



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

DOTTORATO DI RICERCA IN

FISICA

Ciclo 38

Settore Concorsuale: 02/A1 - FISICA SPERIMENTALE DELLE INTERAZIONI
FONDAMENTALI

Settore Scientifico Disciplinare: FIS/01 - FISICA SPERIMENTALE

UPGRADES OF THE SND@LHC DETECTOR DURING LHC RUN 3 AND SEARCH
FOR MUON NEUTRINOS FROM PROTON-PROTON COLLISIONS AT
 $\sqrt{s}=13.6\,\mathrm{TeV}$

Presentata da: Giulia Paggi

Coordinatore Dottorato

Alessandro Gabrielli

Supervisore

Luigi Guiducci

Co-supervisore

Marco Dallavalle



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

DOTTORATO DI RICERCA IN FISICA

CICLO XXXVIII

Settore Concorsuale: 02/A1 - FISICA SPERIMENTALE DELLE INTERAZIONI FONDAMENTALI

Settore Scientifico Disciplinare: FIS/01 - FISICA SPERIMENTALE

Upgrades of the SND@LHC detector during LHC Run 3 and search for muon neutrinos from proton–proton collisions at $\sqrt{s} = 13.6$ TeV

PRESENTATA DA:

Giulia Paggi

SUPERVISORE:

Prof. Luigi Guiducci

COORDINATORE DOTTORATO:

Chiar.mo Prof. Alessandro Gabrielli

CO-SUPERVISORE:

Dott. Gaetano Marco Dallavalle

Esame Finale - Anno 2026

Abstract

This thesis presents the SND@LHC experiment, a forward detector at the CERN LHC. SND@LHC studies neutrinos produced in proton–proton collisions at the LHC, bridging the gap between low-energy reactor experiments and astrophysical neutrino observations. All neutrino flavours are studied; in particular, a small but significant sample of tau neutrino interactions is expected. The experimental setup required to detect such elusive particles also opens sensitivity to feebly interacting particles, offering a window into physics beyond the Standard Model. Moreover, SND@LHC will enable the study of charmed hadrons production in proton–proton interactions and provide precision tests of the Standard Model through the lepton flavour universality of the weak interaction.

Since its commissioning in 2022, SND@LHC has undergone an extensive process of detector characterisation and optimisation, in which this work plays a role spanning both hardware developments and data analysis.

I contributed to the upgrade of the Veto System, crucial for rejecting charged-particle backgrounds in neutrino searches. An additional module was constructed, tested, and installed before the restart of 2024 LHC operation, reducing the system inefficiency by an order of magnitude and achieving uniform performance across the full target acceptance.

I took part in a test-beam campaign in 2023, which was used to calibrate the detector response to hadronic showers. The resulting calibration constants are a crucial step toward neutrino energy reconstruction.

At the same time, I began studying charged-current ν_μ interactions using the 2024 data, which benefited from improved background suppression. Selection strategy, control samples, and preliminary results are presented.

Finally, I contributed to a second upgrade in March 2025: the installation of drift-tube modules to enhance muon identification and tracking. Early performance studies indicate a performance consistent with expectations.

Collectively, these results enhance the SND@LHC detector’s capability for precision neutrino measurements at the LHC and broaden its discovery reach.

Contents

Introduction	4
1 The Physics Case for SND@LHC	6
1.1 Neutrinos at LHC	6
1.1.1 Neutrino Physics at SND@LHC	8
1.2 Feebly interacting particles	15
1.2.1 Feebly interacting particles at SND@LHC	16
2 The SND@LHC Detector	18
2.1 Veto System	20
2.2 Target	21
2.2.1 Emulsion Cloud Chambers	22
2.2.2 Scintillating Fibre tracker	24
2.3 Hadronic calorimeter and Muon System	26
2.3.1 Hadronic calorimeter	26
2.3.2 Muon system	26
2.4 Data acquisition	28
2.4.1 Event Builder	30
3 Veto System Upgrade	33
3.1 Motivation	34
3.2 Design and assembly	35
3.3 Performance with Cosmic Rays	37
3.4 Installation in TI18	40
3.5 Commissioning	40
3.6 Measured Veto Inefficiency	43
4 Calibration for Hadronic Showers	47
4.1 Test Beam 2023: setup and data taking	48
4.2 Shower Tag Algorithm	50
4.3 Energy calibration	54

4.3.1	Ordinary Showers	56
4.3.2	Late Showers	57
4.4	Shower Direction	59
5	Search for muon neutrinos produced in proton-proton collision in 2024 electronic detector data	61
5.1	Dataset and Preselection	62
5.1.1	Neutrino candidates sample preselection	62
5.1.2	Muon background sample preselection	63
5.2	Event Selection	65
5.2.1	Timing studies in TI18 data	65
5.2.2	Custom Shower Tagging	73
5.2.3	Track	81
5.2.4	Validation on 2022-2023 data	88
5.3	Next steps	92
6	Muon System Upgrade	94
6.1	Design	95
6.1.1	CMS Drift Tubes chambers	95
6.1.2	The MiniDTs	96
6.1.3	Application in SND@LHC	98
6.2	Installation in TI18	101
6.2.1	Mechanic support	102
6.2.2	Gas System	102
6.2.3	MiniDTs Readout electronics and integration in SND@LHC DAQ .	102
6.3	MiniDT Performance	107
6.3.1	Event matching	108
6.3.2	MiniDT preliminary efficiency	109
6.4	Next steps	111
	Conclusion	113
A	Additional Figures	115
A.1	Timing studies	115
A.2	Fiducial area	119
A.3	Centroid plots	121

Introduction

In this thesis, I present several works conducted within the SND@LHC experiment, which is dedicated to studying neutrino interactions in an energy range previously unexplored in laboratory experiments. Together with FASER ν , SND@LHC is dedicated to the study of properties of neutrinos produced in proton-proton collisions, at an energy bridging the low part of the spectrum from reactor experiments and neutrinos from astrophysical accelerators. The study of neutrinos in this context is not only motivated by the currently limited observation, especially for ν_τ , but also by the opportunity it provides to investigate the general properties of particles produced in proton-proton collisions and to perform precise tests of the Standard Model, in particular the lepton flavour universality of the weak interaction.

The hardware requirement necessary to identify elusive particles like neutrinos opens the possibility of investigating the presence of feebly interacting particles, allowing for a direct search of physics beyond the Standard Model.

This PhD project started shortly after the beginning of data taking by SND@LHC. As a newly commissioned detector, SND@LHC has undergone an extensive characterisation process, during which its performance was evaluated and potential areas for optimisation or upgrade were identified. This process aimed to ensure optimal operation throughout the data-taking period, which is scheduled to continue until the end of LHC Run 3. For this reason, my work comprises both hardware projects and analysis-related studies.

In Chapter 1, the physics of neutrinos at LHC is presented, as well as the aim of the SND@LHC experiment.

Chapter 2 gives a detailed description of the SND@LHC detector.

Chapter 3 focuses on the Veto System upgrade, which led to the installation of an additional module in March 2024, before the restart of LHC data taking. This intervention proved highly effective, reducing the system inefficiency by an order of magnitude and achieving this level of performance across the entire target area.

Chapter 4 reports the result of the test beam campaign of 2023, aiming to calibrate the energy response of the detector to hadronic showers, crucial to reconstruct the incoming neutrino energy. Thanks to the data collected in this test beam, calibration constants were measured and could be used in the search I focused on.

Chapter 5 summarises the work I have done to search for muon neutrino charged

current interaction with data collected in 2024, exploiting the upgrade of the Veto System. While the search is still ongoing, preliminary results are presented.

Finally, Chapter 6 describes a second upgrade of the detector, done in March 2025, to improve muon identification and tracking. The newly installed modules use drift tube technology, significantly slower than the scintillating material and never employed in SND@LHC. For this reason, the project proved challenging; nevertheless, the modules were successfully installed in March 2025. While final performance tests are still ongoing, the preliminary data indicate that the behaviour is in line with expectations.

Chapter 1

The Physics Case for SND@LHC

Contents

1.1	Neutrinos at LHC	6
1.1.1	Neutrino Physics at SND@LHC	8
1.2	Feebly interacting particles	15
1.2.1	Feebly interacting particles at SND@LHC	16

1.1 Neutrinos at LHC

The Standard Model (SM) of particle physics has demonstrated exceptional success in accurately describing microscopic physical phenomena. However, it is widely recognised that the SM exhibits significant limitations, particularly in the context of cosmological and astrophysical observations. Among the most prominent of these are the existence of dark matter (DM), the baryon-antibaryon asymmetry observed in the Universe, and the presence of neutrino masses and oscillations. These phenomena provide compelling evidence for the necessity of physics beyond the Standard Model (BSM).

Neutrinos allow precise Standard Model tests and are powerful probes for physics beyond the SM. Neutrino interactions - in the form of neutrino-nucleon (νN) cross-section (see Figure 1.1) - have been measured in the energy regime below 350 GeV [1] and astrophysical observations are starting to populate the ultra-high-energy (UHE) regime, from tens of TeV and above [2][3]. The region between 350 GeV and 10 TeV is currently the target of neutrino experiments at LHC, namely the Scattering and Neutrino Detector (SND@LHC) and FASER ν .

The data reported in Figure 1.1 are mostly from measurements of ν_μ , and while extensive data exist for muon neutrinos, electron and tau neutrino measurements remain comparatively limited. The Gargamelle experiment [5] measured electron neutrino-nucleon




Figure 1.1: Available measurements of the neutrino cross-section [4]. The thick dashed curve is a prediction of deep-inelastic scattering (DIS), averaged between ν and $\bar{\nu}$.

cross-sections up to energies of 12 GeV. At higher energies, the E53 experiment [6] confirmed lepton universality in neutrino cross-sections up to 200 GeV using a bubble chamber. Moreover, the ZEUS experiment [7] investigated the process $e^-p \rightarrow \nu_e X$ by selecting events with missing energy in the final state. Interpreting this missing energy as being carried away by an electron neutrino within the Standard Model framework, the results were found to be consistent with the predicted neutrino cross-section.

Concerning tau neutrinos, only a limited number of events have been observed to date, primarily by the DONuT and OPERA experiments. The DONuT experiment at Fermilab achieved the first direct observation of tau neutrinos ν_τ [8], reporting nine tau neutrino candidate events, without distinguishing between neutrinos and antineutrinos, and did not investigate interactions involving other neutrino flavours. The OPERA experiment later identified ten tau neutrino candidates resulting from muon neutrino oscillations and also recorded 35 electron neutrino ν_e interactions [9].

The LHC then offers a unique opportunity to study neutrino interactions in the energy range from several hundred GeV to a few TeV, with a significant contribution from tau neutrinos, whose interactions remain largely uncharted.

The concept of using the LHC as a tau neutrino source was first proposed in the early 1990s by De Rujula [10], Winter [11], and Vannucci [12], who suggested placing a detector to intercept the very forward neutrino flux arising from b and c decays at pseudorapidity $|\eta| > 7$.

Indeed, neutrino production at LHC [13] originates from leptonic decays of W and Z

bosons, and decays of heavy-flavour hadrons, pions, and kaons. Dedicated studies [14] have shown that neutrinos from W decays contribute substantially to the flux at intermediate pseudorapidities ($4 < |\eta| < 5$), with about 30% being of tau flavour. In such events, the associated charged lepton could be detected by central detectors like ATLAS and CMS, enabling flavour identification.

Figure 1.2 shows the neutrino energy versus absolute pseudorapidity for events involving heavy-flavour (b and c) in proton-proton (pp) collisions at a centre-of-mass energy of $\sqrt{s} = 14$ TeV. Neutrinos from heavy quark decays are abundant and highly forward, with tau neutrinos comprising approximately 5% of the flux. Because of the Lorentz boost along the beam axis, their energies can reach the TeV scale at large pseudorapidity, as can be better appreciated in Figure 1.3: in $4 < |\eta| < 5$ the spectrum is confined to low energies while in $|\eta| > 6.5$ it reaches very high energies. Instead, for neutrinos from W production, there is a high-energy contribution mainly from leptonic decays in the $4 < |\eta| < 5$ region.

1.1.1 Neutrino Physics at SND@LHC

In 2020, the SND@LHC Collaboration proposed a new detector [15] positioned off-axis in the secondary LHC tunnel TI18, at ~ 480 m from ATLAS interaction point (the detector position and geometry are better described in Chapter 2), to measure the $pp \rightarrow \nu X$ cross-section for all three neutrino flavours in the pseudorapidity range $7.2 < \eta < 8.6$. The off-axis location is particularly well-suited for studying heavy-quark production in a pseudorapidity range that remains unexplored, as it lies beyond the reach of current LHC experiments. In this setup, ν_τ and $\bar{\nu}_\tau$ originate predominantly from heavy-hadron decays, whereas ν_e and $\bar{\nu}_e$ mostly stem from charmed hadron decays, with minor contamination from kaons at lower energies, where the interaction cross-section with the detector is lower. Meanwhile, ν_μ and $\bar{\nu}_\mu$ receive significant contributions from pion and kaon decays. In this η range, the softer energy spectrum of ν_μ from pion and kaon decays allows differentiation between neutrinos from charm and pion or kaon decays based on their energy, crucial to study charmed-hadron production in pp collisions and constrain the gluon parton distribution function at small values of the *Bjorken* x , the scaling variable identified as the fractional longitudinal momentum carried by the scattered parton [16], one of the goals of SND@LHC, further discussed later in this section. The flux of muon, electron, and tau neutrinos versus their pseudo-rapidity and energy is shown in Figure 1.4.

Charmed-hadron production in pp collisions

Charmed-hadron production in proton-proton (pp) collisions has been extensively studied at the LHC, but only at lower pseudorapidities. For example, the LHCb experiment provided precise measurements of charm and beauty production in the region $|\eta| < 4.5$ [20].




Figure 1.2: Scatter plot of all neutrino flavours energy versus pseudorapidity $|\eta|$ [13]. Top: in b and c decays. Bottom: in pp events with W production.

Figure 1.5 shows the flux of incoming neutrinos and anti-neutrinos in the pseudorapidity region covered by SND@LHC, normalised to an integrated luminosity of 150 fb^{-1} . Between $7.2 < \eta < 8.6$, electron neutrinos and anti-neutrinos are mostly produced in the decays of charmed hadrons. This makes SND@LHC uniquely suited to perform an indirect measurement of charm production at high pseudorapidity via the detection of ν_e and $\bar{\nu}_e$ interactions.

A detailed simulation developed by the CERN EN-STI group, that includes the full propagation of particles through the LHC infrastructure, predicts that approximately 10% of the ν_e and $\bar{\nu}_e$ interactions in the detector arise from kaon decays, particularly K_S^0 ,




Figure 1.3: Energy spectra for neutrinos [13]. Top: from b and c when neutrino has $4 < |\eta| < 5$ (red) and when $|\eta| > 6.5$ (blue). Bottom: from W in $4 < |\eta| < 5$.

with energies below 200 GeV. The contribution from beauty hadron decays produced at ATLAS interaction point (IP1) has been estimated using the PYTHIA event generator and amounts to roughly 3%.

Assuming that the deep-inelastic charged-current cross-section for electron neutrinos follows Standard Model predictions, supported by ZEUS measurements [7] in their SM interpretation, the observed $\nu_e/\bar{\nu}_e$ flux can be used to probe charmed-hadron production in the SND@LHC acceptance, once detector effects are unfolded and the kaon contribution is subtracted. Using the method described in [17], the inclusive $pp \rightarrow \nu_e X$ cross-section can be determined with a total uncertainty of approximately 15%, primarily driven by




Figure 1.4: Neutrino flux as a function of ν energy and pseudorapidity for muon (left), electron (right) and tau (bottom) neutrinos [17].

the systematic uncertainties associated with the unfolding process.

Kaon contributions are modelled using multiple event generators, which show some variation in yield; however, they all agree that these decays populate only the lower energy region below 200 GeV. As a result, considering the kaon component introduces an additional systematic uncertainty in the extracted neutrino flux from heavy-quark production. Neutrinos arising from sequential decays ($D \rightarrow K \rightarrow \nu_e$) are conservatively included in the kaon category, although their contribution is negligible ($\sim 1.5\%$).

Overall, the measurement of charmed-hadron production using neutrino interactions in SND@LHC can be achieved with a statistical uncertainty of approximately 5% and a total systematic uncertainty of about 35%. In addition to providing valuable constraints on charm production at high rapidity, these measurements can also yield information on the gluon parton distribution function (PDF) in a previously unexplored region of Bjorken- x [17]. The measurement of charmed hadrons can be directly related to the production of open charm in the same rapidity region, due to the strong kinematic corre-




Figure 1.5: Energy spectrum of the different types of incoming neutrinos and anti-neutrinos [17] normalised to 150 fb^{-1} . Simulation obtained with DPMJET3 (Dual Parton Model, including charm) [18] embedded in FLUKA [19].

lation between the charm quark and the resulting hadron, as illustrated in Figure 1.6. At the LHC, the dominant partonic process responsible for associated charm production is gluon–gluon scattering. A preliminary analysis [21], summarised in Figure 1.7, reveals the correlation between the momentum fractions x_1 and x_2 of the interacting gluons probed by SND@LHC. The results indicate that the lower of the two momentum fractions typically reaches values around 10^{-6} , a region where the gluon parton distribution function remains entirely unconstrained.

Lepton flavour universality test

The detection of all three types of neutrinos allows for tests of lepton flavour universality using two specific ratios of events: ν_e/ν_μ and ν_e/ν_τ .

In the selected pseudorapidity interval, tau neutrinos are primarily produced through $D_s \rightarrow \tau \nu_\tau$ decays, followed by the subsequent decay of the tau lepton. According to the PYTHIA event generator, approximately 8% of ν_τ originate from beauty hadron decays. Thus, both ν_e and ν_τ can be assumed to arise mainly from the semileptonic and fully leptonic decays of charmed hadrons. Unlike ν_τ , which is produced exclusively in D_s decays, ν_e is generated in the decays of all charmed hadrons, including D^0 , $D^\pm$, D_s , and Λ_c . As a result, the ν_e/ν_τ ratio depends solely on charm hadronization fractions and decay branching ratios. Systematic uncertainties related to the charm-quark production




Figure 1.6: (Left) Correlation between pseudo-rapidities of the parent charm and the charmed hadron. (Right) Only events where the neutrino is in the SND@LHC acceptance are selected [21].




Figure 1.7: Correlation between x_1 and x_2 momentum fractions for events with neutrinos in the SND@LHC acceptance [21].

mechanism cancel out, making this ratio a probe of the ν -nucleon interaction cross-section ratio for these two neutrino species. Consequently, measuring this ratio serves as a test of lepton flavour universality in neutrino interactions.

It is also possible to use ν_e/ν_μ to test lepton flavour universality in neutrino interactions. However, the situation differs significantly for ν_μ compared to ν_τ . While the available ν_μ flux is much more abundant, it is heavily contaminated by pion and kaon decays. This contamination is concentrated at low energies but is predicted to decrease to

about 35% above 600 GeV, remaining stable with energy. Since charmed hadron decays have nearly equal branching ratios into electron and muon neutrinos, the ν_e/ν_μ ratio is unaffected by systematic uncertainties in the weighted branching fractions. However, the yield of muon neutrinos from pion and kaon decays is influenced by two key systematic uncertainties: one arising from pion and kaon production in this pseudorapidity range, and the other from their propagation through the LHC machine elements along the beamline.

Measurement of Neutral Current to Charged Current ratio

At the high energies characteristic of LHC neutrinos, deep-inelastic scattering (DIS) is by far the dominant interaction mechanism. In this regime, neutrino scattering on nuclei can be accurately approximated as an incoherent sum over interactions with individual protons and neutrons. Neutrino flavour is identified in charged-current (CC) interactions by detecting the associated charged lepton in the final state, whereas neutral-current (NC) interactions are flavour-insensitive. When summing over neutrinos and anti-neutrinos, the ratio of NC to CC deep-inelastic cross-sections at a given energy can be expressed as a function of the Weinberg angle [22].

Assuming that the energy spectra of neutrinos and anti-neutrinos are approximately equal, the same relationship holds for the observed interaction rates, since the neutrino flux cancels out in the ratio. Consequently, a measurement of the NC-to-CC ratio can, in principle, be used to extract the value of the Weinberg angle.

At SND@LHC, this ratio will be measured not as a precise determination of the Weinberg angle which is already known with 1% accuracy [23], but rather as an internal consistency check of the detector performance and interaction modelling. The propagation of particles through the LHC machine introduces an asymmetry between neutrino and anti-neutrino fluxes, particularly for muon neutrinos, which predominantly originate from non-prompt decays of pions and kaons. However, this asymmetry is estimated to remain below 20% in all energy bins and is most relevant in the low-energy region, where neutrino cross-sections are relatively small. As a result, the corresponding systematic uncertainty in the NC-to-CC ratio for observed events is estimated to be below 5%.

An additional source of systematic uncertainty arises from the unfolding of instrumental effects, including muon identification efficiency and the subtraction of neutron-induced backgrounds that can mimic NC neutrino interactions. This contribution amounts to approximately 10%. Taking both sources into account, the total systematic uncertainty on the ratio is estimated at the level of 10%, while the statistical uncertainty, driven by the limited number of observed NC events, is expected to be around 5%.

