Eur. Phys. J. C (2023) 83:1126
https://doi.org/10.1140/epjc/s10052-023-12249-5

THE EUROPEAN PHYSICAL JOURNAL C

Regular Article - Theoretical Physics

Towards the optimal beam dump experiment to search for feebly interacting particles

Kyrylo Bondarenko[1,2,3], **Alexey Boyarsky**[4], **Richard Jacobsson**[5], **Oleksii Mikulenko**[4], **Maksym Ovchynnikov**[4,6,a]

[1] IFPU, Institute for Fundamental Physics of the Universe, via Beirut 2, 34014 Trieste, Italy
[2] SISSA, via Bonomea 265, 34132 Trieste, Italy
[3] INFN, Sezione di Trieste, SISSA, Via Bonomea 265, 34136 Trieste, Italy
[4] Instituut-Lorentz, Leiden University, Niels Bohrweg 2, 2333 CA Leiden, The Netherlands
[5] European Organization for Nuclear Research (CERN), Geneva, Switzerland
[6] Institut für Astroteilchen Physik, Karlsruher Institut für Technologie (KIT), Hermann-von-Helmholtz-Platz 1, 76344 Eggenstein-Leopoldshafen, Germany

Received: 25 August 2023 / Accepted: 10 November 2023

Abstract Future searches for new physics beyond the Standard Model are without doubt in need of a diverse approach and experiments with complementary sensitivities to different types of classes of models. One of the directions that should be explored is feebly interacting particles (FIPs) with masses below the electroweak scale. The interest in FIPs has significantly increased in the last ten years. Searches for FIPs at colliders have intrinsic limitations in the region they may probe, significantly restricting exploration of the mass range $m_{\rm FIP} < 5-10\,{\rm GeV/c^2}$. Beam dump-like experiments, characterized by the possibility of extremely high luminosity at relatively high energies and the effective coverage of the production and decay acceptance, are the perfect option to generically explore the "coupling frontier" of the light FIPs. Several proposals for beam-dump detectors are currently being considered by CERN for implementation at the SPS ECN3 beam facility. In this paper, we analyse in depth how the characteristic geometric parameters of a beam dump experiment influence the signal yield. We apply an inclusive approach by considering the phenomenology of different types of FIPs. From the various production modes and kinematics, we demonstrate that the optimal layout that maximises the production and decay acceptance consists of a detector located on the beam-axis, at the shortest possible distance from the target defined by the systems required to suppress the beam-induced backgrounds.

1 Introduction

Despite the success of the Standard Model (SM) of particle physics, evidence for the existence of new physics beyond the Standard Model is already well established because the origin of neutrino oscillations, dark matter, and the baryon asymmetry of the Universe is not known. However, we have no solid predictions of where to search for it. New particles capable of resolving these problems can have masses from sub-eV to Planck scale and coupling constants with SM particles ranging many orders of magnitude. At this crossroad point of particle physics, it is essential to use efficiently available or planned experimental facilities to push forward different frontiers of physics, probing whole classes of models simultaneously.

If the mass of a new particle is below the EW scale, it may be produced at accelerators not only as a resonance but also in decays of SM particles, such as heavy bosons W, Z, h, as well as mesons π, D, B. This makes this range of masses of new particles especially interesting from an experimental point of view. In this mass range, new particles would have escaped detection, not because of the limit on available accelerator energy, but because their creation is extremely rare. Numerous searches at past experiments as well as at the LHC constrain large values of coupling constants, which is why new particles of this type are often called feebly-interacting particles, or FIPs (see e.g. [1,2]).

FIPs can play a direct role in the beyond SM phenomena, like e.g. heavy neutral leptons or HNLs in the sub-EW mass range explain neutrino masses via sea-saw mechanism and matter–antimatter asymmetry via their out-of-equilibrium kinetics in the early Universe at a temperature above 100 GeV.

[a] e-mail: o.maxim@gmail.com (corresponding author)

Published online: 13 December 2023

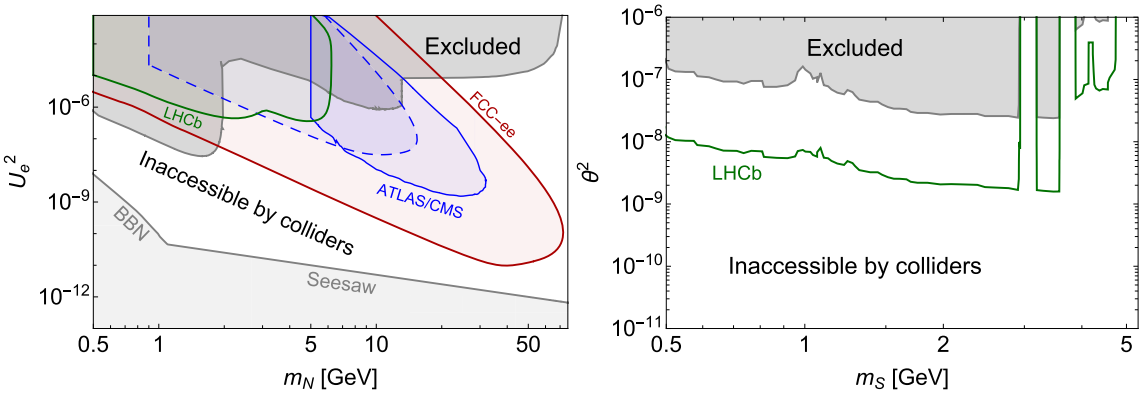

Fig. 1 The potential of future collider searches to probe the parameter space of feebly-interacting particles in the plane FIP mass - FIP coupling to SM particles. The figures demonstrate that colliders cannot efficiently explore the parameter space of FIPs with mass of the order of GeV. Left panel: parameter space of heavy neutral leptons (HNLs) that mix with electron neutrinos. The lower bound (seesaw) is defined by their ability to generate masses for active neutrinos [10]. LHC in high luminosity phase [5,7] and lepton colliders [6] are mainly sensitive to short-lived HNLs, with the typical lifetimes $c\tau_N \lesssim \mathcal{O}(100\,\text{m})$, see text for details. The scaling of the HNL lifetime with the mass is $\tau_N \propto m_N^{-5} U_e^{-2}$. As a result, colliders have poor sensitivity to HNL masses $m_N \lesssim 10$ GeV. Right panel: dark scalars mixing with Higgs bosons. Given the strict event selection to cope with the backgrounds and the available triggers, the event rate with displaced vertices at colliders is insufficient to provide a competitive sensitivity. Instead, scalars may be searched for with prompt events at LHCb [11], even though only relatively large couplings are within reach. The parameter space of these examples and other feebly-interacting particles in the GeV range can instead be more efficiently explored in the coming years with beam dump experiments

They can also be a "portal" that connects the SM sector with a Dark Sector, i.e. the case in which the Dark Sector particles only interact with ordinary matter via the FIP mediator (see e.g. [3]). FIPs, with their tiny coupling constants, form a "coupling frontier" of particle physics, and are part of a whole class of SM extensions.

If there is no particular reason (e.g. symmetry) for making these new particles stable, their lifetime scales with the coupling constant g and mass m_{FIP} as $\tau_{\text{FIP}} \propto g^{-2} m_{\text{FIP}}^{-\alpha}$ where $\alpha = 1-5$ [3]. Depending on the lifetime, different search strategies should be used to probe the FIP parameter space efficiently. In particular, particles with lifetimes $c\tau_{\text{FIP}} \gtrsim 1$ mm can be searched for via displaced-vertices schemes at the LHC and future colliders [4–7]. The limitation on $c\tau_{\text{FIP}}$ from below comes from numerous backgrounds caused by events with SM particles, which occur at small displacements. The limitation from above is FIP-dependent. It is caused by two factors. First, the collider detectors have short decay volumes of the order of $\mathcal{O}(1\,\text{m})$, typically defined by the dimensions of their inner trackers [4]. As a result, long-lived FIPs mainly decay well beyond the fiducial volume. Second, to reduce SM backgrounds, one must impose severe selection criteria on candidate events with FIPs, such as kinematic properties and specific final states, resulting in low signal efficiency. For instance, a typical selection efficiency of recent searches for HNLs at CMS [4] was of the order of 1%. Therefore, even if a long-lived FIP decays inside the tracker, the event will likely be outside the selection acceptance.

Typically, the FIP lifetime and its production rate are controlled by the same coupling, i.e. a significantly long FIP lifetime also means a small production rate.[1] Together with typically low signal efficiencies coming from the trigger and the event-versus-background selection, it is possible that the FIP production rate within the acceptance of the searches is insufficient to provide any sensitivity. This is the case of dark scalars with the coupling through mixing, dark photons, and axion-like particles [1]. Even for FIPs where the production is in principle sufficient (e.g., heavy neutral leptons [5]), future collider searches have a limited potential to probe the parameter space of GeV-scale FIPs. This is explained by the behavior of the FIP decay length $c\tau_{\text{FIP}}\gamma_{\text{FIP}} \propto g^{-2} m_{\text{FIP}}^{-\alpha-1}$ – for a given coupling g, the lifetime rapidly increases with decreasing m_{FIP}, in other words, quickly reducing the FIP decay probability within the fiducial volume and hence the sensitivity. We illustrate these points in Fig. 1, where we show the parameter space of heavy neutral leptons (HNLs) and dark scalars mixing with Higgs bosons. Future development of new trackers [8,9] at ATLAS, CMS, and LHCb may improve the LHC reach for light FIPs, but the remaining explorable region in the parameter space will still be large (see Sect. 6). Altogether, this suggests that we need a special experiment to search for FIPs in the GeV range.

In this paper, we argue that the most suitable experimental setup to search for FIPs with mass below 5 GeV is a beam dump experiment, where an extracted proton beam hits a

[1] A counter-example is the model of dark scalars with the mixing and quartic coupling to the Higgs bosons, $\mathcal{L} = \theta m_h^2 Sh + \alpha h S^2$. Depending on the value of α, the production may be dominated by the Higgs boson decay $h \to SS$, while the decay is mediated by θ.

Table 1 Parameters of the hypothetical experiment used as a reference experiment in our estimates: the longitudinal distance from the target to the beginning of the decay volume; the transverse dimensions of the decay volume and the detector, the longitudinal length of the decay volume; the longitudinal length of the detector; the distance from the centre of the detector in the transverse plane to the beamline. Here and below, we assume that the decay volume is oriented parallel to the beamline, which is motivated by the typical constraints from available space and infrastructure

$l_{\rm min}$	$S_{\rm det}$	$l_{\rm fid}$	$l_{\rm det}$	$r_{\rm displ}$
38 m	$4 \times 6\ {\rm m}^2$	50 m	15 m	0 m

dense target, and where the search is performed in a displaced decay volume. Although having much lower centre-of-mass energy of collisions than at colliders, beam dumps can deliver extremely high luminosity by operating with a more intensive proton beam combined with a high-A/Z target. This means in particular that they are capable of delivering a large number of mesons within a relatively small forward solid angle, in particular B, D, that may further decay into FIPs in the mass range of interest. There is no limitation on the decay volume length. It can easily be several tens of metres to cover a much larger lifetime acceptance than at collider detectors. Finally, backgrounds can be significantly reduced by placing the decay volume behind a chain of components designed to suppress beam-induced particle backgrounds, such as a hadron absorber and a muon deflector, considerably reducing the need to impose strict signal selection criteria. The best placement of such an experiment is the SPS accelerator at CERN, operating with a high-intensity proton beam of energy $E_p = 400$ GeV.

