

Leveraging the ALICE/L3 cavern for long-lived particle searches

Vladimir V. Gligorov,¹ Simon Knapen,^{2,3,4} Benjamin Nachman,² Michele Papucci,^{2,3,5} and Dean J. Robinson^{2,6}

¹*LPNHE, Sorbonne Université et Université Paris Diderot, CNRS/IN2P3, Paris, France*

²*Ernest Orlando Lawrence Berkeley National Laboratory, University of California, Berkeley, California 94720, USA*

³*Department of Physics, University of California, Berkeley, California 94720, USA*

⁴*School of Natural Sciences, Institute for Advanced Study, Princeton, New Jersey 08540, USA*

⁵*Theoretical Physics Department, CERN, CH-1211 Geneva 23, Switzerland*

⁶*Santa Cruz Institute for Particle Physics and Department of Physics, University of California at Santa Cruz, Santa Cruz, California 95064, USA*



(Received 12 November 2018; published 16 January 2019)

Run 5 of the high-luminosity LHC era (and beyond) may provide new opportunities to search for physics beyond the standard model at interaction point 2. In particular, taking advantage of the existing ALICE detector and infrastructure provides an opportunity to search for displaced decays of beyond standard model long-lived particles. While this proposal may well be preempted by ongoing ALICE physics goals, examination of its potential new physics reach provides a compelling comparison with respect to other long-lived particle proposals. In particular, full event reconstruction and particle identification could be possible by making use of the existing L3 magnet and ALICE time projection chamber. For several well-motivated portals, the reach competes with or exceeds the sensitivity of MATHUSLA and SHiP, provided that a total integrated luminosity of approximately 100 fb^{-1} could be delivered to interaction point 2.

DOI: [10.1103/PhysRevD.99.015023](https://doi.org/10.1103/PhysRevD.99.015023)

I. INTRODUCTION

Interaction point 2 (IP2) at the LHC accelerator complex is currently used by the ALICE experiment [1] for the study of the quark-gluon plasma at high temperatures (examples of high temperature QCD discoveries achieved by ALICE can be found in, e.g., Refs. [2,3]). The ALICE Collaboration has firm plans to upgrade its detector and continue running throughout Run 3 and part of Run 4 [4]. However, should the heavy ion program conclude after Run 4 and with the long-term future of the CERN accelerator program now taking shape, it would be remiss not to consider possible new opportunities at IP2 during Run 5 and beyond.

The ALICE experiment comprises in part a gas time projection chamber (TPC) detector housed within the L3 electromagnet [5] and is designed to reconstruct very high multiplicities of tracks from ultrarelativistic ion-ion collisions. The L3 magnet has an interior cylindrical volume of length 12 m and radius 5.9 m and a central field of 0.5 T; the existing ALICE TPC [6] has radius 0.85 m to 2.5 m and a length of 5 m along the beam axis. The combination of a high resolution tracker—the ALICE TPC and/or a larger

one in this volume—and the 0.5 T magnetic field would allow for both particle identification and momentum measurement, which would be tremendously advantageous for establishing an exotic particle discovery. In this study, we investigate the physics reach of a dedicated detector for the decay in flight of long-lived particles in this space. We refer to this hypothetical experiment as A Laboratory for Long-Lived eXotics (AL3X) (pronounced “alex”).

In many well-motivated theoretical frameworks, long lived particles (LLPs) may provide the vestigial signatures through which beyond the Standard Model (SM) physics may be first discovered, in particular through exotic decays of the Higgs boson. Examples include theories of naturalness, extended Higgs sectors, dark matter, baryogenesis, or flavor (see, e.g., Ref. [7] and references therein.). Despite LLPs not being a major design driver for the ATLAS and CMS detectors, they have nevertheless achieved remarkable sensitivity (see, e.g., Refs. [8–11]). On the other hand, there are still important blind spots, some of which can be addressed by LHCb (see, e.g., Refs. [12–15]) or by beam dump experiments such as NA62 [16]. A comprehensive LLP program must, however, have good sensitivity to LLPs produced in Higgs decays, something which is notoriously challenging for all of the above experiments.

The lack of robust coverage for high lifetime LLPs with masses below the weak scale has inspired a number of proposals for dedicated experiments at CERN. The most ambitious along these lines are SHiP [17] and

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

MATHUSLA [18]. SHiP would be a dedicated beam dump experiment at the Super Proton Synchrotron (SPS), roughly ~ 150 m in length, while MATHUSLA would be a detector on the surface above ATLAS or CMS (geometries of 200×200 m, 100×100 m, and 50×50 m are being considered [19]). Other proposals take a more opportunistic approach by trading sensitivity for a smaller size and the advantage of being embedded into existing infrastructure. Following this philosophy, MilliQan [20,21] aims to search for millicharged particles in a drainage tunnel above CMS; CODEX-b [22] proposes to make use of the soon-to-be vacated data acquisition space next to LHCb; and FASER [23–26] would consist of a small detector volume in a service tunnel in the far forward regime of the ATLAS interaction point.

In this paper, we consider a LLP detector constructed inside the L3 magnet that is screened from SM backgrounds by heavily shielding the interaction point, located outside the magnet, as shown in Fig. 1. The proximity of the proposed detector to a LHC interaction point, with a considerable geometric acceptance, permits sensitivity not only to LLPs generated by high center-of-mass energy portals such as the Higgs invisible width but also from low scale vector, scalar, or fermion mixing portals, thereby covering all possible renormalizable couplings of the SM to exotic sectors in one detector concept. In this proof-of-concept study, we examine the AL3X reach for a LLP produced in an exotic Higgs or B decay as well as for the production of a kinetically mixed dark photon. For an integrated luminosity of order 100 fb^{-1} , we find that the AL3X reach meets, exceeds, or complements the combined reach of other LLP proposals. Much of our discussion will be informed by those applicable to the MATHUSLA [18] and CODEX-b [22] proposals, though the challenges from

backgrounds will be significantly different from the former and somewhat different from the latter.

II. UPGRADING IP2

Before further motivating and elaborating on the detector concept, we discuss up front some of the potential challenges as they relate to delivering $\mathcal{O}(100) \text{ fb}^{-1}$ luminosity to IP2 in the AL3X configuration. There are at least four main concerns: (i) moving the IP, (ii) beam quality, (iii) luminosity sharing, and (iv) cost.

For LHC collisions at 40 MHz, an IP can only be moved by multiples of $12.5 \text{ ns} \times c \simeq 3.75$ m. For this reason, we envision moving the IP by 11.25 m from its current location at the center of the magnet, which should provide sufficient room for shielding the detector from the IP. However, moving this distance with a reasonably low β^* would require changing the layout of the quadrupole magnets in addition to general modifications to the optics. The fact that the injection of one of the proton beams is located near IP2 is a possible additional complication. The current luminosity delivered to IP2 is also so low that it has little impact on the beam quality and lifetime. Increasing the instantaneous luminosity to be a non-negligible fraction of the ATLAS and CMS collision rate would make beam preservation more challenging. Another consequence of the higher luminosity is that the magnets (triplets and probably also the beam separator magnets) would need additional shielding from forward-going radiation. In addition to adding absorbers in front of the magnets, one may need to cool the absorbers. The final concern is the cost. It is too early to give a reliable price tag of configuring IP2 for AL3X, but given the known feasibility for a similar upgrade at IP8 for LHCb, the cost may not be prohibitive

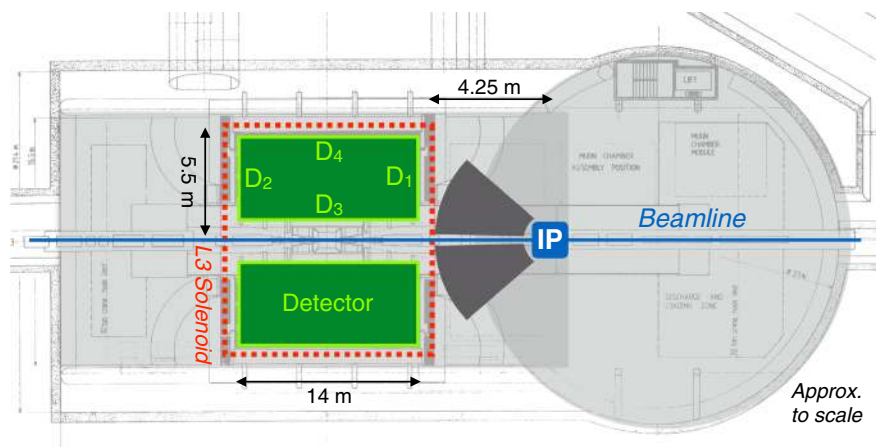


FIG. 1. Schematic representation of the proposed detector layout. Cavern layout information is from Refs. [1,27]; the cavern diagram is reproduced from Ref. [27]. The current L3 magnet is shown in dashed red for reference. The four surfaces bounding the detector volume are labeled $D_{1\dots 4}$ (see Sec. III for details).

at this time, especially in the context of other ambitious proposals for LLP detectors at CERN.

At this stage, none of these issues appears to preclude an efficient, robust, and cost effective implementation of AL3X, but further engineering studies are required to establish a realistic configuration of the ALICE cavern and the surrounding LHC tunnel and beam line. We further emphasize that the 100 fb^{-1} target is somewhat arbitrary and is chosen to roughly balance the physics reach against the challenges mentioned above and anticipated limitations from backgrounds. To give the reader a sense of how the various projections scale with the luminosity, we will therefore also show 250 fb^{-1} projections. With the above caveats in mind, we now proceed to present the nominal detector concept.