1.2 Feebly interacting particles

One of the main challenges in particle physics today is determining the microscopic identity and the cosmological origin of dark matter (DM). The theoretical landscape is broad and spans many orders of magnitude in the mass/coupling parameter space.

A widely studied candidate in the framework where dark matter is a thermal relic from the early universe is the Weakly Interacting Massive Particle (WIMP), a hypothetical particle with a mass in the GeV–TeV range that interacts with the visible sector via forces of weak interaction strength. Searches for WIMPs are actively ongoing [24], and there is a growing interest in the search for Feebly Interacting Particles (FIPs) [25] driven by the absence of discoveries in various direct detection dark matter experiments, which have reached unprecedented precision, as well as in collider experiments that have now large datasets and are being analyzed with increasingly sophisticated techniques. Most of these efforts focus on masses in the 10 GeV to 1 TeV range, leaving the lower mass region largely unexplored. Recently, significant attention has shifted toward light dark matter (LDM) in the keV–GeV mass range [26]. The limited sensitivity at lower masses stems from the challenge of detecting the extremely soft recoils expected due to the non-relativistic nature of galactic dark matter.

While direct-detection experiments searching for elastic nuclear recoils lose sensitivity rapidly for dark matter particles with masses below a few GeV, accelerator-based experiments operating at the intensity frontier offer a promising and complementary approach in this mass regime [26]. Indeed, in a fixed-target experiment at an accelerator, the candidate dark matter particle is ultrarelativistic and therefore the experimental challenges in its detection are different and less demanding than for experiments searching for dark matter originating in the galaxy. The most stringent constraints on light feebly interacting particle searches currently come from electron beam-based experiments, which benefit from precise control of backgrounds and accurate measurement of final-state particle energies. These include b-factories such as Belle II experiment [27], which employs the missing-mass technique, and experiments at fixed-target facilities such as NA64 [28], operating at the CERN SPS, and the proposed LDMX [29].

Collisions delivered by accelerators offer a general source of feebly coupled particles, but the main LHC experiments, such as ATLAS, CMS, and LHCb, have limited sensitivity to searches for light dark matter or feebly interacting particles. In particular, the visible decay modes of a light DM mediator often occur outside the detector volume due to the mediator’s long lifetime, or are obscured by overwhelming backgrounds in the low-mass region most relevant to the relic abundance parameter space. Invisible decays, on the other hand, typically yield low missing transverse momentum, since the dominant production mechanisms at the LHC do not impart a significant boost to the DM particles. Strategies involving additional handles—such as the production of hard jets or vector bosons in association—tend to increase model dependence and significantly reduce the accessible cross section. Instead, SND@LHC can investigate their scattering off electrons

or nucleons, making it possible for a direct search of LDM candidates.

1.2.1 Feebly interacting particles at SND@LHC

SND@LHC is designed primarily to study high-energy neutrino interactions in the far-forward region of the ATLAS interaction point but it is well suited to search for new light particles such as LDM, which may interact with SM particles via weakly coupled mediators [30].

A key challenge in probing LDM-nucleon scattering lies in distinguishing signal events from the dominant neutral current neutrino background. Although not background-free, the signal can be extracted statistically by identifying an excess in the reconstructed energy spectra. Elastic and inelastic scattering channels provide complementary sensitivity, depending on the mediator mass. Elastic scattering dominates for light mediators $\lesssim 1$ GeV, while inelastic processes become more relevant at higher mediator masses. In the inelastic case, an excess in the NC/CC interaction ratio could signal the presence of new physics, since this ratio is well-predicted within the SM.

However, if systematic uncertainty quoted in Section 1.1.1 cannot be reduced exploiting data-driven analyses, the uncertainty on this measurement is expected to be too large for the required sensitivity. An additional handle on separating LDM from neutrino events is time-of-flight measurement. Given the 200 ps timing resolution of SND@LHC, massive LDM particles ($\gamma \lesssim 100$) could be distinguished from relativistic neutrinos, even if this requires LDM masses in the tens of GeV range (see Figure 1.8), where production and scattering probabilities are typically suppressed.

SND@LHC may also be sensitive to the decays of mediators into visible final states, such as di-lepton or meson pairs. While some of these decay modes suffer from neutrino-induced backgrounds, they can often be distinguished kinematically, especially at high energies. However, to produce observable scattering signals large couplings are necessary, resulting in the mediator's decay length too short to allow decay within the detector volume. This opens the possibility for complementary coverage: large couplings probed through scattering, and smaller couplings through displaced decays.




Figure 1.8: Sensitivity curves for a Time-of-flight measurement, illustrated from 1σ up to 5σ in the plane $(M_{\text{NP}}, p_{\text{NP}})$ of mass and momentum of the New Physics candidates to be detected in the SND@LHC, assuming a time resolution of ~ 200 ps [21].

Chapter 2

The SND@LHC Detector

Contents

2.1	Veto System	20
2.2	Target	21
2.2.1	Emulsion Cloud Chambers	22
2.2.2	Scintillating Fibre tracker	24
2.3	Hadronic calorimeter and Muon System	26
2.3.1	Hadronic calorimeter	26
2.3.2	Muon system	26
2.4	Data acquisition	28
2.4.1	Event Builder	30

The Scattering and Neutrino Detector at LHC [15] is a compact and stand-alone experiment to perform measurements with neutrinos produced at the LHC in an unexplored pseudo-rapidity region of $7.2 < \eta < 8.6$, complementary to all the other experiments at the LHC. The experiment is located 480 m downstream of IP1 in the unused TI18 tunnel (see Figure 2.1). The detector, shown in Figure 2.2, comprises a hybrid system based on an 800 kg target mass of tungsten plates, interleaved with emulsion and electronic trackers, followed downstream by a calorimeter and a muon system. The configuration efficiently distinguishes between all three neutrino flavours, opening a unique opportunity to probe the physics of heavy flavour production at the LHC in the region inaccessible to ATLAS, CMS and LHCb.

The detector, described in the following sections, occupies the full available volume within the TI18 tunnel to achieve the required pseudo-rapidity coverage.

The top and side views of the detector installed inside the tunnel are shown in Figure 2.3. Several aspects of the tunnel geometry, such as its dimensions, the sloped floor,




Figure 2.1: The location of the SND@LHC detector in the TI18 tunnel, accessible from the LHC one via UJ18, 480 m from the ATLAS interaction point (IP1).

and the relative distances between the tunnel walls, floor, and the nominal collision axis, posed significant constraints on the detector design.

As a result, the layout of the detector was optimised to balance these geometrical limitations with key physics objectives: achieving a precise calorimetric energy measurement (requiring approximately 10 interaction lengths λ_0), ensuring high muon identification efficiency (which demands sufficient material to absorb hadrons), and covering the necessary azimuthal angular acceptance in the target region. The requirements for energy measurement and muon identification impose a lower limit on the detector's length. However, given the tunnel constraints, this competes with the need for a wide azimuthal angular coverage, which directly affects the total neutrino flux intercepted and, consequently, the number of observed interactions. Additionally, the tunnel geometry constrains the detector position with respect to the collision axis, as shown in Figures 2.3 and 2.4.

The detector was installed in 2021 and has been running successfully since 2022.

An upgraded version of the detector [33] is foreseen to take data during the operation of HL-LHC. The foreseen increase in instantaneous luminosity is incompatible with the use of nuclear emulsion films as high-resolution tracking detectors in the neutrino target region (better described in Section 2.2.1), motivating their replacement with silicon trackers. In addition to improved statistical precision, significant upgrades to the experimental setup are foreseen, including the installation of a magnetic spectrometer. This will enable the separate identification of neutrino and antineutrino interactions for both muon and tau flavours, allowing in particular the first direct observation and study of tau antineutrinos, extending tests of lepton flavour conservation, and enhancing sensitivity to potential exotic phenomena.




Figure 2.2: Layout of the SND@LHC detector.

2.1 Veto System

In this section, the original SND@LHC detector configuration will be described, while Chapter 3 is dedicated to the upgrade work and the new configuration of the veto system in place since 2024.

The Veto system is crucial to veto charged particles crossing the instrumented volume. Given its role, it is located upstream of the rest of the detector. The technology used for this detector is EJ200 [34] scintillating bars with dimension $42 \times 6 \times 1 \text{ cm}^3$ read by silicon photomultipliers (SiPMs).

The Veto system was designed to consist of two planes, each made of seven bars, for a total area of $42 \times 42 \text{ cm}^2$. Each bar is wrapped in Mylar foil to avoid light sharing between adjacent bars. The Mylar foil, even if $20 \mu\text{m}$ thick, creates a small inefficient region in between bars as it can be seen in single-plane efficiency plots shown in Section 3.6. Each bar is read out on both sides by 8 Hamamatsu S14160-6050HS SiPMs [35] with a dimension of $6 \times 6 \text{ mm}^2$, $50 \mu\text{m}$ pitch. The SiPMs are placed on a single printed circuit board (PCB), as shown in Figure 2.5, for each side of a detector plane and read out by TOFPET2 ASICS [36], a 64-channel readout and digitisation ASIC which incorporates signal amplification circuitry, discriminators, charge integrators, charge-to-digital converters (QDC) and time-to-digital converters (TDC).

The stacked scintillator bars for each plane are enclosed within an aluminium frame featuring 4 mm thick walls. Readout PCBs are mounted on 4 cm wide rectangular flanges at both ends of the frame, serving as structural end caps. A scheme of the mechanical




Figure 2.3: Top and side views of the SND@LHC detector in the TI18 tunnel [31].




Figure 2.4: Frontal and side views of the SND@LHC detector with the collision axis position. The side view includes an schematic representation of a ν_μ charged-current interaction in the target with a hadronic shower sampled in the emulsion target, target trackers, and hadronic calorimeter, and a muon track visible in the muon system [32].

structure is in Figure 2.6. Each end is sealed with an aluminium cover to ensure light tightness and to function as a heat sink for the front-end (FE) board. The FE board is accommodated in a dedicated groove on the side opposite the PCB. The two frames comprising the veto system are connected and supported by a compact mechanical structure, and to ensure complete coverage, the two planes are aligned to be vertically staggered by 2 cm, corresponding to $1/3$ of a bar.

2.2 Target

The target region uses a hybrid technology. The target mass of ~ 800 kg is divided into five modules, called walls, of emulsion cloud chambers (ECCs), and each module is followed by a Scintillating Fibre (SciFi) plane.




Figure 2.5: PCB hosting the 56 SiPMs to read out a veto plane side. It is possible to see the subset of SiPMs dedicated to the readout of a bar.




Figure 2.6: Left: View of a Veto plane showing the placement of SiPMs and read-out PCB. Centre: Frontal and top view of a plane. Right: Side view of a plane.

The region of the target is enclosed in a box made of 4 cm-thick 30% borated polyethylene and 5 cm-thick acrylic layers, which serves mainly to shield from low-energy - from a few MeV to a few hundreds of MeV - neutrons produced in the interaction of proton beams with the residual gas inside the LHC beam pipe. It provides a background reduction of a factor of 10^{-7} . At the same time, this box acts as an insulation chamber to maintain humidity and temperature constant in the correct condition for the emulsion and SciFi tracker to work.

2.2.1 Emulsion Cloud Chambers

In an emulsion cloud chamber, layers of emulsion films, capable of reaching sub-micrometric resolution, are alternated to layers of a passive material, in this case tungsten, which is the neutrino target.

Each ECC wall, shown in Figure 2.7, covers an area of $384 \times 384 \text{ mm}^2$ but it is further

divided into four units called bricks. A brick consists of 60 emulsion film layers with a transverse size of $192 \times 192 \text{ mm}^2$ and a thickness of $\sim 300 \mu\text{m}$, interleaved by 59 tungsten layers with the same area and 1 mm thick.




Figure 2.7: Scheme of an emulsion cloud chamber wall, made of four bricks, composed of 60 layers of emulsion film interleaved by 59 tungsten layers.

Nuclear emulsion films are the most compact and lightest tracking detectors with sub-micrometric position and milliradian angular resolution. Nuclear emulsions consist of AgBr crystals in a gelatin binder, sensitive to minimum ionising particles (MIPs). As a MIP passes through, it leaves a trail, which is enhanced through chemical development, forming visible silver grains. The emulsion film has two $70 \mu\text{m}$ -thick sensitive layers on either side of a $170 \mu\text{m}$ -thick plastic base. By matching the hits left on both active layers, the track's slope can be measured with milliradian accuracy.

Tungsten was chosen as the target material to maximise the interaction rate per unit volume. Its small radiation length of 3.5 mm allows good performances in the electromagnetic shower reconstruction in the ECC, and the low intrinsic radioactivity makes tungsten a suitable material for an emulsion detector.

Each brick has a total thickness of $17 X_0$ or $0.5 \lambda_{\text{int}}$, so the whole target accounts to $\sim 85 X_0$ or $2.5 \lambda_{\text{int}}$.

Emulsion films have a maximum luminosity they can integrate, given that analysis is possible up to a track density not larger than $\sim 10^7 \text{ cm}^{-2}$ [13]. The optimal value was determined to be $\sim 10^5 / \text{cm}^2$, that given the background measurements in the detector location [13] [37] corresponds to $\sim 20 \text{ fb}^{-1}$. This led to multiple sets of emulsion in each year of data taking. All the irradiated emulsions are then chemically developed in a process of five steps:

- Development (Fujifilm Developer + Fujifilm Starter): to turn a cluster made of a

Table 2.1: Summary of scanning microscope available for SND@LHC emulsion scan.

Laboratory	CERN	Napoli	Bologna	LNGS	Salerno	Santiago	Nagoya
Scanning station	4	3	2	2	1	1	1

few silver atoms into a visible metallic silver grain

- Stop (Acetic acid): to immediately stop the development process
- Fix (Fujifilm UR-F1): to dissolve all residual silver halide crystals, leaving the metallic silver to form the image
- Wash (Water): to remove all the silver thiosulphate complexes in the emulsion
- Thickening (Glycerine + Fujifilm Driwell): to re-inflate the emulsions to their original thickness

The chemical treatment of SND@LHC emulsion films is performed in the dark room at CERN, instrumented to develop 1200 emulsion films in one week.

After preparing the emulsions, readout is performed in dedicated laboratories equipped with automated optical microscopes. The collaboration made an effort to set up 14 scanning stations in 7 different laboratories, summarised in Table 2.1. The laboratory located in Bologna hosts two microscopes, one shown in Figure 2.8, whose operation I have significantly contributed to over the past three years: I commissioned the initial microscope; I then assembled and commissioned the second, while coordinating shifter organisation and training.

2.2.2 Scintillating Fibre tracker

Scintillating Fibre modules are between the ECCs walls. SciFi contribution is crucial both in the case where its timing and spatial information is used to discriminate events reconstructed in the emulsions layers, but also standalone, as it provides key measurements about the development of electromagnetic and hadronic showers through the target volume. Moreover, SciFi is essential for reconstructing isolated tracks, such as muon tracks, necessary to disentangle piled-up events in the emulsion and to perform event reconstruction between emulsion walls.

The main components employed in the target tracker, fibre mats and photodetectors, were developed by the EPFL group for the LHCb SciFi Tracker [38].

The detection layers are based on the blue light-emitting Kuraray SCSF-78MJ scintillating fibres of $250\,\mu\text{m}$ diameter. The layers are made of six densely packed staggered fibre layers glued together, shown in Figure 2.9 forming fibre mats of $133 \times 390\,\text{mm}^2$ length. Three mats are integrated into a plane with less than $500\,\mu\text{m}$ dead zones between mats.




Figure 2.8: One of the two microscopes in the Bologna scanning laboratory.




Figure 2.9: A SciFi fibre mat is composed of six layers of fibres glued with a titan oxide-loaded epoxy glue to suppress cross-talk between fibres.

The readout consists of the photo-detector (S13552 SiPM multichannel arrays by Hamamatsu) at the end of the fibre module, a short Kapton flex PCB holding the photo-detector and signal connectors and the front-end electronics board. The light tightness of the module is ensured by a seal on the flat Kapton flex, the aluminium module frame and an opaque Tedlar sheet on both sides of the module. The SiPM multichannel array is optimised for low-intensity light detection. For application in SND@LHC, the SciFi performance has to be tuned to maximise the hit detection efficiency at an acceptable noise rate. The final threshold chosen for operation produces a noise rate of 25 Hz per channel, which poses no problem to the event builder (described in Section 2.4.1) and can be efficiently suppressed by the online noise filter.

The SciFi tracker readout is also performed through the TOFPET2 ASIC, as for the Veto System, to read out not only the time information but also an accurate amplitude

measurement required for shower shape information, and to improve energy measurement, electron-hadron separation, and shower centre-of-gravity measurement.

2.3 Hadronic calorimeter and Muon System

The hadronic calorimeter and muon system are located downstream of SciFi and are necessary to identify passing-through muons and, together with SciFi, they serve as a sampling hadronic calorimeter, enabling measurement of the energy of hadronic jets. These two systems consist of eight modules, interleaved by iron slabs 20 cm thick, corresponding to $9.5 \lambda_{\text{int}}$. The active planes are made out of the same scintillating material of the Veto planes, described in Section 2.1.

2.3.1 Hadronic calorimeter

The first five modules, called UpStream (US) chambers, are dedicated to hadronic calorimetry and time tagging of the passing particles. The bars have a dimension of $82.5 \times 6 \times 1 \text{ cm}^3$. Each plane is made of 10 bars, which add up to a total active area of $82.5 \times 60 \text{ cm}^2$. Each bar is wrapped in aluminised Mylar and read out by 8 SiPMs: six Hamamatsu S14160-6050HS [35] - described in the Veto section - and two Hamamatsu S14160-3010PS ($3 \times 3 \text{ mm}^2$, with a $10 \mu\text{m}$ pitch) SiPMs. The readout PCB is shown in Figure 2.10.



Figure 2.10: PCB hosting the SiPMs to read out a US plane side.

The mechanical structure of the upstream layers closely resembles that of the veto system and can be seen in Figure 2.11. The sides flare out at the ends to accommodate the size of the PCB. The overall width is fixed to match the width of the iron blocks, with added tolerance. As a result, the PCB end cap extends beyond the iron blocks, allowing the thickness of the detector plane between the iron walls to be minimized to that of the scintillator bar plus the aluminum frame wall. The total available space between the iron blocks for the detector plane is 26 mm, comprising 16 mm for the detector itself and two 5 mm gaps on either side to account for uncertainties in the iron wall thickness.

2.3.2 Muon system

The final modules, known as DownStream (DS), are dedicated to muon tracking. The muon system includes three modules, each composed of two layers of thin bars - one




Figure 2.11: Left: View of the mechanical structure of a US plane. Centre: Frontal and top view. Right: Side view.

horizontal and one vertical - providing a hit spatial resolution better than 1 cm. An additional module with only vertical bars is included to enhance redundancy. Horizontal planes are made of 60 bars, each measuring $82.5 \times 1 \times 1 \text{ cm}^3$, visible in Figure 2.12, and each bar is read out by a single Hamamatsu S14160-6050HS SiPM, same as the one employed in the Veto system, on both sides. The vertical planes have 60 bars $63.5 \times 1 \times 1 \text{ cm}^3$, read out by a single SiPM from the top. An example of PCB to readout the DS module is in Figure 2.13. To ensure optimal readout, the bar edges are placed in close contact with the SiPMs without using optical gel, while the other side is carefully closed with flat aluminized Mylar pieces to act as efficient mirrors.

The structure of the downstream planes is illustrated in Figure 2.14. The space available between the iron blocks for these planes is 34 mm, comprising 29 mm for the scintillator bars and aluminium frame, plus a 2.5 mm gap on each side to account for tolerances between the iron walls. Since the corner of the most downstream layer would interfere with the tunnel wall, the entire structure is laterally shifted by 22 mm, without changing its dimensions.

Given the similarities in all technology employed in the hadronic calorimeter and muon system detectors with the veto system planes, the readout is done in the same way using the TOFPET2 ASIC.

The muon system was upgraded in 2025 with the addition of two drift chambers. This




Figure 2.12: DS bars aligned in their frame. The module will be completed by adding the readout PCB and closing the frame.




Figure 2.13: Close up of PCB hosting the SiPMs to read out a DS plane side.

upgrade is further discussed in Chapter 6.

2.4 Data acquisition

As described in the previous sections, the SND@LHC detector includes two active sub-systems: scintillator bars read out by SiPMs, used in the veto (see Section 2.1), hadron calorimeter and muon systems (see Section 2.3), and scintillating fibres read out by SiPMs, used in the SciFi tracker (see Section 2.2).

All active systems share the same Data Acquisition (DAQ) architecture [17]. As mentioned above, to read the signals produced by SiPM, front-end boards are used, based on the TOFPET2 ASIC developed by PETsys [36]. A single FE board, shown in Figure 2.15, hosts two TOFPETs that read up to 64 channels each, for a total of 128 for the whole board.

