



Expression of interest for the CODEX-b detector

Giulio Aielli¹, Eli Ben-Haim³, Roberto Cardarelli², Matthew John Charles³, Xabier Cid Vidal⁴, Victor Coco⁵, Biplab Dey⁶, Raphael Dumps⁵, Jared A. Evans⁷, George Gibbons⁸, Olivier Le Dortz³, Vladimir V. Gligorov³, Philip Ilten^{8,17}, Simon Knapen⁹, Jongho Lee^{5,10}, Saul López Soliño⁴, Benjamin Nachman¹¹, Michele Papucci^{11,12}, Francesco Polci³, Robin Quessard¹³, Harikrishnan Ramani^{11,14}, Dean J. Robinson^{11,a}, Heinrich Schindler⁵, Michael D. Sokoloff⁷, Paul Swallow⁸, Riccardo Vari¹⁵, Nigel Watson⁸, Mike Williams¹⁶

- ¹ Università e INFN Sezione di Roma Tor Vergata, Rome, Italy
- ² INFN Sezione di Roma Tor Vergata, Rome, Italy
- ³ LPNHE, Sorbonne Université, Paris Diderot Sorbonne Paris Cité, CNRS/IN2P3, Paris, France
- ⁴ Instituto Galego de Física de Altas Enerxías (IGFAE), Universidade de Santiago de Compostela, Santiago de Compostela, Spain
- ⁵ European Organization for Nuclear Research (CERN), Geneva, Switzerland
- ⁶ ELTE Eötvös Loránd University, Budapest, Hungary
- ⁷ Department of Physics, University of Cincinnati, Cincinnati, OH 45221, USA
- ⁸ University of Birmingham, Birmingham, UK
- ⁹ School of Natural Sciences, Institute for Advanced Study, Princeton, NJ 08540, USA
- ¹⁰ Kyungpook National University (KNU), Daegu, South Korea
- ¹¹ Ernest Orlando Lawrence Berkeley National Laboratory, University of California, Berkeley, CA 94720, USA
- ¹² Walter Burke Institute for Theoretical Physics, California Institute of Technology, Pasadena, CA 91125, USA
- ¹³ Département de Physique de l'École Normale Supérieure, LPENS, Paris, France
- ¹⁴ Department of Physics, University of California, Berkeley, CA 94720, USA
- ¹⁵ INFN Sezione di Roma La Sapienza, Rome, Italy
- ¹⁶ Laboratory for Nuclear Science, Massachusetts Institute of Technology, Cambridge, MA 02139, USA
- ¹⁷ Present address: University of Cincinnati, Cincinnati, OH 45221, USA

Received: 13 April 2020 / Accepted: 25 November 2020 / Published online: 21 December 2020
© The Author(s) 2020

Abstract This document presents the physics case and ancillary studies for the proposed CODEX-b long-lived particle (LLP) detector, as well as for a smaller proof-of-concept demonstrator detector, CODEX- β , to be operated during Run 3 of the LHC. Our development of the CODEX-b physics case synthesizes ‘top-down’ and ‘bottom-up’ theoretical approaches, providing a detailed survey of both minimal and complete models featuring LLPs. Several of these models have not been studied previously, and for some others we amend studies from previous literature: In particular, for gluon and fermion-coupled axion-like particles. We moreover present updated simulations of expected backgrounds in CODEX-b’s actively shielded environment, including the effects of shielding propagation uncertainties, high-energy tails and variation in the shielding design. Initial results are also included from a background measurement and calibration campaign. A design overview is presented for the CODEX- β demonstrator detector, which will enable background calibration and detector design studies. Finally, we lay out brief studies of various design drivers

of the CODEX-b experiment and potential extensions of the baseline design, including the physics case for a calorimeter element, precision timing, event tagging within LHCb, and precision low-momentum tracking.