III. DETECTOR CONCEPT

For LLPs with relatively long lifetimes, the reach of any particular detector scales with the angular coverage and the size of the detector. This is the main reason for the rather large size of the two proposed experiments with the highest sensitivity: SHiP and MATHUSLA. Since SHiP would operate in beam dump mode off the SPS, the LLPs are necessarily fairly boosted no matter the portal, requiring a long fiducial decay volume. MATHUSLA would be located $\sim 150 \text{ m}$ from the IP due to space restrictions and thus would require a very large detector volume to ensure good geometric coverage. Consequently, it is not feasible to instrument MATHUSLA with precision tracking or calorimetry; nor is a magnetic field possible. A similarly ambitious detector, like AL3X, installed near a LHC interaction point, could have the best of both worlds, that is, moderately boosted LLPs and access to high center-of-mass energy—e.g., for Higgs portal production—but with good geometric coverage in a relatively small fiducial volume. Such a more modest volume might be instrumented with a dedicated TPC and potentially a calorimeter and in the case of IP2 make use of an already existing magnet.¹

The configuration in Fig. 1 is informed by considerations of both signal acceptance as well as the need to control backgrounds, that is, to look for LLP decays in flight in a heavily shielded environment. The proposed 11.25 m shift of IP2 from the center of the L3 magnet provides 4.25 m of remaining space that can be used for shielding the IP: the L3 magnet half length comprises 6 m plus a set of 1 m thick solid iron doors that serve as a return yoke for the magnet. As we show in the next section, approximately 40λ of shielding suffices to suppress the primary hadron and lepton backgrounds to acceptable levels, where λ is a nuclear interaction length.

¹If IP2 continues to be used for heavy ion physics after Run 4, it may be that the old ALEPH (IP4) or OPAL (IP6) caverns could be used for a proposal similar to what is described in this paper.

Although the iron doors provide some shielding already, for the sake of simplicity, we will model the shield by 40λ of tungsten, corresponding to 4 m of material. We leave a further optimization of the shield configuration for future work.² To veto backgrounds from secondaries produced in the shield itself, an active shield veto is included, embedded inside the shield volume, discussed below.

The nominal detector geometry is a 12 m long cylinder, with inner radius 0.85 m and outer radius 5 m centered on the beam line, leaving $\sim 1 \text{ m}$ between the inner (outer) cylindrical detector surface and the beam line (L3 magnet). (This extra space is included to allow for support structures and trigger layers, as well as to mitigate some of the forward backgrounds.) The detector geometry corresponds to a pseudorapidity acceptance $0.9 \leq \eta \leq 3.7$.

With appropriate shielding of the IP, the occupancy of the detector is expected to be relatively low, even with 40 MHz collisions. A gas TPC could therefore be a plausible choice for the detector technology because of its excellent tracking resolution and the possibility of reusing the existing ALICE TPC. In our new physics (NP) sensitivity estimates below, we will consider the reach for the ALICE TPC as well as for a larger TPC filling the entire detector volume. In a realistic design, the size and shape of the volume needed to be instrumented can likely be optimized to an interpolation between these two configurations; we leave this for future studies.

A “time stamp” to enable calibration of the TPC drift time can be achieved by including a trigger layer on the outer surface (D_4) and back face (D_2) of the cylindrical detector volume, as shown by the light green strips in Fig. 1. This trigger layer could, e.g., be composed of a scintillator. The flux of charged tracks, mostly muons, originating from the beam line and the shield is expected to be large. In order to suppress the triggering rate to manageable levels for a TPC readout ($1\text{--}10 \text{ kHz}$), veto layers on the front (D_1) and inner surface (D_3) of the detector complement the outer trigger layers, as discussed below in Sec. IV.

Before elaborating on the background estimates of this hypothetical detector, it is worth briefly estimating its fiducial efficiency, as it compares to other proposals. Concretely, in the limit where $\beta\gamma c\tau$ is much larger than the distance of the detector from the IP, the probability for a particle to decay in the detector volume is approximately

²A more realistic and affordable configuration would make use of a tungsten and steel or lead hybrid shield; in addition to the shielding already provided by the 1 m (6λ) thick iron doors, one could consider 2.5 m (25λ) of tungsten next to the IP, followed by 1.5 m (9λ) of additional steel or lead. One could also move the IP a farther 3.75 m away at a mild cost in geometric acceptance, providing enough space for a solely lead or steel shield.

$$\epsilon_{\text{fid}} \simeq \frac{\Delta\phi}{2\pi} \int_{\eta_0}^{\eta_1} d\eta d\gamma f(\eta, \gamma) \frac{\ell}{\beta\gamma c\tau} \quad (1)$$

with $\Delta\phi$ the azimuthal angular coverage of the detector, η_0 (η_1) the lower (upper) end of the pseudorapidity coverage of the detector, $f(\eta, \gamma)$ the distribution of the signal as a

function of boost γ and pseudorapidity η , and ℓ the typical path length of the LLP in the detector. For an example, we compute $f(\eta, \gamma)$ with PYTHIA8.2 [28,29] for a LLP with mass 1 GeV produced in an exotic Higgs decay (see Sec. V), such that $\gamma \sim \mathcal{O}(100)$. This results in the following comparison:

$$\begin{aligned} \text{AL3X: } 0.9 < \eta < 3.7; & \quad \frac{\Delta\phi}{2\pi} = 1; \quad \epsilon_{\text{fid}} = \frac{3.2 \times 10^{-2}}{c\tau/m}, \\ \text{MATHUSLA: } 0.9 < \eta < 1.4; & \quad \frac{\Delta\phi}{2\pi} = 0.15; \quad \epsilon_{\text{fid}} = \frac{6.9 \times 10^{-3}}{c\tau/m}, \\ \text{CODEX-b: } 0.2 < \eta < 0.6; & \quad \frac{\Delta\phi}{2\pi} = 0.06; \quad \epsilon_{\text{fid}} = \frac{1.1 \times 10^{-3}}{c\tau/m}. \end{aligned}$$

Being the closest the IP, AL3X has the largest angular coverage of all proposals, though the typical path length of a LLP is a factor of ~ 2 less than in MATHUSLA (~ 12 vs ~ 25 m). Since AL3X is more forward than MATHUSLA and CODEX-b, the LLPs tend to be a bit more boosted as well. In the long-lifetime limit, AL3X ends up having somewhat larger fiducial efficiency than MATHUSLA, for which we have assumed the $100 \text{ m} \times 100 \text{ m}$ configuration [19]. The efficiency for MATHUSLA's $200 \times 200 \text{ m}$ configuration is roughly a factor of 3 larger and similar to that of AL3X. The relative sensitivity between both detector concepts will therefore largely be driven by the luminosity that could be delivered to IP2 in Run 5.

In the short life-time regime, this scaling does not apply, and instead the main driver of the sensitivity is the distance of the IP to the detector, in the frame of the LLP. As we will see in Sec. V C, for a kinetically mixed dark photon, an experiment like AL3X could be competitive with SeaQuest, FASER, and SHiP in this regime, despite the lower number of collisions. The reason is the relatively short baseline of AL3X, as compared to FASER, and its access to very highly boosted LLPs, as compared to SeaQuest and SHiP.

IV. BACKGROUNDS AND SHIELDING

There are two classes of backgrounds for AL3X: those which are attenuated by the shield and those that are not. We discuss both in turn, as well as the necessary specifications for the shield to achieve the desired low background regime. The shielding analysis is similar in spirit to the CODEX-b proposal [22], though the background analysis and shield design for AL3X is complicated by the detector surrounding the beam line itself.

A. Shield configuration

The proof-of-concept background shield configuration is taken to be a 40λ (4 m) spherical shell segment of tungsten, centered on the IP with an inner and outer radius of 0.25

and 4.25 m, respectively, covering a pseudorapidity range $0.8 < \eta < 4.3$, as shown in Fig. 2. As for the CODEX-b shield, the prompt primary background fluxes pass through the S_1 surface and are then attenuated by the shield. The primary fluxes, in particular muons, may produce secondary backgrounds via scattering inside the shield. These secondary backgrounds can be reduced by a judiciously located active veto inside the shield itself. However, the extremely large forward production of pions near the beam line means that the cavity around the beam line itself can also be a source of a large flux of daughter muons, that instead transect the S_3 inner surface—“shield-clipping” muons—producing copious secondary neutrons and kaons or miss the shield entirely—“shield-evading” muons. To control these backgrounds, the shield coverage is extended beyond the angular acceptance of the detector, and an additional, radially oriented active veto is included, as shown in Fig. 2. Except when explicitly stated below, effects of shield-clipping muons are found to always be highly subleading compared to the background fluxes from muons traversing the full shield and are therefore hereafter neglected.

Control of the background processes in AL3X are determined by detector and signal-specific interplay between three different rates:

- (i) the detector trigger rate,
- (ii) the shield veto rate,
- (iii) the “potentially irreducible” background rate.