This ASIC was well suited for the SND@LHC application since it records all the hits that exceed a predefined threshold and provides energy and timestamp information




Figure 2.14: Left: View of the mechanical structure of a US plane. Centre: Frontal and top view. Right: Side view.

with a resolution of 30 ps. Moreover, it allows per-channel energy and time threshold configuration, channel disabling for noise control, and time and energy calibration through pulse injection.

Up to four FE boards are connected to a DAQ board that collects the digitised data. The DAQ board, in Figure 2.16, features a Cyclone V FPGA [39] and has:

- Four high-speed connectors for FE boards;
- TTCrx ASIC [40] with optical fibre input for receiving clock and sync signals;
- 1 Gb Ethernet interface for data/control communication;
- LEMO connector for SiPM bias voltage.

A total of 36 DAQ boards operate synchronously with the LHC bunch crossing clock. A TTC crate receives and distributes LHC clock and orbit signals via the Beam Synchronous Timing (BST) system. This is crucial since it keeps SND@LHC readout aligned to the LHC machine by providing the bunch crossing (BX) frequency (40.08 MHz), the orbit signal (11.245 kHz), beam mode, collision energy, and GPS time.




Figure 2.15: The FE board. The two TOFPET2 ASICs (centre) and the SiPM connectors (left) are visible.

These signals are received by a BST-TTC board, which extracts the clock/orbit signals, applies jitter cleaning via a PLL, and distributes the signals to the detector electronics using the TTC system.

In the absence of a valid BST signal, the board can generate substitute clocks and transmit beam metadata over Ethernet to the DAQ server.

The readout system is trigger-less, meaning all hits are continuously recorded and transmitted to a DAQ server for storage and analysis. Noise suppression is handled at the front-end level through configurable thresholds or channel disabling.

The DAQ system, illustrated schematically in Figure 2.17, is completed by:

- Control PC: located in TI18, manages the TTC crate and environmental sensors to ensure the correct humidity and temperature in the cold box;
- DAQ server: located in a surface building, collects and processes data from all boards. If communication with the DAQ server is lost, data are temporarily stored on local microSD cards in the DAQ boards: a 1 GB card can buffer several hours of data.

2.4.1 Event Builder

Each DAQ board transmits all recorded hits to the DAQ computer server, where the event building is performed. Hits are grouped into events based on their timestamps and saved to disk in the form of ROOT files. In addition to the hit data, the DAQ boards also receive and forward periodic triggers from the Timing, Trigger and Control (TTC) system. These so-called heartbeat triggers are used by the event-building software on the server to monitor synchronization across boards and to verify correct operation, even in the absence of data.




Figure 2.16: The DAQ board. The four FE board connectors are visible on the left, the TTCrx and optical receiver on the bottom-right, the Enclustra Mercury SA 1 module [41] in the centre.

The event building process consists of two main steps, as illustrated in Figure 2.18. In the first step, hits from different DAQ boards that occur within a time window of four clock cycles, so 25 ns, are grouped into a single event. The timestamp assigned to the event corresponds to the earliest hit within that window. Events are then filtered and processed online before being written to disk. The exact processing pipeline depends on the selected configuration, but always includes an online noise suppression filter.

Noise filtering is implemented in two stages:

1. events must satisfy a minimum requirement on the number of DAQ boards registering a given number of hits: this is a fast pre-selection that eliminates most spurious events generated by isolated noise hits;
2. hits are then grouped by the plane in which they occurred, enabling more advanced selection criteria based on event topology.




Figure 2.17: Scheme of the SND@LHC DAQ system.




Figure 2.18: Scheme of the SND@LHC DAQ system.

Chapter 3

Veto System Upgrade

Contents

3.1	Motivation	34
3.2	Design and assembly	35
3.3	Performance with Cosmic Rays	37
3.4	Installation in TI18	40
3.5	Commissioning	40
3.6	Measured Veto Inefficiency	43

As discussed in 2.1, the Veto system plays a critical role in rejecting backgrounds from charged particles crossing the instrumented volume. The measured inefficiency of the Veto system during the 2022-2023 data-taking was $(7.8 \pm 2.8) \times 10^{-8}$ on a fiducial area of about 64% of the whole target.

This value is not satisfactory, not only due to the limited coverage area, but also because of the expected background rates. Given the muon flux measured by SND@LHC [42] is $\sim 2.07 \times 10^4 \text{ fb/cm}^2$, and the veto area is $42 \times 42 \text{ cm}^2$, the 2022–2023 dataset of 68.6 fb^{-1} yields an expected number of $\sim 2 \times 10^9$ muons traversing the detector. To ensure adequate background rejection, considering the expectations of 250 fb^{-1} integrated in LHC Run 3 (2022-2025), the Veto inefficiency should therefore be at the level of 10^{-10} .

To enhance the system’s performance, both in terms of efficiency and coverage, a major upgrade was implemented during the Year-End Technical Stop (YETS) of 2023–2024. This included the installation of a third Veto plane and the repositioning of the system to fully cover the active target volume.

As a result of this intervention, the inefficiency was significantly reduced to $(4.9 \pm 1.9) \times 10^{-9}$, close to the expectation, now measured over the entire target area [43].

This chapter provides a detailed account of the motivations for the upgrade, the design of the new Veto 3 module, and the performance improvements obtained.

3.1 Motivation

Using data collected in 2022, the SND@LHC Collaboration reported the observation of neutrino interactions from muon neutrinos originating at the LHC IP1 [32]. The main point of the analysis was observing the rare neutrino signal over the prevailing background implied adopting a selection with strong rejection power, designed to yield a clean set of events.

To obtain the required rejection power, the two most upstream planes of the SciFi tracker were included in the Veto logic. While this reduced background contamination, it also severely limited the fiducial volume, rejecting approximately 92.4% of neutrino charged-current interactions. This effect was particularly severe in the lower part of the detector, where the neutrino flux is highest. Figure 3.1 illustrates this issue by showing reconstructed tracks projected onto the Veto planes.




Figure 3.1: The extrapolated position of the reconstructed SciFi track at Veto 1 (left) and Veto 2 (right) for events with less than 13 fired Veto channels. The red square encloses the fiducial area used for the observation of neutrino interactions.

The initial Veto configuration, described in Section 2.1, exhibited an overall inefficiency of the order of 10^{-4} in 2022. This non-optimal performance was mainly due to the lack of time alignment between planes during a large fraction of the data-taking period.

Following timing improvements, the inefficiency decreased to 10^{-6} in 2023, but was still constrained by the detector dead time [43].

Only selecting the smaller area enclosed in the red squares in Figure 3.1, it was possible to achieve the value quoted before of $(7.8 \pm 2.8) \times 10^{-8}$.

To improve the efficiency of the Veto system, during the Year End Technical Stop (YETS) of 2023-2024, the collaboration decided to:

- lower the Veto by ~ 2.7 cm to provide a better coverage of the bottom edge of the neutrino tungsten target,

- add a third layer (Veto 3) of scintillating bars aligned vertically to improve the Veto system efficiency

as shown in Figure 3.2.




Figure 3.2: 2022-2023 Veto system layout with two planes with horizontal bars (left). The upgraded Veto system with a third plane with vertical bars (right).

3.2 Design and assembly

The Veto 3 module, made of scintillating bars read out by SiPMs, uses the same technology as Veto 1 and 2, described in Section 2.1.

While Veto 1 and 2 have horizontally disposed bars, the Veto 3 bars are aligned vertically.

The vertical segmentation added by the new plane is crucial to recover efficiency since it strongly reduces the area over which all three planes overlap. This configuration effectively lowers the probability of two particles crossing the veto system so close in time and space that the second one goes across undetected due to the readout dead time.

For this new plane, each bar dimension is $46 \times 6 \times 1 \text{ cm}^3$, extending by 2 cm above and below the Veto 1 and 2 range (see Figure 3.2). Seven bars are framed side by side between two 4 mm thick aluminium sheets, to constitute a vertical module of 42 cm width

and 46 cm height. The bars are pushed against each other to minimise the dead area, as can be seen in figure 3.3.

The Veto 3 bars are read out at the top edge by 56 Hamamatsu Photonics MPPC S14160-6050HS SiPMs ($6 \times 6\text{mm}^2$, $50\mu\text{m}$ pitch) [35] hosted on a printed circuit board and subdivided into seven equally spaced bunches, which are carefully aligned to be centred between bar edges, so that each bar is uniquely read out by 8 SiPMs (see Figure 2.5).

As for DS vertical planes (described in 2.3), to achieve optimal readout performance, the bar edges are placed in direct contact with the SiPMs without optical gel, while the opposite ends are sealed with flat aluminised Mylar to serve as mirrors.

I contributed to the assembly of the Veto 3 module, spending three months at CERN working within a small team. My activities included including the alignment of the SiPMs, verification of light tightness, and initial tests of the readout electronics. Particular attention was devoted to monitoring the power consumption as an early diagnostic tool for identifying potential issues.




Figure 3.3: The seven scintillating bars wrapped in aluminised Mylar and aligned side by side.

3.3 Performance with Cosmic Rays

After the construction and verification of the absence of shorts, abnormal currents, and other obvious defects, I conducted the initial performance tests of the Veto 3 module with a dedicated laboratory setup. The primary objectives were to verify that all readout channels were operational and to investigate potential asymmetries in efficiency along the scintillator bars, particularly between the end directly coupled to the SiPMs and the opposite end, which relies on a reflective surface.

The test setup, shown in Figure 3.4, consisted of the Veto 3 module placed on top of a Downstream chamber. The DS was most suited for this test since its bars are arranged in two perpendicular layers to provide both horizontal (x) and vertical (y) position information. These layers are referred to as DS H (horizontal) and DS V (vertical), respectively. As described in section 2.3, the DS V layer is read out on one side, with a single SiPM at the top of each bar, while DS H bars are read out at both ends, labelled DS R and DS L.




Figure 3.4: Initial tests setup. The Veto plane is on the top of the DS chamber.

The Veto 3 module and DS chamber were aligned horizontally to facilitate the detection of vertical cosmic ray tracks. The DS chamber served both as a trigger and as a reference tracker for determining the location of cosmic ray impacts within the Veto 3 module.

To identify clean cosmic ray events suitable for the efficiency analysis, the following selection criteria were applied:

- Signals were required from DS R, L, and V. In case of multiple hits, the one with the highest integrated charge (QDC) was retained.

- The DS R and L signals had to originate from the same bar to ensure correct position reconstruction in the horizontal layer.
- All signals had to exceed a predefined QDC threshold to further reject noise hits.

Once suitable events were selected, the performance of the Veto 3 module was evaluated by examining the hit distribution and QDC response for each channel. All channels were found to be functional. As illustrated in Figure 3.5, a peak in the hit distribution was observed in the third bar, deviating from the uniform distribution expected for cosmic rays. This effect was attributed to the front-end gain being set too high, as further evidenced by the QDC distribution in Figure 3.6, where a saturation effect is visible for the same bar. This problem was solved by modifying the readout threshold and gain settings as described in Section 3.5.




Figure 3.5: Hit distribution per Veto channel in the selected cosmic rays events. The red lines delimit the scintillator bars.

To reduce background and noise contributions in the efficiency computation, the Veto hit was required to be spatially consistent with the reconstructed DS track, and a QDC threshold was applied to the Veto signal as well.

To assess the impact of single-sided readout, the Veto 3 module was divided into three zones of 12 cm along the bar length, based on the distance from the SiPM side. The first and last 3 cm of each bar were excluded from the analysis to avoid edge effects.

The resulting efficiency as a function of position is shown in Figure 3.7. The hit rate is found to be uniform along the bar, and the relative efficiency is computed as:

$$\text{Relative efficiency} = \frac{\text{Efficiency at far end}}{\text{Efficiency at SiPM end}} = 0.95 \pm 0.02$$




Figure 3.6: QDC distribution per Veto channel in the selected subset of cosmic ray events. The red lines delimit the bars.




Figure 3.7: Efficiency Veto 3 bars as a function of the hit position. Each bin is the average of 12 cm. Bin 0 is the side furthest from the readout SiPM. Bin 1 is the central region. The last group is closest to the readout side.

This result confirms that the single-sided readout does not lead to a significant loss in efficiency at the far end of the bars.

No performance issues were observed during the laboratory cosmic ray tests, and the Veto 3 module was approved for installation in the TI18 experimental area.

3.4 Installation in TI18

With the addition of the third plane, the overall thickness of the Veto system increased and its position had to be modified accordingly. A 3D integration model revealed that this new configuration was incompatible with the existing floor layout of TI18. Consequently, a local excavation was required to accommodate the upgraded Veto system.

Before the civil engineering work, the area was prepared by dismantling the previous Veto system and the SciFi modules. Local cabling and piping were either disconnected or relocated, and protective films were installed to shield surrounding equipment and surfaces from dust generated during the excavation.

Upon completion of the excavation, the area was cleaned and restored in compliance with CERN Radiation Protection requirements. The SND@LHC components were then reassembled and reconnected. The upgraded Veto system was installed on a redesigned support structure and repositioned upstream of the detector. SciFi cooling pipes were rerouted to enable the integration of the new configuration.

I was actively involved during all installation period before and after the excavation period.

The final layout of the SND@LHC detector, incorporating the upgraded Veto system and associated modifications, is shown in Figure 3.8.

3.5 Commissioning

After installation, recabling, and integration into the SND@LHC DAQ system, the third Veto plane underwent a commissioning phase in TI18.

I conducted the first part of this commissioning during the final months of the YETS, using cosmic ray data. The initial objective was to verify that the Veto plane had not sustained any damage during transportation and installation, and to ensure that all channels were operational. This was confirmed by inspecting the hit distribution per channel, as shown in Figure 3.9. Figures 3.9 and 3.10 also demonstrate that the thresholds and gain settings used in TI18, identical to those employed for the other two Veto planes, successfully resolved the noise and QDC saturation issues observed during the laboratory tests described in Section 3.3.

Following these preliminary checks, I assessed the spatial efficiency of the Veto in its final configuration.

The efficiency of Veto 3 was measured by selecting cosmic ray events that crossed the plane. These events were identified by requiring coincident signals in Veto 2 and the first two SciFi modules, as illustrated in Figure 3.11. These conditions enabled the reconstruction of the particle trajectory using the SciFi modules, allowing precise determination of the expected impact point on the Veto 3 plane.

To minimise angular bias in the efficiency determination, an additional cut was applied




Figure 3.8: Final configuration of SND@LHC detector with upgraded Veto system.




Figure 3.9: Hit distribution in Veto 3 from cosmic ray events selected in TI18. The red lines delimit the scintillator bars.




Figure 3.10: Veto 3 QDC distribution in cosmic ray events selected in TI18. The red lines delimit the scintillator bars.




Figure 3.11: Schematic representation of the selection of cosmic ray events. The Veto planes are highlighted in green. The SciFi modules are identified in blue.

on the incident angle of the muons. Specifically, the displacement between hits in the two SciFi planes was required to be less than 1.5 cm in both the x and y directions.

These selection criteria can also be applied to muons originating from LHC collisions, which typically traverse the detector horizontally.

The Veto plane was considered efficient when a signal was recorded in the expected bar or in one of its immediate neighbours.

The resulting efficiency map is shown in Figure 3.12, and confirms a uniform and near-unity response across the entire detector surface. The relative efficiency along the bars, as defined in Section 3.3, was measured to be close to 100%. This can be appreciated in Figure 3.13, where the efficiency along the bars is shown. In this figure, the final value takes into account the contribution of all bars. The lower efficiency at $y \sim 40$ cm, the bar edge closer to the readout SiPM, is due to border effect and linked to the computation of the expected position. Thus, the commissioning confirmed that the single-ended readout does not degrade performance, consistent with the results from earlier performance tests conducted during construction.




Figure 3.12: Veto 3 efficiency in TI18 measured during commissioning.

3.6 Measured Veto Inefficiency

After the start of the proton-proton run of the LHC in 2024, the inefficiency of the upgraded Veto System was measured using passing-through muons [44]. The inefficiency of individual Veto planes was estimated by extrapolating reconstructed tracks from the SciFi and Downstream detectors to each Veto plane and checking for the presence of a corresponding signal. A plane is considered inefficient if there are no hits where one is expected because of a passing track.

The (x, y) positions of SciFi tracks for which there is no corresponding Veto plane hit




Figure 3.13: Veto 3 efficiency along the bar, averaged for all bars, in TI18. The lower efficiency at $y \sim 40$ is expected and due to border effects.

are shown in Figure 3.14. The single plane and whole system inefficiencies are reported in Table 3.1. The dead areas between scintillating bars are visible. This was observed also in previous years (Figure 3.1): regions outside the acceptance of a single Veto plane, particularly the horizontal band below Veto 1, still show a high density of tracks. However, these areas are covered by the sensitive regions of other Veto planes, as demonstrated in the combined inefficiency map in Figure 3.15.

The green rectangle in the inefficiency plots highlights the fiducial area of best system performance, for which an upper limit on the Veto system inefficiency of $(4.9 \pm 1.9) \times 10^{-9}$ is obtained.

Table 3.1: Veto plane and system inefficiencies calculated with passing-through muons using TI18 data collected between May and August 2024. The values are calculated for a fiducial area of $40 \times 48 \text{ cm}^2$ in ranges $-47 \text{ cm} \leq x \leq -7 \text{ cm}$ and $10 \text{ cm} \leq y \leq 58 \text{ cm}$, shown as green rectangles in Figures 3.14 and 3.15.

Veto 1	Veto 2	Veto 3	system
$(1.90 \pm 0.01) \times 10^{-3}$	$(2.65 \pm 0.01) \times 10^{-5}$	$(4.1 \pm 0.005) \times 10^{-4}$	$(4.9 \pm 1.9) \times 10^{-9}$

This fiducial area for neutrino interactions is now 1.6 times larger than the area used in the 2022-2023 analysis and fully covers the target region. Monte Carlo simulation studies (see Figure 3.16) indicate that this extension could lead to a $\sim 56\%$ increase in the number of observable neutrino interactions.




Figure 3.14: The extrapolated position of the reconstructed SciFi tracks at Veto 1 (left), Veto 2 (middle) and Veto 3 (right) for events without any fired Veto channel in that plane. The red square encloses the fiducial area used for the 2022-2023 analyses, while the green rectangle represents the fiducial area used in the 2024 Veto inefficiency study.




Figure 3.15: The extrapolated position of the reconstructed SciFi tracks at Veto 1 for events without any fired Veto channel in all three planes. The red square encloses the fiducial area used for the 2022-2023 analyses, while the green rectangle represents the fiducial area used in the 2024 Veto inefficiency study.




Figure 3.16: Position of muon neutrino interactions from Monte Carlo simulation. The dashed line represents the fiducial area used in 2022-2023, while the solid one encloses the acceptance of the upgraded Veto system.

Chapter 4

Calibration for Hadronic Showers

Contents

4.1	Test Beam 2023: setup and data taking	48
4.2	Shower Tag Algorithm	50
4.3	Energy calibration	54
4.3.1	Ordinary Showers	56
4.3.2	Late Showers	57
4.4	Shower Direction	59

In neutrino-nucleon (νN) collisions of all flavours, energetic interactions produce hadronic showers, as summarised in Figure 4.1. The hadronic shower energy, along with the direction of the shower itself and of the outgoing charged lepton, constitutes a key set of observables required to estimate the energy of the incoming neutrino in Charged Current (CC) interactions.




Figure 4.1: Scheme of different neutrino flavour interaction with a nucleon. Left: electron neutrino. Centre: tau neutrino. Right: muon neutrino.

Being such crucial information to reconstruct, a replica of the detector was exposed to

hadron beams with 100 to 300 GeV in the CERN SPS H8 test beam line (see Figure 4.2) in summer 2023 to calibrate the SND@LHC response.




Figure 4.2: Map of the CERN accelerator complex. The H8 line is located inside the North Area.

In this chapter, the most important results obtained from the test beam campaign are reported.

4.1 Test Beam 2023: setup and data taking

The HCAL system of the test beam detector was made of five Upstream (US) modules and one Downstream (DS) module, constructed using spare materials to assemble a system as similar as possible to the one installed in TI18.

I was involved in all stages of the assembly, from the preparation of the aluminized mylar, to the wrapping of the bars, to their arrangement and alignment inside the frames and with the readout PCBs, to ensuring the light tightness of each module, to test the readout electronics, to installing it in the dedicated location in H8 and taking data, and finally in analysing the collected data.

Upstream of the hadronic calorimeter (HCAL), the setup included a target composed of three iron walls, each 10 cm thick or $1\lambda_{int}$, instrumented with four SciFi tracking stations, each made of an x module and a y module, covering an area of $13 \times 13 \text{ cm}^2$,

provided by the Losanna group. Although this target had a reduced size compared to that of the SND@LHC experiment, the configuration allowed for a comprehensive study of hadronic shower development.

The setup, installed in H8, is visible in Figure 4.3.

The data-taking was designed to investigate the detector response as a function of both the hadron energy and the depth at which the interaction occurred within the target.