Contents

Executive summary	2
1 Introduction	2
1.1 Motivation	2
1.2 Experimental requirements	4
1.3 Baseline detector concept	4
1.4 Search power, complementarities and unique features	5
1.5 Timeline	6
2 Physics case	6
2.1 Theory survey strategies	6
2.2 Novel studies	7
2.3 Minimal models	7
2.3.1 Abelian hidden sector	8
2.3.2 Scalar-Higgs portal	9
2.3.3 Axion-like particles	10

^ae-mail: d Robinson@lbl.gov (corresponding author)

2.3.4 Heavy neutral leptons	13
2.4 Complete models	14
2.4.1 R-parity violating supersymmetry	16
2.4.2 Relaxion models	16
2.4.3 Neutral naturalness	16
2.4.4 Inelastic dark matter	18
2.4.5 Dark matter coscattering	19
2.4.6 Dark matter from sterile coannihilation	21
2.4.7 Asymmetric dark matter	21
2.4.8 Other dark matter models	21
2.4.9 Baryogenesis	22
2.4.10 Hidden valleys	23
3 Backgrounds	23
3.1 Overview	24
3.2 Simulation	25
3.2.1 Primary fluxes	25
3.2.2 Simulated shield propagation	25
3.2.3 Neutrinos	26
3.2.4 Shielding marginal performance	27
3.2.5 Simulated track production	29
3.3 Measurement campaign	29
4 CODEX- β	32
4.1 Motivation	32
4.2 Technical description and timetable	33
4.3 New physics reach	34
5 Detector case studies	35
5.1 Tracking	35
5.1.1 Design drivers	35
5.1.2 Studies performed	37
5.2 Timing	38
5.3 Calorimetry	38
5.4 Tagging of events at LHCb	40
6 Outlook	41
References	41

Executive summary

The Large Hadron Collider (LHC) provides unprecedented sensitivity to short-distance physics. Primary achievements of the experimental program include the discovery of the Higgs boson [1, 2], the ongoing investigation of its interactions [3], and remarkable precision Standard Model (SM) measurements. Furthermore, a multitude of searches for physics beyond the Standard Model (BSM) have been conducted over a tremendous array of channels. These have resulted in greatly improved BSM limits, with no new particles or force carriers having been found.

The primary LHC experiments (ATLAS, CMS, LHCb, ALICE) have proven to be remarkably versatile and complementary in their BSM reach. As these experiments are scheduled for upgrades and data collection over at least another 15 years, it is natural to consider whether they can be further

complemented by one or more detectors specialized for well-motivated but currently hard-to-detect BSM signatures. A compelling category of such signatures are long-lived particles (LLPs), which generally appear in any theory containing a hierarchy of scales or small parameters, and are therefore ubiquitous in BSM scenarios.

The central challenge in detecting LLPs is that not only their masses but also their lifetimes may span many orders of magnitude. This makes it impossible from first principles to construct a single detector which would have the ultimate sensitivity to all possible LLP signatures; multiple complementary experiments are necessary, as summarized in Fig. 1.

In this expression of interest we advocate for CODEX-b (“Compact Detector for EXotics at LHCb”), a LLP detector that would be installed in the DELPHI/UXA cavern next to LHCb’s interaction point (IP8). The approximate proposed timeline is given in Fig. 2: Here “CODEX- β ” refers to a smaller proof-of-concept detector with otherwise the same basic geometry and technology as CODEX-b.

The central advantages of CODEX-b are:

- Very competitive sensitivity to a wide range of LLP models, either exceeding or complementary to the sensitivity of other existing or proposed detectors;
- An achievable zero background environment, as well as an accessible experimental location in the DELPHI/UXA cavern with all necessary services already in place;
- Straightforward integration into LHCb’s trigger-less readout and the ability to tag events of interest with the LHCb detector;
- A compact size and consequently modest cost, with the realistic possibility to extend detector capabilities for neutral particles in the final state.

We survey a wide range of BSM scenarios giving rise to LLPs and demonstrate how these advantages translate into competitive and complementary reach with respect to other proposals. We furthermore detail the experimental and simulation studies carried out so far, showing that CODEX-b can be built as planned and operate as a zero background experiment. We also discuss possible technology options that may further enhance the reach of CODEX-b. Finally, we discuss the timetable for the construction and data taking of CODEX- β , and show that it may also achieve new reach for certain BSM scenarios.