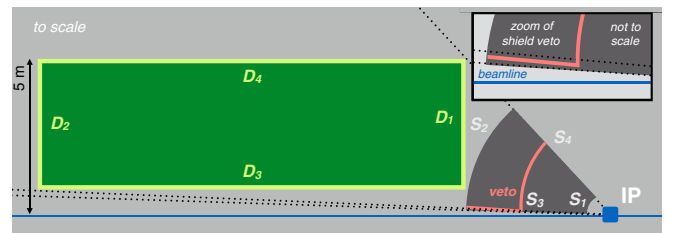


FIG. 2. Configuration of the shield, to scale.

The first of these is limited by the capabilities of the TPC readout, while being driven by the total charged flux through the detector, which can be large. An appropriate triggering strategy will reduce this rate to acceptable levels.

The shield veto rate is driven by the requirement that neutral secondary backgrounds produced downstream in the shield—mainly muoproduction of K_L 's and neutrons—can be vetoed down to acceptable levels by tagging the primary muons. This veto rate must not be so large that a significant fraction of all events is vetoed. The trigger rate sets an upper bound for the rate at which the shield veto needs to be read out, so that a fast shield veto readout need not be required.

Finally, the irreducible background rate sets the sensitivity to low rate signals. These backgrounds may arise mainly from the above-mentioned secondary K_L 's and neutrons or primary neutrinos. The extent of the detector's ability to reduce these backgrounds is both detector technology and LLP signal dependent; signals with no missing energy will be easier to distinguish from these backgrounds than missing energy LLP decay signatures, since one can require the vertex to point back to the IP. (This will be the case for all signal benchmarks we consider in Sec. V.) It is likely that signals and backgrounds can be somewhat well characterized and separated by using the TPC and B -field of the L3 magnet, though a full examination of these capabilities is beyond the scope of this proof-of-concept study. We therefore refer to the remaining backgrounds as potentially irreducible rather than irreducible. To eliminate some of the backgrounds which are soft and more difficult to estimate, we will hereafter impose a 3 GeV cut on the scalar sum of the momenta of the tracks of a candidate vertex. The effect of this cut on the signal efficiency for the benchmark models considered in Sec. V is 10% or less.

B. Shield-attenuated backgrounds

To estimate the backgrounds, we simulate minimum bias production of pion, kaon, muon, neutron, proton, and neutrino fluxes with PYTHIA8 [28,29]. Leptons produced from pion decay vertices at $\eta < 0.8$ and $r > 0.75$ m are neglected, under the assumption that their parent pions can be suppressed with a moderate amount additional shielding close to the IP, external to the geometric acceptance of the primary shield.

The propagation of primary backgrounds and production of secondary backgrounds inside tungsten is simulated with GEANT4 10.3 with the SHIELDING 2.1 physics list for which high energy interactions are modeled with the FTFP_BERT physics list based on the Fritiof [30–33] and Bertini intranuclear cascade [34–36] models and the standard electromagnetic physics package [37]. Propagating a large amount of events through the full shield is computationally prohibitive, so we instead use a “particle gun,” binned in energy and particle species, applied to a 5λ shield subelement (see Appendix). The resulting map between the incoming and outgoing fluxes is then applied recursively

to obtain the attenuation and response of the full 40λ shield. Neutrino production of neutral hadrons occurs at a prohibitively small rate and is not included in this analysis; these backgrounds are discussed in Sec. IV C below.

An active veto layer is located at a depth of 20λ inside the shield, with a rejection factor $\epsilon = 10^{-8}$, achievable, e.g., with several redundant layers of scintillator. The purpose of this “shield veto” is to detect charged tracks that may produce neutral secondary fluxes— K_L 's or neutrons—downstream in the shield itself, that may then enter the detector and produce a LLP-like event by decay or scattering. The location of the veto is determined by a balance between detecting charged particles before they create secondaries, not having too large a shield veto rate, and having sufficient material downstream of the veto to suppress neutral primary or secondary fluxes through the veto. The expected correlation between primary charged fluxes and neutral secondary fluxes within the shield—when a charged particle produces a secondary, it is typically not fully stopped—in principle permits vetoing some of these neutral secondaries produced upstream of the veto layer, so that the veto might be located deeper in the shield with a correspondingly lower shield veto rate. To be conservative, we have assumed the charged and neutral fluxes are instead fully decorrelated. The corresponding shield veto rate derived from this analysis, as well as the amount of shielding material required, is therefore expected to be an overestimate.

In Table I, we show the efficiencies (background flux/pp collision) for each relevant primary and secondary background entering the detector volume after propagation through the shield and application of the shield veto, integrated over energy above a minimal threshold. Also shown are the “shield veto rates,” corresponding to the flux of charged particles through the veto itself, relevant for an estimate of the event rejection rate by the shield veto. We divide the background fluxes, as appropriate, into those that transit the full shield, i.e., S_1 – S_2 , and those that are produced by shield-evading muons. Effects of shield-clipping muons are negligible for all backgrounds, with the exception of the muon rate itself, for which they comprise approximately 50% of the S_1 – S_2 rate, i.e., a flux/pp collision of 0.3%. We also provide, for the purpose of estimating the maximum required detector-trigger rate, the net background flux into the detector volume after propagation through the shield, but without application of the shield veto, in the second-to-last column. In order to characterize the sensitivity of the background rates to the 3 GeV cut, in Fig. 3, we show the neutron, K_L and neutrino kinetic energy spectra.

The shield veto rate is driven mainly by the muon flux and is approximately $\sim(0.6 + 0.3)\%$. Including a pileup factor of ~ 10 , this implies a 10% event rejection rate by the shield veto. As discussed above, this event veto rate is an overestimate, expected to be reduced once correlations

TABLE I. Results from the preliminary GEANT4 background simulation for $(20 + 20)\lambda W$ shield, i.e., with an active shield veto at 20λ , applying a veto efficiency of $\epsilon = 10^{-8}$. For outgoing neutrons and neutrinos, a cut on their kinetic energy was applied, as indicated in the first column. Background (BG) fluxes per pp collision (pp) are shown for fluxes entering the detector by traversing the full shield (S_1 – S_2) or by missing the shield entirely (Evade), together with veto rate for charged BG fluxes passing through the veto itself. Also shown are (upper bounds for) the net background fluxes that enter the detector, i.e., without the application of the veto rejection factor, relevant for the trigger veto rate. Actual potential background rates for 100/fb, shown in the final column, are obtained from the BG fluxes/pp, by folding in the decay or scattering probabilities, which are detector dependent, and assuming a minimum bias cross section of 100 mb (see the text for details).

BG species	Full shield (S_1 – S_2)		Evade shield	Net BG flux/pp into detector (no cuts)	BG rate per 100 fb $^{-1}$
	Shield veto rate	BG flux/pp	BG flux/pp		
$n + \bar{n}$ (>3 GeV)	...	$4. \times 10^{-16}$...	$3. \times 10^{-6}$	$\lesssim 0.2$
$p + \bar{p}$	$2. \times 10^{-6}$	$1. \times 10^{-14}$...	$5. \times 10^{-7}$...
μ	0.006	$3. \times 10^{-11}$	0.007	0.01	...
e	$5. \times 10^{-7}$	$3. \times 10^{-15}$...	$3. \times 10^{-7}$...
K_L^0	...	$1. \times 10^{-15}$...	$6. \times 10^{-8}$	$\lesssim 1$
K_S^0	...	$4. \times 10^{-16}$...	$3. \times 10^{-8}$	$\ll 1$
γ	...	$1. \times 10^{-15}$...	$1. \times 10^{-7}$...
π^\pm	$2. \times 10^{-6}$	$5. \times 10^{-15}$...	$4. \times 10^{-7}$...
K^\pm	$2. \times 10^{-7}$	$9. \times 10^{-16}$...	$8. \times 10^{-8}$...
$\nu + \bar{\nu}$ (>3 GeV)	...	0.01	$3. \times 10^{-4}$	0.2	$\lesssim 10$

between charged and neutral fluxes in the shield are included. There is also likely a spatial correlation between the charged primary and neutral secondary velocities, such that detector and veto might be segmented. This would permit vetoing only part of the detector, with the remainder open to detect a signal.

The net background flux into the detector is overwhelmingly dominated by muons, which have a flux/event of $\sim(0.6 + 0.3 + 0.7)\% \times$ pileup. While this background can easily be eliminated by detecting a track in the front of the detector, it can contribute to the trigger rate. In particular, muons would induce a trigger rate on D_2 and D_4 of $\mathcal{O}(\text{MHz})$, which is far too high for the TPC readout. However, one can use all four trigger/veto layers on the detector surfaces to significantly reduce the trigger rate from muons, by computing the number of hits in D_2 and D_4 less the number of hits in D_1 and D_3 :

$$T = \#\text{hits}(D_2 + D_4) - \#\text{hits}(D_1 + D_3). \quad (2)$$

Nominally, $T = 0$ for background muons, and $T > 0$ for a LLP decay, up to small instrumentation inefficiencies; the probability of tracks from a LLP decay hitting D_1 or D_3 is low, since they are expected to be mostly forward moving, and the average projective size of D_3 is only $\sim 10\%$ of D_4 . A triggering strategy that requires $T > 0$ may then reduce the trigger rate to acceptable levels. If needed, one may further segment the T variable azimuthally. The flux of other backgrounds will also contribute to the trigger rate, either as charged particles directly transiting the trigger layers—i.e., the protons, electrons, and pions—or by decaying or scattering to charged tracks in the detector volume—i.e., from the kaons and neutrons. However, these induce trigger rates of $\mathcal{O}(\text{Hz})$, well within TPC readout capabilities. Finally, cosmogenic muons would also trigger the detector with a rate of $\mathcal{O}(\text{kHz})$ [38], though if needed, this rate can be reduced substantially with a timing cut.