Several proposals of beam dump experiments at the SPS have been made [12–14]. They differ in geometric parameters, including the placement with respect to the beam axis, the decay volume size, and the detector angular coverage, both in terms of production acceptance and decay acceptance. The experiments also differ in the choice of material for the proton target. Maximum production of FIPs, and simultaneously maximum suppression of background from pion and kaon decays to muons and neutrinos, are achieved with a target of the highest possible atomic mass and atomic number, as well as minimized internal cooling for density. However, to discuss the optimal experimental layout, we assume in this paper the same target material for all experiments.

In this paper, we perform an analysis of the sensitivity to FIPs with respect to the geometric parameters of a beam dump experiment, in a maximally model-independent way, in order to find the optimal configuration. We start with an on-axis experiment specified in Table 1. Its close analog is the SHiP experiment [12,15]. We then study how the FIP sensitivity is affected by changing the parameters from the table. To be model-independent, we consider a few FIP models covering a wide class of production mechanisms and decay modes.

Our main results are summarized in Figs. 5, 6, 7, 8. We demonstrate that the setup in Table 1 is optimal for searching for FIPs independently of their lifetime, being also compatible with the absence of backgrounds. One of the main reasons for providing the largest signal acceptance is the on-axis placement. The off-axis location leads to a substantial loss of acceptance, and significantly worsens the ability to reconstruct properties of FIPs such as mass, spin, and decay modes, see Fig. 8.

The paper is organized as follows. In Sect. 2, we start from the expression for the number of events and discuss the FIP phenomenology (Sect. 3), describing, in particular, their production and decay modes. In Sect. 4, we analyze how the number of events at the lower bound of the sensitivity varies with the experimental configuration, considering separately its placements on-axis (Sect. 4.1) and off-axis (Sect. 4.2) relatively to the beamline. In Sect. 4.3, we study the impact of the configuration on the potential to probe FIPs at the upper bound of the sensitivity. In Sect. 5, we apply our findings to compare the sensitivity of the experiments proposed at SPS, SHiP, SHADOWS, and HIKE. Finally, in Sect. 6, we make conclusions. Appendices contain all the relevant technical information on the phenomenology of FIPs and the calculations.

2 Number of signal events

Let us start with the expression for the number of events in the regime of large lifetimes, where the typical decay length of FIPs, $c\tau_{\rm FIP}\langle\gamma_{\rm FIP}\rangle$, is much larger than the characteristic scale of the experiments $\simeq$ 100 m (the opposite cases of short lifetimes $c\tau_{\rm FIP}\langle\gamma_{\rm FIP}\rangle \lesssim l_{\rm min}$ is discussed in Sect. 4.3). In this case, the number of events may be represented in the following schematic form (see Appendix A):

$$N_{\rm events} \approx N_{X,\rm prod} \times \frac{l_{\rm fid}\langle p_{\rm FIP}^{-1}\rangle}{c\tau_{\rm FIP} m_{\rm FIP}} \times \epsilon_{\rm geom} \times \epsilon_{\rm rec} \times {\rm Br}_{\rm vis} \quad (1)$$

Here, $N_{X,\rm prod} = N_{\rm PoT} \times {\rm Br}(pp \to X)$ is the total number of the FIPs produced in the collisions, with $N_{\rm PoT}$ being the number of protons on target, and ${\rm Br}(pp \to X)$ the probability of producing of FIPs by any mechanism; the second factor is the decay probability in the regime of large lifetimes, with $\langle p_{\rm FIP}^{-1}\rangle$ being the mean inverse momentum for the FIP at the experiment; $\epsilon_{\rm geom}$ is the overall geometric acceptance, following from geometric limitations of the decay volume and detectors (discussed in details below); ${\rm Br}_{\rm vis}$ is the branching ratio of the decay of FIPs into states visible at the given experiment; finally, $\epsilon_{\rm rec}$ is the total reconstruction efficiency – the fraction of events within the geometric acceptance that may be reconstructed.

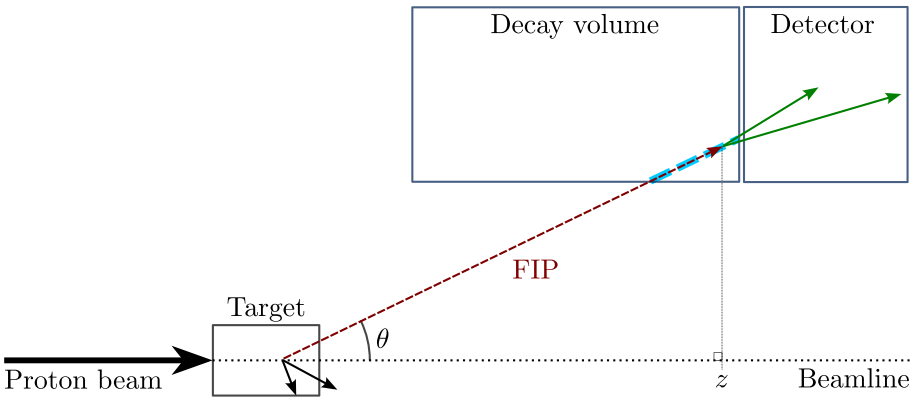

Fig. 2 Illustration of the impact of different contributions to the geometric acceptance defined by Eq. (3). First, the FIPs produced by collisions of the proton beam with the fixed target must point to the detector (the red arrow). The fraction of such events is given by ϵ_{FIP}. The effective length inside the decay volume passed by decaying FIPs (the dashed blue line) may differ significantly from the nominal decay volume length l_{fid}. This results in the factor $l_{\text{fid,eff}}/l_{\text{fid}}$. Finally, the decay products of FIPs (the green arrows) also have to point to the detector, which is incorporated by ϵ_{dec}

Independently of the configuration, we will assume $\epsilon_{\text{rec}} = 1$, and that Br_{vis} includes all decays with at least two electrically charged particles or two photons.

Considering two experiments located at the same facility (such that the beam energy and target configuration are the same), for the ratio of the number of events (1) one has

$$\frac{N_{\text{events},1}}{N_{\text{events},2}} \approx \frac{l_{\text{fid},1}}{l_{\text{fid},2}} \times \frac{\epsilon_{\text{geom},1}}{\epsilon_{\text{geom},2}} \times \frac{\langle E_{\text{FIP}}^{-1}\rangle_1}{\langle E_{\text{FIP}}^{-1}\rangle_2} \tag{2}$$

Therefore, to compare the lower bounds of sensitivity of two experiments, we need to understand the behavior of ϵ_{geom} and $\langle E_{\text{FIP}}^{-1}\rangle$.

2.1 Geometric acceptance

Schematically, the geometric acceptance is given by

$$\epsilon_{\text{geom}} \simeq \epsilon_{\text{FIP}} \times l_{\text{fid}}^{\text{eff}}/l_{\text{fid}} \times \epsilon_{\text{dec}}, \tag{3}$$

see Fig. 2.

The first factor is the FIP acceptance, i.e., the fraction of FIPs with trajectories pointing to the cross-section of the end of the detector.[2] $l_{\text{fid}}^{\text{eff}}$ is effective fiducial volume length – the mean length inside the decay volume passed by FIPs pointing to the end of the detector. If the considered experiment is located on-axis, then $l_{\text{fid}}^{\text{eff}} \approx l_{\text{fid}}$. However, in case of the off-axis placement, parallel to the beamline, $l_{\text{fid}}^{\text{eff}}$ gets effectively reduced. Finally, the last factor is the decay product acceptance, i.e., the fraction of decays of FIPs within ϵ_{FIP} with at least two of their decay products pointing to the end of the detector. Roughly, the decay products of FIPs with the gamma factor γ_{FIP} have the opening angle[3]

$$\Delta\theta_{\text{dec}} \sim 2/\gamma_{\text{FIP}} \tag{4}$$

If this angle becomes comparable with the angle covered by the detector as seen from the FIP's decay point, then ϵ_{dec} would significantly reduce the event rate.

3 Kinematic distributions of FIPs

To understand the behavior of ϵ_{FIP} and ϵ_{dec}, we need to study in a model-independent way how FIPs may be produced in proton-target collisions and how they decay. To this extent, we consider different types of FIPs: dark photons, dark scalars, HNLs, and ALPs with photon coupling [1]. By considering all of them, we can perform the analysis in a maximally model-independent fashion. The dominant production mechanisms and decay modes are shown in Fig. 3 and in Table 2. The production channels are decays of mesons (light unflavoured mesons π, η, η', ρ^0, as well as heavy flavored mesons B, D), or the direct production through proton bremsstrahlung, Drell–Yan process, coherent proton-nucleus and photon-nucleus scattering. We generate the distribution of the light mesons at the SPS using the approach of [16] (see also [17]), and use the distribution of B, D mesons from [18]. We follow the description of the bremsstrahlung process from [19] and the coherent scattering from [16]. For

[2] Naively, the mentioned definition of the FIP acceptance looks too restrictive, as it does not take into account the FIPs decaying inside the decay volume but not pointing to the detector. However, because of the 4-momentum conservation, FIPs that do not point to the detector typically cannot decay into particles pointing to the detector. Therefore, instead of considering FIPs decaying in any direction inside the decay volume, we considered only the FIPs that already point to the detector.

[3] For simplicity, we considered here 2-body decays; however, for 3-body decays, the situation is qualitatively similar.