1 Introduction

1.1 Motivation

New Physics (NP) searches at the LHC and other experiments have primarily been motivated by the predictions of various

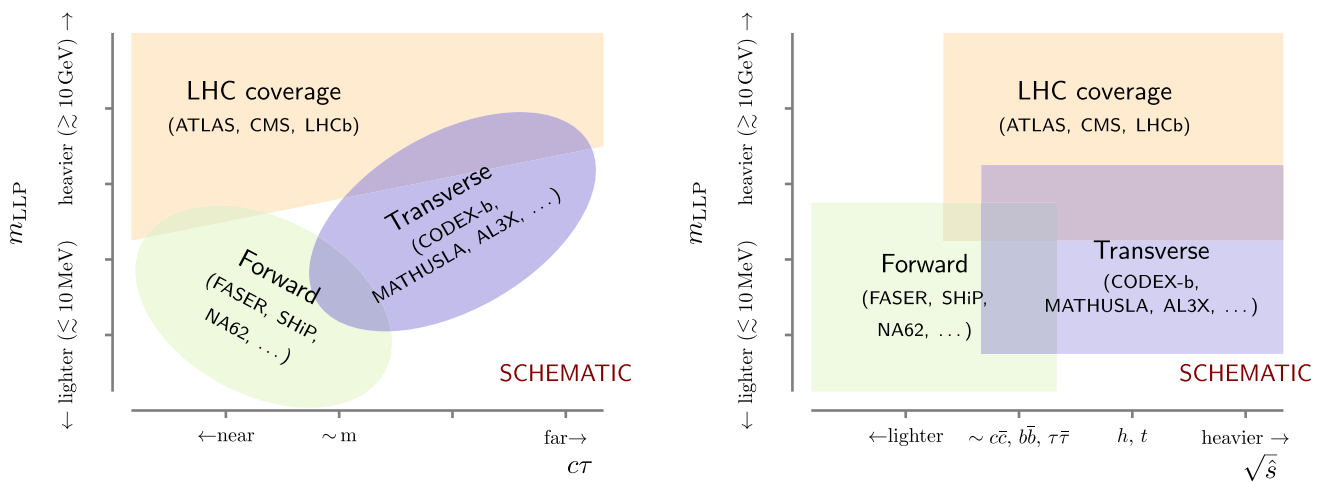


Fig. 1 Schematic summary of reach and coverage of current, planned or proposed experiments in terms of the LLP mass, lifetime and the required parton center-of-mass energy, $\sqrt{\hat{s}}$

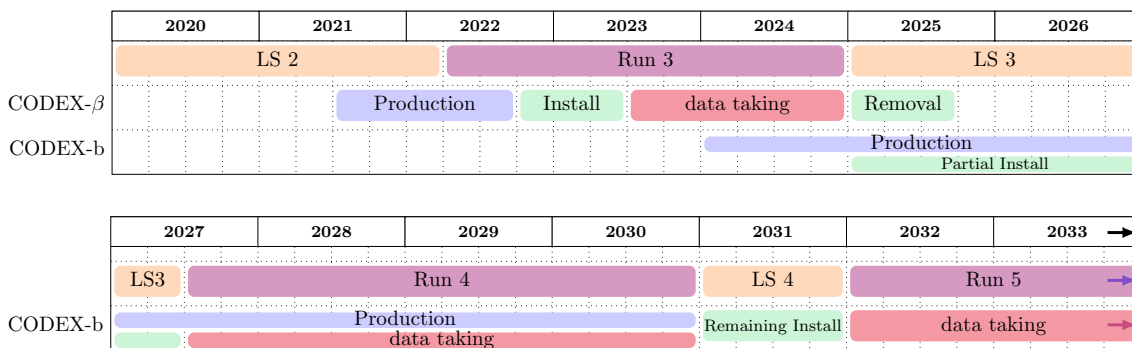


Fig. 2 Approximate production, installation and data-taking timelines for demonstrator (CODEX- β) and full (CODEX-b) detectors. (Updated according to COVID constraints.)

extensions of the SM, designed to address long-standing open questions. These include e.g. the origin and nature of dark matter, the detailed dynamics of the weak scale, the mechanism of baryogenesis, among many others. However, in the absence of clear experimental NP hints, the solutions to these puzzles remain largely mysterious. Combined with increasing tensions from current collider data on the most popular BSM extensions, it has become increasingly imperative to consider whether the quest for NP requires new and innovative strategies: A means to diversify LHC search programs with a minimum of theoretical prejudice, and to seek signatures for which the current experiments are trigger and/or background limited.

