germany*
- ¹⁷ *Sezione INFN di Bari, Bari, Italy*
- ¹⁸ *Sezione INFN di Bologna, Bologna, Italy*
- ¹⁹ *Sezione INFN di Cagliari, Cagliari, Italy*
- ²⁰ *Sezione INFN di Napoli, Napoli, Italy*
- ²¹ *Laboratori Nazionali dell'INFN di Gran Sasso, L'Aquila, Italy*
- ²² *Aichi University of Education, Kariya, Japan*
- ²³ *Kobe University, Kobe, Japan*
- ²⁴ *Nagoya University, Nagoya, Japan*
- ²⁵ *College of Industrial Technology, Nihon University, Narashino, Japan*
- ²⁶ *Toho University, Funabashi, Chiba, Japan*
- ²⁷ *Institute of Nuclear Physics, Almaty, Kazakhstan*
- ²⁸ *Physics Education Department & RINS, Gyeongsang National University, Jinju, Korea*
- ²⁹ *Gwangju National University of Education^d, Gwangju, Korea*
- ³⁰ *Sungkyunkwan University^d, Suwon-si, Gyeong Gi-do, Korea*
- ³¹ *University of Leiden, Leiden, The Netherlands*
- ³² *Laboratory of Instrumentation and Experimental Particle Physics(LIP),Portugal*
- ³³ *Joint Institute for Nuclear Research, Dubna, Russia*
- ³⁴ *Institute of Physics, University of Belgrade, Serbia*
- ³⁵ *European Organization for Nuclear Research (CERN), Geneva, Switzerland*
- ³⁶ *Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland*
- ³⁷ *Physik-Institut, Universität Zürich, Zürich, Switzerland*
- ³⁸ *Middle East Technical University (METU), Ankara, Turkey*
- ³⁹ *Ankara University, Ankara, Turkey*
- ⁴⁰ *H.H. Wills Physics Laboratory, University of Bristol, Bristol, United Kingdom*
- ⁴¹ *University of Liverpool, Liverpool, United Kingdom*
- ⁴² *Imperial College London, London, United Kingdom*
- ⁴³ *University College London, London, United Kingdom*
- ⁴⁴ *Taras Shevchenko National University of Kyiv, Kyiv, Ukraine*
- ^b *Università di Bari, Bari, Italy*
- ^c *Università di Cagliari^j, Cagliari, Italy*
- ^d *Università di Napoli “Federico II“, Napoli, Italy*
- ^e *Associated to Gyeongsang National University, Jinju, Korea*
- ^f *Consorzio CREATE^j, Napoli, Italy*
- ^g *Università della Basilicata^j, Potenza, Italy*
- ^h *Università di Napoli Parthenope^j, Napoli, Italy*
- ⁱ *Università degli Studi del Sannio di Benevento^j, Benevento, Italy*
- ^j *Associated to SAPHIR, Chile*
- ^k *Associated to Università di Napoli “Federico II“, Napoli, Italy*
- ^l *Currently at the University of Bergamo, Bergamo, Italy*
- ^m *Also at the University of Pisa, Pisa, Italy*
- ⁿ *Individuals who made crucial contributions to the LoI*
- ^o *Currently at KEK, Tsukuba, Japan*
- ^p *Also at Pegaso University, Napoli, Italy*

We acknowledge contributions from A. Anokhina, E. Atkin, T. Barbe, A. Bagulya, A.Y. Berdnikov, Y.A. Berdnikov, M. Chernyavskiy, T. Colin, C. Chatron, V. Dmitrenko, T. Enik, A. Etenko, O. Fedin, K. Filippov, G. Gavrilov, V. Giuliaeva, V. Golovtsov, D. Golubkov, S. Gorbunov, M. Gorshenkov, V. Grachev, V. Grichine, N. Gruzinskii, O. Id Bahmane, Yu. Guz, D. Karpenkov, M. Khabibullin, E. Khalikov, A. Khotyantsev, V. Kim, N. Konovalova, I. Korol'ko, Y. Kudenko, P. Kurbatov, V. Kurochka, L. Le Mao, E. Kuznetsova, V. Maleev, R. Mauny, A. Malinin, A. Mefodev, O. Mineev, P. Moyret, S. Nasybulin, B. Obinyakov, N. Okateva, A. Petrov, D. Podgrudkov, M. Prokudin, T. Roganova, V. Samsonov, E.S. Savchenko, A. Shakin, P. Shatalov, T. Shchedrina, V. Shevchenko, A. Shustov, M. Skorokhvatov, S. Smirnov, N. Starkov, P. Teterin, S. Than Naing, S. Ulin, A. Ustyuzhanin, Z. Uteshev, L. Uvarov, D. Valencon, K. Vlasik, A. Volkov, R. Voronkov, A. Zelenov.