Table 2 Dominant production and decay modes of GeV-scale FIPs in proton-proton collisions at the SPS. We consider the dark photon, the ALP with the photon coupling, the Higgs-like scalar with mixing coupling, and HNLs with the dominant coupling to electron neutrinos. For more details, see [20]

FIP	Prod. modes	Decay modes
DP V	$\begin{cases}\pi^0/\eta \to V,\ m_V < m_\eta \\ \text{Brem/DIS},\ m_V > m_\eta\end{cases}$	$V \to ll$ $V \to 2\pi, 3\pi, KK, m_V \lesssim 1$ GeV $V \to qq, m_V \gtrsim 1$ GeV
ALP$_\gamma$ a	$\gamma + Z \to a + Z$ $p + Z \to p + Z + a$	$a \to \gamma\gamma$
Scalar S	$\begin{cases}K \to S + \pi,\ m_S < m_K - m_\pi \\ B \to S + X,\ m_S > m_K + m_\pi\end{cases}$	$S \to ll$ $S < \pi\pi/KK, m_S < 2$ GeV $S \to qq, m_S > 2$ GeV
HNL N	$\begin{cases}K \to N + X,\ m_N < m_K, \\ D \to N + X, m_N < m_{D_s}, \\ B \to N + X,\ m_N > m_{D_s}\end{cases}$	$N \to l l \nu$ $N \to \pi^0\nu, \eta\nu, \pi l, m_N \lesssim 1$ GeV $N \to qq\nu, qql, m_N \gtrsim 1$ GeV

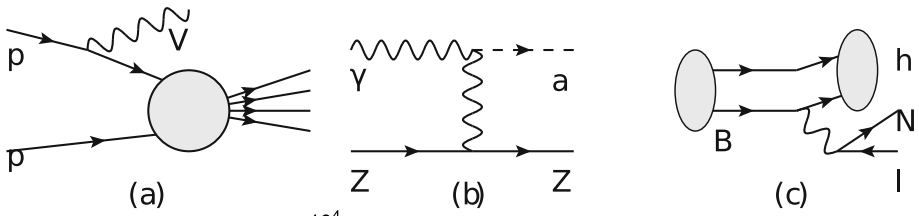

Fig. 3 Examples of production processes for various FIPs: **a** proton bremsstrahlung (dark photon V), **b** coherent scattering off nuclei (ALP with the photon coupling a), **c** decays of B mesons (HNLs N, dark scalars)

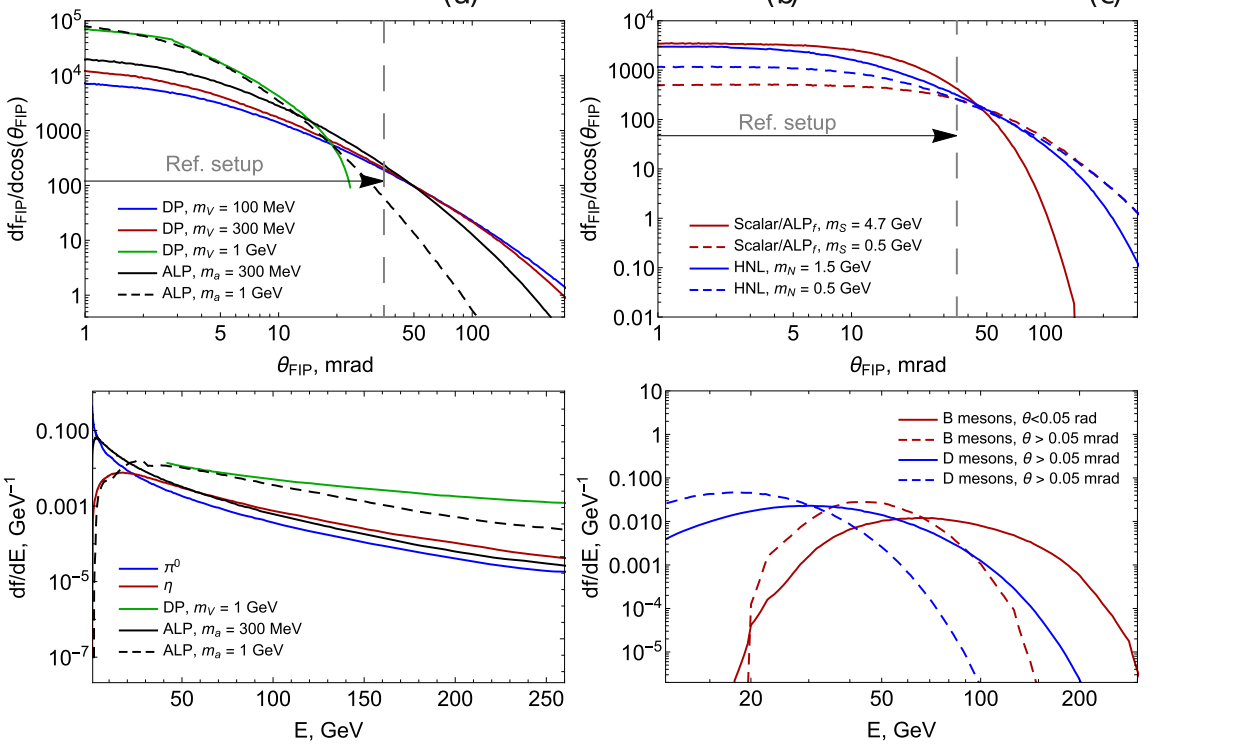

Fig. 4 Kinematics of FIPs produced in proton-target collisions at the SPS. Molybdenum target is considered. Top panels: solid angle distributions $df_{\rm FIP}/d\Omega_{\rm FIP} \sim df_{\rm FIP}/d\cos(\theta_{\rm FIP})$ of various FIPs. Different masses are considered, corresponding to different production channels (Table 2). Note that the distribution of heavy HNLs with $m_N > 3$ GeV is very similar to the distribution of scalars, because of the same mother particle and decay kinematics. The polar angle coverage of the detector of the reference setup (Table 1) is indicated with arrows and the vertical dashed line. Bottom panels: energy spectra of the mesons producing FIPs, dark photons produced by the proton bremsstrahlung, and ALPs with photon coupling. For the case of heavy mesons B, D, the distribution is shown assuming two different angular coverage: "on-axis" $\theta < 0.05$ rad, and "off-axis", $\theta > 0.05$ rad, to demonstrate how the spectrum gets softer off-axis. See text and Ref. [20] for details

the details on the derivation of the distributions of FIPs from these production channels, see Appendix A.

The solid angle distributions

$$df_{\rm FIP}/d\Omega_{\rm FIP} \sim df_{\rm FIP}/d\cos(\theta_{\rm FIP})) \tag{5}$$

of the FIPs produced by these mechanisms, as well as their energy distributions are shown in Fig. 4. Because of the kinematics of the collisions with a fixed target, the bulk of the distributions of FIPs are contained within a relatively small forward solid angle around the beam axis, being flat up to and quickly dropping at large angles $\theta > \theta_{\rm flat}$. The direct production processes are characterized by very small typical transverse momentum compared to the momentum of the incoming proton. In the case of Drell–Yan and proton bremsstrahlung, it is of the order of $m_p \sim 1$ GeV [21]. Given very large typical energies of the FIPs produced by these mechanisms, $E_{\rm FIP} \sim 100 - 200$ GeV, the angle is $\theta_{\rm flat} \simeq 10$ mrad. For the Primakov production off nuclei and nucleons, $\theta_{\rm flat}$ is determined by typical transverse momenta carried by virtual photons, being $\theta_{\rm flat} \simeq m_p/E_{\rm beam} = 2.5 \times 10^{-3}$ rad [22].

Light mesons produced in collisions have characteristic p_T of the order of $\Lambda_{\rm QCD}$, and relatively small mean energy of order of $\langle E\rangle \simeq 20 - 30$ GeV, which leads to $\theta_{\rm flat}$ being of the order of a few tenths of mrad [16].

Heavy mesons D, B have large p_T of the order of their mass, but, at the same time, much larger characteristic energies. As a result, their distributions start dropping at even smaller angles $\theta_{\rm flat} \lesssim 10$ mrad. For lighter masses, the FIPs produced by decays of B, D may have broader angular distribution. The reason is an additional transverse momentum of the order of the energy of the FIP at the rest frame of the decaying meson, which may be as large as the meson mass for FIPs with $m_{\rm FIP} \ll m_{\rm meson}$.

The main uncertainty for the FIP production (which translates to the main uncertainty in the signal rate estimates) comes from the thickness of the target, which gives rise to a cascade production of the SM particles that may potentially decay into FIPs. The cascade population affects both the total yield of the particles and the shape of the angle-energy distribution: the angular distribution of these extra particles is broader, while the energy distribution is shifted to lower energies.

While for the thin target and SPS beam energies the fluxes of mesons have been measured [23,24], observing the cascade flux in the case of thick targets is more complicated (as the mesons or their decay products may be absorbed inside the target) and requires a dedicated experiment. Therefore, we may only estimate it via simulations. Such estimates have already been performed by the SHiP collaboration [18,19]. They suggest that for heavy flavor production, the cascade enhancement gives a factor of 2 to the total meson yield and is important for both on-axis and off-axis experiments (see also [25]); we account for these cascade mesons in our calculations. For lighter mesons, the enhancement may be a factor of 10 or larger [19]. However, these mesons are very soft, with energies of order of 1 GeV, and decay into a very low-energy FIP produced in a broad range of polar angles. Even if such FIPs would decay inside the decay volume, the products would not be properly reconstructed because of low energies.[4] As a result, only a tiny fraction of such cascade mesons contributes to the signal yield. Therefore, we may safely neglect them.

The decay modes of FIPs may differ by the number of the decay products and their phase space if the number of products is fixed, which comes from the ratio of the FIP-to-decay product masses and the matrix element of the decay. The main decay modes of the FIPs are provided in Table 2.

Another uncertainty in calculating the number of events may come from theoretical uncertainties in the description of the phenomenology of FIPs. Namely, for the particles such as Higgs-like scalars, various ways to describe hadronic decays may change the decay width by 1–2 orders of magnitude (see, e.g., [26,27] and references therein), which affects both the lower and upper bounds of the sensitivity. While it may be important for the full sensitivity region, it is not relevant when computing the ratio of the number of events in the regime of long lifetimes, as we do in this paper – just the parameter space gets shifted.

4 Effect on the number of signal events from varying parameters of the experiment

In this section, we study the impact of varying the geometric parameters of the experiment on the number of events. To illustrate the effect of the decay products acceptance and the FIP angular distribution, we consider the following FIPs and properties:

- HNLs with masses 0.5 GeV, 1.5 GeV produced by decays of D mesons, and 4 GeV produced by decays of B.
- Dark scalars with masses 0.5 GeV and 4 GeV produced by decays of B.
- Dark photons with masses $m_V = 300$ MeV produced by decays of η and $m_V = 1$ GeV by proton bremsstrahlung.
- ALPs with photon coupling with masses $m_a = 300$ MeV and 1 GeV.

[4] For instance, selection criterion at SHiP for the decay products is $E > 1$ GeV.