Data were collected using hadron beams of five different energies and three target configurations, as summarised in Table 4.1. The beams, produced colliding the 400 GeV protons from SPS onto Beryllium targets [45], consisted of positively charged hadrons (a mixture of protons and positive pions, referred to as the π^+ beam) with energies of 100, 140, and 180 GeV, as well as negatively charged pions (π^-) at 240 and 300 GeV, delivered at the CERN SPS H8 beam line. The beam energy spread was optimised to remain below 2%.

Beam intensities were tuned to approximately $\sim 5 \times 10^3$ particles per spill, the period during which particles are extracted from the synchrotron and delivered to the fixed-target experimental areas, counted using beam counter scintillators located before the target region. Higher intensities were avoided, as the large iron mass of the HCAL would have necessitated the installation of additional absorbers around the detector to mitigate the radiation background generated by hadronic interactions. To ensure optimal containment of showers even in the smaller SciFi modules, the beam was aligned with the centre of the detector.

Table 4.1: Number of events collected (in multiples of 10^6 events), and beam energies and target configurations used.

	π^+ 100 GeV	π^+ 140 GeV	π^+ 180 GeV	π^- 240 GeV	π^- 300 GeV
3 target walls	51	50	83	98	121
2 target walls	15	15	30	61	61
1 target walls	15	21	27	64	50

The DAQ system of the test beam was based on spare DAQ boards, as the ones described in Section 2.4. The readout system is asynchronous and triggerless, as in SND@LHC. However, since in the test beam setup no LHC timing signal was available to synchronise the DAQ boards, the internal system clock was used instead. It had a frequency of 160.316 MHz, corresponding to a clock cycle of about 6.25 ns.

Events were collected if:

- at least five out of eight SciFi planes had hits;
- all signals in HCAL were accepted, to guarantee no bias on shower shape analysis.




Figure 4.3: Detector setup in the SPS H8 test beam line with the beam coming from the right. From right to left: small-size Beam Counter scintillators (green), three target walls interleaved with four SciFi detectors (light blue), and the hadronic calorimeter (HCAL), consisting of 5 US planes (red) and 1 DS (violet) interspaced with iron walls.

The online event builder, described in Section 2.4.1, was modified specifically for the test beam application to allow events spanning up to 16 clock cycles or ~ 100 ns.

4.2 Shower Tag Algorithm

In the SND@LHC experiment, hadronic showers in the HCAL are observed downstream of an interaction target with a depth of approximately 2.5 interaction lengths (λ_{int}), to be compared to $1.5 \lambda_{\text{int}}$ in the test beam setup. The shower development observed in the US planes is strongly dependent on the position of the primary hadronic interaction within the target. Therefore, the first step in calibrating the response to hadronic showers is the development of a shower tagging algorithm that not only detects the presence of a hadronic shower and effectively distinguishes it from the muon background, but also provides an estimate of the shower's origin.

An essential preliminary study to the shower tagging was the selection of in-time hits. In the LHC data from 2022 and 2023 [46], a “ringing effect” was observed: delayed

hits and even entirely new events appeared shortly after genuine ones. Spatial patterns of SiPM activity in these secondary events are correlated with those of the preceding valid events. A similar phenomenon was identified in the test beam data. As discussed in Section 4.1, the event duration in the test beam configuration was 100 ns. Analysis revealed an anomalously high number of consecutive events separated by roughly this time interval (see Figure 4.4), prompting a more detailed investigation.




Figure 4.4: Left: time difference between two consecutive events. Right: multiplicity of hits, both SciFi and US, in consecutive events as a function of the time difference between two consecutive events.

Figure 4.5 shows the distribution of hit times within the 100 ns window for the HCAL DAQ boards. Most hits cluster around well-defined time-stamps for each DAQ board, peaking between 10 ns and 20 ns. Since the same DAQ board read out the first two US planes, its timing distribution displays two partially overlapping peaks. Notably, all distributions exhibit a broad shoulder extending beyond 30 ns. This long tail persists even after the nominal event window ends, and could lead to the creation of subsequent HCAL events immediately after the original one, contributing to the excess of closely spaced events observed in Figure 4.4.

To suppress noise-induced events, a filter condition is applied to the SciFi data, requiring more than ten hits across the full system since noise hits tend to be sparse and randomly distributed, unlike those produced by actual particles traversing the detector. Although this filter efficiently removes most spurious events, the periodic nature of the hit time distributions - with peaks recurring at intervals close to the DAQ clock period of 6.25 ns - allows some of the ringing events to pass the selection criteria. To mitigate this, only events with no prior event in the preceding 150 ns are selected for the analysis.

Within a selected event, SciFi hits are classified as in time if they lie within ± 0.5 clock cycles of the most probable value of the time distribution, denoted t_{ref} . The hit distribution and the in-time interval for one event are shown in Figure 4.6. US hits are




Figure 4.5: Timestamp distribution of hits in the DAQ boards collecting HCAL and Beam Counter scintillators data.

considered in time if recorded within $+3$ clock cycles (~ 19 ns) from t_{ref} (Figure 4.7). For the US stations, only hits from large SiPMs are used, as signals from small SiPMs were found to be significantly affected by cross-talk from adjacent large SiPMs.

The shower tagging algorithm operates exclusively on in-time hits. Multiple algorithmic strategies were evaluated, with the shower tagging via the hit density method proving the most effective. In this implementation, the algorithm searches for in-time hits in at least 35 channels within a 128-channel sliding window, in both the x and y directions of each SciFi station in the target. The most upstream SciFi station that satisfies this condition is identified as the location of the iron wall where the shower originated.

This tagging approach showed stable and consistent performance across varying run conditions, beam energies, and target configurations.

The possibility of modifying the test beam setup by adding and removing target walls placed between SciFi stations was key to developing the hit density-based tagging algorithm. Without any iron walls, the SciFi hit distributions in all four stations reflect the transverse profile of the incident beam. As iron walls are inserted into the target, this beam profile becomes progressively distorted in the downstream stations, due to the development of showers induced by hadronic interactions.

Since only a subset of incoming particles interacts within each target wall, the beam shape observed when no iron is present should still be observable by selecting events in which no shower has yet developed.

Based on this idea, the optimal threshold for the hit density could then be determined by comparing the beam profile obtained under two conditions: (1) when no iron wall is




Figure 4.6: Typical time distribution of SciFi hits in one event. SciFi hits are considered in time if they are within a ± 0.5 clock cycle window around the most probable value of the SciFi hit time distribution (in this plot, the range is bounded by the dashed lines). One clock cycle equals ~ 6.25 ns.




Figure 4.7: Scatter plot of the digitised charge (QDC) versus time for all SiPM signals of US stations; the time is considered with respect to the SciFi reference time t_{ref} . US signals are considered in time if earlier than 3 clock cycles (in this plot, the range is bounded on the right by the dashed line). One clock cycle equals ~ 6.25 ns.

located upstream of the SciFi plane, and (2) when the full target is in place, but only events without shower activity are selected, aiming to recover the original beam profile in this second case.

This analysis was carried out using the default timing window for SciFi hits, defined as $[-0.5; +0.5]$ clock cycles relative to t_{ref} . To ensure that showers are sufficiently developed when sampled by SciFi, a hit density working point of 35 hits in a 128-channel sliding window was chosen [46].

4.3 Energy calibration

Energy calibration studies were done on the dataset acquired with the three-wall target configuration. To ensure an unbiased energy calibration, the data samples collected in the target configuration with three walls were split into *calibration* and *test* samples.

For every event, the main quantities considered were: origin of shower as measured with the shower tagging algorithm, the digitized integral charge (QDC) of the SiPM signals summed over in time SciFi hits (QDC_{SciFi}) and QDC sum of in time US hits (QDC_{US}).

As shown in Figure 4.8, both the mean QDC_{SciFi} and QDC_{US} increase with beam energy. However, while the mean QDC_{SciFi} decreases the more downstream the shower originated along the target, the mean QDC_{US} increases the more downstream the shower started within the target. This behaviour is expected and can be modelled as:

$$E_{\text{hadron}} = k \times QDC_{\text{SciFi}} + \alpha \times QDC_{\text{US}} \quad (4.1)$$

where E_{hadron} is the hadron energy, k and α two calibration constants to be determined.

Data are further divided into calibration subsamples, each representing one of the 15 combinations of beam energy and the iron wall where the shower originated. For each subsample, it is possible to see a clear linear anti-correlation between the event distribution of QDC_{SciFi} and QDC_{US} . Two methods were used to extract the parameters k and α :

- a linear fit performed globally on the scatter plots, such as those shown in the top panel of Figure 4.9;
- a Principal Component Analysis (PCA): an iterative, event-by-event procedure that identifies the axis minimising the spread of the data, corresponding to the major axis of the elliptical distribution [47].

The fitting range is chosen by selecting the linear region of the scatter plot profiles (bottom of Figure 4.9).

The correlation between QDC_{SciFi} and QDC_{US} is strong when showers start in the first or second iron wall, and both methods reliably determine the principal axes of the ellipses. However, when the shower originates late (i.e. in the third iron wall), the




Figure 4.8: Top figure: mean QDC_{SciFi} observed in events with different beam energies and start of shower. The error bars represent the standard deviations of the corresponding QDC_{SciFi} distributions. As expected, the mean QDC_{SciFi} increases with the beam energy, but it shows a decreasing trend with increasing depth of the shower origin. Bottom figure: mean QDC_{US} observed in events with different beam energies and start of shower. The error bars represent the standard deviations of the corresponding QDC_{US} distributions. As expected, the mean QDC_{US} increases with the beam energy and also grows the more downstream the shower originates.

correlation weakens. In this case, only the last SciFi station (SciFi 4) samples the shower in the target, while the energy deposit peaks in the first US plane, potentially causing SiPM saturation, as discussed later. Thus, it was necessary to analyse ordinary and late showers separately.




Figure 4.9: Top: QDC_{SciFi} vs QDC_{US} scatter plot for 300 GeV beam energy and shower origins tagged in the second SciFi wall (SciFi 2). Bottom: the corresponding profile along the x axis.

4.3.1 Ordinary Showers

Ordinary showers are defined as those that develop sufficiently early to be sampled in the target by at least two SciFi stations in both the x and y projections. In the test beam setup, these correspond to showers originating in the first two iron walls. The calibration subsample is used to compute the constants, resulting in the average values of $k = (0.059 \pm 0.006)$ GeV/QDC and $\alpha = (0.0145 \pm 0.0010)$ GeV/QDC, where the uncertainties include a $\pm 5\%$ systematic component accounting for the slight energy dependence observed across the tested range.

These constants are then used to reconstruct the energy in the test data samples. The standard deviation of a Gaussian fit to the core of each reconstructed energy distribution provides an estimate of the energy resolution (Figure 4.10). For energies above 140 GeV, the distributions develop a low-energy tail, attributed to saturation effects in the first US station when large energy deposits occur. This saturation contribution increases with energy, remaining below 1% at 140 GeV and reaching approximately 10% at 300 GeV. A

more detailed study of this effect is presented in Section 4.3.2.




Figure 4.10: Reconstructed energy E_{hadron} for showers starting in SciFi 3, with different beam energies.

Figure 4.11 summarises the relative offset between the reconstructed and nominal beam energies for each sample. The error bars represent the standard deviation of the reconstructed energy divided by the nominal beam energy. For showers beginning in SciFi 2 and SciFi 3, the reconstructed energy agrees with the true hadron energy within $\pm 10\%$ across the full energy range. The figure also shows that the achieved energy resolution improves from 22% at 100 GeV to 12% at 300 GeV.

4.3.2 Late Showers

Late showers originate near the end of the target depth and are sampled by only one SciFi station, SciFi 4 in the test beam setup. Two main factors contribute to the large uncertainties obtained in this case:

- the single SciFi measurement exhibits large shower-to-shower fluctuations;
- the US signal is significantly affected by SiPM saturation, which distorts the energy response.

Regarding the SciFi calibration constant k , it is important to note that SciFi 4 is also used in the reconstruction of ordinary showers, alongside SciFi 2 and SciFi 3. Therefore, the conversion factor from QDC counts to energy must remain consistent. The same value of k derived from ordinary showers is then used for late showers.




Figure 4.11: Top: Relative offset between reconstructed energy E_{hadron} and beam energy E_{true} . The error bars are the standard deviation of the reconstructed energy divided by the beam energy. Bottom: Energy resolution (standard deviation of the reconstructed energy divided by the beam energy) as a function of the beam energy.

In contrast, the calibration constant α for the US system must be corrected. Due to SiPM saturation, the QDC response in the US layers reaches an upper limit, particularly for high-energy showers. The optimal value of α for late showers deviates from the value computed in ordinary showers, decreasing with increasing hadron energy. At 300 GeV, the decrease in QDC_{US} can reach up to 40%.

While a hadron-energy-dependent correction to α is feasible in test beam conditions, it cannot be applied to SND@LHC data where the hadron energy is not known a priori. Nevertheless, saturation is expected to be most severe in US 1, the first US station located roughly $1 \lambda_{int}$ downstream of the target. At this stage, the shower is still compact and tends to deposit its energy in only a few scintillating bars, saturating the corresponding SiPMs. By US 2, one interaction length further downstream, the shower has spread out, distributing its energy over more bars and reducing the likelihood of saturation.

This observation suggests a possible correction to the α value for US 1: using the energy measured in the rest of the detector, excluding US 1, as a rough estimator of the true shower energy. In exploring this approach, significant variability among SiPMs is observed, as they saturate at different thresholds. Therefore, any correction must be tailored specifically for the detector configuration.

It is important to emphasise that such a correction would only restore the average energy response. On an event-by-event basis, the reconstructed energy would still carry substantial uncertainty. For showers developing late in the target, only a lower bound on the hadron energy can be reliably determined.

4.4 Shower Direction

Reconstructing the direction of hadronic showers is crucial for studying neutrino interactions. Precise knowledge of the shower position and direction enables the correlation of the observed hadronic activity with the incident neutrino, thereby allowing an accurate determination of the neutrino energy.

The energy of the incoming neutrino, E_ν , assuming the masses of the produced particles are negligible with respect to their energy, so that $m_\nu \simeq m_\mu \simeq m_{\text{shower}} \simeq 0$, can be derived from the conservation of energy to be:

$$E_\nu \simeq E_{\text{shower}} \left(1 + \frac{\sin(\theta_{\text{shower}})}{\sin(\theta_\mu)} \right),$$

where θ_{shower} and θ_μ are the directions of the shower and the muon relative to the incoming neutrino direction, shown in Figure 4.12. Under the small-angle approximation, this relation can be further simplified to be:

$$E_\nu \simeq E_{\text{shower}} \left(1 + \frac{\theta_{\text{shower}}}{\theta_\mu} \right),$$

The achievable resolution depends on both the energy measurement of the hadronic shower and the precision in determining the muon direction.

The shower direction is estimated by performing a linear fit of the centroid positions across the detector planes as a function of depth. Separate linear fits are carried out for the x and y directions with respect to the z coordinate. The fit is applied only when more




Figure 4.12: Schematic representation of a muon neutrino CC interaction. The vector S represents the average shower direction to be reconstructed.

than two centroids are identified and both x and y coordinates are available in the same plane.

This method was developed and validated using test beam data, exploiting the measurement of incoming particle position [48].

The final validation must be performed on Monte Carlo samples, which are currently only partially available and require a detailed study before they can be effectively used, as discussed in Section 5.3.

Chapter 5

Search for muon neutrinos produced in proton-proton collision in 2024 electronic detector data

Contents

5.1	Dataset and Preselection	62
5.1.1	Neutrino candidates sample preselection	62
5.1.2	Muon background sample preselection	63
5.2	Event Selection	65
5.2.1	Timing studies in TI18 data	65
5.2.2	Custom Shower Tagging	73
5.2.3	Track	81
5.2.4	Validation on 2022-2023 data	88
5.3	Next steps	92

The observation of neutrino interactions at a collider marks a significant milestone in experimental high-energy physics.

Data taking at the LHC restarted in 2022, and the SND@LHC and FASER ν experiments have already independently achieved the first direct observations.

In particular, FASER ν reported the observation of muon neutrinos [49] and the first measurement of electron and muon neutrino cross sections using an emulsion detector [50].

The SND@LHC collaboration also reported a direct observation of muon neutrino charged current interactions based on a data sample of 36.8 fb^{-1} collected in 2022 [32], and later published a measurement on neutrino interactions with no muons in the final state, thus sensitive to CC interactions of either ν_e or ν_τ flavours, and neutral current neutrino interactions [51].

The ongoing analysis of muon neutrino interactions in a subset of the 2024 SND@LHC dataset aims to characterise neutrino-induced events in the updated detector configuration and to study hadronic shower observables by using the new algorithm developed with the 2023 test beam data, described in Sections 4.2 and 4.3.

As in previous searches conducted by the SND@LHC collaboration, this analysis focuses on identifying the most suitable set of selection cuts. The goal is to isolate events with the topology of charged-current muon neutrino interactions, to obtain a clean sample that remains statistically significant. This will enable a detailed study of the properties of neutrino events, which was not possible with data from the first observation due to the limited statistics available after the final selection.

5.1 Dataset and Preselection

The analysis of muon neutrino interactions in the 2024 SND@LHC dataset is ongoing, with the current study relying on a partial unblinding of the collected data. The total luminosity integrated during 2024 by the SND@LHC detector is $\sim 120 \text{ fb}^{-1}$, as can be seen in Figure 5.1. Given that the dataset is very large, the collaboration decided to limit the available sample to approximately $\sim 20 \text{ fb}^{-1}$, keeping the rest blinded. The runs selected for unblinding are distributed across the entire data-taking period to promptly identify and account for any possible systematics due to operational variations throughout the year.

5.1.1 Neutrino candidates sample preselection

The total number of events collected in the available runs is 1.14×10^{10} , so we decided to proceed with a two-step selection: an initial preselection with looser constraints to reduce the dataset to a more manageable size, followed by the actual selection applied to the preselected events.

The general strategy is to select events with no incoming charged particles detected by the Veto system, a hadronic shower developing within the target region, and enough hits in the DS stations, which can later be used to reconstruct a muon track exiting the detector.

This first set of cuts is composed of:

- no more than a veto hit, to allow for some inefficiency;
- at least two hits in both x and y of the DS system;
- a shower tagged in SciFi with the algorithm, better described in Section 5.2.2.




Figure 5.1: Luminosity delivered by LHC (solid green line) and recorded by SND@LHC (dashed green line) in 2024. The luminosity delivered and integrated by each emulsion target is also shown. Eleven targets were necessary to cover the 2024 data-taking period. The twelfth one, during the heavy ions campaign of LHC, was entirely passive as it had no emulsion films since the track density would be too high to properly profit from emulsions.

At this stage, there is no cut on the timing of hits inside a single event, as described also in Section 4.2, to replicate the study and determine the most suitable time window for the 2024 data.

A summary of the unblinded runs and their respective luminosity, as well as the preselection efficiency and number of selected events, can be found in Table 5.1.

5.1.2 Muon background sample preselection

The two background sources of this analysis are muon-induced with two different mechanisms:

- muons entering the detector undetected and producing showers due to bremsstrahlung or deep inelastic scattering (DIS);
- muons interacting in the surroundings of the detector and producing both charged particles that can reach the instrumented volume from the sides or neutral particles which simulate the neutrino signal.

With the applied analysis cuts, only muon bremsstrahlung, DIS, and charged particles produced around the instrumented volume can be addressed. The contribution of neutral particles can only be estimated through Monte Carlo simulations.

To study muon-induced DIS background, events were preselected with a similar set of properties as for a neutrino event, described in Section 5.1, by requiring:

- at least two veto hits;
- at least two DS stations, each with hits at least one hit in both the x and y planes;
- a tagged shower.

When relevant, the reconstructed tracks for these events were also required to pass the same quality cuts described in Section 5.2.3 to ensure a consistent comparison with the ν_μ candidate sample.

The only difference in the preselection of muon and neutrino events lies in the veto requirement to ensure the muon entrance into the instrumented volume.

This sample is used to study the properties of muon DIS events to better reject them.

The number of events passing this selection made the file size unmanageable, even for one entire run. Thus, the analysis was restricted to 0.1 fb^{-1} from Run 8323 data, corresponding to 1.34×10^6 events.

Table 5.1: Summary of runs used in the 2024 data analysis.