Folding the postshield background fluxes against their relevant decay or scattering probabilities allows estimation

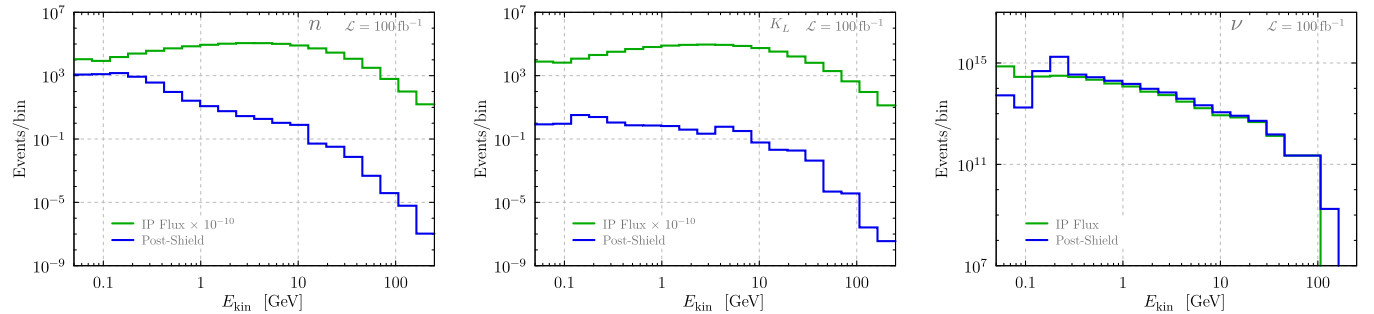


FIG. 3. Kinetic energy spectra of outgoing neutrons, K_L 's and neutrinos (blue), compared to their fluxes from the IP (green). For the sake of visual clarity, both the neutron and K_L fluxes from the IP are scaled by a factor of 10^{-10} .

of the total potential neutral background rates into the detector volume, shown in the final column of Table I. These are estimated assuming a minimum bias cross section of 100 mb, and for the estimation of scattering probabilities, we assume most of the TPC gas is neon at standard temperature and pressure. For the neutron background, only events which produce at least two tracks are relevant, notably reactions of the form $nn \rightarrow nn\pi^+\pi^-$ and $nn \rightarrow np^+\pi^-$. Using isospin symmetry, we estimate both processes from the analogous pp reactions [39], which results in a combined rate of ~ 5 mb per nucleon for $E_{\text{neutron}} \gtrsim 3$ GeV. Accounting for the TPC's size and gas pressure, this implies that roughly 5% of the neutrons entering the detector will create two or more tracks. This leads to an upper bound of ~ 0.2 neutron-induced background events per 100 fb $^{-1}$. It should be emphasized that this is still a substantial overestimate of the actual background rate, since these reactions predominantly occur at very low momentum transfer (see, e.g., Ref. [40]), resulting in mostly soft tracks. In reality, the neutron must thus be substantially harder than 3 GeV for the tracks to satisfy our 3 GeV cut on the scalar sum of their momenta, further reducing the background rate. We also note that a large control sample for this background can be acquired by simply inverting the shield veto, and it should be possible to characterize it carefully with data-driven methods. Moreover, the ongoing data-driven calibration for the CODEX-b shield will likely assist in calibration of expected background rates at AL3X.

C. Other backgrounds

Backgrounds induced by atmospheric neutrinos have been shown to be negligible for a detector of this size [22], but neutrinos originating from decays of primary pions must be considered; the kinetic energy spectrum of primary and secondary neutrinos entering the detector is shown in Fig. 3. We have verified that the contribution from pion decays dominates that of prompt neutrino production from $c\bar{c}$ and $b\bar{b}$ production by roughly an order of magnitude. The main process of concern is either coherent particle production of the neutrinos on the gas of the TPC, e.g., $\nu n \rightarrow \mu p\pi$, $\nu n \rightarrow p\mu$, or deep inelastic scattering. We require the neutrino to have energy $E_\nu \geq 3$ GeV and conservatively assume that all the remaining scattering events satisfy our 3 GeV cut on the scalar sum of the track momenta. The neutrino-nucleon cross section is taken from Ref. [41]; as a rough guide, for neutrinos with $E_\nu \gtrsim 1$ GeV, the cross section per nucleon $\sim 0.01(E_\nu/\text{GeV})$ pb. Integrating the cross section against the primary neutrino flux implies $\mathcal{O}(1)$ ν -Ne scattering events per 100 fb $^{-1}$. This background is threshold dominated [42], which means that for a charged- (neutral-)current interaction the majority of the energy is carried away by a relatively hard muon (neutrino), accompanied with a few soft hadronic tracks, which do not point to the IP. It should therefore be possible to further reduce this background by placing additional

cuts on the minimal track momentum or by applying an impact parameter cut.

The neutrino flux may also interact inelastically with the shield and create secondary neutral hadrons, in particular neutrons or K_L 's. If this occurs in the last few interaction lengths, these secondary hadrons may reach the detector. The MINERvA Collaboration measured the inclusive, neutral-current K^+ production cross section for a neutrino beam centered around $E_\nu \approx 3.5$ GeV to be $\sim 10^{-40}$ cm 2 [43], which we take as an approximation for the K_L production cross section. Charged-current scatters can be vetoed effectively by tagging the associated muon, but we do not exploit this here. Since the neutrino-nucleon cross section is expected to rise linearly with energy, we weight the MINERvA result with $E_\nu/3.5$ GeV and fold this against the neutrino spectrum obtained from PYTHIA8 (described above). As a result, we find $\mathcal{O}(3)K_L$'s which are produced in the last interaction length of the shield. These residual K_L should be further attenuated by an $\mathcal{O}(1)$ number under propagation through the last interaction length. Since they are expected to be soft, they are further likely reducible by a cut on track momenta and/or a requirement that the tracks point to the IP.

Using the total inelastic cross section [41], we bound the amount of neutrons produced in the last interaction length of the shield to be less than $\mathcal{O}(300)$. With a $\sim 5\%$ scattering probability in the TPC gas (see the previous section), this implies an upper bound of $\mathcal{O}(15)$ events. This number is conservative in two ways: (i) the produced neutrons should have substantially lower energy than the incoming neutrino, further softening the spectrum, and (ii) in the majority of the events, one or more charged states will be created along with the neutron. These charged states are likely to reach the detector as well and can therefore be used to veto the event. A full simulation, including a realistic detector response, is beyond the scope of this work, but for the time being, it appears plausible that this handful of background events can be fully reduced by the kinematic cuts described above.

Finally, there are additional sources of background, such as cavern backgrounds, beam-gas backgrounds, and hot-shield-induced backgrounds. We expect that none of these is a major source of background, but they would need to be carefully considered, possibly with *in situ* studies, should the AL3X proposal move forward. The cavern and hot-shield-induced backgrounds are likely eliminated by the cuts on the track momenta but are expected to contribute to the noise levels in the detector. Beam-gas events can produce neutral hadrons, which may bypass the shield. These events are very boosted in the lab frame, and for the tracks to hit the TPC, the beam-gas vertex must be located either far behind the IP, in which case the hadrons would pass through the full shield, or in the far forward region of the beam line well beyond the detector volume. In the latter case, any vertex made in the detector would generate tracks pointing toward, rather than from, the IP, which will not

occur for a LLP vertex. A hadron from a beam-gas event could in principle also deflect off the beam pipe or create more secondaries, and for a realistic design, it may be therefore be advisable to clad the beam pipe with a layer of tungsten, as is the case for the current forward muon detector of ALICE [1]. The beam-induced backgrounds have been measured in the hotter ATLAS cavern and were found to be small ($\ll 40$ MHz) [44] and should not affect the trigger strategy laid out in the previous section.

In summary, our baseline configuration has a trigger rate, veto rate, and irreducible background rate that are compatible with a close-to-zero background search for 100 fb^{-1} of luminosity. A trigger rate well below the TPC bandwidth can be achieved with a relatively simple algorithm using the T trigger variable, while the fraction of events vetoed by the shield is $\sim 1\% \times$ pileup. This veto can be applied offline, as the detector trigger rate is sufficiently slow.

V. REACH

In this section, we present three example benchmark models, representing high, medium, and low mass portals. For our reach estimates, we require three signal events, which roughly corresponds to a 95% C.L. exclusion, assuming zero background.

A. Exotic Higgs decays

Searching for exotic Higgs decays is a top priority for the high-luminosity LHC (HL-LHC) program. The small width of the SM Higgs means that relatively small couplings to exotic states, with mass $< m_h/2$, can lead to an appreciable Higgs exotic branching ratio. Combined with the large sample of Higgs bosons expected from HL-LHC—approximately 10^8 Higgs—this leads to a powerful portal for probing new physics. For a benchmark, we consider an exotic decay of $h \rightarrow XX$, with X a long-lived particle decaying to two or more charged SM particles. X could for instance be a kinetically mixed dark photon (e.g., Refs. [45–48]) or another (pseudo)scalar of an extended Higgs sector (e.g., Ref. [49]).