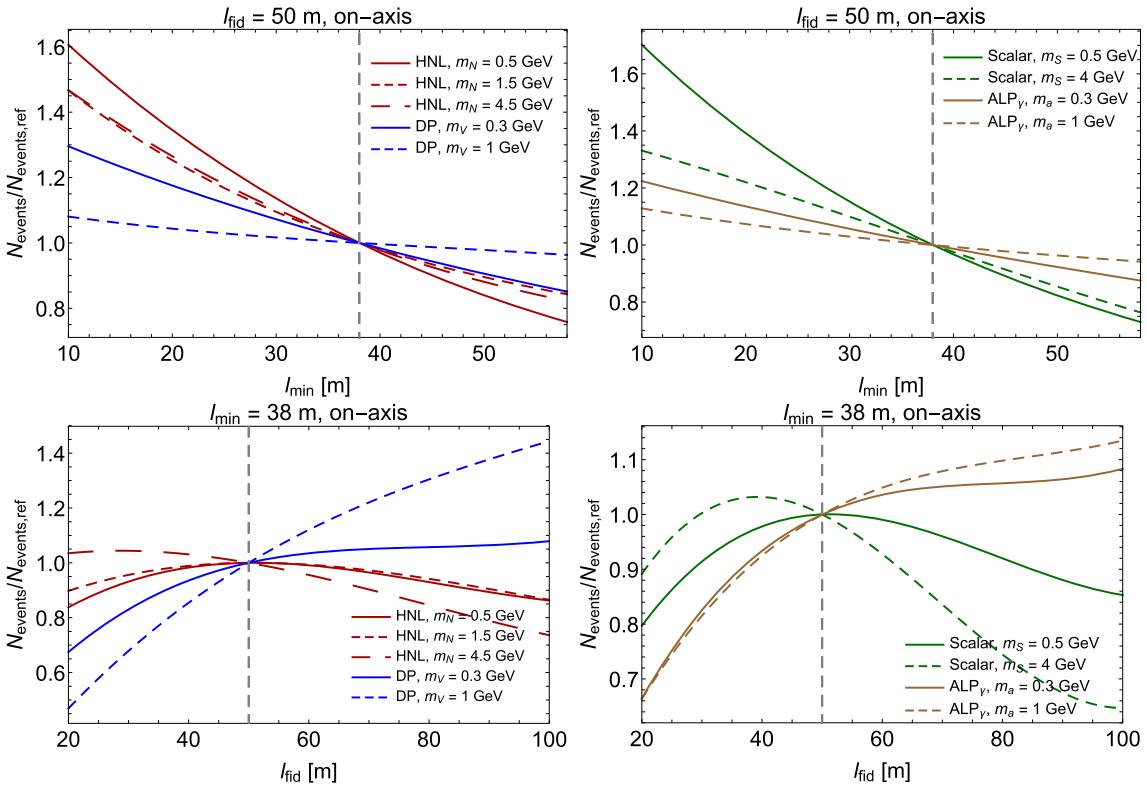

Fig. 5 The behavior of the number of signal events of a beam dump on-axis experiment at the SPS at the lower bound of the sensitivity ($c\tau_{\rm FIP}\langle\gamma_{\rm FIP}\rangle \gg 100$ m, see Sect. 2 for details) under change of the distance to the decay volume $l_{\rm min}$ (top panels) and its length $l_{\rm fid}$ (bottom panels) for different models of FIPs. On the one hand, changing these parameters may have a significant impact for the backgrounds to be removed, the complexity of the setup, and costs. On the other hand, the maximal impact of these parameters on the number of events is small, $< \mathcal{O}(2)$, see text for details. Therefore, we conclude that the optimization of these parameters should be a subject of background considerations and costs rather than the maximization of the number of FIP events. The other parameters defining the experimental setup – the transverse size of the decay volume and the detector dimensions – are summarized in Table 1. For convenience, we normalize the number of events to the one for the configuration from Table 1

4.1 On-axis location

Let us first assume the on-axis placement of the experiment. We analyze how the number of long-lived FIPs is affected by changing the distance to the decay volume $l_{\rm min}$, and its length $l_{\rm fid}$ (the case of short-lived FIPs with $c\tau_{\rm FIP}\gamma_{\rm FIP} \lesssim l_{\rm min}$ is discussed in Sect. 4.3).

The dependence of the number of events on $l_{\rm min}$ is shown in Fig. 5 (top panels). We normalize all the values to the corresponding values of the SHiP-like experiment from Table 1. The main impact of $l_{\rm min}$ is in defining the solid angle covered by the detector as seen from the target,

$$\Omega_{\text{det-target}} = S_{\rm det}/(l_{\rm min} + l_{\rm fid} + l_{\rm det})^2, \tag{6}$$

and hence the fraction of FIPs $\epsilon_{\rm FIP}$ pointing to the detector. If the FIP's angular distribution $df_{\rm FIP}/d\Omega$ (Fig. 4) is flat within the angles covered by the detector (the case of light HNLs, dark scalars, and ALPs), then $\epsilon_{\rm FIP} \propto \Omega_{\text{det-target}}$. From Eq. (6), it follows that to increase $N_{\rm events}$ for the SHiP-like configuration by a factor of two, it is necessary to decrease $l_{\rm min}$ from $l_{\rm min} = 38$ m to $l_{\rm min} \approx 8$ m. At the same time, if $df_{\rm FIP}/d\Omega$ is collimated and falls at the boundaries (dark photon, ALPs, and heavy HNLs/dark scalars), the effect on $\epsilon_{\rm FIP}$ is even smaller. On the other hand, the price of such a close placement would be a significant increase in the SM background.

The effect of $l_{\rm fid}$ is less trivial, see the bottom panels of Fig. 5. It affects the product $\epsilon_{\rm FIP} \times l_{\rm fid} \times \epsilon_{\rm dec}$. The impact on $\epsilon_{\rm FIP}$ is similar to the $l_{\rm min}$ case (Eq. (6)). The second factor comes from the decay probability of long-lived FIPs. In addition, by increasing $l_{\rm fid}$ and maintaining the aperture of the detector constant, $\epsilon_{\rm dec}$ decreases, and vice versa. Indeed, $l_{\rm fid}$ enters the solid angle covered by the detector as seen from the beginning of the decay volume,

$$\Omega_{\text{det-fid}} = S_{\rm det}/(l_{\rm fid} + l_{\rm det})^2 \tag{7}$$

If $l_{\rm fid}$ is too large, the opening angle between the FIP's decay products $\Delta\theta_{\rm dec}$ (Eq. (4)) becomes comparable with the detector size, and such events do not contribute. Hence, $l_{\rm fid} \times \epsilon_{\rm dec}$ remains constant under further increase of $l_{\rm fid}$.

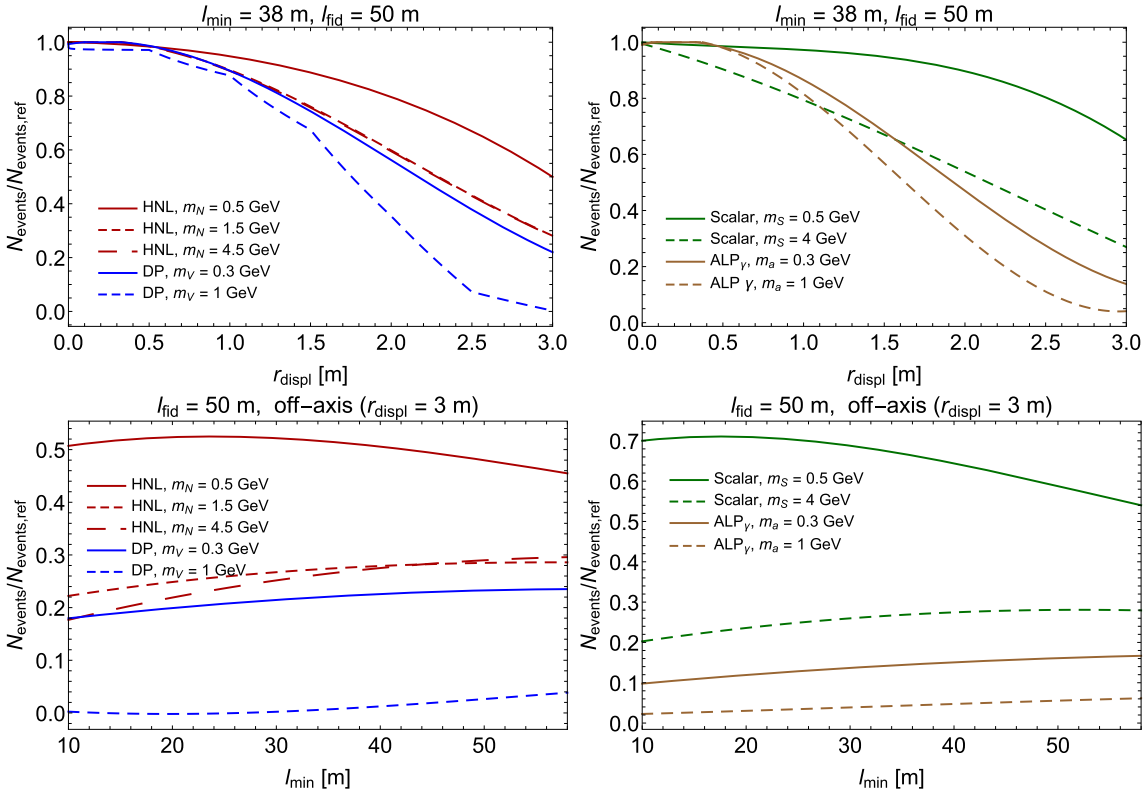

Fig. 6 The behavior of the number of signal events of a beam dump experiment at the SPS at the lower bound of the sensitivity ($c\tau_{\rm FIP}\langle\gamma_{\rm FIP}\rangle \gg 100$ m, see Sect. 2) assuming an off-axis placement of the centre of its decay volume parametrized in terms of the displacement $r_{\rm displ}$. From the figures, we see that independently on the FIP type, by increasing $r_{\rm displ}$, the number of events decreases. This results from the very forward-pointing FIP angular distribution that falls at large polar angles (Fig. 4). Depending on the FIP, the decrease may be an order of magnitude or larger (top panel). It is impossible to compensate for this decrease by placing the experiment closer to the target (bottom panel): despite the increase of the solid angle covered by the detector, the minimal covered polar angle increases, which again results in a decrease of the FIP flux. The other parameters defining the experiment – transverse dimensions of the decay volume, and detector dimensions – are fixed as specified in Table 1. Note that $r_{\rm displ}$ does not equal the off-axis displacement of the side of the decay volume. In particular, for the configuration considered, this displacement becomes non-zero only if $r_{\rm displ} > 2$ m. The displacement $r_{\rm displ} = 3$ m corresponds to 1 m gap between the side of the decay volume and beam axis. For convenience, we normalize the number of events to the one for the configuration specified in Table 1

As a result, the product $\epsilon_{\rm FIP} \times \epsilon_{\rm dec} \times l_{\rm fid}$ first scales linearly with $l_{\rm fid}$, then reaches maximum where $\Delta\theta_{\rm dec}$ becomes comparable with the detector size, and then decreases as a result of decreasing $\Omega_{\text{det-target}}$.

For the FIPs produced by decays of heavy mesons, the maximum is around $l_{\rm fid} = 50$ m. The situation is different for the FIPs produced by decays of π, η, bremsstrahlung, or Primakov process. In these cases, either the FIPs have small masses or very large energies (Fig. 4), which leads to a very large $\gamma_{\rm FIP}$. Hence, the suppression by $\epsilon_{\rm dec}$ is not so severe, meaning that the number of events increases with $l_{\rm fid}$ over a very large range.

4.2 Off-axis location

Let us now analyze the impact on the number of signal events when displacing the detector off-axis. As in the previous subsection, we first consider the configuration with the same dimensions and distance from the beam dump as in Table 1. We will consider only the parallel orientation of the decay volume and detector relative to the beamline, motivated by the limitations typically imposed by the infrastructure. Let us start with increasing the transverse displacement of the centre of the detector relative to the beamline, $r_{\rm displ}$. We vary $r_{\rm displ}$ from 0 to 5 m. Note that $r_{\rm displ}$ is not the same as the off-axis displacement of the side of the decay volume. In particular, for the considered configuration, the latter is > 0 only if $r_{\rm displ} > 2$ m. For $r_{\rm displ} = 3$ m, the gap between the side of the decay volume and the beam axis is 1 m.