Run	Emulsion target	Events	Luminosity [fb^{-1}]	Preselected ν events	Preselected events %
8285	1	4.12×10^8	0.6945	1.17×10^4	0.0028
8315	1	4.23×10^8	0.7400	1.19×10^4	0.0028
8320	2	4.48×10^8	0.7827	1.25×10^4	0.0028
8323	2	4.50×10^8	0.7900	1.22×10^4	0.0027
8638	3	6.05×10^8	1.0690	1.76×10^4	0.0029
8724	3	5.43×10^8	0.9128	1.58×10^4	0.0029
9015	4	4.76×10^8	0.8447	1.36×10^4	0.0029
9094	4	5.44×10^8	0.9650	1.56×10^4	0.0029
9258	5	5.07×10^8	0.8990	1.47×10^4	0.0029
9262	5	5.24×10^8	0.9320	1.51×10^4	0.0029
9288	6	5.14×10^8	0.9125	1.48×10^4	0.0029
9361	6	5.38×10^8	0.9420	1.55×10^4	0.0029
9411	7	5.32×10^8	0.9311	1.51×10^4	0.0028
9436	7	6.14×10^8	0.9080	1.76×10^4	0.0029
9562	8	5.33×10^8	0.9420	1.56×10^4	0.0029
9569	8	5.27×10^8	0.9350	1.51×10^4	0.0029
9622	9	7.04×10^8	1.2320	2.02×10^4	0.0029
9685	9	5.48×10^8	0.9771	1.57×10^4	0.0029
9715	10	3.36×10^8	0.6102	9.04×10^3	0.0027
9880	10	5.36×10^8	0.8814	1.55×10^4	0.0029
9885	11	5.31×10^8	0.9594	1.53×10^4	0.0029
9913	11	5.18×10^8	0.9361	1.48×10^4	0.0029
Total	—	1.14×10^{10}	19.7965	3.25×10^5	0.0029

5.2 Event Selection

The first part of the analysis focuses on determining the most suitable set of cuts to select a clean sample of muon neutrino CC events. The workflow is summarised here and better described in the rest of this section. The cuts applied are:

1. no hits in the veto system, to enhance the purity of the selected subsample;
2. event synchronised with interaction in IP1;
3. hit timing, studying the QDC vs time distribution of hits, reported in 5.2.1;
4. all US planes have some activity, to ensure the muon produced in the interaction reaches the DS planes;
5. a shower tagged in the event with the strict version of the algorithm, described in 5.2.2;
6. a *good* track reconstructed, as defined in 5.2.3.

Once the most appropriate set of cuts is identified, the second part of the analysis will aim to reconstruct the neutrino energy spectrum and give an estimate of the cross section.

5.2.1 Timing studies in TI18 data

Although the preselection already addresses the issue of events triggered by the afterglow of previous ones, discussed in Section 4.2 for the test beam setup, the analysis of hit time distribution within each event was repeated using the TI18 dataset.

Figure 5.2 shows, for SciFi (left) and US (right), the distribution of signal times relative to t_{ref} (see Section 4.2), defined as the most probable value for hits in SciFi, in the neutrino candidates sample. The two distributions exhibit distinctly different shapes.

In the SciFi case, a prominent cluster of hits appears around t_{ref} .

On the positive side of the x-axis, a double-tailed structure is visible. Since these hits occur later than t_{ref} , they can be attributed to cross-talk and afterglow effects induced by the large number of particles produced in showers. Dividing the hits in walls before the shower starts, on the first wall in which the shower is visible, and after this wall confirms this hypothesis. Figure 5.3 shows, as an example, the distributions obtained when the shower starts on the third wall, located in the middle of the instrumented volume.

Hits recorded before the shower can originate from background muons bypassing the veto-instrumented area and entering from the sides before interacting, from the early stages of the shower that are too small to be tagged, or from backscattering, which contributes only marginally to the tails of the distributions. The plots for showers starting in the other SciFi modules are available in Appendix A.1.




Figure 5.2: Neutrino sample. Distribution of QDC as a function of the time difference between the hit timestamp and t_{ref} for SciFi (left) and US (right).

In the US distribution, the impact of saturation is evident with a major cluster at low QDC values and a minor one around ~ 180 .

As in the SciFi case, it's possible to identify the main cluster of hits due to the signal centred around 0 and a longer tail. This distribution is substantially different from the one shown in Section 4.2 because TI18 data are calibrated so that a particle coming from IP1 should leave a signal at $t = t_{\text{ref}}$ in all modules of the detector. The tail also has two distinct components: one at low values, primarily attributed to noise, and another near the saturation threshold, likely associated with physical signals of large amplitude, which may induce significant cross-talk and afterglow.

It is also important to note that the dataset resulting from the preselection is still largely dominated by muons undergoing Deep Inelastic Scattering. The muon component influences the QDC distributions, particularly by contributing to the population of hits at low QDC values, as illustrated in Figure 5.4. In SciFi hits, the central cluster is accompanied by a tail extending toward increasingly negative QDC values, more pronounced compared to the neutrino candidates sample. For the US modules, the cluster at high QDC values is strongly reduced, leading to the conclusion that US saturation is predominantly associated with large showers not induced by muons.

For both SciFi and US, in both samples, on the negative side, a well-separated cluster corresponds to hits arriving significantly earlier than the reference time. The only exception is for US signals in the neutrino sample, which, although exhibiting a significant contribution of early hits, do not display a clear separation.

The first hypothesis was that this distribution was due to residual afterglow of earlier




Figure 5.3: Neutrino candidates sample. Distribution of QDC as a function of the time difference between the hit timestamp and t_{ref} for SciFi 3. Top left: hits in SciFi stations before the shower starts. Top right: hits in the SciFi station where the shower develops. Bottom: hits in SciFi in the stations after the one in which the shower starts.

events that were rejected during the preselection process.

To test this assumption, a subset of one million events from an unfiltered data file was analysed. The first requirement imposed was that no other event be recorded in SND@LHC during the 1000 clock cycles preceding the event under study. To ensure a consistent comparison, the same criteria as for the preselection were applied. To retain a




Figure 5.4: Muon sample. Distribution of QDC as a function of the time difference between the hit timestamp and t_{ref} for SciFi (left) and US (right).

sufficiently large sample, the muon preselection described in Section 5.1.2 was used.

Figure 5.5 shows the distribution of QDC as a function of the distance of the hit time from t_{ref} in such a sample. The cluster of hits arriving much earlier than the reference time is visible, ruling out the possibility of them being hits from the preceding event, broken down into two parts by the maximum duration of an event due to the online event builder settings, described in Section 2.4.1.

A second test was performed using a Monte Carlo (MC) sample of muons undergoing DIS, generated with the 2022 detector geometry. Although this sample was not specifically produced for this study, it was the only one suitable for the test. However, the distributions shown in Figure 5.6 differ substantially from those obtained from data. In particular, the US timing does not appear to be aligned with the detector data for particles produced at IP1.

One more hypothesis was that the hit arrival time depended on the position of the hit itself. Using the selected muon sample, the distribution of the difference from t_{ref} was investigated for both SciFi (Figure 5.7 and Figure 5.8) and US (Figure 5.9).

For SciFi hits, shown in Figure 5.7, the distribution is centred around zero and does not show any particular dependence on the position. The larger number of entries at $y \gtrsim 40$ cm depends on the muon flux distribution, which has a gradient, increasing towards the top of the detector [42].

It can be observed that the hit time was calibrated to average to zero at the centre of the module. As a result, small deviations are present: towards negative values in the first half of the detector ($x \gtrsim -25$ cm and $y \lesssim 38$ cm), and towards positive values in the




Figure 5.5: No preselection sample. Distribution of QDC as a function of the time difference between the hit timestamp and t_{ref} for SciFi (left) and US (right) in data not preprocessed.




Figure 5.6: Monte Carlo sample. Distribution of QDC as a function of the time difference between the hit timestamp and t_{ref} for SciFi (left) and US (right) in the same range as the plots shown before.

other half.

To check for dependencies of this behaviour on the position of the hit along the fibre,

the track extrapolated position was used to estimate the missing coordinate, so x for SciFi Y and y for SciFi X. The result is shown in Figure 5.8. Using the extrapolated point from a track reconstructed with DS hits leads to a worse resolution than a SciFi hit, as highlighted by the much broader distribution of hits, which, as a quick comparison with Figure 5.7 reveals, extends beyond the active area. More information on how tracking is performed is given in Section 5.2.3. No clear dependence is visible.




Figure 5.7: Muon sample. Distribution of the hit position as a function of the time difference between the hit timestamp and t_{ref} for SciFi X (left) and Y (right).

Regrettably, also in the case of US, reported in Figure 5.9, no clear explanation is possible. In this detector, both coordinates are extrapolated from the track, due to the large size of the US bars. The plots show that the timing of hits in the US has not yet been calibrated, unlike in the SciFi module, to account for light propagation and other effects introduced by the electronics.

Nonetheless, none of these distributions sheds any light on the origin of early hits, as for both SciFi and US, these seem to be evenly distributed along all directions, making it impossible to identify a clear reason.

To further investigate this issue, the ratio of early hits, defined as those occurring at least one clock cycle (6.25 ns) earlier than the reference time, was studied, with particular attention to the US, where a clearer separation is observed around this value. Figure 5.10 shows that, when early hits are present, in most cases the majority of them fall into the early category. This suggests that the problem is related to the entire event and, potentially, to the choice of reference time. To explore this hypothesis, the SciFi hits time distribution was studied for a subset of events characterised by very early US hits. Figure 5.11 shows one of these events.




Figure 5.8: Muon sample. Distribution of the hit position along the fibre as a function of the time difference between the hit timestamp and t_{ref} for SciFi X (left) and Y (right) .




Figure 5.9: Muon sample. Distribution of the hit position as a function of the time difference between the hit timestamp and t_{ref} for US x projection (left) and y projection (right).

The double-peaked time distribution suggests a problem linked to the readout electronics. The excess of late signals was observed in the same plane where the shower had been tagged at the preselection stage, suggesting that the issue originates from the large




Figure 5.10: Muon sample. Ratio of US early hits over the total number of US hits in the event.




Figure 5.11: Muon sample. SciFi time distribution in a single event.

number of developing particles and the consequent light production in the fibres, which leads to multiple channels firing. Nonetheless, this effect is neither caused by individual channels firing multiple times nor by specific channels firing in every event. While the issue is still under investigation by the collaboration, the shape of this distribution hinders

the identification of a reliable reference time for the interaction. To minimise the impact of such bimodal events, a peak search is performed using the ROOT TSpectrum class [52]. If two or more peaks are found separated by more than one clock cycle, the reference time is taken to be the first peak.

This procedure ensures the selection of the appropriate reference time for the crossing particle, which in turn suppresses the cluster of hits that had been previously classified as early, both for SciFi and US.

After the rejection of such events, the distribution is shown in Figure 5.2 and 5.4 are now shown respectively in Figure 5.12 and 5.13.




Figure 5.12: Neutrino sample. Distribution of QDC as a function of the time difference between the hit timestamp and t_{ref} for SciFi (left) and US (right) after the new selection of t_{ref} .

A cut is still applied to reject as many spurious hits as possible and limit their contribution to the overall QDC estimation, directly linked to the energy reconstruction. The final cuts used in the analysis are $[-0.5, 1]$ clock cycles ($[-3.125, 6.25]$ ns) around t_{ref} for SciFi and $[-1, 1]$ clock cycles ($[-6.25, 6.25]$ ns) for US.

5.2.2 Custom Shower Tagging

To enable the effective application of the shower tagging algorithm developed and validated on the test beam dataset, as described in Section 4.2, to the current analysis, adaptations were implemented to account for differences in detector conditions.

The first difference to account for is the target material.




Figure 5.13: Muon sample. Distribution of QDC as a function of the time difference between the hit timestamp and t_{ref} for SciFi (left) and US (right) after the new selection of t_{ref} .

Hadronic shower development depends on the average interaction length, which in turn depends on the atomic number of the material. The test beam setup used an iron target ($Z = 26$), with interaction length $\lambda_{Fe} = 16.77$ cm [53], while the target in TI18 is made of tungsten ($Z = 74$), $\lambda_W = 9.946$ cm [54]. The ratio between the interaction lengths is ~ 0.59 , thus we expect more compact showers in TI18, so we decided to reduce the minimum number of hits required from 36 to 10. This reduction allows us to successfully tag the shower even at an early stage of development.

For the preselection of events, as stated in Section 5.1, at least one SciFi plane had to meet this requirement.

In the stricter search for neutrino candidates, a more stringent constraint was applied to enhance background rejection. Requiring that two consecutive SciFi stations satisfy the shower tagging criteria significantly reduces the contribution from electromagnetic showers. This approach exploits the differing longitudinal development of hadronic and electromagnetic showers. A hadronic shower typically extends over $6 - 9 \lambda_{\text{int}}$ to reach 95% containment, while a single tungsten wall in the target corresponds to $0.5 \lambda_{\text{int}}$. On the other hand, the average longitudinal containment for an electromagnetic shower at 95% is $15 - 20 X_0$, comparable to the $17 X_0$ of a tungsten wall. As a result, electromagnetic showers are unlikely to be sampled in more than one SciFi station, making the requirement of two consecutive SciFi stations an effective discriminator.

A third requirement was added to reduce the contribution of muons that are not detected by the veto but enter the target region.

The distribution of hits in the SciFi plane preceding the shower was studied. Ideally, no signal should be present in the detector before a neutrino interaction. In practice, an unavoidable contribution arises from electronic noise, plus hits originating from charged particles that are not tagged by the Veto System, either due to detector inefficiency or to their incoming direction.

After the time filtering of hits, the number of hits due to noise is expected to be minimal. The contribution was further studied.

In the muon sample, the distribution of the total number of hits in the x and y stations of this plane, shown in Figure 5.14, exhibits a distinct peak at around four hits, consistent with the expected response of approximately two hits per SciFi module for a crossing muon. The distribution also features a long tail corresponding to events with a larger number of hits in the plane immediately upstream of the shower, originating from a combination of noise and backscattering effects, particularly pronounced in cases where large showers develop.




Figure 5.14: Muon sample. Number of hits in the SciFi module preceding the shower. The contributions from the x and y modules are summed. Right: zoom on the peak region.

The same distribution for the neutrino sample is reported in Figure 5.15. As expected, a fraction of events display no activity before the shower. However, the contribution of muons remains significant, as indicated by the pronounced peak analogous to that in Figure 5.14.

The cut applied to reject these events, but still allow the presence of hits due to backscattering, was not to require no activity before the shower, but to discard those




Figure 5.15: Neutrino sample. Number of hits in the SciFi module preceding the shower. The contributions from the x and y modules are summed.

events with a number of hits smaller than one and larger than seven in which the contribution of muons is most significant.

As the last two requirements described depend on the SciFi modules preceding and following the one in which the shower was identified at the preselection stage, dedicated criteria need to be applied to the first and last SciFi modules to obtain a similar performance.

Special case: SciFi 1

SciFi 1 proved to be the most difficult plane to reject noise.

If the shower is tagged in the first SciFi plane, the requirement on the hit before the shower cannot be applied, but the request on the veto system does not have the same performance, as can be seen in Figure 5.16.

In the muon sample, a smaller secondary shower is observed in many muons undergoing DIS in the first plane, following the main hadronic shower. This secondary shower is likely due to bremsstrahlung photons.

A preliminary check done on the neutrino interaction MC in the 2024 configuration leads to expecting a fraction of muons produced from neutrino interactions to have sufficient energy to undergo bremsstrahlung. Since no reliable criterion could be established to distinguish these events from high-energy incoming muons, the cut was discarded, and the mitigation of this background was deferred to a later phase following a more detailed




Figure 5.16: Neutrino sample. Distribution of shower starts in the five SciFi planes, with the stricter algorithm but no special cuts. These events satisfy the track requirements described in Section 5.2.3.

Monte Carlo analysis.

Special case: SciFi 5

The request of two consecutive SciFi modules tagging the shower cannot be applied to SciFi 5, since it is the last SciFi wall, and, as Figure 5.16 shows, it has a significant impact on the final result.

A dedicated condition must be defined for US planes to meet, aiming to achieve a background rejection power comparable to that of a single SciFi plane shower tag. The challenge arises from the larger dimensions of bars and from the fact that, in front of US 1, the absorber material is no longer the ~ 6 cm tungsten but 20 cm of iron. The information extrapolated from the SciFi must then be adapted to account for shower development in a different medium, while considering that a single US bar, 6 cm-wide, can fully contain the shower. Indeed, a shower starting in SciFi 5 may generate a signal in only one US 1 bar, while an inclined muon can leave signals in more than one bar. Six SiPMs per side read out each bar, but light diffusion inside the bar and intrinsic SiPM cross-talk make the number of hit SiPMs an unreliable discriminator. To reject any residual electronic noise hits, a bar is considered if at least 5 out of 12 large SiPMs have a signal.

The initial strategy was therefore to require either more than one bar with a signal, or, in the case of a single bar, to apply a threshold on the QDC value, chosen by looking at the events selected by the 2022/23 analysis, better discussed in Section 5.2.4. In Figure 5.17,

the number of showers tagged in SciFi 5 highlights how suboptimal the applied condition is, being strongly reduced but still nearly three times larger than in the central walls.




Figure 5.17: Neutrino sample. Distribution of shower starts in the five SciFi planes, using an algorithm asking in US 1 for either more than one bar with a signal, or, in the case of a single bar, to apply a threshold on the QDC value. These events satisfy the track requirements described in Section 5.2.3.

A second possibility was investigated: analysing the energy deposited in the SciFi modules and in US 1 by muons, to find a suitable threshold to discriminate showers from passing muons. This approach leverages the detector calibration to allow for a direct comparison between the SciFi and US modules, despite their otherwise considerably different QDC scales.

The first step was to evaluate the energy deposition in SciFi. The neutrino sample results proved difficult to interpret due to the small statistics available. The distributions are shown in Figure 5.18, normalised to compare their shape regardless of the number of events. Even if it is difficult to find an average value for the energy deposition in the SciFi planes in which the shower develops, the distributions of the shower wall and the following one are similar, suggesting that the main contribution is from showers. Instead, the energy deposit is clearly peaked around zero in the module preceding the shower, as expected.

At this stage, the goal was to discriminate between hadronic and electromagnetic showers. For this purpose, the muon sample was used, as it provides a much larger number of events and includes hadronic showers originating from DIS. The energy distributions for the muon sample are shown in Figure 5.19. In this case, a clear shape is visible, showing a peak at low energies but shifted away from zero.




Figure 5.18: Neutrino sample. Energy deposited in the SciFi module after the wall in which the shower starts (top left), after one more wall (top right), and in the wall before (bottom).

Moving to US 1, the first check done was on the energy deposited in this station for a shower tagged to start in the first and last walls. In the first case, only the outgoing muon, or a minimal part of the shower, should reach US 1, while in the second case, the shower should be in its developing phase. The result reported in Figure 5.20 does not show the expected behaviour, as the two distributions are very similar, making it not very useful to set a reasonable threshold.




Figure 5.19: Muon sample. Energy deposited in the SciFi module after the wall in which the shower starts (top left), after one more wall (top right), and in the wall before (bottom).

Indeed, the shower tagged in SciFi 5, setting an energy threshold at 3 GeV in US 1, still results in 66 events accepted in the neutrino sample, essentially matching the combined contribution of all other SciFi stations, which highlights the need for a different method.

The larger value of the interaction length of iron, λ_{Fe} , can be exploited to set a threshold on the overall detector response, assuming that the shower lateral dimension increases across the iron walls. Therefore, relying solely on US 1 limits the discrimination power.




Figure 5.20: Muon sample. Energy deposit in US 1 for showers starting in SciFi 1 and in SciFi 5.

The requirement chosen to validate a SciFi 5 shower tag is to have > 6 bars with signal across the entire US subsystem. In addition, the condition that each module must independently register a signal implies that either at least three bars are hit within a single module, or at least two planes register a minimum of two hits each.

The number of events tagged in this case is reported in Figure 5.21. With this condition, the number of candidate interactions flagged in SciFi 5 is coherent with the previous planes.

The validity of this cut in selecting the desired events is verified by studying the energy deposit in US 1. The distribution previously shown in Figure 5.20, in events selected with this requirement, is reported in Figure 5.22.

In this case, there is a clear bump for energies $\gtrsim 3$ GeV, as expected, now visible since a large fraction of noise is being rejected.

The result was considered satisfactory, and no further selection criteria were applied.

5.2.3 Track

After the preselection, track reconstruction is attempted to extract and exploit the muon information. Tracking was performed using the algorithm implemented in SND@LHC [42]. Two algorithms are available: the Simple Tracking (ST) and the Hough Transform (HT) methods [55]. Both use the Kalman Filter [56] for track fitting, as implemented in the GENFIT package [57]. Tracking is performed separately in the horizontal (x-z) and




Figure 5.21: Neutrino sample. Distribution of shower starts in the five SciFi planes, using an algorithm requiring at least 7 bars in the US subsystem. These events satisfy the track requirements described in Section 5.2.3.




Figure 5.22: Muon sample. Energy deposit in US 1 for showers starting in SciFi 1 and in SciFi 5 with the final shower selection algorithm.

vertical (y - z) projections for both the SciFi and DS systems, since their planes and readout schemes provide two-dimensional spatial information. These 2D tracks are later combined

to form a 3D trajectory. In events with multiple candidates per projection, associating the correct pair of projections becomes nontrivial and is handled differently by the two algorithms. This analysis employed the HT method, using DS hits as seeds and limiting the reconstruction to a maximum of one track per event.

Hough Transform Tracking

The HT method implemented in `sndsw` [58] is optimised for high-occupancy environments, such as those resulting from muon neutrino charged-current (CC) interactions, where a hadronic shower develops partially superimposed onto the outgoing muon. HT is designed to detect linear patterns in complex hit maps. When multiple track candidates are found in a single projection, the algorithm verifies whether their slopes fall within a range consistent with the detector’s angular resolution. If so, the track whose slope is closest to the median is selected; otherwise, the event is discarded.