We estimate the fiducial efficiency for this benchmark with PYTHIA8 and show the resulting reach for 95% exclusion in Fig. 4, assuming negligible irreducible backgrounds. We see that AL3X can reach $h \rightarrow XX$ branching ratios as low as $\sim 5 \times 10^{-6}$, which is close to the best reach that is achievable with a 100 fb^{-1} data sample, corresponding to 6×10^6 Higgs bosons. In the large lifetime limit, the AL3X reach falls in between the reach for CODEX-b and MATHUSLA. For comparison, we also show the (optimistic) reach for 250 fb^{-1} and the reach assuming only the existing ALICE TPC for 100 fb^{-1} . Naturally, AL3X would have much better reach at low $c\tau$, regardless of the delivered luminosity, since it is much closer to the IP.

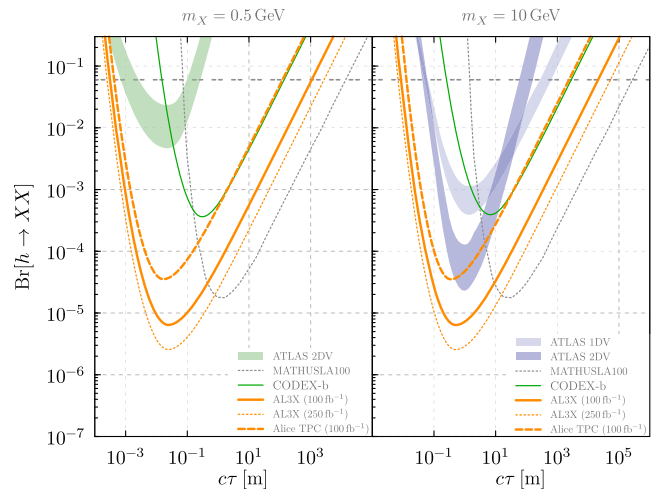


FIG. 4. Projected reach for AL3X, CODEX-b, MATHUSLA, and ATLAS (see the text) for $h \rightarrow XX$. For MATHUSLA, the 100×100 m configuration was assumed [19]. The reach for $h \rightarrow$ invisibles is also shown (horizontal gray dashed) [50].

The projected ATLAS reach for a Higgs decay to a pair of displaced dijets is also included in Fig. 4; the shaded bands attempt to indicate the uncertainty on these projections. For ATLAS and CMS, the mass of the LLP is a crucial parameter, as the number of tracks associated with the vertex is a key discriminant between signal and background. This is why the ATLAS reach for the 0.5 GeV benchmark [51] in Fig. 4 is substantially lower than for the 10 GeV mass benchmark [52,53]; the lower edge of the green ATLAS band is obtained by rescaling the current expected limit in Ref. [51], assuming that the systematic uncertainties could be lowered with a factor of 5 (for more details, see Sec. III of Ref. [22] and see Refs. [54,55] for the analogous searches by CMS.) To further reduce the backgrounds, ATLAS and CMS often require two displaced vertices, as indicated by the “2DV” label in Fig. 4, as compared to higher background searches with a single displaced vertex (“1DV”). In the latter case, both the current and future limits are merely projections [53], and it is conceivable that innovations in future analyses may substantially improve on this, e.g., by making use of the timing detectors [56]. For these reasons, the largest increase in sensitivity from proposals such as AL3X, CODEX-b, and MATHUSLA over ATLAS or CMS will be for LLPs with mass $\lesssim 10$ GeV, though there is gain for higher mass LLPs as well. Especially for low mass LLPs, the reach on the Higgs branching ratio can be improved by several orders of magnitude.

B. Exotic B decays

A new scalar state ϕ , lighter than the B meson, can also be produced in a neutral-current $B \rightarrow X_s \phi$ decay, even if the coupling of ϕ to the SM satisfies the ansatz of minimal flavor violation. A canonical example is the case where ϕ

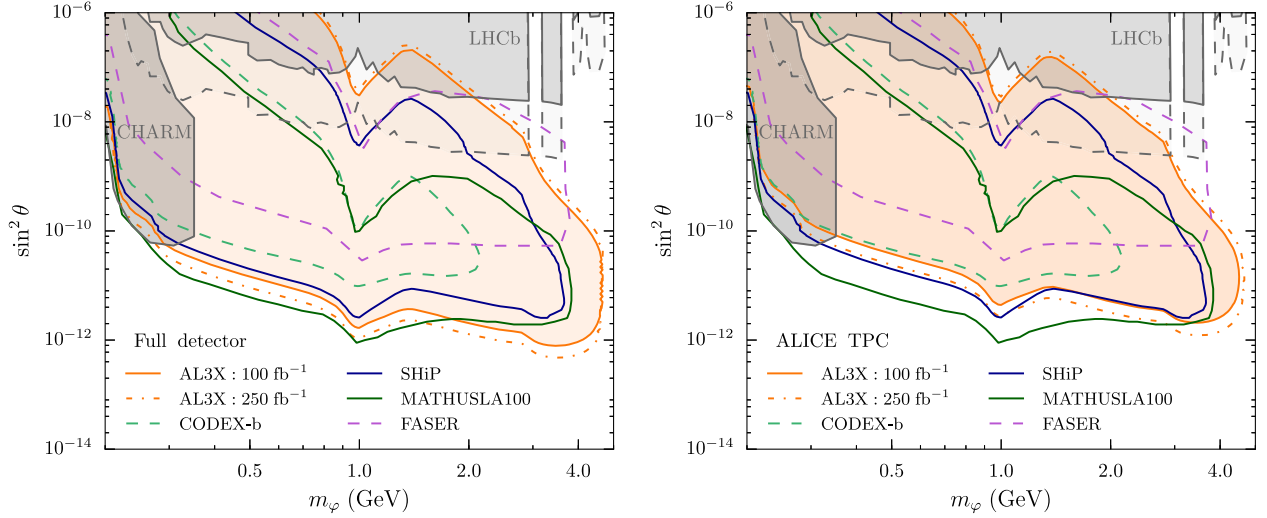


FIG. 5. AL3X reach for $B \rightarrow X_s \varphi$ in the s_θ^2 - m_φ plane for a full detector (left) and with just the existing ALICE TPC (right). The remaining solid (dashed) curves indicate the various existing (projected) constraints from the various existing, past, or proposed experiments, as described in the text.

mixes with the SM Higgs and thus obtains a coupling to the SM fermions proportional to their masses. The inclusive branching ratio for this process is [57–59]

$$\frac{\text{Br}[B \rightarrow X_s \varphi]}{\text{Br}[B \rightarrow X_c e \nu]} \simeq \frac{27 g^2 s_\theta^2}{256 \pi^2} \frac{m_t^4}{m_b^2 m_W^2} f(m_c/m_b) \left| \frac{V_{ts} V_{tb}}{V_{cb}} \right|^2 \approx 6. \times s_\theta^2 \times \lambda^2(m_\varphi/m_B, m_K/m_B), \quad (3)$$

where

$$\lambda(x, y) \equiv \sqrt{(1 - (x - y)^2)(1 - (x + y)^2)}, \quad (4)$$

and $s_\theta \equiv \sin \theta$ parametrizes the Higgs- φ mixing.

The φ must decay back to the SM through its induced Yukawa couplings with the SM fermions. Its lifetime is therefore also determined by s_θ^2 but is affected substantially by hadronic resonances for $m_\varphi \gtrsim 1$ GeV, as well as threshold effects. The theory uncertainties in this region are rather large, and we make use of the data-driven result from Refs. [60,61]. (This result is in good agreement with another, more recent calculation of the lifetime [62].) For our estimates, we assume a $b\bar{b}$ production cross section of $500 \mu\text{b}$ and compute the boost and pseudorapidity distributions with PYTHIA8.

Figure 5 shows the reach in the Higgs mixing portal s_θ^2 - m_φ parameter space, assuming 95% C.L. exclusion and negligible expected backgrounds. The existing constraints are from CHARM [63] and LHCb [12]; we also show the projected reach for LHCb, SHiP [64], MATHUSLA [65], CODEX-b [22], and FASER [24] for comparison. The LHCb reach was estimated by rescaling the $B \rightarrow K(\varphi \rightarrow \mu\mu)$ limit [12], optimistically assuming zero background. When

needed, the reach of other proposals was recast to match the assumptions for the lifetime of φ in Refs. [60,61].

C. Exotic π^0 and η decays

The final LLP production scenario we consider is via exotic decays of unflavored mesons. A popular benchmark for this type of process is a light $U(1)$ gauge boson that kinetically mixes with the SM photon through the operator $\epsilon/2F^{\mu\nu}F'_{\mu\nu}$. The π^0 and η branching ratios are [66]

$$\text{Br}(\pi^0 \rightarrow \gamma A') = 2\epsilon^2 \left(1 - \frac{m_{A'}^2}{m_{\pi^0}^2}\right)^3 \text{Br}(\pi \rightarrow \gamma\gamma) \quad (5)$$

$$\text{Br}(\eta \rightarrow \gamma A') = 2\epsilon^2 \left(1 - \frac{m_{A'}^2}{m_\eta^2}\right)^3 \text{Br}(\eta \rightarrow \gamma\gamma). \quad (6)$$

The $\eta' \rightarrow \gamma A'$ and $\omega \rightarrow \pi^0 A'$ processes may be also considered, but we have verified that they do not contribute substantially to the sensitivity. The A' width is given by

$$\Gamma = \frac{1}{3} \epsilon^2 \alpha m_{A'} \sqrt{1 - \frac{4m_e^2}{m_{A'}^2}} \times \frac{1}{\text{Br}(A' \rightarrow e^+e^-)}, \quad (7)$$

where we take the branching ratio to electrons from Ref. [67]. The lifetime of the dark photon is therefore only prolonged by the smallness of the mixing parameter ($\epsilon \ll 1$). The long lifetime regime is therefore only accessible at the LHC because of the huge numbers of π^0 and η mesons that are produced with a relatively high boost.