The dependence of the number of events is shown in Fig. 6, top panels. For all FIP models considered, the number of events decreases with $r_{\rm displ}$. The reason for this is that the FIP angular distribution decreases at the larger polar angles (Fig. 4) covered with the off-axis displacement $r_{\rm displ}$. In the

case that the whole detector is placed off-axis ($r_{\rm displ} > 2$ m, such that the side of the decay volume is entirely away from the beam axis), we would not only shift the detector in the domain of large θ, but also decrease the azimuthal coverage in the domain of small polar angles, which further reduces the FIP acceptance. Finally, the acceptance gets further suppressed by the shortening of the effective decay volume length for FIPs that enter from the side (Eq. (3)).

For heavy FIPs, the drop is more significant, which is explained by more forward-pointing angular distributions and smaller decay acceptance. The decrease in $\epsilon_{\rm dec}$ is due to the softer energy spectrum at large polar angles. It leads to an increase of the typical opening angle between the decay products $\theta \simeq 1/\gamma_{\rm FIP}$ (Eq. (4)), and hence a decrease in $\epsilon_{\rm dec}$.

It is also interesting to compare the effect of shortening $l_{\rm min}$ in the case of off-axis and on-axis placements (Fig. 5 and discussion therein). The behavior of the number of events is illustrated in Fig. 6, bottom panels. Unlike the on-axis case, the number of events tends to grow if *increasing* $l_{\rm min}$, which is again a result of the very forward-pointing angular distributions and decreasing $\epsilon_{\rm dec}$ at the off-axis locations.

4.3 Upper bound of the sensitivity for on-axis and off-axis

Finally, let us also examine the role of the geometric parameters in determining the potential of the experiment to search for FIPs with large couplings, for which the typical FIP decay length is smaller than the distance to the decay volume, $c\tau_{\rm FIP}\langle\gamma_{\rm FIP}\rangle \lesssim l_{\rm min}$.

At the upper bound, the number of events is proportional to the following integral (see Appendix A):

$$N_{\rm events}^{\rm upper\ bound} \propto \epsilon = \int dL \int dE_{\rm FIP} \times \frac{1}{c\tau_{\rm FIP}\gamma_{\rm FIP}} \exp\left[-\frac{L}{c\tau_{\rm FIP}\gamma_{\rm FIP}}\right] f_{E_{\rm FIP},L}^{l_{\rm min}}, \quad (8)$$

where L is the modulus of the FIP decay position, and $f_{E_{\rm FIP},L}^{l_{\rm min}} \equiv \left\langle \frac{df_{\rm FIP}}{dE_{\rm FIP}}\epsilon_{\rm dec}\right\rangle_\theta$ is the FIP distribution averaged over the angular coverage of the detector. The integral (8) effectively plays the role of geometric acceptance in the case of short-lived FIPs. It is sensitive to the high-energy tail of the FIP distribution ($E_{\rm FIP} > \langle E_{\rm FIP}\rangle$).

We will start with the simpler case of the on-axis configuration. In this case, $f_{E_{\rm FIP},L}^{l_{\rm min}}$ depends weakly on $l_{\rm min}$ and on L: $f_{E_{\rm FIP},L}^{l_{\rm min}} \approx f_{E_{\rm FIP}}$. Indeed, independently of $l_{\rm min}$, the detector covers the far-forward domain, which determines the high-energy tail of the distribution function. As a result, the only impact of decreasing $l_{\rm min}$ comes from decreasing $L_{\rm min} \approx l_{\rm min}$. Namely, the whole integral is saturated around its value $L_{\rm min}$:

$$\epsilon \approx \int dE_{\rm FIP} \exp\left[-\frac{l_{\rm min}}{c\tau_{\rm FIP}\gamma_{\rm FIP}}\right] f_{E_{\rm FIP}} \quad (9)$$

Up to logarithmic corrections, at a fixed mass, the upper bound of the sensitivity, i.e., the smallest lifetimes that may be probed, scales as $\tau_{\rm FIP}^{\rm upper} \propto l_{\rm min}^{-1}$.

Let us now consider the off-axis case, concentrating on the case where the whole detector lies off-axis. For the detector dimensions in Table 1, this case would correspond to the transverse displacement of the centre of the detector $r_{\rm displ} > 3$ m. In this case, the situation is less trivial. First, the FIP energy spectrum becomes softer (remind Fig. 4) at large polar angles. As a result, the value of the $E_{\rm FIP}$ integral in Eq. (8) at fixed L decreases compared to the off-axis case. Second, the FIPs pointing to the closest (to the beamline) part of the detector enter the decay volume from the side (Fig. 2). This also decreases $f_{E_{\rm FIP},L}^{l_{\rm min}}$. This behavior is in tension with the exponent in Eq. (8): small values of L close to $L_{\rm min}$ correspond to maximal polar angle and hence smaller energies/acceptance. This destructive interplay destroys the positive impact of decreasing $l_{\rm min}$ on the potential to probe short FIP lifetimes. As a result, the fully off-axis configurations have a lower number of signal events than the on-axis, even given much smaller $l_{\rm min}$.

To illustrate these qualitative arguments, let us consider three setups: the on-axis experiment from Table 1, the same on-axis experiment but with $l_{\rm min} = 10$ m, and finally an off-axis experiment with $l_{\rm min} = 10$ m and $r_{\rm displ} = 3$ m. For the FIPs, we will consider dark scalars with mass $m_S = 3$ GeV and lifetime $c\tau_S = 5$ cm, which corresponds to the upper bound of the sensitivity.

The behavior of the integrand of (8),

$$\frac{d\epsilon}{dE_{\rm FIP}} = \int dL \frac{1}{c\tau_{\rm FIP}\gamma_{\rm FIP}} \exp\left[-\frac{L}{c\tau_{\rm FIP}\gamma_{\rm FIP}}\right] f_{E_{\rm FIP},L}^{l_{\rm min}}, \quad (10)$$

for the three setups is shown in Fig. 7. Obviously, the on-axis setup with $l_{\rm min} = 10$ m has the largest flux of FIPs. However, despite the fact that the off-axis decay volume is located ~ 4 times closer to the beam dump than the on-axis setup with $l_{\rm min} = 38$ m, the off-axis setup has a $\simeq 30$ times lower value of the total integral (8), following from the factors described above.

5 Comparison between the experiment proposals under study at the CERN SPS

In this section, we make a comparison of the physics yields between the three experiment proposals that are currently being considered for implementation in the ECN3 beam facility at CERN's SPS accelerator, HIKE [13], SHADOWS [14], SHiP [12]. All proposals are based on a similar detector setup

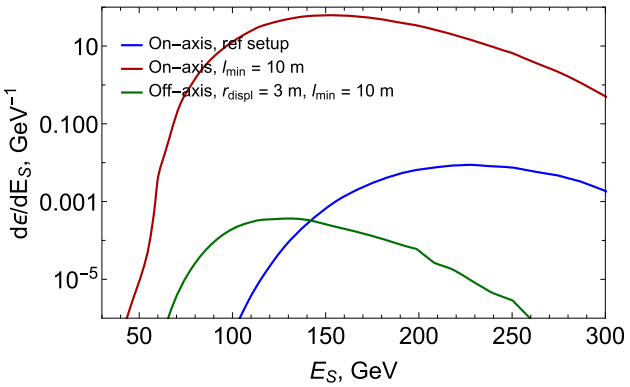

Fig. 7 The energy distribution of the function (10) determining the number of events at the upper bound of the sensitivity ($c\tau_{\rm FIP}\langle\gamma_{\rm FIP}\rangle \lesssim l_{\rm min}$, where $l_{\rm min}$ is the distance from target to the decay volume). Its shape and normalization depend on the value of the distance from the target to the decay volume $l_{\rm min}$ and the high-energy tail of FIPs within the acceptance for the given experiment. If decreasing $l_{\rm min}$ for the on-axis placement of the experiment, the high-energy tail would remain unchanged; as a result, we would increase the event rate. This may not be the case for an off-axis placement: the energy spectrum may become much softer and compensate for the decrease of $l_{\rm min}$ (see text for details). To illustrate these points, we consider a dark scalar with mass $m_S = 3$ GeV and lifetime $c\tau_S = 0.05$ m at three various experimental setups at SPS: the configuration from Table 1; the same configuration but with $l_{\rm min} = 10$ m; and the off-axis experiment with $l_{\rm min} = 10$ m, the displacement of the lower edge of its decay volume from the beamline of 1 m, and the decay volume length $l_{\rm fid} = 20$ m. The number of events at the closer on-axis experiment is larger, while at the off-axis experiment, it is smaller since scalars have smaller energies

in that they conceptually consist of large decay volumes followed by spectrometers and particle identification, together with various veto systems. HIKE is primarily a kaon experiment located on-axis. It requires a specialised beam setup with a kaon target, a secondary beam line for the kaon selection, and an absorber of copper/iron for the remaining proton beam and secondary hadrons from the kaon target. HIKE's distance to the kaon target is defined by the optimisation for the kaon physics programme, resulting in a relatively large distance between HIKE's decay volume and the absorber. The kaon physics optimisation imposes limitations on the maximum beam intensity that is due to both the secondary beam line setup and the detector, effectively four times lower than it would be for SHiP. HIKE is also proposed to partially operate in beam-dump mode for FIP physics. In this mode, the kaon target is moved aside to let the proton beam of the same intensity be directly dumped on the absorber. With the very small solid angle coverage, HIKE has mainly sensitivity to dark photons and ALPs with photon coupling. SHADOWS is an off-axis experiment that would be located alongside HIKE's kaon beam, downstream of the absorber, with the detector covering angles $\theta \gtrsim 30$ mrad (Fig. 4). SHADOWS' distance to the proton absorber is defined by the infrastructure around the absorber, shielding requirements, a muon sweeper, and the subsequent beam line elements. SHADOWS would operate together with HIKE in beam-dump mode, with SHADOWS searching for the FIPs flying off-axis and HIKE for those produced in the far-forward region. In this respect, it is expected that the beam time for HIKE and SHADOWS is split between periods of kaon physics and beam-dump physics. SHiP is instead a dedicated on-axis experiment with the detector located as close as possible to a compact target station housing a target of molybdenum/tungsten that is optimised for FIP physics. SHIP's distance to the target is defined by a hadron stopper with minimum depth, and a specialised magnetic muon deflector that sweeps the muon flux away from the fiducial volume. SHIP's location allows it to cover all the FIP production modes at the SPS.

In Fig. 8 we show the 90% CL sensitivities of HIKE$_{\rm dump}$ + SHADOWS and of SHIP to HNLs, dark scalars, dark photons, and ALPs with the coupling to photons. In addition, for SHiP, we include the iso-contours corresponding to $N_{\rm events} = 100$. Such a large number of events allows not only to establish the existence of a new particle but also identify its properties such as branching ratios of various decay channels, precise mass, etc. Details of the sensitivity estimates are described in Appendix A. We see that the lower bound of the sensitivity of SHADOWS+HIKE$_{\rm dump}$ is close to the 100 events line of SHiP. This may be easily understood if using Eq. (1). Namely, for SHADOWS, the product $N_{\rm PoT} \times l_{\rm fid}$ is 10 times smaller than at SHiP. The rest of the suppression in the number of events at SHADOWS comes from the off-axis placement parallel to the beamline, significantly decreasing $\epsilon_{\rm FIP} \times \langle p_{\rm FIP}^{-1}\rangle$ – by a factor $1/3-1/20$ depending on the FIP mass. For HIKE, the product $N_{\rm PoT} \times l_{\rm fid}$ is only twice smaller than at SHiP. However, its detector covers ~ 20 times smaller solid angle as seen from the target and the beginning of the decay volume, which results in a much lower overall acceptance.