After a track candidate is identified, its trajectory has to pass through hit detector elements within a predefined tolerance of 0.1 cm for SciFi and 1 cm for DS. These values correspond to about four SiPM channels in SciFi and one bar width in the DS. In low-occupancy events (planes with fewer than 4 hits), the tolerances are relaxed to 0.5 cm for SciFi and 3 cm in the DS to enhance tracking efficiency without significantly increasing the risk of ghost tracks.

To attempt a track reconstruction, the algorithm requires that the HT line intersect at least three horizontal and three vertical planes within the defined tolerance. This constraint is stricter than the one used in the preselection, but allows DS 4, which is only vertical, to recover tracking efficiency.

A good track

Each event that passes the first four cuts listed at the beginning of this section is processed in the next analysis step, which focuses on the reconstructed muon track information.

To ensure track quality, the following requirements are applied:

- a track must be successfully reconstructed, meaning the fit has converged;
- reduced $\chi^2 \leq 5$ [42].
- computed $\frac{1}{v} - \frac{1}{c} > -0.07 (\text{cm/ns})^{-1}$, where v is the particle velocity and c is the speed of light, as this value is associated to tracks of particles moving from DS to SciFi, produced in IP2 instead of IP1 and crossing the detector backwards [59].

If these conditions are met, the reconstructed track is extrapolated back to the target region. Since no magnetic field is present in the instrumented volume, the extrapolation is performed with a straight line. The effects of multiple scattering must be accounted

for when matching positions in the vertex region, which can be achieved by allowing a sufficiently large window.

This set of criteria aims to ensure that the reconstructed muon track is consistent with the interaction vertex and not produced by a muon traversing the detector at a large angle, unrelated to the event in the target.

The first fiducial area cut excludes muon interactions with the tunnel wall or floor that result in the hadronic shower reaching the target region and the muon being recorded in the US and DS chambers, and the absence of a corresponding signal in the Veto system mimics the desired event topology.

The first step involves extrapolating the track to the SciFi station where the shower was tagged. The distribution of the extrapolated points is in Figure 5.23 for the first SciFi plane (the others can be found in Appendix A at A.5, A.6, A.7, A.8).

These plots show a visible accumulation of extrapolated muon tracks on the left and right sides of the SciFi plane. These correspond to muons entering the instrumented volume at angles that allow them to bypass the veto system and thus remain undetected. The majority of these events occur on the small- x side, which lies closer to both the collision axis and the tunnel wall. In this region, the probability of particle interactions with the surrounding material is higher, resulting in an increased production of showers and muons detected within the instrumented area. The veto upgrade strongly mitigates the same effect from the small- y region, corresponding to the tunnel floor, thanks to its new position (see Figure 3.2), which provides improved coverage of the lower part of the target.

To reject this class of events, an inspection of the extrapolated position distribution motivates the requirement that the track intersection lie within a fiducial region. Based on the features observed in the distribution, in particular the high number of points extrapolated near the edges of the active area, this region is defined by excluding 2, cm from the top and bottom edges of the SciFi module and 3, cm from both sides in the x direction. This results in a total fiducial area of $33 \times 35 \text{ cm}^2$, corresponding to approximately 76% of the full SciFi plane, significantly larger than the $25 \times 26 \text{ cm}^2$ fiducial region used in the previous analysis [32].

An example of this kind of discarded events is shown in Figure 5.24.

Shower Centroid

To further constrain the muon's position within the SciFi plane, a proxy for the interaction vertex was defined as the centroid of the shower. This centroid is computed as the QDC-weighted mean position, using only signals with positive QDC values to avoid unphysical results. Additionally, a minimum of five hits is required in both the x and y projections to ensure a reliable estimate.

Validation was performed using the DIS muon sample, with the additional requirement that the incoming muon had exactly 1 hit in the x and y SciFi modules in the wall




Figure 5.23: Neutrino sample. Reconstructed muon track position in SciFi 1. The solid red line shows the position of the tracker module, while the red dashed line shows the fiducial area.

immediately preceding the shower start.

This subsample was the most suitable, as it includes both hadronic showers and the SciFi hits associated with the muon entering the instrumented volume. This information enabled an estimation of the distribution of the position differences between the muon entry point and the shower centroid, using data from two consecutive planes to reduce the influence of multiple scattering.

This distribution is shown for the x and y coordinates and the different SciFi planes in the Figure 5.25.

It is important to note that for this study, the first wall of SciFi could not be used as the veto planes segmentation does not provide a sufficiently precise measurement of the incoming particle position. Nonetheless, given the homogeneity of results across the remaining four walls, there is no indication that the first module would behave differently.

A check that the centroid distribution is not biased by showers not fully contained, so the distribution of the residual value as a function of the incoming muon position, measured with the SciFi wall immediately in front of the one in which the shower started, is visible in Figure 5.26 (and Figures A.9, A.10, A.11) for both coordinates.

The slight increase in occupancy toward larger values of y is not an artefact of the centroid algorithm but reflects the muon background distribution at the SND@LHC location. As shown in the muon flux study [42], a gradient is present, with a higher muon rate observed in the upper part of the detector.

Once the centroid computation was validated, it was used to refine the definition of a




Figure 5.24: Event display from the preselected sample. A muon producing a shower near the top edge of the SciFi modules, leaving no hits in the veto.

good track: the extrapolated point must lie within the fiducial area of the SciFi module where the shower originates, and additionally, it must be within 3 cm [42] of the computed centroid position in both the x and y directions. This threshold value was chosen because it corresponds to the resolution of the DS track extrapolated to the SciFi region, taking into account the degradation of the nominal resolution caused by multiple scattering.

If this condition is satisfied, indicating that the muon track is consistent with the shower location, the centroid is incorporated as an additional point in the track reconstruction, and the track is subsequently refitted with a line. This added point provides a further constraint with a resolution determined by the total QDC of the signals contributing to the centroid, proportional to $1/\sqrt{\text{sum of QDC}}$. Then, the centroid resolution is significantly superior to that of the DS tracking point, especially considering the uncertainty introduced by extrapolating backwards through the entire instrumented volume and the substantial amount of passive material.

The improvement introduced by this correction is illustrated in Figure 5.27, which show the difference between the muon hit position measured in the SciFi and the position




Figure 5.25: Distribution of the difference between the incoming muon position measured in the SciFi wall before shower development and the centroid position computed in the first shower wall. Left: SciFi X. Right: SciFi Y.




Figure 5.26: Distribution of the centroid residual as a function of the muon position measured in the wall preceding shower development. Left: SciFi 2X. Right: SciFi 2Y.

extrapolated using the DS-only track (in black) and the corrected track including the centroid point (in red). The distributions demonstrate not only a narrower peak around zero but also the disappearance of the multi-peak structure caused by the coarse granularity

of the DS bars.




Figure 5.27: Difference between the measured hit position before shower development and the reconstructed position using DS only (black) or DS with centroid correction (red). Left: SciFi X. Right: SciFi Y.

The improved track is then used as a proxy for the incoming neutrino direction. This allows for verifying whether the extrapolated direction intersects the active area of the Veto System, ensuring that the absence of a signal corresponds to the passage of a neutral particle rather than a charged particle entering outside the instrumented volume.

The points obtained by extrapolating the corrected track onto each veto plane should lie within the active region, with a margin of 1 cm from each edge to exclude border effects.

5.2.4 Validation on 2022-2023 data

The selection steps described above have been applied to the neutrino events identified with an independent method in the combined 2022–2023 dataset, which serve as a benchmark for validating the performance. This independent method is based on and aims to be an upgraded version of the reconstruction and selection procedure used in [49], and includes the events selected for the observation claim. In this dataset, 28 candidates are available, which is still too few to allow for a detailed study of observables such as the ν_μ energy spectrum.

Surely, the different veto efficiencies correspond to different rejection powers for incoming muons; however, this was the only sample available for the initial validation of

the chosen selection cuts. The issues related to the Monte Carlo samples, mentioned in Section 5.2.1, are further discussed in Section 5.3.

Out of 28 candidate events, the current analysis developed for 2024 data rejects 13 at different stages of the selection.

US hits

One event was initially rejected because US 1 shows no remaining signal after the applied time filtering, leading to the requirement that all US planes contain a signal, intended to ensure that the muon originates from the vertex region, to fail. An inspection of the hit timing, shown in Figure 5.28, suggests that the most probable cause is a suboptimal timing alignment. A slight relaxation of the time filtering window in the US, from $[-1, 1]$ to $[-1.2, 1]$, enables the recovery of this event. This relaxation has no impact on the 2024 dataset, as the final number of selected candidates remains the same in both cases.




Figure 5.28: Event 19118383 of run 6018 (2023). Hit time difference from the reference time as a function of the US station. A clear diagonal structure is visible, pushing values in the first 2 planes of US toward the negative part of the distribution.

Shower tag

The most important contribution to rejected events comes from the shower tagging algorithm.

Figure 5.29, as an example in which a shower is tagged, summarises the hit distributions in SciFi X and Y as functions of QDC and time. The two plots display the SciFi X and Y views. In each plot, the red lines indicate the SciFi planes, ordered from bottom to

top; each band is further subdivided into five bins, each corresponding to one clock cycle, while the colour scale encodes the QDC information.




Figure 5.29: Event 9764224 of run 6596 (2023). Hit distribution in SciFi X and Y, as a function of time and QDC.

In the 2022-2023 dataset:

- One event is not tagged by the shower identification used in the preselection. In this case, the distribution of hits in the SciFi never satisfies the minimum requirements. This result is independent of the time filtering. The SciFi hit distribution is visible in Figure 5.30.
- Two events, tagged by the preselection shower requirement in SciFi 5, fail the final condition on the number of US bars with signal. One of them could be recovered by removing the time filtering of hits, although both are later rejected because the reconstructed track extrapolates outside the SciFi area.
- Four events, tagged by the preselection shower requirement in SciFi 3 and SciFi 4: the hit distribution does not permit tagging two consecutive walls, causing the strict selection to fail. This behaviour is independent of the time filtering. An example of such events is in Figure 5.31.

In this case, the shower tag algorithm was not modified to recover such events. Further studies will be performed when a Monte Carlo sample is available.




Figure 5.30: Event 38637101 of run 6268 (2023). Hit distribution in SciFi X and Y, as a function of time and QDC. No shower found by the shower tag algorithm.

Reduced χ^2

Of the remaining events, five fail to satisfy the requirement on the maximum reduced χ^2 of the reconstructed track. In the previous analysis, track quality was instead enforced by discarding events with more than 10 DS hits per projection. This cut intends to reduce the probability of including noise hits in the reconstructed track, which could otherwise lead to an incorrect determination of the particle direction.

Although both approaches share the same objective, they clearly exhibit different rejection powers. Once again, a Monte Carlo sample could help in discriminating the best approach.

Centroid position

Two more events are rejected because the reconstructed centroid and the extrapolated track position are not compatible. In both cases, one coordinate lies within the acceptance window while the other exceeds the applied cut. Since the rejection occurs once in the x-coordinate and once in the y-coordinate, there is no indication of a systematic bias in the centroid reconstruction algorithm.




Figure 5.31: Event 49645511 of run 5888 (2023). Hit distribution in SciFi X and Y, as a function of time and QDC. The initial shower tag identified a shower in SciFi 4; however, the requirement was already not met in SciFi 5, which rejected the event.

While acknowledging that two analyses may select slightly different subsets of events, the overlap should be substantial, as both are designed to achieve the same goal. Given the current situation of 13 events out of 28 rejected, a more detailed investigation is therefore required on both sides to understand the residual differences.

5.3 Next steps

The analysis has progressed substantially, yet it remains far from completion.

In particular, the work is still focusing on optimising the event selection, exploiting the three available data samples: the preselected data from the 2024 dataset and the 2022/23 candidate neutrinos. The detailed study of events, beginning with their fundamental features such as hit timing, required a significant investment of time but has proven valuable for advancing the overall understanding of the detector and will surely help in the next step of the analysis.

Unluckily, it was not possible to validate the selection flow on Monte Carlo samples.

The neutrino sample in the 2024 configuration became available too late to yield meaningful results within the time limit imposed by this thesis, despite the efforts of the Collaboration and of the SY-STI colleagues who provided the FLUKA simulations

underlying the Monte Carlo production. This first step simulates the proton interactions at the LHC and the subsequent particle production, which are then propagated into SND@LHC. Correctly estimating the flux of all particles reaching the detector is crucial; thus, a careful calibration, guided by measured data, was required before the simulated sample could be used. Indeed, the LHC configuration underwent significant changes with respect to previous years, particularly in the beam optics, which directly affects the particle flux reaching SND@LHC. Consequently, before the correction, the discrepancy between data and Monte Carlo simulations was considered too large to proceed with the large-scale Monte Carlo production intended for analyses, as the simulations failed to reproduce the measured distributions.

The background sample of muons undergoing DIS in the detector region is not yet available, as the collaboration is updating the DIS simulation. Starting from the FLUKA sample, particles must be propagated into the detector region, with their interactions and decays handled by PYTHIA, and their interactions within the active volume modelled using Geant4. The collaboration is focusing on optimising the production of interactions in the active volume, aiming to reduce the computational time required to generate a statistically significant sample.

These Monte Carlo samples will be used to validate the current choices adopted in the selection flow, albeit with significant limitations at the present stage, as already discussed in Section 5.2.1 for the hit-timing selections and for the detector response in QDC. A careful matching between collected data and simulation will be required, as this comparison has already proven to be a significant challenge. Once the Monte Carlo description of the detector response is fully validated, it will enable a more detailed justification of the timing cuts and shower-tagging criteria, together with a quantitative evaluation of their associated systematic uncertainties. In this framework, a summary of the dominant detector- and analysis-related uncertainties will be derived, and the impact of operational effects will be systematically assessed. Ultimately, a validated simulation will allow for a precise estimation of the expected number of interactions and their contributions to the overall uncertainty at each stage of the selection, as well as a reliable assessment of the background, including not only muons but also neutral hadrons, which will require dedicated Monte Carlo samples that are not yet available.

Nevertheless, although the selection cut still requires some optimisation, it already correctly identifies the desired topology. Moreover, the available tools, in particular the energy calibration of the detector (Section 4.3) and the reconstruction of the shower direction (Section 4.4), allow the reconstruction of both the shower energy and direction, enabling the reconstruction of the incoming neutrino event and paving the way for a cross-section measurement in an energy region where no data currently exist.

Chapter 6

Muon System Upgrade

Contents

6.1	Design	95
6.1.1	CMS Drift Tubes chambers	95
6.1.2	The MiniDTs	96
6.1.3	Application in SND@LHC	98
6.2	Installation in TI18	101
6.2.1	Mechanic support	102
6.2.2	Gas System	102
6.2.3	MiniDTs Readout electronics and integration in SND@LHC DAQ	102
6.3	MiniDT Performance	107
6.3.1	Event matching	108
6.3.2	MiniDT preliminary efficiency	109
6.4	Next steps	111

The muon system plays a central role in the identification and precise tracking of muons, which are key signatures in a wide range of physics analyses. To enhance its performance, an upgrade has been implemented with the addition of two new modules based on drift chambers, which provide accurate position measurements and robust tracking capabilities. By integrating these additional modules into the existing muon system, the detector benefits from improved spatial resolution, redundancy in track reconstruction, and an increased efficiency in muon identification.

In this chapter, the design and motivation for the upgrade are discussed, along with the integration into the overall muon detection system.

6.1 Design

The two modules added to the muon system are called MiniDTs, as they are small replicas of CMS Drift Tube (DT) chambers.

In drift chambers, charged particles ionise the gas, and the resulting electrons drift toward the anode wire, where an avalanche is produced. By measuring the drift time and knowing the electron drift velocity in the chosen gas mixture, the distance of the track from the wire can be determined. The left–right ambiguity of this position is resolved during reconstruction using hits from multiple drift tube layers. This technology is well-suited for covering large areas, such as the CMS outer barrel region.

6.1.1 CMS Drift Tubes chambers

Drift Tube chambers in CMS cover the barrel region, integrated into the magnetic field return yoke, and are a crucial part of the CMS muon system [60]. A CMS DT chamber is composed of three independent units, called SuperLayers (SL). Each SL contains four layers of rectangular drift cells, staggered by half a cell to allow correlation of drift times across layers for the determination of track coordinates and angles. The mechanical structure of a SL is made of five 2 mm thick aluminium sheets, separated by insulated aluminium I-beams that also define the cell boundaries.

The fundamental building block of the DT chamber is the drift tube, shown schematically in Figure 6.1. Each cell has a cross section of $13 \times 42 \text{ mm}^2$. Inside each cell, there are five electrodes:

- the anode wire ($V_{wire} = +3600 \text{ V}$), which collects the signal, a $50 \mu\text{m}$ diameter stainless steel wire, gold-plated to avoid oxidation;
- two aluminium strips ($V_{strip} = +1800 \text{ V}$), insulated via a mylar layer, which shape the drift field to improve the linearity of the time and space relationship and thus the resolution of each cell;
- and aluminium cathodes ($V_{cathode} = -1200 \text{ V}$), also insulated via a mylar layer, placed on the I-beams that separate adjacent cells.

This design is optimised to achieve a good performance even if stray magnetic fields are present in some regions of the chambers, necessary in the CMS application.

The number of cells per layer, as well as the tube length, depends on its position in the CMS Barrel, in particular, on how far it is from the interaction point. The largest chambers have more than 90 cells per layer and are up to $\sim 2.5 \text{ m}$ long, limited by mechanical constraints.

The gas mixture is made of 85% Ar and 15% CO_2 . Argon is a noble gas with a high ionisation yield and a low electronegativity, allowing electrons to drift through the




Figure 6.1: Schematic view of a single drift tube cell [60]. The shape of the nominal electric field lines is also shown.

chamber without being captured. Carbon dioxide, on the other hand, acts as a quencher by absorbing most of the photons produced during the avalanche process. Its presence increases the stability of the detector, extending the efficiency plateau before discharges occur.

The chosen ratio between Ar and CO₂ allows for good quenching while saturating the drift velocity of electrons to preserve the space-time relationship linearity.

Given the voltages of the electrodes in a cell, the nominal drift velocity is about $54.5 \mu\text{m}/\text{ns}$, resulting in a maximum drift time of about 380 ns. A single DT cell has a detection efficiency of $\sim 98\%$ and an intrinsic spatial resolution of $\sim 200 \mu\text{m}$.

6.1.2 The MiniDTs

MiniDTs are smaller-sized Super Layers, constructed at INFN Legnaro Laboratory in Padova (Italy) in 2017-2021, using spare material from the original CMS DT chambers. Like the standard SLs, each MiniDT consists of four layers of parallel drift tubes, with 16 cells per layer, for a total of 64 drift tubes per SuperLayer. The wire length of the MiniDTs is ~ 0.6 m.

The differences between CMS DTs and MiniDTs are minimal and related to the manual assembly operations [61] and the intended use in a simpler environment than CMS.

Some of the chamber elements were redesigned to facilitate detector construction while ensuring precise alignment of the cathodes, wires and strips with external positioning tools.

Custom Mechanics

MiniDTs have a custom version of the cells' endcap holding the crimping block (Figure 6.2) designed to exploit the preprocessing of aluminium sheets to have precisely located holes. This allows for better alignment of the cathodes in the chamber and between each other. Figure 6.3 shows how the endcaps are positioned onto the aluminium layer.




Figure 6.2: Schematic representation of the MiniDTs' endcaps.

In the MiniDT case, the wire plugs are truncated cones, developed to fit the new endcap design. These are made of insulating material, namely Noryl resin reinforced with glass fibres, and have a small hole to insert the metallic pipe onto which the wire is crimped. This part has a wider section to maintain the correct position with respect to each wire plug. In this way, the crimping pipe also serves as a pin to connect the front-end circuitry.

The crimping of the wire is done manually after positioning the wire and tensioning it with a dedicated tool shown in Figure 6.4.

Custom electronics

The High Voltage (HV) and Front End (FE) boards used for MiniDTs are those used in the original CMS chambers.

The MiniDTs have a custom HV connector and an HV splitter box for each chamber, to interface standard HV power supplies with the on-chamber connector.

The powering and control of the Front End electronics in CMS Drift Tubes chamber is based on a system of so-called *Splitter Boards*, developed to implement the complex, distributed powering of multiple large chambers in CMS. To simplify the operation of the MiniDTs, a custom board, called Low Voltage Front End (LVFE), was designed in Legnaro and developed and implemented by the electronics service of INFN Bologna.

The MiniDTs readout chain begins at the front-end electronics, where signals from the anode wires are amplified and discriminated by Front End Boards (FEBs), installed within the gas volume. FEBs outputs, one asynchronous LVDS [62] pair per channel, are collected and transmitted to the "On detector Board for Drift Tubes" (OBDT) [63],




Figure 6.3: Positioning of a group of four cathodes onto the aluminium sheet.



Figure 6.4: Tool used to tension the wires during the MiniDTs' assembly.

shown in Figure 6.5. The OBDT receives the analogue signals from the FE electronics and performs the time-to-digital conversion with a coarse count in steps of the LHC bunch crossing (BX), which corresponds to 25 ns, and a fine bin size with a least significant bin of $1/32$ of a BX, or 0.781 ns.