Searches for a kinematically mixed dark photon have been conducted for several decades, leading to a considerable list of existing constraints from a variety of probes.

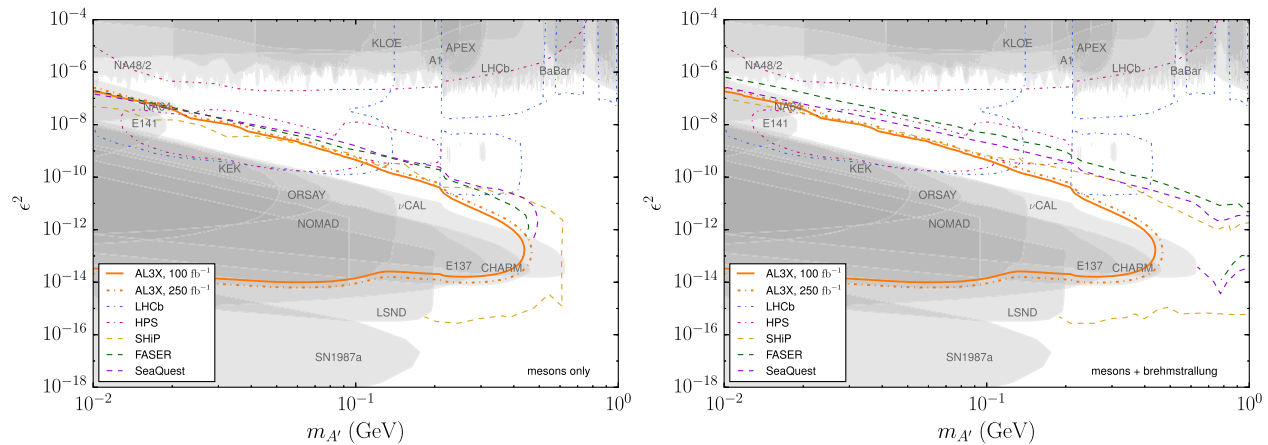


FIG. 6. (left) Existing constraints and projected reach for a kinetically mixed dark photon, including only meson decays. (right) Same projections and constraints, but for SeaQuest, SHiP, and FASER, the bremsstrahlung has been included.

In the high ϵ , short lifetime regime, the dominant constraints come from high intensity lepton colliders and B-factories: A1 [68], APEX [69], BABAR [70], KLOE [71–74], and LHCb [75]. (Several of the existing constraints were computed with the DARKCAST package [76]). The low ϵ , high $c\tau$, part of the dark photon parameter space has been probed by a range of beam dump and neutrino experiments: LSND [66,77,78], CHARM [79], SLAC beam dumps [80–82], KEK [83], NA48 [84], NA64 [85], NOMAD [86], ν CAL [87,88], and ORSAY [89]. Finally, the very low ϵ regime is constrained by limits on the anomalous cooling of supernova SN1987a [90].

We compute the boost distribution of the A' and the geometric acceptance of AL3X, using a minimum bias sample generated with PYTHIA8, using the measured inelastic cross section of 68 mb [91]. To model the tail of the A' boost distribution, we also include several weighted dijet samples with increasingly stringent cuts on the parton level process. Specifically, we demand that at least one hard parton satisfies $\eta < 4$ and $p_T > 30$ GeV, as well as a lower bound on its energy, where the latter is varied over different samples. The PYTHIA level cross sections are corrected with a κ -factor of 1.1, by comparing with the corresponding measurements [92]. To compute the detector efficiency, we add the efficiencies obtained from each of these samples, weighted by the appropriate fiducial cross section. The resulting reach is shown in Fig. 6. Also shown are the aforementioned existing constraints as well as the projected limits from planned or proposed experiments like SeaQuest [93,94], LHCb [95,96], SHiP [97], HPS [98], and FASER [23].

We find that for exotic π^0 and η decays AL3X can probe new parameter space in the high ϵ , low $c\tau$, regime but does not exceed the reach of dedicated forward experiments like FASER, SHiP, and SeaQuest. This can be understood as follows: though AL3X would likely have much fewer proton collisions than FASER, SeaQuest, and SHiP, it

would have the shortest effective baseline. In other words, since AL3X is somewhat forward from the IP, the tail of the kinetic energy distribution of the A' extends as high as several TeV. This is higher than SeaQuest and SHiP and somewhat lower than FASER, though FASER would be located ~ 400 m from the IP, compared to ~ 5 m for AL3X.

For the kinetically mixed dark photon model in particular, FASER, SeaQuest, and SHiP also have reach beyond the η mass due to A' production through bremsstrahlung of the protons,³ which is not available to AL3X. It should be mentioned that at AL3X a large flux of charged hadrons impinges on the shield, and it is possible this could result in a bremsstrahlung contribution as well. We save this calculation for future work. For reference, the FASER, SeaQuest, and SHiP reach, including this additional contribution, is shown in the right-hand panel of Fig. 6.

We conclude that AL3X is competitive with the dark photon reach obtainable by HPS, LHCb, SeaQuest, and FASER, though these either exist already in some form or are substantially smaller in scale than AL3X. While AL3X is then likely not the most ideal configuration for the specific case of the kinetically mixed dark photon, these results nonetheless demonstrate that AL3X has non-negligible reach for low mass LLP portals.

VI. DISCUSSION

The potential strength of an experiment like AL3X lies in its versatility over a wide mass range of LLP portals; thanks to its short baseline near a LHC interaction point, it can

³There is also an inelastic contribution from A' plus jet(s) production, which is difficult to reliably quantify in the very forward regime (FASER, SeaQuest, and SHiP). AL3X is located centrally enough that this process can be computed with standard tools, and we have verified that this contribution is negligible.

probe a range of LLP production modes all the way from exotic pion decays up to high mass portals like Higgs or top decays, or even heavier exotic states. We have seen that the AL3X concept can significantly increase the LLP reach over a broad range of portals, with respect to ATLAS, CMS, and LHCb. The AL3X configuration is in essence a tracking detector behind a heavy shield, which can be thought of as analogous to a calorimeter that is solely an absorber. While ATLAS, CMS, and LHCb do have tracking stations (muon spectrometers) following their calorimeters, their hadronic calorimeters are much thinner than the shield in the AL3X design. This permits AL3X to search for much rarer signals in a very low background environment compared to ATLAS and CMS, and in this sense, AL3X would be complementary to the existing (and proposed upgraded) multipurpose detectors.

The AL3X configuration shown in Fig. 1 only makes use of collision debris for $\eta > 0.8$. This leaves the entire region with $\eta < 0.8$ available for another experiment to coexist. One such possibility could be a dedicated detector to study τ -neutrinos, along the lines of what is proposed for SHiP [99]. We leave the exploration of such a companion experiment to future studies. To avoid challenges of moving the IP, an alternative configuration for AL3X could be to keep IP2 fixed and build a cylindrical shield and detector around the interaction point within the L3 magnet itself. Such a configuration would have smaller, lower energy backgrounds as it is more central, and its signal acceptance would be roughly comparable to the configuration in Fig. 1. However, the restricted radial size inside the L3 magnet—at most, 5.9 m—would require a tungsten-only shield to preserve enough space for a viable TPC detector, and the larger angular size would require a much larger, and heavier, amount of tungsten than in Fig. 1. There are also significant engineering challenges with this concept due to the mechanical weakness of the L3 magnet; both the shield and the detector would effectively need to be suspended inside the magnet.

To fully establish the feasibility and cost estimate of an experiment like AL3X at IP2, a number of things should be studied further. First and foremost, a study of the required beam configuration is needed to establish the technical feasibility and cost of delivering more luminosity, needed for virtually any new experiment looking for beyond the Standard Model physics at IP2. Next, a more detailed simulation of the shield configuration and the detector response is required, possibly by making use of some of the existing ALICE reconstruction software. Including the background rejection power of the detector itself, such a simulation will inform a more realistic design and size for the shield, compared to our conservative estimates in this analysis. Simultaneously, on the theory side, the reach of a

number of additional benchmark models—e.g., dark matter model(s), axionlike particles, and heavy neutral leptons—can be evaluated [100].