Moreover, despite the fact that SHADOWS is placed much closer to the target, which could naively mean that it should be able to probe larger couplings of FIPs, Fig. 8 shows that this is not the case. One of the reasons for this is that FIPs flying off-axis have a much softer energy spectrum and hence shorter typical decay length (see Fig. 6 and discussions therein). An additional reason is that the most energetic FIPs enter the decay volume from the side, resulting in a shorter flight path within the decay volume and thus suppressing the fraction of their decays inside the decay volume.

6 Conclusions

Future searches for new physics beyond the Standard Model are without a doubt in need of a diverse approach and experiments with complementary sensitivities to different classes

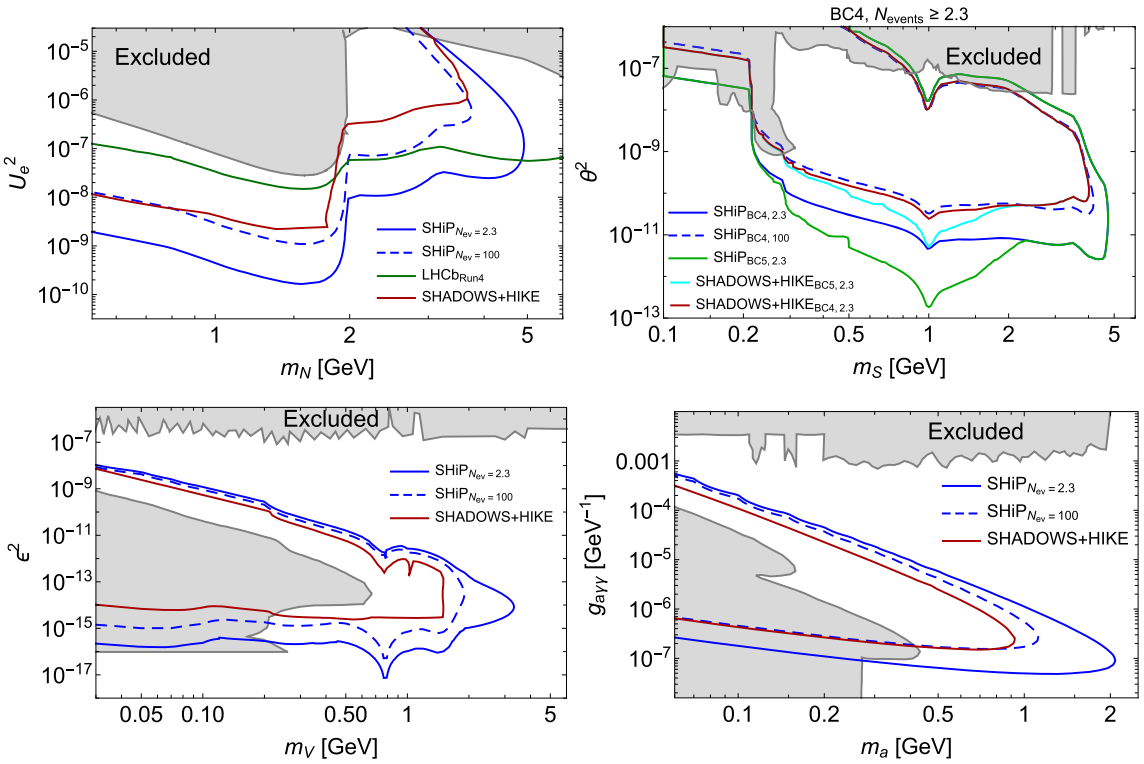

Fig. 8 Comparison of the potential of the HIKE, SHADOWS, and SHiP proposals to explore the parameter space of FIPs in beam-dump mode. The FIP physics is here illustrated with HNLs (top left panel), dark scalars with the mixing coupling and with the quartic coupling fixed by $\mathrm{Br}(h \to SS) = 10^{-3}$ (top right panel), dark photons (bottom left), and ALPs with the coupling to photons (bottom right). For all the models except for dark scalars, we use the 90% CL combined sensitivity of SHADOWS and $\mathrm{HIKE}_{\mathrm{dump}}$ from the LoIs [13,14], while for dark scalars we report our own estimates based on the SHADOWS configuration specified in the LoI (see Appendix A for detail). For SHiP, we show two curves: the 90% CL sensitivity, and the domain $N_{\mathrm{events}} > 100$, where it may be possible to determine properties of FIPs such as their mass and decay branching ratios. The combined impact of the lower number of protons on the target and the non-optimal placement for SHADOWS and HIKE lead to a significant limitation on their physics potential. Their exclusion domain lies within the FIP identification domain of SHiP. Moreover, in the case of FIPs produced by decays of B mesons (such as e.g. HNLs, dark scalars, and ALPs coupled to fermions), the reach of SHADOWS and HIKE may be overcome by future searches at LHCb [9] with triggers allowing the use of the muon stations as trackers. We show the sensitivity of such types of searches in the case of HNLs

of models, ranging both mass and coupling scales. A theoretically and experimentally attractive case for new physics that is largely unexplored but within reach at current facilities, consists of particles with masses below the electroweak scale that may be produced at accelerators in decays of SM particles. Their couplings should be small in order to avoid existing experimental constraints. Hence, they are often called feebly-interacting particles, or FIPs.

FIPs are now being actively searched for at the LHC and will be searched for at future collider experiments. However, this type of search has limitations in probing the parameter space of long-lived FIPs in the mass range $m_{\mathrm{FIP}} \lesssim 5-10$ GeV, see Fig. 1, mainly because of strong backgrounds and too short decay volumes as compared to the typical decay length of light FIPs. Beam dump-like experiments, characterized by the possibility of extremely high luminosity at relatively high energies, and effective coverage of the production and decay acceptance, are the perfect setup to generically explore the "coupling frontier" and the case of FIPs with mass below the B meson mass. Beam dump experiments can be equipped with long decay volumes and be located at some distance from the target to accommodate absorbers and deflectors of SM particles, as well as veto systems, to reduce the backgrounds. The best placement for such a FIP facility is the SPS accelerator at CERN, where the existing infrastructure and the currently available proton yield of up to 4×10^{19} protons per year at 400 GeV make it possible to implement and operate a-state-of-the-art experiment at relatively low costs. There are currently three experiment proposals being considered for implementation in the ECN3 beam facility at CERN's SPS accelerator, HIKE [13], SHADOWS [14], SHiP [12]. Their respective objectives and layouts are briefly summarised in Sect. 5.

In order to determine the optimal experimental geometry, we have made an in-depth study of the dependence of the number of FIP signal events at the lower (Sects. 4.1, 4.2) and

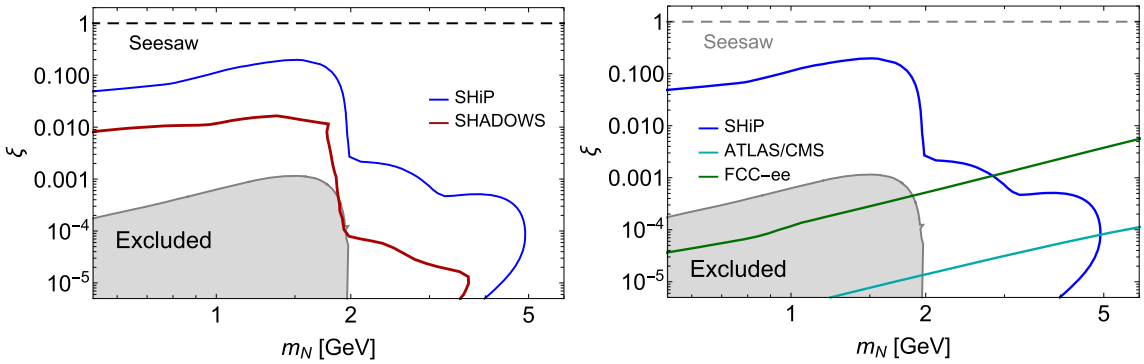

Fig. 9 The left panel: the HNL parameter space reach of SHiP and SHADOWS in the plane $\xi = U^2_{\text{seesaw}}/U^2$-HNL mass. The departure of ξ from the seesaw line $\xi = 1$ shows the scale of fine-tuning needed to explain neutrino masses by two HNLs with large couplings. The right panel: the same figure but with the reach of colliders included

the upper (Sect. 4.3) bounds of the sensitivity, as a function of the length of the decay volume, the distance from the target, and the transverse displacement with respect to the beamline, and other parameters. We have performed the analysis for several classes of models of FIPs ("portals") with different production channels and decay modes, see Sect. 3. Given that all the proposals claim to reach a background-free regime, we do not consider backgrounds specifically other than the constraints that such assumptions impose on the geometric parameters.

In particular, we analyzed the effect of displacing the detector off-axis for different portals. Generically, it leads to a decrease in the number of events for two specific reasons: first, the angular flux of FIPs decreases at large angles (Fig. 4, and second, it causes geometric shortening of the effective length of the decay volume along trajectories pointing to the detector (Fig. 2). The impact of the off-axis placement depends on the dominant FIP production channel, i.e., by the decays of mesons (such as for HNLs, ALPs, and dark scalars), or directly in proton-target collisions (dark photons). In the former case, shifting the decay volume entirely to the side of the beam axis such that it covers polar angles $\theta \gtrsim 10$ mrad, leads to a loss of up to a factor of five in the number of events at the lower bound, depending on the FIP mass, as compared to the reference configuration in Table 1. In the latter case, the same decrease in yield is seen as soon as the side of the decay volume wall gets off-axis. Practically, this means that the off-axis configurations have no sensitivity to dark photons.

In contrast, we have seen that changing the distance from the target to the decay volume, and its length, for the on-axis configuration does not affect the number of events at the lower bound by more than a factor of two (Fig. 5) over a broad range of parameters. On the other hand, significantly decreasing the distance to the target affects the complexity of the experiment and background suppression, and hence the cost. Therefore, optimizing the on-axis configuration is rather a subject of minimizing background and cost.

Finally, we have applied the analysis to the three ECN3 proposals (Sect. 5). In Fig. 8, we show the 90% CL sensitivities of $\text{HIKE}_{\text{dump}}$ + SHADOWS and of SHIP to HNLs, dark scalars, dark photons, and ALPs with the photon coupling. We also include the projection of the sensitivity of the LHC experiments, taking LHCb as an example, for the case that new triggers will be developed, allowing, e.g., the use of the muon chambers as trackers for FIPs independently of their production channel [9]. In this case, the sensitivity of SHADOWS+HIKE may be limited in the domain of heavy FIPs with $m \gtrsim 2$ GeV.