HI-ECN3 Project Team³⁵

C. Ahdida, M. Aversa, B. Balhan, D. Banerjee, H. Bartosik, D. Belohrad, E. Berthome, J. Braken, K.E. Buchanan, T.A. Bud, F. Butin, M. Calviani, G.P. Cnudde, J. Coleman-Smith, A. Colinet, P.M. Curran, J.J.S. Currie, P. Santos Diaz, M. Dole, F. Dragoni, A. Dutruel, L. Esposito, V. Ferrentino, M.A. Fraser, Y. Gaillard, H.G. Gavela, L. Gentini, S.D. Giovannantonio, B. Goddard, D.R. Gomez, E.C. Gonzalez, A. Gorn, J. Grenard, D. Grenier, C.D. Gutierrez, A. Harrison, G.S. Humphreys, A. Huncikova, R. Jacobsson, R.K. Janardhan, L.O. Jorat, I. Josifovic, N. Kahn, S.W. Kamugasa, T. Kolstad, M. Lazzaroni, K.B. Li, M. Liebsch, M.D. Liu, G. Mazzola, F. Metzger, L.J. Nevay, E. Nowak, M. Parkin, F. Baltasar Dos Santos Pedrosa, O. Pinto, P.G. Pisano, T. Prebibaj, N. Quinquis, J. Ridewood, R. Rinaldesi, G. Romagnoli, F. Roncarolo, O.R. Rubiras, M.L.D.D. Santos, P. Schwarz, S. Sgobba, P.A. Sota, F.W. Stummer, B.M. Sutil, I.A. Vaquero, F.M. Velotti, C. Vendeuvre, P. Wojtyla, D.O. Wasik, W. Weterings, R.F. Ximenes, C. Zamantzas, N. Zaric, D. Zemkova

Acknowledgements

The SHiP Collaboration wishes to acknowledge the support from the National Research Foundation of Korea with grant numbers of 2018R1A2B2007757, 2018R1D1A3B07050649, 2018R1D1A1B07050701, 2017R1D1A1B03036042, 2017R1A6A3A01075752, 2016R1A2B4012302, and 2016R1A6A3A11930680, from the FCT - Fundação para a Ciencia e a Tecnologia of Portugal with grant number CERN/FIS-PAR/0030/2017, and from the TAEK of Turkey.

Outside of the SHiP collaboration and the HI-ECN3 Project, we acknowledge in particular, for their contribution to:

- HI-ECN3 project: F.L. Aberle, O. Aberle, R.G. Alia, C. Bernard, J. Bernhard, P. Ber-

treix, C. Biot, J. Blanc, P. Boisseaux-Bourgeois, J. Borburgh, R.A. Bozzi, A.-M. Branza, J. Bremer, S. Burger, M.D. Castro, N. Charitonidis, R. Charousset, H. Danielsson, G. Dumont, Y. Dutheil, J.C.H. Emonds-Alt, B.M.S. Ferrando, M. Ferro-Luzzi, A.T.P. Fontenla, A. Funken, F. Galleazzi, S. Girod, E. Grenier-Boley, R. Haddad, T. Ladzinski, J. Lendaro, G. Lerner, K. Leventaki, S. Marsh, C. Moelck, R. Mompot, C. Mucher, M. Munos, G. Papotti, Y.P. Pira, B. Pittet, S. Prodon, O.B. Prouteau, M. Remta, J. Tan, M. Zamkovsky, T. Zickler

- superconducting spectrometer magnet: T. Barbe, H. Bajas, C. Chatron, T. Colin, O. Id Bahmane, L. Le Mao, R. Mauny, P. Moyret, D. Tommasini, D. Valencon,
- and to the development of the SHiP detectors: M. Andreini, L. Di Giulio.