ACKNOWLEDGMENTS

We thank Anthony Barker, Asher Berlin, David Curtin, Nathaniel Craig, Jared Evans, Maurice Garcia-Sciveres, Andy Haas, Peter Jacobs, Spencer Klein, Zoltan Ligeti, Peilian Liu, Harikrishnan Ramani, Simone Pagan-Griso, Mateusz Ploskon, Scott Thomas, Jessie Shelton, David Stuart, and Jose Zurita for useful discussions. We are particularly grateful to Sergey Antipov, John Jowett, Gianluigi Arduini, and Burkhard Schmidt for helping us understand some of the possible accelerator related challenges; to Peter Jacobs for supplying us with details on the ALICE cavern and setup; and to David Curtin, Maurice Garcia-Sciveres, Peter Jacobs, and Spencer Klein for providing valuable comments on the manuscript. V. V. G. acknowledges funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme under Grant No. 724777 “RECEPT.” S. K. and M. P. are supported in part by the LDRD program of LBNL under Contract No. DE-AC02-05CH11231 and by the National Science Foundation (NSF) under Grants No. PHY-1002399 and No. PHY-1316783. S. K. also acknowledges support from DOE Grant No. DE-SC0009988 and from the Kavli Institute for Theoretical Physics, supported in part by the National Science Foundation under Grant No. NSF PHY-1748958, where part of this work was performed. The work of D. R. was supported in part by NSF Grant No. PHY-1720252. D. R. also acknowledges support from the University of Cincinnati and thanks the Aspen Center of Physics, supported by the NSF Grant No. PHY-1607611, where parts of this work were completed. This research used resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the DoE under Contract No. DE-AC02-05CH11231.

APPENDIX: PARTICLE GUN GRIDS

Sample grids are shown in Fig. 7 for muons, K_L^0 , and neutrons. The full background estimation includes similar grids for each incoming particle species and energy, for a total of 18×20 grids to account for all 18 incoming particle species and energies logarithmically spaced in incoming particle energy between 50 MeV and 300 GeV. The transfer probability for K_L^0 ’s for various sample grid points was validated against the measured values in Ref. [101] using a Pb shield, while the muon propagation is consistent with continuous-slowing-down-approximation estimates [102].

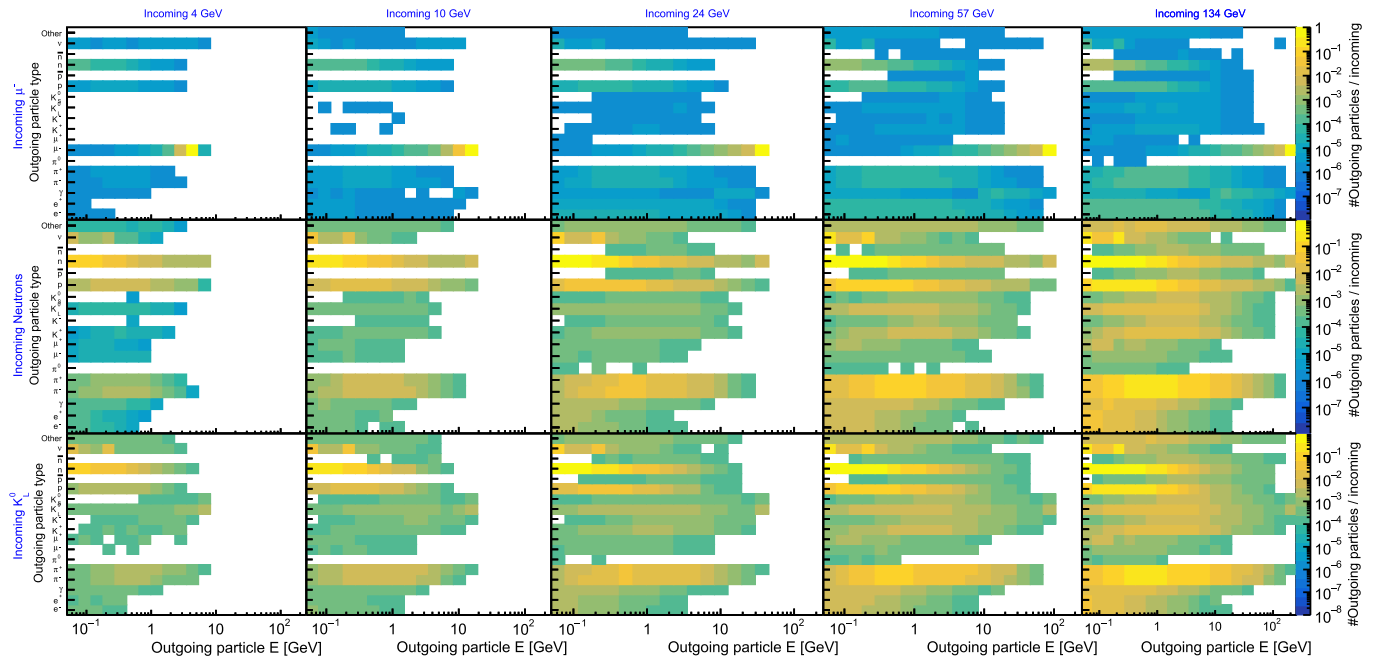


FIG. 7. The outgoing particle spectrum after 5λ of W for an incoming μ^- (top row), K_L^0 (middle row), and neutrons (bottom row) at various energies. These are a selection of 15 grids from a total of 18×20 : 18 species, logarithmically spaced in energy from 50 MeV to 300 GeV in 20 bins.

-
- [1] K. Aamodt *et al.* (ALICE Collaboration), *J. Instrum.* **3**, S08002 (2008).
- [2] J. Adam *et al.* (ALICE Collaboration), *Nat. Phys.* **13**, 535 (2017).
- [3] J. Adam *et al.* (ALICE Collaboration), *Nat. Phys.* **11**, 811 (2015).
- [4] B. Abelev *et al.* (ALICE Collaboration), Upgrade of the ALICE Experiment: Letter of intent, CERN, Technical Reports No. CERN-LHCC-2012-012, LHCC-I-022, ALICE-UG-002, 2012.
- [5] F. Wittgenstein, A. Hervé, M. Feldmann, D. Luckey, and I. Vetlitsky, Construction of the 13 magnet, in *11th International Conference on Magnet Technology (MT-11)*, edited by T. Sekiguchi and S. Shimamoto (Springer, Dordrecht, Netherlands, 1990), Vol. I, pp. 130–135.
- [6] Upgrade of the ALICE time projection chamber, CERN, Technical Reports No. CERN-LHCC-2013-020, ALICE-TDR-016, 2013.
- [7] D. Curtin *et al.*, [arXiv:1806.07396](https://arxiv.org/abs/1806.07396).
- [8] Search for heavy stable charged particles with 12.9 fb^{-1} of 2016 data, CERN, Technical Report No. CMS-PAS-EXO-16-036, 2016.
- [9] M. Aaboud *et al.* (ATLAS Collaboration), *Phys. Rev. D* **93**, 112015 (2016).
- [10] Inclusive search for new particles decaying to displaced jets at $\sqrt{s} = 13 \text{ TeV}$, CERN, Technical Report No. CMS-PAS-EXO-16-003, 2017.
- [11] Search for long-lived neutral particles decaying in the hadronic calorimeter of ATLAS at $\sqrt{s} = 13 \text{ TeV}$ in 3.2 fb^{-1} of data, CERN, Technical Report No. ATLAS-CONF-2016-103, 2016.
- [12] R. Aaij *et al.* (LHCb Collaboration), *Phys. Rev. D* **95**, 071101 (2017).
- [13] R. Aaij *et al.* (LHCb Collaboration), *Phys. Rev. Lett.* **115**, 161802 (2015).
- [14] A. Pierce, B. Shakya, Y. Tsai, and Y. Zhao, *Phys. Rev. D* **97**, 095033 (2018).
- [15] S. Antusch, E. Cazzato, and O. Fischer, *Phys. Lett. B* **774**, 114 (2017).
- [16] E. Cortina Gil *et al.* (NA62 Collaboration), *J. Instrum.* **12**, P05025 (2017).
- [17] S. Alekhin *et al.*, *Rep. Prog. Phys.* **79**, 124201 (2016).
- [18] J. P. Chou, D. Curtin, and H. J. Lubatti, *Phys. Lett. B* **767**, 29 (2017).
- [19] C. Alpigiani *et al.*, A letter of intent for MATHUSLA: A dedicated displaced vertex detector above ATLAS or CMS, CERN, Technical Reports No. CERN-LHCC-2018-025, LHCC-I-031, 2018.
- [20] A. Haas, C. S. Hill, E. Izaguirre, and I. Yavin, *Phys. Lett. B* **746**, 117 (2015).
- [21] A. Ball *et al.*, [arXiv:1607.04669](https://arxiv.org/abs/1607.04669).
- [22] V. V. Gligorov, S. Knapen, M. Papucci, and D. J. Robinson, *Phys. Rev. D* **97**, 015023 (2018).