For SHiP, we also include the iso-contours corresponding to $N_{\text{events}} = 100$ in Fig. 8. Such a large number of events allows not only to establish the existence of a new particle but also identify its properties, such as branching ratios of various decay channels, precise mass, etc. We conclude that the combination of lower beam intensity and non-optimal geometric placement for HIKE and SHADOWS worsens their potential to explore the "coupling frontier" compared to SHiP. In particular, assuming that all the experiments operate in background-free regime, in the domain of small couplings, where SHADOWS and HIKE may detect 1 event at 90% CL, SHiP would be able to reconstruct their parameters, such as mass, spin, and the probability of various decays.

To stress the importance of the maximization of the FIP event yield, let us look closer at the model of HNLs. The main motivation for the HNLs is that they provide a very simple explanation of the observed active neutrino masses, in a way very similar to other fermions of the SM, i.e., by the mixing of left-handed and right-handed states via a Dirac mass term. For this, however, (i) at least two HNLs are required in order to explain to observed neutrino mass differences $\Delta m^2_{\text{solar}}$, Δm^2_{atm}; (ii) their coupling constants should be at the level presented in Fig. 1 as the "sea-saw line". As we see,

such couplings are very far from what will be accessed by the current experiments. It is, however, possible that the couplings of each of the HNLs are orders of magnitude larger, but that their contributions to the active neutrino masses cancel with high precision, due to a fine-tuning or a symmetry. The fine-tuning scale associated with this cancellation is $\xi = U^2_{\text{seesaw}}/U^2$, where $U^2_{\text{seesaw}} \sim 5 \times 10^{-11}(1\,\text{GeV}/m_N)$ is the see-saw line which does not require the fine-tuning. As shown in Fig. 9, beam dump experiments allow probing the part of the parameter space that is unreachable to collider experiments. In addition, we find that the layout of the SHiP experiment at the SPS ECN3 beam facility is close to optimal, assuming that the background is negligibly small.

Acknowledgements AB is supported by the European Research Council (ERC) Advanced Grant "NuBSM" (694896). KB is partly funded by the INFN PD51 INDARK grant. OM is supported by the NWO Physics Vrij Programme "The Hidden Universe of Weakly Interacting Particles" with project number 680.92.18.03 (NWO Vrije Programma), which is (partly) financed by the Dutch Research Council (NWO). MO received support from the European Union's Horizon 2020 research and innovation program under the Marie Sklodowska-Curie grant agreement No. 860881-HIDDeN.

Data Availability Statement This manuscript has no associated data or the data will not be deposited. [Authors' comment: The data for this study has been generated using the public code SensCalc accessible to anyone].

Funded by SCOAP[3]. SCOAP[3] supports the goals of the International Year of Basic Sciences for Sustainable Development.

A. Technical details of the sensitivity estimates

A.1. The number of events

The expression for the number of events that we use looks as follows:

$$N_{\text{events}} = \sum_{\text{mother}} N_{X,\text{prod}} \times \int dE_X \int_{\text{decay volume}} d\theta_X dz_X \times f_X(\theta_X, E_X) \times \epsilon_{\text{az}}(\theta_X, z_X) \times \frac{dP_{\text{dec}}}{dz_X} \times \epsilon_{\text{dec}}(\theta_X, E_X, z_X) \times \text{Br}_{\text{vis}}(m_X) \quad (11)$$

Here, $N_{X,\text{prod}} = N_{\text{mother}} \times P_{\text{mother}\to X}$ is the total number of particles X produced by a mother particle, with $P_{\text{mother}\to X}$ being the probability of production. z_X, θ_X, E_X are, correspondingly, the longitudinal position, polar angle, and energy of the decaying FIP of type X, $f_X(\theta_X, E_X)$ is the differential distribution of FIPs in the polar angle and energy, and $\epsilon_{\text{az}}(\theta_X, z_X)$ is the azimuthal acceptance:

$$\epsilon_{\text{az}} = \frac{1}{2\pi}\min\big[\Delta\phi_{\text{decay volume}}(\theta_X, z_X), \Delta\phi_{\text{end of spectrometer}}(\theta_X))\big] \quad (12)$$

i.e. fraction of azimuthal angle for which FIPs (1) decays inside the decay volume and (2) have a trajectory intersecting the whole spectrometer. The differential decay probability $\frac{dP_{\text{dec}}}{dz_X}$ is:

$$\frac{dP_{\text{dec}}}{dz_X} = \frac{\exp[-l(z,\theta)/l_{\text{X,decay}}]}{l_{\text{X,decay}}} \frac{dl(z_X, \theta_X)}{dz_X}, \quad (13)$$

with $l = z_X/\cos(\theta_X)$ being the displacement. The decay product acceptance $\epsilon_{\text{dec}}(\theta_X, E_X, z_X)$ is the fraction of FIPs within the azimuthal acceptance whose decay products point to the detector. Finally, $\text{Br}_{\text{vis}}(m_X)$ is the branching ratio of the FIP decay into final states visible at the given experiment.

Further details may be found in the works [28–30], where the approach has been used for various experiments located at different facilities. In addition, the publicly available code `SensCalc`, which allows computing sensitivities using this approach, will be provided in a paper [20] soon to appear.

A.1.1. Lower bound of the sensitivity

In the regime of long FIP lifetimes, the exponent in the expression for the decay probability is

$$\exp[-l/l_{\text{X,decay}}] \approx 1, \quad (14)$$

and the number of events becomes

$$N_{\text{events}} \approx N_{X,\text{prod}} \times \frac{l_{\text{fid}}}{c\tau_X \langle E_X^{-1}\rangle^{-1} m_X} \times \epsilon, \quad (15)$$

where ϵ is the overall acceptance:

$$\epsilon \equiv \frac{1}{l_{\text{fid}}} \int dE_X \int_{\text{decay volume}} d\theta_X dz_X \times f_X(\theta_X, E_X) \times \frac{\langle E_X^{-1}\rangle^{-1}}{E_X} \times \epsilon_{\text{dec}}(\theta_X, E_X, z_X) \times \epsilon_{\text{az}}(\theta_X, z_X) \quad (16)$$

A.2. Sensitivity to dark scalars

In Fig. 8, we report the 90% CL sensitivity of SHADOWS to dark scalars. The sensitivity has been obtained assuming the setup described in the SHADOWS LoI [14] and requiring $N_{\text{events}} > 2.3$.

Let us discuss the setup in detail. The distance of the decay volume from the target is considered to be $l_{\rm min} = 10$ m,[5] and the decay volume is defined to be located entirely off-axis with the closest decay volume wall being parallel to the beam-axis at a distance of 1 m. We consider the given box geometry of the decay volume with length 20 m and the cross-section equal to 2.5×2.5 m^2, and use the detector cross-section equal to 2.5×2.5 m^2. In addition, we require at least two charged decay products that point to the beginning of the dipole magnet, located $\simeq 2.5$ m downward the end of the decay volume. We note that the last specification may significantly overestimate the sensitivity of SHADOWS as the LoI specifies that the spectrometer and its downstream particle identification system, consisting of a calorimeter and muon system, have a total length of 12–14 m. For proper identification of the decay products and reconstruction of their kinematics, their trajectories have to be contained within the acceptance of the entire detector, also requiring that the spectrometer magnet field is accounted for.

As for the scalar phenomenology, we follow [26]. For the probability of producing scalars, we use

$$\mathrm{P}(B \to S) \approx \sum_{B_x = B^{0,+}} f_{b\to B_x} \times \max\left[\mathrm{Br}_{\rm incl}(B \to S), \mathrm{Br}_{\rm excl}(B \to S + X_{d/s})\right] \quad (17)$$

Here, $f_{b\to B_x}$ is the fragmentation fraction ($f_{b\to B^+} \approx f_{b\to B^0} \approx 0.411$). $\mathrm{Br}_{\rm incl}(B \to S)$ is the inclusive branching ratio, which is equal to ≈ 6.6 in the domain $m_S \ll m_b$. $\mathrm{Br}_{\rm excl}$ is the sum over all exclusive channels considered in [26], and $X_{d/s}$ is the final state appearing from the hadronization of d or s quarks. The inclusive description breaks down at large scalar masses [26]; in addition, PBC recommended using the exclusive description [1]. Nevertheless, in [14], a purely inclusive estimate has been used. As a result, the branching ratio (17) matches the value used in [14] for scalars with masses $m_S \lesssim 2$ GeV, but is larger for larger masses.

For the distribution of mother particles, B mesons, we use the angle-energy distribution and the production probability $\chi_{b\bar{b}} = 2.7 \times 10^{-7}$ from [18], which also includes the cascade enhancement from the amount of Bs produced in secondary interactions in the thick target.

Finally, we assume that the detectable decays of scalars are those containing at least two "visible" particles (photons or charged particles). In this way, we allow not only fully reconstructable states (where the kinematics may be fully reconstructed) but also partially reconstructed decays, containing particles such as neutrinos among the other decay products.

[5] Note that $l_{\rm min}$ reported in the LoI in the section about experiment specifications is different, being 14 m. The reason for this discrepancy is unclear to us.

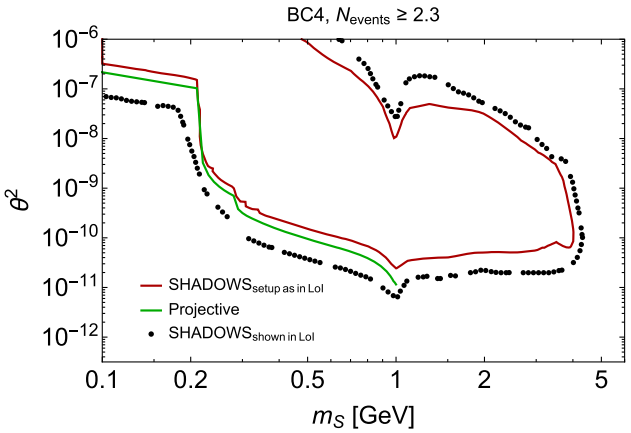

Fig. 10 Sensitivity of SHADOWS to dark scalars with the mixing coupling. The red line shows the sensitivity obtained from our computation (Eq. (11)) for the experimental setup used in the SHADOWS LoI [14]. The black dots show the sensitivity reported in the LoI. Finally, the green line shows the sensitivity of the projective setup where scalars are required to point to the SHADOWS detector acceptance window and with the decay products acceptance set to 1. The sensitivity of SHADOWS cannot be better than the sensitivity of the projective setup (see text for details)

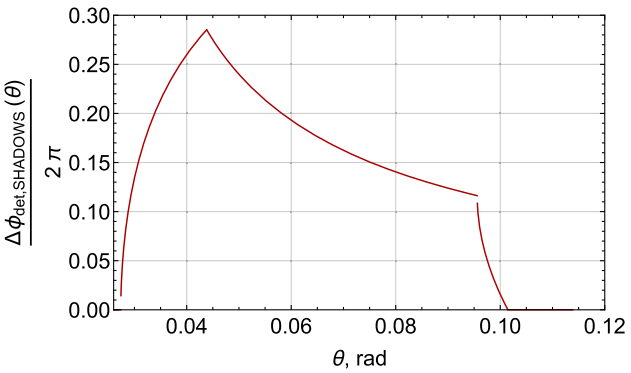

Fig. 11 The fraction of the azimuthal angle covered by the SHADOWS detector window located 2.5 m downwards the end of the decay volume. The solid angle covered by the detector, $\Omega_{\rm det} = \int \sin(\theta)\Delta\phi_{\rm det,SHADOWS} d\theta$, matches the simple formula $S_{\rm det}/L^2_{\rm to\ det} \approx 4.6 \times 10^{-3}$ sr

Examples of the latter are decays $S \to \tau\bar{\tau} \to \nu_\tau \bar{\nu}_\tau +$visible, $S \to D\bar{D} \to \nu + \bar{\nu} +$ visible, which dominate the scalar decays above di-τ mass threshold. In contrast, in the LoI, only fully reconstructable states have been required. We do not require any other cut, such as the minimum opening angle or the energy cuts.