The digitised data are then sent to the VCU118 FPGA, which computes the trigger primitives (TPGs) track segments built with hits from 3 or 4 layers of each MiniDT and stores all TDC hits and TPG information into temporary buffers, accessible for readout through a 1 Gbps Ethernet interface. In addition, a Technical Trigger (TT) bit is asserted when a TPG is built; the use of this bit is further described in Section 6.2.3.

The system is synchronised with the LHC clock via the TTCrq [64] and an HPTC PLL card [65], which provides the required 40 and 320 MHz clocks to the VCU118.

Hits and trigger segments from the MiniDTs can then be read out, connecting the VCU118 to a readout server.

6.1.3 Application in SND@LHC

As described in Section 2.3, until the end of the 2024 data-taking period, the SND@LHC muon system consisted of scintillating modules using 1 cm-thick bars for muon tracking. The limited bar thickness, combined with multiple scattering in the iron blocks of the hadron calorimeter acting as passive material, results in a non-optimal determination of the points where reconstructed muon tracks cross the target region. This directly affects the achievable performance, e.g. the requirement for matching the reconstructed track




Figure 6.5: The OBDT board.

and the centroid discussed in Section 5.2.3. Acknowledging the limit imposed by multiple scattering, which depends on the crossing particle energy, many physics analyses would benefit from more precise tracking.

A preliminary study indicates that adding two MiniDT modules—one for the x coordinate and one for the y —could not only improve the resolution of reconstructed tracks by a factor of two, but also smooth out the multi-peak structure observed in DS bar tracking.

In this preliminary study, the MiniDTs enter the tracking process as 4 additional points in each coordinate, one for each measured point in each module. Even if this simple simulation does not account for detector inefficiencies, it is evident that the redundancy of 8 additional points will positively affect the tracking efficiency, considering the current requirement of 3 points per coordinate. The resolution associated with each point is $250\,\mu\text{m}$, the single hit resolution expected from DTs, to introduce a safety margin. The Monte Carlo sample used was a muon particle gun with energy ranging from 100 GeV to 1 TeV.

Figure 6.6 shows the position difference between the track reconstructed with the `sndsw` [58] HT algorithm, described in Section 5.2.3, using SciFi information (considered as the true position) and the extrapolation of the track reconstructed with the same algorithm but from the muon system points, with (red line) and without (black line) the inclusion of the MiniDTs. The additional information from the MiniDTs not only improves precision but also mitigates the multi-peak structure observed in DS bar tracking.

Figure 6.7 summarises the overall improvement, showing in the two cases the distribution of $\Delta R = \sqrt{(\Delta x)^2 + (\Delta y)^2}$ where Δx and Δy are the event-by-event residuals computed as described above.

Finally, Figure 6.8 shows the average of the residual distribution as a function of the incoming particle momentum. It is worth noting that the MiniDT information also reduces




Figure 6.6: Distribution of residuals for DS-only tracks (black) or DS+MiniDT tracks (red). The difference is computed between the reconstructed SciFi track and the track extrapolated from the muon system to the target region. Left: x coordinate. Right: y coordinate.




Figure 6.7: Distribution of $\Delta R = \sqrt{(\Delta x)^2 + (\Delta y)^2}$ for DS-only tracks (black) or DS+MiniDT tracks (red).

the momentum dependence, as the distribution remains stable in both coordinates, while the DS-only tracks show a worsening of the resolution approaching 1 TeV.




Figure 6.8: Distribution of resolution for DS-only tracks (black) or DS+MiniDT tracks (red) as a function of momentum. Left: x coordinate. Right: y coordinate.

Moreover, each MiniDT chamber provides enough information to produce a stand-alone track segment, as is done in the CMS trigger system. In CMS, each SL is used to produce a trigger primitive, i.e. a tentative track segment, which is then matched locally with other primitives at the DT level and then with tracker and calorimeter information to form a global track [66].

Even if the SND@LHC detector is smaller and simpler, the MiniDTs track segments could be exploited to provide seeds for the tracking algorithm, or to enable matching between DS and MiniDTs at the track level rather than at the level of measured points. A study is currently underway to identify the optimal strategy for fully leveraging the new subdetector capabilities. The status of the work is better discussed in Section 6.3.

6.2 Installation in TI18

The MiniDTs were installed in TI18 during the YETS 2024-2025. Preliminary work was necessary, from the reparation and commissioning of two MiniDTs to the design of the mechanical structure to hold every module in place, to the installation of the gas system, and to the integration into the SND@LHC data acquisition system.

6.2.1 Mechanic support

To install the two MiniDTs at the downstream end of the SND@LHC experiment, several design options for their placement were considered. The position of SND@LHC in TI18, near the curved tunnel wall, along with the limited available space, imposed strict constraints on integrating the chambers into the muon system. The final configuration places the MiniDTs between DS 3 and DS 4: the module closest to the SND@LHC target is dedicated to y-position measurements (MiniDT Y), followed by the one for x-position measurements (MiniDT X). This design also required the removal of the last iron block, which was not part of the HCAL but served only as a mechanical support for DS 3 and 4. Nevertheless, the best position available offers only a partial overlap of the target area ($\sim 83\%$) and of the DS modules ($\sim 73\%$).

The installed detectors are in Figure 6.9: on the left, the installed MiniDTs and the DS 4, relocated to its new position, while on the right, the installed electronics, described in Section 6.2.3.

6.2.2 Gas System

The MiniDTs are the first detector employing gas-based technology in SND@LHC. Before their installation, it was necessary to obtain the safety clearance and to install a dedicated gas system. The ideal location to store the gas bottle was identified in the UJ18 cavern, from which the TI18 tunnel departs, as Figure 2.1 shows. To mitigate the impact of the system on the underground area of LHC, it was decided to install the gas bottle support using pre-existing holes in the tunnel wall and to use an already in place metallic rail to support the gas pipes. Figure 6.10 shows the chosen location before and after installation of the gas bottle.

The installation of the gas system was done in parallel with the installation of the MiniDTs.

6.2.3 MiniDTs Readout electronics and integration in SND@LHC DAQ

The readout setup developed for the MiniDTs, described in Section 6.1.2, was operated in triggerless mode. In this configuration, all signals from the detector, filtered only by the front-end thresholds, set at 50 mV, are continuously streamed and recorded on a dedicated server located in a surface building. This threshold value was determined through preliminary studies aimed at optimising the signal-to-background ratio, summarised in Figure 6.11. While the HV value for the wire $V_{wire} = +3600$ V was the best for the efficiency, the choice for threshold was at first to employ 100 mV, trying to limit the overall rate, specifically from MiniDTX, but after the installation the noise rate proved to be




Figure 6.9: The MiniDTs installed. Left: side view of the updated muon system. It is possible to see the MiniDTs inserted between the DS 3 and DS 4 modules. Right: MiniDTs electronics.

much smaller with respect to this preliminary studies thus, since the rate became not a limiting factor, to lower the threshold to maximize the efficiency.

While a wire high voltage of $V_{wire} = +3600\text{ V}$ was identified as optimal in terms of efficiency, the discriminator threshold was initially set to 100 mV to limit the overall rate, in particular that of MiniDT X. After installation, however, the measured noise rate was found to be significantly lower than in the preliminary studies. As a consequence, the rate no longer represented a limiting factor, allowing the threshold to be lowered to maximise the detection efficiency. This choice is further motivated by the requirement of reconstructing tracks with four hits: even a small gain in single-hit efficiency has a non-negligible impact on the overall tracking performance.

During the installation, all MiniDTs electronics boards, except the HV and SND-PULSE boards, described later, were arranged in an *electronics box* built from $50 \times 50\text{ cm}^2$




Figure 6.10: The UJ18 tunnel wall position for the MiniDTs gas bottle. On the left, the two pre-existing holes chosen to fix the gas bottle and a metallic rail to run the gas pipes that connect the bottle to the MiniDTs. On the right, the gas bottle installed on its support.

aluminium grid plates (see Figure 6.9 right). This box was installed directly on their mechanical support downstream of the MiniDTs. The electronics box interfaces the MiniDTs with the SND@LHC experiment and the external systems through FE, LV, and slow-control cables, one TTC fibre for the LHC clock, one connection to the SNDPULSE board, and two Ethernet cables for readout and slow control.

The MiniDT data stream is independent of the SND@LHC data stream due to the different detector technology employed. The SND@LHC online event builder, described in Section 2.4.1, is optimal for full scintillator-based detectors as it collects events within a 25 ns-wide time window, matching LHC interaction rate. The drift time of the MiniDTs, up to ~ 380 ns, summed with the readout electronics latency, makes the timing of their signals incompatible with the rest of the detector.

The synchronisation of the two data streams is key to performing an offline match, associating MiniDTs hits to the correct SND@LHC event. The SND@LHC DAQ asserts a TTC command at the start of each run, configured to be synchronous with the orbit signal received from the LHC machine. This command is used to generate the MiniDT internal Bunch Crossing zero (BC0) signal, which, together with internal orbit and bunch crossing counters, is used by the electronics to ensure synchronisation.

The synchronisation test was done before installation, with a dedicated test setup, shown in Figure 6.12.

This figure shows the SNDPULSE board and its associated TOFPET, connected to an SND@LHC DAQ board. This was a proposed solution to simplify the offline matching.




Figure 6.11: Summary plot of the preliminary studies, varying the voltage applied to the wire (x axis) and the threshold applied at the FE level (marker shapes). Top: MiniDT Y. Bottom: MiniDT X.




Figure 6.12: The setup for the synchronisation tests of MiniDTs and SND@LHC DAQ systems.

The aim was to receive a signal in the SND@LHC data stream to be included in the event as a flag for MiniDTs activity. The TT bits asserted by the VCU118 when a TPG is reconstructed are sent to the SNDPULSE board, which reshapes the signal to make it compatible with the TOFPET, originally designed for analogue rather than digital inputs. The TOFPET is then connected to the SND DAQ board, where the information is integrated with the main detector readout and can be included in an SND event.

The synchronisation test demonstrated that it is possible to start the two data acquisition streams simultaneously and maintain their alignment throughout the entire data-taking period, both with respect to each other and to an external reference clock. In the test setup, this was provided by an external clock generator, while in the detector, it corresponds to the LHC clock.

The test also highlighted significant challenges in making the additional hit in the SND@LHC DAQ useful, due to a relatively high ($\sim 85\%$), but still insufficient, efficiency in detecting the signal converted by the SNDPULSE. This limitation arose from a combination of the non-optimal design of the SNDPULSE board and the difficulty in tuning the TOFPET parameters to accept the signal. Despite these discouraging results, the upgraded SNDPULSE board, with improved performance, together with the MiniDT DAQ board, were installed in their final configuration, as mentioned before, to enable further

studies during the first phase of the LHC recommissioning, before data taking at nominal energy and luminosity. However, the attempt carried out in this configuration also proved inconclusive, leading to the decision to rely only on time information for event matching.

6.3 MiniDT Performance

The first runs after installation were dedicated to verifying the operational status of the MiniDTs by checking the presence of signals across all channels and monitoring the overall signal rate from the chambers, as indicators of proper detector performance. These initial runs were performed with the MiniDTs operating in standalone mode, before finalising all the details required for joint operation with the SND@LHC detector, since a few modifications were needed to reproduce the synchronisation test setup.

The first test was to compute the rate of the MiniDTs. To do so, each hit was assigned a global timestamp calculated as

$$t_{hit} = n_{OC} \times T_{OC} + O \times T_O + BX \times T_{BX} + TDC \times \frac{25}{32} \quad (6.1)$$

where n_{OC} is the number of times the Orbit Counter (OC) counter has reset, $T_{OC} \approx 46.6$ s is the OC counter period, O is the hit orbit number, $T_O \approx 88.9 \mu\text{s}$ is the orbit period, BX is the hit bunch crossing number, $T_{BX} = 25$ ns is the single bunch slot, TDC is the hit fine TDC value, which is multiplied by the TDC resolution (32 bit, $\frac{25}{32}$ least significant bit).

Figure 6.13 shows the hit rate measured for each MiniDT layer, while Figure 6.14 reports the rate for individual cells. Consistent with preliminary studies, MiniDT X exhibits a higher noise level and an increased hit rate.




Figure 6.13: Hit rate per layer in a cosmic run for MiniDT X (left) and MiniDT Y (right).




Figure 6.14: Hit rate per channel in a cosmic run for MiniDT X (left) and MiniDT Y (right).

The rate of hits and TPGs provides an effective monitoring tool. For example, Figure 6.15 shows on the left the rate of TPGs during a stable beam run, while on the right the ratio of high-quality (H-quality) TPGs, reconstructed with one hit in each layer, to low-quality (L-quality) TPGs, made from three hits. The TPG distribution follows the beam intensity profile, and the H/L ratio remains constant over the entire run. But the plots also reveal a performance discrepancy between the two MiniDTs, which led to further investigation and the identification of a gas-related issue in the MiniDT system.

Continuous monitoring of hit and TPG rates revealed a progressive decrease in the H-quality TPG rate of MiniDT X, while MiniDT Y remained stable. The behaviour excluded HV or LV issues, indicating a significant gas leak. Limited tunnel access initially delayed the intervention, allowing the problem to persist until the LHC Technical Stop 1 (TS1). At the first available access, tests localised the leak to MiniDT Y, specifically at the slit between the LVFE box and the aluminium cover, which reduced the gas supply to MiniDT X. The problem was solved during TS1 by replacing the faulty MiniDT with a spare chamber and restoring the gas flow.

Following this repair, the same monitoring plots were employed to confirm that the gas leak had been fully resolved, and the MiniDTs performance, quantified by the number of hits per inverse nanobarn, shows a good stability over the data-taking period, as shown in Figure 6.16.

6.3.1 Event matching

At the start of each run, the MiniDT and SND@LHC DAQ systems are synchronised, as described in 6.2.3. To verify the synchronization, an offline study based on the presence




Figure 6.15: TPGs rate plot, part of the monitoring tools. Left: TPGs rate as a function of time for 3 and 4 hits TPGs for both MiniDT X and MiniDT Y. Right: ratio of high and low quality TPGs as a function of time.

of one H-quality TPG in each chamber within $[-15, 15]$ ns and corresponding hits in DS 3 and DS 4 on the SND@LHC side, with a timestamp difference $|t_{MiniDT} - t_{SND}| < 1 \mu s$. The time difference between TPGs and SND hits distribution in Figure 6.17 shows a narrow peak of about 70 ns, confirming that the two systems detect the same events.

The use of TPGs was viable for an initial verification, but incorporating MiniDT information into the SND data file requires the use of the single hits. At this stage, it was crucial to avoid introducing detector inefficiencies through restrictive data selection criteria, leaving more stringent cuts for the analysis of the collected data. Therefore, the final matching window for the difference was set to $[-200, 1000]$ ns. The widened interval also accounts for the fact that, on the SND@LHC side, the reference corresponds to the global event timestamp rather than the hit times in the DS stations. After matching the events, the selected ones must be included in the data files. For this purpose, it was necessary to implement in `sndsw` a dedicated class for the MiniDTs geometry and hit properties.

6.3.2 MiniDT preliminary efficiency

After matching the information from the MiniDTs with SND@LHC data, preliminary measurements of the MiniDT efficiency were performed. The goal is to measure the fraction of DS tracks that, when extrapolated to the MiniDT average position, correspond to events with at least three hits in a MiniDT.




Figure 6.16: The MiniDTs hits per nb^{-1} after TS1. The distribution shows a flat profile, as expected. No data were recorded at the beginning of September due to an LHC machine development period with no collision runs.

In this case, to reject noise, the chosen interval is different from the one described in Section 6.3.1: MiniDT hits are accepted within $[-150, 650]$ ns of the SND@LHC event timestamp. This choice was driven by the need to reject noise hits that could mistakenly be accounted for in the final efficiency value. On the SND@LHC side, the event selection required exactly one reconstructed muon track in the DS stations, with reduced $\chi^2 < 5$, and at least one hit in each horizontal and vertical plane of all DS stations except DS 4. The reconstructed track is then extrapolated to the MiniDT midplane, the plane positioned at the geometric centre of the MiniDT, equidistant from its two external surfaces.

The efficiency of a MiniDT was defined as

$$\epsilon = \frac{(\text{DS track in } xy \text{ bin}) \wedge (\geq 3 \text{ MiniDT hits})}{(\text{DS track in } xy \text{ bin})}$$

which does not rely on spatial matching between DS extrapolated tracks and MiniDT hit positions. The contribution of accidental coincidences is expected to be negligible. The resulting efficiency maps are shown in Figure 6.18, where the blue profile reflects the partial DS–MiniDT overlap already discussed in Section 6.2.1.

Both MiniDTs reach $\sim 97\%$, in agreement with expectations from CMS DT performance.




Figure 6.17: Time difference between MiniDTs and SND@LHC in matched events.




Figure 6.18: Efficiency as a function of position for both MiniDT X (left) and MiniDT Y (right). In red, the DS system position, in green, the position of the MiniDT.

6.4 Next steps

The first performance studies confirmed the proper operation of the MiniDTs: no dead channels were observed, the measured rate is consistent with the muon rate in SND, and

the preliminary efficiency evaluation further validates that their performance is consistent with expectations.

The MiniDTs information can be matched to the corresponding SND@LHC event, exploiting the timestamp of the events, given the synchronicity of the two data acquisition streams. The necessary updates in `sndsw` to handle the MiniDT geometry and hits have also been implemented.

All this preliminary work makes it possible to focus on the main objective of the upgrade: the production of tracks incorporating MiniDT information and the investigation of the most effective strategies to exploit them for muon track reconstruction, to improve future analyses.

Indeed, muon identification and tracking are essential in SND@LHC. An accurate measurement of the total muon flux is necessary for the correct operation of the detector, enabling the determination of the most suitable emulsion exposure time and programming the target replacement. An enhanced muon tracking resolution will significantly improve the identification of physics processes with muons in the final state, including charged-current muon neutrino interactions, muon deep inelastic scattering, and potential Beyond the Standard Model signatures, while also providing a more robust determination of background contributions in a wide range of analyses. As an example, muon tracks play a central role in the current muon neutrino analysis of the 2024 dataset, which is described in detail in Chapter 5. They are used not only to identify the correct interaction topology but also as one of the most effective handles for background rejection. In this analysis, the reconstructed muon track is extrapolated back to the target region and matched to the tracker modules to associate the muon with a candidate neutrino interaction (see Section 5.2.3). Given the currently limited spatial resolution of the muon tracking system, this matching requires relatively large acceptance windows compared to the dimensions of the tracker modules, which in turn reduces the discriminating power against background events: halving the size of this matching window would lead to an improvement in both signal efficiency and purity.

Conclusion

SND@LHC pursues ambitious objectives: from collecting data on neutrino interactions in previously unexplored energy ranges, to refining the theoretical understanding of neutrinos, to performing precise tests of the Standard Model, and ultimately to probing potential signatures of new physics. One drawback of being a recently deployed experiment is the ongoing process of identifying areas that could be improved; however, this also creates opportunities for upgrade projects aimed at maximising performance.

The first significant limitation identified during the initial analyses concerned the Veto System. To address this, the collaboration decided to upgrade the system by adjusting its placement with respect to the target region to improve the coverage and by adding one more module to enhance efficiency. I contributed to this effort by participating in the assembly of the new module and its initial tests. In particular, I carried out a study to validate that the single-sided readout configuration does not compromise the overall module efficiency. Once this was confirmed, I also took part in the installation of the new plane in its updated position. A dedicated study based on the 2024 dataset, corresponding to an integrated luminosity of 102 fb^{-1} , collected in this upgrade configuration, confirmed the positive outcome of the intervention. The upper limit on the total system inefficiency was computed to be $(4.9 \pm 1.9) \times 10^{-9}$ over the whole target area. The previous best estimate from 2023 data analysis was 6.6×10^{-6} , optimised to reach $(7.8 \pm 2.8) \times 10^{-8}$ on a smaller fiducial area covering 64% of the target area.

An equally important step toward achieving more precise analyses was the calibration of the hadronic calorimeter. This was achieved through a dedicated test beam campaign, where I contributed to every stage, from assembling and testing the hardware to data acquisition and the subsequent analysis that produced the calibration constants. Hadronic showers are crucial for the SND@LHC physics program, as they occur in interactions of all neutrino flavours, and their energy reconstruction is essential both for estimating the incoming neutrino energy and thus for performing spectrum and interaction cross-section measurements.

Being directly involved in both the hardware upgrade and the calibration of the detector, I then started the search for charged current muon neutrino interaction in the data collected in 2024, after the installation of the additional veto plane. The idea was not to extend the previous analysis done on 2022 data, but to develop an entirely new strategy.

To this end, I studied the detector response in detail to gain a deeper understanding of its behaviour and to define a set of selection cuts based not only on global event properties or overall quantities.

While the aim is to compute the neutrino energy spectrum and interaction cross-section as a function of energy, the analysis is still in the phase of selecting neutrino interaction candidates. Most of the cuts have been studied in depth and fully understood; however, the discrepancy with the 2022-23 dataset suggests the need for further investigation using Monte Carlo samples, which are not yet fully available.