- [23] J. L. Feng, I. Galon, F. Kling, and S. Trojanowski, *Phys. Rev. D* **97**, 035001 (2018).
- [24] J. L. Feng, I. Galon, F. Kling, and S. Trojanowski, *Phys. Rev. D* **97**, 055034 (2018).
- [25] F. Kling and S. Trojanowski, *Phys. Rev. D* **97**, 095016 (2018).
- [26] J. L. Feng, I. Galon, F. Kling, and S. Trojanowski, *Phys. Rev. D* **98**, 055021 (2018).
- [27] Technical proposal: L3, CERN, Technical Reports No. CERN-LEPC-83-5, LEPC-P-4, 1983.
- [28] T. Sjostrand, S. Mrenna, and P. Z. Skands, *J. High Energy Phys.* **05** (2006) 026.
- [29] T. Sjostrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, and P. Z. Skands, *Comput. Phys. Commun.* **191**, 159 (2015).
- [30] B. Andersson, G. Gustafson, and B. Nilsson-Almqvist, *Nucl. Phys.* **B281**, 289 (1987).
- [31] B. Andersson, A. Tai, and B.-H. Sa, *Z. Phys. C* **70**, 499 (1996).
- [32] B. Nilsson-Almqvist and E. Stenlund, *Comput. Phys. Commun.* **43**, 387 (1987).
- [33] B. Ganhuyag and V. Uzhinsky, *Czech. J. Phys.* **47**, 913 (1997).
- [34] M. P. Guthrie, R. G. Alsmiller, and H. W. Bertini, *Nucl. Instrum. Methods* **66**, 29 (1968).
- [35] H. W. Bertini and M. P. Guthrie, *Nucl. Phys.* **A169**, 670 (1971).
- [36] V. A. Karmanov, *Nucl. Phys.* **B166**, 378 (1980).
- [37] H. Burkhardt, V. M. Grichine, P. Gumplinger, V. N. Ivanchenko, R. P. Kokoulin, M. Maire, and L. Urban, in *IEEE Symposium Conference Record Nuclear Science 2004* (IEEE, Rome, Italy, 2004), Vol. 3 pp. 1907–1910.
- [38] V. Drollinger, Simulation of beam halo and cosmic muons, CERN Technical Report No. CMS-NOTE-2005-012, 2005.
- [39] G. Alexander, O. Benary, and U. Maor, *Nucl. Phys.* **B5**, 1 (1968).
- [40] A. P. Colleraine and U. Nauenberg, *Phys. Rev.* **161**, 1387 (1967).
- [41] J. A. Formaggio and G. P. Zeller, *Rev. Mod. Phys.* **84**, 1307 (2012).
- [42] P. A. Rodrigues *et al.* (MINERvA Collaboration), *Phys. Rev. Lett.* **116**, 071802 (2016).
- [43] C. M. Marshall *et al.* (MINERvA Collaboration), *Phys. Rev. Lett.* **119**, 011802 (2017).
- [44] G. Aad *et al.* (ATLAS Collaboration), *J. Instrum.* **11**, P05013 (2016).
- [45] R. M. Schabinger and J. D. Wells, *Phys. Rev. D* **72**, 093007 (2005).
- [46] S. Gopalakrishna, S. Jung, and J. D. Wells, *Phys. Rev. D* **78**, 055002 (2008).
- [47] D. Curtin, R. Essig, S. Gori, and J. Shelton, *J. High Energy Phys.* **02** (2015) 157.
- [48] M. J. Strassler, arXiv:0801.0629.
- [49] D. Curtin *et al.*, *Phys. Rev. D* **90**, 075004 (2014).
- [50] Proceedings, 2013 Community Summer Study on the Future of U.S. Particle Physics: Snowmass on the Mississippi (CSS2013): Minneapolis, MN, USA, 2013, CERN, Technical Report No. CMS-NOTE-13-002, 2013.
- [51] Search for long-lived neutral particles decaying into displaced lepton jets in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector, CERN, Technical Report No. ATLAS-CONF-2016-042, 2016.
- [52] G. Aad *et al.* (ATLAS Collaboration), *Phys. Rev. D* **92**, 012010 (2015).
- [53] A. Coccaro, D. Curtin, H. J. Lubatti, H. Russell, and J. Shelton, *Phys. Rev. D* **94**, 113003 (2016).
- [54] S. Chatrchyan *et al.* (CMS Collaboration), *J. High Energy Phys.* **02** (2013) 085.
- [55] A search for beyond standard model light bosons decaying into muon pairs, CERN, Technical Report No. CMS-PAS-HIG-16-035, 2016.
- [56] J. Liu, Z. Liu, and L.-T. Wang, arXiv:1805.05957.
- [57] R. S. Willey and H. L. Yu, *Phys. Rev. D* **26**, 3287 (1982).
- [58] R. Chivukula and A. V. Manohar, *Phys. Lett. B* **207**, 86 (1988).
- [59] B. Grinstein, L. J. Hall, and L. Randall, *Phys. Lett. B* **211**, 363 (1988).
- [60] A. Fradette and M. Pospelov, *Phys. Rev. D* **96**, 075033 (2017).
- [61] F. Bezrukov and D. Gorbunov, *J. High Energy Phys.* **05** (2010) 010.
- [62] M. W. Winkler, arXiv:1809.01876.
- [63] F. Bergsma *et al.*, *Phys. Lett. B* **157**, 458 (1985).
- [64] G. Lanfranchi (SHiP Collaboration), <https://cds.cern.ch/record/2243034>.
- [65] J. A. Evans, *Phys. Rev. D* **97**, 055046 (2018).
- [66] B. Batell, M. Pospelov, and A. Ritz, *Phys. Rev. D* **80**, 095024 (2009).
- [67] M. Buschmann, J. Kopp, J. Liu, and P. A. N. Machado, *J. High Energy Phys.* **07** (2015) 045.
- [68] H. Merkel *et al.*, *Phys. Rev. Lett.* **112**, 221802 (2014).
- [69] S. Abrahamyan *et al.* (APEX Collaboration), *Phys. Rev. Lett.* **107**, 191804 (2011).
- [70] J. P. Lees *et al.* (BABAR Collaboration), *Phys. Rev. Lett.* **113**, 201801 (2014).
- [71] A. Anastasi *et al.*, *Phys. Lett. B* **750**, 633 (2015).
- [72] A. Anastasi *et al.* (KLOE-2 Collaboration), *Phys. Lett. B* **757**, 356 (2016).
- [73] D. Babusci *et al.* (KLOE-2 Collaboration), *Phys. Lett. B* **720**, 111 (2013).
- [74] D. Babusci *et al.* (KLOE-2 Collaboration), *Phys. Lett. B* **736**, 459 (2014).
- [75] R. Aaij *et al.* (LHCb Collaboration), *Phys. Rev. Lett.* **120**, 061801 (2018).
- [76] P. Ilten, Y. Soreq, M. Williams, and W. Xue, *J. High Energy Phys.* **06** (2018) 004.
- [77] C. Athanassopoulos *et al.* (LSND Collaboration), *Phys. Rev. C* **58**, 2489 (1998).
- [78] R. Essig, R. Harnik, J. Kaplan, and N. Toro, *Phys. Rev. D* **82**, 113008 (2010).
- [79] F. Bergsma *et al.* (CHARM Collaboration), *Phys. Lett.* **157B**, 458 (1985).
- [80] J. D. Bjorken, S. Ecklund, W. R. Nelson, A. Abashian, C. Church, B. Lu, L. W. Mo, T. A. Nunamaker, and P. Rassmann, *Phys. Rev. D* **38**, 3375 (1988).
- [81] E. M. Riordan *et al.*, *Phys. Rev. Lett.* **59**, 755 (1987).
- [82] A. Bross, M. Crisler, S. H. Pordes, J. Volk, S. Errede, and J. Wrbanek, *Phys. Rev. Lett.* **67**, 2942 (1991).

- [83] A. Konaka *et al.*, *Phys. Rev. Lett.* **57**, 659 (1986).
- [84] J. R. Batley *et al.* (NA48/2 Collaboration), *Phys. Lett. B* **746**, 178 (2015).
- [85] D. Banerjee *et al.* (NA64 Collaboration), *Phys. Rev. Lett.* **120**, 231802 (2018).
- [86] P. Astier *et al.* (NOMAD Collaboration), *Phys. Lett. B* **506**, 27 (2001).
- [87] J. Blumlein *et al.*, *Z. Phys. C* **51**, 341 (1991).
- [88] J. Blumlein *et al.*, *Int. J. Mod. Phys. A* **07**, 3835 (1992).
- [89] M. Davier and H. Nguyen Ngoc, *Phys. Lett. B* **229**, 150 (1989).
- [90] J. H. Chang, R. Essig, and S. D. McDermott, *J. High Energy Phys.* **01** (2017) 107.
- [91] M. Aaboud *et al.* (ATLAS Collaboration), *Phys. Rev. Lett.* **117**, 182002 (2016).
- [92] G. Aad *et al.* (ATLAS Collaboration), *Phys. Rev. D* **86**, 014022 (2012).
- [93] S. Gardner, R. J. Holt, and A. S. Tadepalli, *Phys. Rev. D* **93**, 115015 (2016).
- [94] A. Berlin, S. Gori, P. Schuster, and N. Toro, *Phys. Rev. D* **98**, 035011 (2018).
- [95] P. Ilten, J. Thaler, M. Williams, and W. Xue, *Phys. Rev. D* **92**, 115017 (2015).
- [96] P. Ilten, Y. Soreq, J. Thaler, M. Williams, and W. Xue, *Phys. Rev. Lett.* **116**, 251803 (2016).
- [97] E. Graverini (SHiP Collaboration), <https://cds.cern.ch/record/2214092>.
- [98] O. Moreno, in *Meeting of the APS Division of Particles and Fields (DPF 2013) Santa Cruz, California, USA, 2013* (2013).
- [99] M. Anelli *et al.* (SHiP Collaboration), [arXiv:1504.04956](https://arxiv.org/abs/1504.04956).
- [100] J. A. Evans, S. Knapen, H. Ramani, and D. J. Robinson (to be published).
- [101] A. Gsponer, J. Hoffnagle, W. R. Molzon, J. Roehrig, V. L. Telegdi, B. Winstein, S. H. Aronson, G. J. Bock, D. Hedin, and G. B. Thomson, *Phys. Rev. Lett.* **42**, 9 (1979).
- [102] C. Patrignani *et al.* (Particle Data Group), *Chin. Phys. C* **40**, 100001 (2016).