Therefore, given the discussion above, our sensitivity estimate must match the sensitivity from the LoI in the mass range $m_S \lesssim 2$ GeV. In contrast, for larger masses, it should be more optimistic.

The comparison of the sensitivities of SHADOWS obtained by our calculations and by the calculations in the LoI is shown in Fig. 10.

Table 3 The values of the parameters entering Eq. (18) for the particular scalar mass $m_S = 150$ MeV

Parameter	$\mathrm{Br}(B \to S)$	ϵ_S	$m_S\langle p_S^{-1}\rangle$	$c\tau_S$, m	N_{events}
Value	≈ 6.6	0.09	3.7×10^{-3}	$7.7 \times 10^{-3}/\theta^2$	$2.3\left(\frac{\theta^2}{1.4\times 10^{-7}}\right)^2$

We notice a significant difference at the lower and upper bounds of the sensitivity. Namely, at small masses, the number of events at the lower bound reported by the LoI is larger by up to a factor of 15 compared to our predictions. The situation is similar for large masses, where the sensitivity reported in the LoI should be much more conservative.

While we cannot identify the reason for such discrepancy, we can present a simple qualitative argument suggesting that the sensitivity presented in the LoI could be overestimated. Namely, let us consider an experiment with the projective decay volume along the lines pointing from the target to the SHADOWS detector window. For the latter, we choose the same plane with cross-section $2.5 \times 2.5\ \mathrm{m}^2$ located 2.5 m downward the end of the decay volume.

It would not be possible to build such an experiment because of the limitations imposed by the infrastructure. However, what is important for us is that it must have better sensitivity than the real SHADOWS configuration. Indeed, its detector covers the same polar angles as the detector of SHADOWS, but, there would be no geometric shortening of the length of the decay volume for the scalars pointing to the detector, and there would be less azimuthal suppression.

Further, let us consider the scalar mass $m_S \lesssim 1$ GeV, for which the decay products fly roughly in the same direction as the decaying scalar, and we may approximate the decay products acceptance by its upper bound $\epsilon_{\text{dec}} \approx 1$. Finally, we may consider the lower bound of the sensitivity only. Under these assumptions, the number of decays is approximately

$$N_{\text{events}} \approx N_{\text{PoT}} \times \chi_{b\bar{b}} \times 2(f_{b\to B^+} + f_{b\to B^0}) \times \mathrm{Br}_{B\to S} \times \epsilon_{\text{geom}} \times \frac{l_{\text{fid}} \cdot m_S\langle p_S^{-1}\rangle}{c\tau_S}, \tag{18}$$

where ϵ_{geom} is the fraction of the scalars pointing to the detector, and $\langle p_S^{-1}\rangle$ is the mean inverse momentum of such scalars. These quantities may be obtained just from knowing the scalar distribution function, which may be easily derived given the simple kinematics of the scalar production in 2-body decays $B \to S + X_s$, and the azimuthal coverage of the SHADOWS detector window, which is shown in Fig. 11.

The values of all the relevant parameters for the particular point $m_S = 150$ MeV are reported in Table 3.

Comparing the lower bound predicted by this formula with the sensitivity from the SHADOWS LoI, we find that the sensitivity of the projective experiment lies *well above* the latter, see Fig. 10. In contrast, it lies below our estimate. At small masses, it is very close to our estimate. The increasing difference at large masses is explained by the decay products acceptance which is assumed to be 1 for the projective setup.

References

1. J. Beacham et al., Physics beyond colliders at CERN: beyond the Standard Model Working Group report. J. Phys. G **47**(1), 010501 (2020). arXiv:1901.09966 [hep-ex]
2. P. Agrawal et al., Feebly-interacting particles: FIPs 2020 workshop report. Eur. Phys. J. C **81**(11), 1015 (2021). arXiv:2102.12143 [hep-ph]
3. S. Alekhin et al., A facility to search for hidden particles at the CERN SPS: the SHiP physics case. Rep. Prog. Phys. **79**(12), 124201 (2016). arXiv:1504.04855 [hep-ph]
4. CMS Collaboration, A. Tumasyan et al., Search for long-lived heavy neutral leptons with displaced vertices in proton-proton collisions at $\sqrt{s}$ =13 TeV. JHEP **07**, 081 (2022). arXiv:2201.05578 [hep-ex]
5. I. Boiarska, K. Bondarenko, A. Boyarsky, S. Eijima, M. Ovchynnikov, O. Ruchayskiy, I. Timiryasov, Probing baryon asymmetry of the Universe at LHC and SHiP. arXiv:1902.04535 [hep-ph]
6. A. Blondel et al., Searches for long-lived particles at the future FCC-ee. Front. Phys. **10**, 967881 (2022). arXiv:2203.05502 [hep-ex]
7. A.M. Abdullahi et al., The present and future status of heavy neutral leptons. J. Phys. G **50**(2), 020501 (2023). arXiv:2203.08039 [hep-ph]
8. D. Acosta et al., Review of opportunities for new long-lived particle triggers in Run 3 of the large hadron collider. arXiv:2110.14675 [hep-ex]
9. L. Shchutska, M. Ovchynnikov, Prospects for LLP searches at the LHC in Run 3 and HL-LHC (2022)
10. T. Asaka, S. Blanchet, M. Shaposhnikov, The nuMSM, dark matter and neutrino masses. Phys. Lett. B **631**, 151–156 (2005). arXiv:hep-ph/0503065
11. LHCb Collaboration, R. Aaij et al., Search for long-lived scalar particles in $B^+ \to K^+\chi(\mu^+\mu^-)$ decays. Phys. Rev. D **95**(7), 071101 (2017). arXiv:1612.07818 [hep-ex]
12. SHiP Collaboration, O. Aberle et al., BDF/SHiP at the ECN3 high-intensity beam facility, Tech. Rep. CERN-SPSC-2022-032, SPSC-I-258 (CERN, Geneva, 2022). http://cds.cern.ch/record/2839677
13. HIKE Collaboration, E. Cortina Gil et al., HIKE, High Intensity Kaon Experiments at the CERN SPS: Letter of Intent, Tech. Rep. CERN-SPSC-2022-031, SPSC-I-257, SPSC-I-257 (CERN, Geneva, 2022). arXiv:2211.16586. https://cds.cern.ch/record/2839661. Letter of Intent submitted to CERN SPSC. Address all correspondence to hike-eb@cern.ch
14. SHADOWS Collaboration, M. Alviggi et al., SHADOWS Letter of Intent, Tech. Rep. CERN-SPSC-2022-030, SPSC-I-256 (CERN, Geneva, 2022). https://cds.cern.ch/record/2839484
15. SHiP Collaboration, M. Anelli et al., A facility to Search for Hidden Particles (SHiP) at the CERN SPS. arXiv:1504.04956 [physics.ins-det]
16. B. Döbrich, J. Jaeckel, T. Spadaro, Light in the beam dump—ALP production from decay photons in proton beam-dumps. JHEP **05**,

213 (2019). arXiv:1904.02091 [hep-ph]. [Erratum: JHEP 10, 046 (2020)]
17. J. Jerhot, B. Döbrich, F. Ertas, F. Kahlhoefer, T. Spadaro, ALPINIST: axion-like particles in numerous interactions simulated and tabulated. JHEP **07**(TTK-22-04, CP3-22-02), 094 (2022). arXiv:2201.05170 [hep-ph]
18. SHiP Collaboration, H. Dijkstra, T. Ruf, Heavy Flavour Cascade Production in a Beam Dump (2015). https://cds.cern.ch/record/2115534. CERN-SHiP-NOTE-2015-009
19. SHiP Collaboration, C. Ahdida et al., Sensitivity of the SHiP experiment to dark photons decaying to a pair of charged particles. Eur. Phys. J. C **81**(5), 451 (2021). arXiv:2011.05115 [hep-ex]
20. M. Ovchynnikov, J.-L. Tastet, O. Mikulenko, K. Bondarenko, Sensitivities to feebly interacting particles: public and unified calculations. arXiv:2305.13383 [hep-ph]
21. D. Gorbunov, A. Makarov, I. Timiryasov, Decaying light particles in the SHiP experiment: signal rate estimates for hidden photons. Phys. Rev. D **91**(3), 035027 (2015). arXiv:1411.4007 [hep-ph]
22. B. Döbrich, J. Jaeckel, F. Kahlhoefer, A. Ringwald, K. Schmidt-Hoberg, ALPtraum: ALP production in proton beam dump experiments. JHEP **02**, 018 (2016). arXiv:1512.03069 [hep-ph]
23. M. Adamovich et al., Results on charm physics from WA82. Nucl. Phys. B Proc. Suppl. **27**, 212–218 (1992)
24. M. Aguilar-Benitez et al., Inclusive particle production in 400-GeV/c p p interactions. Z. Phys. C **50**, 405–426 (1991)
25. B. Döbrich, F. Ertas, F. Kahlhoefer, T. Spadaro, Model-independent bounds on light pseudoscalars from rare B-meson decays. Phys. Lett. B **790**, 537–544 (2019). arXiv:1810.11336 [hep-ph]
26. I. Boiarska, K. Bondarenko, A. Boyarsky, V. Gorkavenko, M. Ovchynnikov, A. Sokolenko, Phenomenology of GeV-scale scalar portal. JHEP **11**, 162 (2019). arXiv:1904.10447 [hep-ph]
27. D. Gorbunov, E. Kriukova, O. Teryaev, Scalar decay into pions via Higgs portal. arXiv:2303.12847 [hep-ph]
28. I. Boiarska, A. Boyarsky, O. Mikulenko, M. Ovchynnikov, Constraints from the CHARM experiment on heavy neutral leptons with tau mixing. Phys. Rev. D **104**(9), 095019 (2021). arXiv:2107.14685 [hep-ph]
29. A. Boyarsky, O. Mikulenko, M. Ovchynnikov, L. Shchutska, Exploring the potential of FCC-hh to search for particles from B mesons. JHEP **01**, 042 (2023). arXiv:2204.01622 [hep-ph]
30. M. Ovchynnikov, V. Kryshtal, K. Bondarenko, Sensitivity of the FACET experiment to heavy neutral leptons and dark scalars. JHEP **02**, 056 (2023). arXiv:2209.14870 [hep-ph]