The analysis of the improved data set is still ongoing, but in parallel, my work has also turned back to the hardware side, focusing on the upgrade of the muon system. Precise muon tracking in SND@LHC is important for many reasons: muons are the main background in neutrino searches, muon flux is crucial to determine the right moment for emulsion replacement to optimise exposure without exceeding what is reconstructable, and for physics searches of rare events with muons in the final state, such as muon trident production or the dark Higgs dimuon decay. To enhance performance, it was proposed to complement the DS modules with two drift chambers, called MiniDTs, based on the design of CMS DT chambers.

These two modules were installed in March 2025, following the initial phase of the project, which addressed the key technical challenges: how to integrate the MiniDTs in the SND detector, how to install a gas system in the LHC underground area, and how to prove synchronisation with this new subsystem was possible despite the very different technology employed. During this phase, I made significant contributions to the project's organisational aspects, particularly by supporting the acquisition of the necessary safety clearances for the gas system, coordinating and participating in the DAQ synchronisation tests, and assisting in organising the installation activities.

After the installation, performance was assessed and proved to be in line with expectations. To achieve these results, a major challenge was tackled: merging the MiniDT information, collected through a dedicated data acquisition flow, into SND@LHC data. The matching was successfully done, and dedicated classes were added to `sndsw`. The next phase of work involves integrating MiniDT information into the tracking process to reach the final goal of an enhanced tracking resolution.

Appendix A

Additional Figures

A.1 Timing studies




Figure A.1: Neutrino candidates sample. Distribution of QCD as a function of the time difference between the hit timestamp and t_{ref} (see Section 4.2) for SciFi 1. Since this is the first module, there are no hits in SciFi stations before the shower starts. Left: hits in the SciFi station where the shower develops. Right: hits in SciFi stations after the shower.




Figure A.2: Neutrino candidates sample. Distribution of QCD as a function of the time difference between the hit timestamp and t_{ref} (see Section 4.2) for SciFi 2. Top left: hits in SciFi stations before the shower starts. Top right: hits in the SciFi station where the shower develops. Bottom: hits in SciFi stations after the shower.




Figure A.3: Neutrino candidates sample. Distribution of QCD as a function of the time difference between the hit timestamp and t_{ref} (see Section 4.2) for SciFi 4. Top left: hits in SciFi stations before the shower starts. Top right: hits in the SciFi station where the shower develops. Bottom: hits in SciFi stations after the shower.




Figure A.4: Neutrino candidates sample. Distribution of QCD as a function of the time difference between the hit timestamp and t_{ref} (see Section 4.2) for SciFi 5. Left: hits in SciFi stations before the shower starts. Right: hits in the SciFi station where the shower develops. Since this is the last wall, there are no hits in SciFi stations after the shower.

A.2 Fiducial area




Figure A.5: Distribution of the reconstructed muon track position in SciFi 2. The solid red line shows the position of the tracker module, while the red dashed line shows the fiducial area.




Figure A.6: Distribution of the reconstructed muon track position in SciFi 3. The solid red line shows the position of the tracker module, while the red dashed line shows the fiducial area.




Figure A.7: Distribution of the reconstructed muon track position in SciFi 4. The solid red line shows the position of the tracker module, while the red dashed line shows the fiducial area.




Figure A.8: Distribution of the reconstructed muon track position in SciFi 5. The solid red line shows the position of the tracker module, while the red dashed line shows the fiducial area.

A.3 Centroid plots




Figure A.9: Distribution of the centroid residual as a function of the muon position measured in the wall preceding shower development. Data for SciFi 3X (left) and SciFi 3Y (right).




Figure A.10: Distribution of the centroid residual as a function of the muon position measured in the wall preceding shower development. Data for SciFi 4X (left) and SciFi 4Y (right).




Figure A.11: Distribution of the centroid residual as a function of the muon position measured in the wall preceding shower development. Data for SciFi 5X (left) and SciFi 5Y (right).

Bibliography

- [1] Particle Data Group. “Review of Particle Physics”. In: *Progress of Theoretical and Experimental Physics* 2020.8 (Aug. 2020), p. 083C01. ISSN: 2050-3911. DOI: 10.1093/ptep/ptaa104. eprint: <https://academic.oup.com/ptep/article-pdf/2020/8/083C01/34673722/ptaa104.pdf>. URL: <https://doi.org/10.1093/ptep/ptaa104>.
- [2] The IceCube collaboration. “Measurement of the multi-TeV neutrino interaction cross-section with IceCube using Earth absorption”. In: *Nature* 551.7682 (Nov. 2017), pp. 596–600. ISSN: 1476-4687. DOI: 10.1038/nature24459. URL: <http://dx.doi.org/10.1038/nature24459>.
- [3] *Observation of an ultra-high-energy cosmic neutrino with KM3NeT - Nature — nature.com*. <https://www.nature.com/articles/s41586-024-08543-1>. [Accessed 14-07-2025].
- [4] Mauricio Bustamante et al. “Extracting the Energy-Dependent Neutrino-Nucleon Cross Section above 10 TeV Using IceCube Showers”. In: *Phys. Rev. Lett.* 122 (4 Jan. 2019), p. 041101. DOI: 10.1103/PhysRevLett.122.041101. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.122.041101>.
- [5] Gargamelle Collaboration. “Total cross sections for ν_e and $\bar{\nu}_e$ interactions and search for neutrino oscillations and decay”. In: *Nuclear Physics B* 133.2 (1978), pp. 205–219. ISSN: 0550-3213. DOI: [https://doi.org/10.1016/0550-3213\(78\)90299-7](https://doi.org/10.1016/0550-3213(78)90299-7). URL: <https://www.sciencedirect.com/science/article/pii/0550321378902997>.
- [6] N. J. Baker et al. “ $\nu_\mu - \nu_e$ universality in charged-current neutrino interactions”. In: *Phys. Rev. D* 41 (9 May 1990), pp. 2653–2666. DOI: 10.1103/PhysRevD.41.2653. URL: <https://link.aps.org/doi/10.1103/PhysRevD.41.2653>.
- [7] *Measurement of high- Q^2 charged current cross sections in $e + p$ deep inelastic scattering at HERA - The European Physical Journal C — doi.org*. <http://doi.org/10.1140/epjc/s2003-01363-5>. [Accessed 14-07-2025].
- [8] The DONuT Collaboration. “Final tau-neutrino results from the DONuT experiment”. In: *Physical Review D* 78.5 (Sept. 2008). ISSN: 1550-2368. DOI: 10.1103/physrevd.78.052002. URL: <http://dx.doi.org/10.1103/PhysRevD.78.052002>.

- [9] OPERA Collaboration. *Final results of the search for $\nu_\mu \rightarrow \nu_e$ oscillations with the OPERA detector in the CNGS beam*. 2018. arXiv: 1803.11400 [hep-ex]. URL: <https://arxiv.org/abs/1803.11400>.
- [10] Alvaro De Rújula et al. “Neutrino fluxes at future hadron colliders”. In: *Nucl. Phys. B* 405 (1993), pp. 80–108. DOI: 10.1016/0550-3213(93)90427-Q. URL: <https://cds.cern.ch/record/239864>.
- [11] Klaus Winter. “Detection of the tau-neutrino at the LHC”. In: (1990). DOI: 10.5170/CERN-1990-010-V-2.37. URL: <https://cds.cern.ch/record/220304>.
- [12] François Vannucci. *Neutrino physics at LHC/SSC*. Tech. rep. Paris: Paris 6. Lab. Phys. Nucl. Théor. Hautes Energ., 1993. URL: <https://cds.cern.ch/record/253670>.
- [13] N Beni et al. “Physics potential of an experiment using LHC neutrinos”. In: *Journal of Physics G: Nuclear and Particle Physics* 46.11 (Oct. 2019), p. 115008. ISSN: 1361-6471. DOI: 10.1088/1361-6471/ab3f7c. URL: <http://dx.doi.org/10.1088/1361-6471/ab3f7c>.
- [14] S. Buontempo et al. *CMS-XSEN: LHC Neutrinos at CMS. Experiment Feasibility Study*. Tech. rep. (1) INFN, Napoli, Italy, (2) INFN, Bologna, Italy, (3) Univ. Federico II and INFN, Napoli, Italy, (4) Boston University, USA, (5) Dipartimento di Fisica and INFN, Bologna, Italy CERN CMS NOTE-2018/001 , 9 pages with 3 figures and 1 table, March 2018. Geneva: CERN, 2018. arXiv: 1804.04413. URL: <https://cds.cern.ch/record/2309742>.
- [15] G. Acampora et al. “The scattering and neutrino detector at the LHC”. In: *Journal of Instrumentation* 19.05 (May 2024), P05067.
- [16] S. Navas et al. “Review of Particle Physics”. In: *Phys. Rev. D* 110.3 (2024). See section “Structure Functions” for the definition of Bjorken x , p. 030001. DOI: 10.1103/PhysRevD.110.030001. URL: <https://pdg.lbl.gov/2025/reviews/rpp2024-rev-structure-functions.pdf>.
- [17] The SND@LHC collaboration. *Scattering and Neutrino Detector at the LHC*. Tech. rep. Geneva: CERN, 2020. URL: <https://cds.cern.ch/record/2729015>.
- [18] S. Roesler et al. “The Monte Carlo Event Generator DPMJET-III”. In: *Advanced Monte Carlo for Radiation Physics, Particle Transport Simulation and Applications*. Springer Berlin Heidelberg, 2001, pp. 1033–1038. ISBN: 9783642182112. DOI: 10.1007/978-3-642-18211-2_166. URL: http://dx.doi.org/10.1007/978-3-642-18211-2_166.
- [19] T T Böhlen et al. “The FLUKA Code: Developments and Challenges for High Energy and Medical Applications”. In: *Nucl. Data Sheets* 120 (2014), pp. 211–214. DOI: 10.1016/j.nds.2014.07.049. URL: <https://cds.cern.ch/record/2109973>.

- [20] The LHCb collaboration et al. “Measurements of prompt charm production cross-sections in pp collisions at $\sqrt{s} = 13$ TeV”. In: *Journal of High Energy Physics* 3 (2016), p. 159. DOI: 10.1007/JHEP03(2016)159. URL: [https://doi.org/10.1007/JHEP03\(2016\)159](https://doi.org/10.1007/JHEP03(2016)159).
- [21] C Ahdida et al. *SND@LHC - Scattering and Neutrino Detector at the LHC*. Tech. rep. Geneva: CERN, 2021. URL: <https://cds.cern.ch/record/2750060>.
- [22] Sergey Alekhin et al. “A facility to search for hidden particles at the CERN SPS: the SHiP physics case”. In: *Reports on Progress in Physics* 79.12 (Oct. 2016), p. 124201. DOI: 10.1088/0034-4885/79/12/124201. URL: <https://dx.doi.org/10.1088/0034-4885/79/12/124201>.
- [23] A. Blondel et al. “Electroweak parameters from a high statistics neutrino nucleon scattering experiment”. In: *Z. Phys. C - Particles and Fields* 45 (1990), pp. 361–379. DOI: 10.1007/BF01549665. URL: <https://doi.org/10.1007/BF01549665>.
- [24] Marc Schumann. “Direct detection of WIMP dark matter: concepts and status”. In: *Journal of Physics G: Nuclear and Particle Physics* 46.10 (Sept. 2019), p. 103003. DOI: 10.1088/1361-6471/ab2ea5. URL: <https://dx.doi.org/10.1088/1361-6471/ab2ea5>.
- [25] European Strategy for Particle Physics Preparatory Group. *Physics Briefing Book*. 2020. arXiv: 1910.11775 [hep-ex]. URL: <https://arxiv.org/abs/1910.11775>.
- [26] Marco Battaglieri et al. *US Cosmic Visions: New Ideas in Dark Matter 2017: Community Report*. 2017. arXiv: 1707.04591 [hep-ph]. URL: <https://arxiv.org/abs/1707.04591>.
- [27] E et al. Kou. “The Belle II Physics Book”. In: *Progress of Theoretical and Experimental Physics* 2019.12 (Dec. 2019). ISSN: 2050-3911. DOI: 10.1093/ptep/ptz106. URL: <http://dx.doi.org/10.1093/ptep/ptz106>.
- [28] D. Banerjee et al. “Dark Matter Search in Missing Energy Events with NA64”. In: *Physical Review Letters* 123.12 (Sept. 2019). ISSN: 1079-7114. DOI: 10.1103/physrevlett.123.121801. URL: <http://dx.doi.org/10.1103/PhysRevLett.123.121801>.
- [29] Torsten Åkesson et al. *Light Dark Matter eXperiment (LDMX)*. 2018. arXiv: 1808.05219 [hep-ex]. URL: <https://arxiv.org/abs/1808.05219>.
- [30] Alexey Boyarsky et al. “Searches for new physics at SND@LHC”. In: *Journal of High Energy Physics* 2022.3 (Mar. 2022). ISSN: 1029-8479. DOI: 10.1007/jhep03(2022)006. URL: [http://dx.doi.org/10.1007/JHEP03\(2022\)006](http://dx.doi.org/10.1007/JHEP03(2022)006).

- [31] G. Acampora et al. “SND@LHC: The Scattering and Neutrino Detector at the LHC”. In: *JINST* 19.05 (2024). 57 pages. To be published in JINST, P05067. DOI: 10.1088/1748-0221/19/05/P05067. arXiv: 2210.02784. URL: <https://cds.cern.ch/record/2838901>.
- [32] R. Albanese et al. “Observation of Collider Muon Neutrinos with the SND@LHC Experiment”. In: *Phys. Rev. Lett.* 131 (3 July 2023), p. 031802. DOI: 10.1103/PhysRevLett.131.031802. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.131.031802>.
- [33] The SND@LHC collaboration. *Addendum to the AdvancedSND LoI*. Tech. rep. Geneva: CERN, 2024. URL: <https://cds.cern.ch/record/2909524>.
- [34] ELJen. *EJ-200*. <https://eljentechnology.com/products/plastic-scintillators/ej-200-ej-204-ej-208-ej-212>.
- [35] Hamamatsu. *S14160-6050HS*. https://www.hamamatsu.com/eu/en/product/optical-sensors/mppc/mppc_array/S14160-6050HS.html.
- [36] PETsys. *TOFPET2 ASIC*. <https://www.petsyselectronics.com/web/public/products/1>.
- [37] N Beni et al. “Further studies on the physics potential of an experiment using LHC neutrinos”. In: *Journal of Physics G: Nuclear and Particle Physics* 47.12 (Nov. 2020), p. 125004. ISSN: 1361-6471. DOI: 10.1088/1361-6471/aba7ad. URL: <http://dx.doi.org/10.1088/1361-6471/aba7ad>.
- [38] The LHCb Collaboration. *LHCb Tracker Upgrade Technical Design Report*. Tech. rep. 2014. DOI: 10.17181/CERN.000E.909T. URL: <https://cds.cern.ch/record/1647400>.
- [39] Altera Corporation. *Cyclone V GT FPGA Development Board Reference Manual*. Technical Reference Manual. Document ID: 654538. Intel (formerly Altera). Sept. 2014. URL: <https://www.intel.com/programmable/technical-pdfs/654538.pdf>.
- [40] CERN TTC Project. *Introduction to the TTC System*. <https://ttc.web.cern.ch/intro.html>. 2025.
- [41] Enclustra GmbH. *Mercury SA1 System-on-Chip Module*. <https://www.enclustra.com/en/products/system-on-chip-modules/mercury-sa1/>. 2025.
- [42] LHC Collaboration. *Measurement of the muon flux at the SND@LHC experiment*. 2023. arXiv: 2310.05536 [hep-ex]. URL: <https://arxiv.org/abs/2310.05536>.
- [43] The SND@LHC collaboration. “Installation and performance of the 3rd Veto plane at the SND@LHC detector”. In: *Journal of Instrumentation* 20.07 (July 2025), P07011. DOI: 10.1088/1748-0221/20/07/P07011. URL: <https://dx.doi.org/10.1088/1748-0221/20/07/P07011>.

- [44] Federico Cindolo et al. “Installation of the 3rd Veto plane”. In: (2024). URL: <https://cds.cern.ch/record/2923932>.
- [45] Luke Aidan Dyks et al. “Proton beam delivery study for the NA60+ experiment in the H8 beamline at the CERN SPS North Area”. In: (2025). URL: <https://cds.cern.ch/record/2923184>.
- [46] SND@LHC Collaboration. *Studies of Hadronic Showers in SND@LHC*. 2025. arXiv: 2504.01716 [physics.ins-det]. URL: <https://arxiv.org/abs/2504.01716>.
- [47] H Wind. *Principal component analysis, Pattern recognition for track finding, Interpolation and functional representation*. Tech. rep. Geneva: CERN, 1982. URL: <https://cds.cern.ch/record/139228>.
- [48] Alex Ferrara. “Reconstruction of hadronic showers in the SND@LHC experiment”. URL: <https://amslaurea.unibo.it/id/eprint/34233/>.
- [49] Henso Abreu et al. “First Direct Observation of Collider Neutrinos with FASER at the LHC”. In: *Phys. Rev. Lett.* 131 (3 July 2023), p. 031801. DOI: 10.1103/PhysRevLett.131.031801. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.131.031801>.
- [50] Roshan Mammen Abraham et al. “First Measurement of ν_e and ν_μ Interaction Cross Sections at the LHC with FASER’s Emulsion Detector”. In: *Phys. Rev. Lett.* 133 (2 July 2024), p. 021802. DOI: 10.1103/PhysRevLett.133.021802. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.133.021802>.
- [51] SND@LHC Collaboration. *Observation of collider neutrinos without final state muons with the SND@LHC experiment*. 2024. arXiv: 2411.18787 [hep-ex]. URL: <https://arxiv.org/abs/2411.18787>.
- [52] ROOT Team et al. *TSpectrum Class Reference*. <https://root.cern.ch/doc/master/classTSpectrum.html>. CERN, 2024.
- [53] Particle Data Group. *Atomic and Nuclear Properties of Materials: Iron (Fe)*. https://pdg.lbl.gov/2019/AtomicNuclearProperties/HTML/iron_Fe.html. 2019.
- [54] Particle Data Group. *Atomic and Nuclear Properties of Materials: Tungsten (W)*. https://pdg.lbl.gov/2019/AtomicNuclearProperties/HTML/tungsten_W.html. 2019.
- [55] P. V. Hough. “Method and Means for Recognizing Complex Patterns”. US3069654A. U.S. Patent, 3069654. Dec. 1962.
- [56] R. E. Kalman. “A New Approach to Linear Filtering and Prediction Problems”. In: *J. Fluids Eng.* 82.1 (1960), pp. 35–45. DOI: 10.1115/1.3662552.

- [57] Johannes Rauch et al. “GENFIT — a Generic Track-Fitting Toolkit”. In: *Journal of Physics: Conference Series* 608 (May 2015), p. 012042. ISSN: 1742-6596. DOI: 10.1088/1742-6596/608/1/012042. URL: <http://dx.doi.org/10.1088/1742-6596/608/1/012042>.
- [58] SND-LHC Collaboration. *sndsw: SND@LHC experiment software framework*. Software framework for simulation and reconstruction of the SND@LHC detector. 2025. URL: <https://github.com/SND-LHC/sndsw>.
- [59] Simona Ilieva Ilieva. “Measurement of the muon flux at SND@LHC”. In: (2023). URL: <https://cds.cern.ch/record/2859193>.
- [60] S. Chatrchyan et al. “The CMS Experiment at the CERN LHC”. In: *JINST* 3 (2008), S08004. DOI: 10.1088/1748-0221/3/08/S08004.
- [61] Giulia Paggi. “Construction and test of a cosmic ray telescope based on CMS Drift Tube chambers and data acquisition prototypes of the Phase-2 upgrade”. URL: <https://amslaurea.unibo.it/id/eprint/26567/>.
- [62] *Electrical Characteristics of Low Voltage Differential Signaling (LVDS) Interface Circuits*. Telecommunications Industry Association, 2001.
- [63] Javier Sastre Alvaro. *The OBDT board: A prototype for the Phase 2 Drift Tubes on-detector electronics*. Tech. rep. Geneva: CERN, 2020. URL: <https://cds.cern.ch/record/2797780>.
- [64] Paulo Moreira et al. *TTCrx and QPLL Mezzanine Card (TTCrq)*. Technical Report Rev. 0.1. CERN, Dec. 2003. URL: <https://ttc.web.cern.ch/TTCrqSpec.pdf>.
- [65] HPTD collaboration. *HPTC – High-Precision Timing Clock generator board for timing characterization*. GitLab repository (created 26 Sept 2019) :contentReference[oaicite:3]index=3. Sept. 2019. URL: <https://gitlab.cern.ch/HPTD/hptc>.
- [66] A Tapper et al. *CMS Technical Design Report for the Level-1 Trigger Upgrade*. Tech. rep. Additional contacts: Jeffrey Spalding, Fermilab, Jeffrey.Spalding@cern.ch Didier Contardo, Universite Claude Bernard-Lyon I, didier.claude.contardo@cern.ch. 2013. URL: <https://cds.cern.ch/record/1556311>.