A central component of this program will be the ability to probe ‘dark’ or ‘hidden’ sectors, comprised of degrees of freedom that are ‘sterile’ under the SM gauge interactions. Hidden sectors are ubiquitous in many BSM scenarios, and typically may feature either suppressed renormalizable couplings, heavy mediator exchanges with SM states, or both.¹

¹ This heavy mediator may itself be an exotic state or part of the SM electroweak sector, such as the Higgs. The possibility of exploring

The sheer breadth of possibilities for these hidden sectors mandates search strategies that are as model-independent as possible.

Suppressed dark–SM couplings or heavy dark–SM mediators may in turn give rise to relatively long lifetimes for light degrees of freedom in the hidden spectrum, by inducing suppressions of their total widths via either small couplings, the mediator mass, loops and/or phase space. This scenario is very common in many models featuring e.g. Dark Matter (Sects. 2.4.4, 2.4.5, 2.4.6, 2.4.7 and 2.4.8), Baryogenesis (Sect. 2.4.9), Supersymmetry (Sect. 2.4.1) or Neutral Naturalness (Sect. 2.4.3). The canonical examples in the SM are the long lifetimes of the K_L^0 , π^\pm , neutron and muon, whose

Footnote 1 continued

the ‘Higgs portal’ is particularly compelling, because our theoretical understanding of Higgs interactions is likely incomplete, and new states might interact with it. In addition, the Higgs itself may have a sizable branching ratio to exotic states since its SM partial width is suppressed by the b -quark Yukawa coupling. Experiments capable of leveraging large samples of Higgs bosons are then natural laboratories to search for NP. Understanding the properties of the Higgs sector will be central to ongoing and future particle physics programs.

widths are suppressed by the weak interaction scale required for flavor changing processes, as well as phase space. Vestiges of hidden sectors may then manifest in the form of striking morphologies within LHC collisions, in particular the *displaced decay-in-flight* of these metastable, light particles in the hidden sector, commonly referred to as ‘long-lived particles’ (LLPs). Surveying a wide range of benchmark scenarios, we demonstrate in this document that by searching for such LLP decays, CODEX-b would permit substantial improvements in reach for many well-motivated NP scenarios, well beyond what could be gained by an increase in luminosity at the existing detectors.

1.2 Experimental requirements

In any given NP scenario, the decay width of an LLP may exhibit strong power-law dependencies on a priori unknown ratios of various physical scales. As a consequence, theoretical priors for the LLP lifetime are broad, such that LLPs may occupy a wide phenomenological parameter space. In the context of the LHC, LLP masses from several MeV up to $\mathcal{O}(1)$ TeV may be contemplated, and proper lifetimes as long as $\lesssim 0.1$ seconds may be permitted before tensions with Big Bang Nucleosynthesis arise [4–6].

Broadly speaking, the ability of any given experiment to probe a particular point in this space of LLP mass and lifetimes will depend strongly not only on the center-of-mass energy available to the experiment, but also on its fiducial detector volume, distance from the interaction point (IP), triggering limitations, and the size of irreducible backgrounds in the detector environment [7]. The latter is large for light LLP searches, requiring a *shielded, background-free* detector. Further, LLP production channels involving the decay of a heavy parent state – e.g. a Higgs decay – require sufficient partonic center-of-mass energy, $\sqrt{\hat{s}}$, to produce an abundant sample of heavy parents. Such channels are thus probed most effectively *transverse* to an LHC interaction point. Taken together, these varying requirements prevent any single experimental approach from attaining comprehensive coverage over the full parameter space.

Experimental coverage of LLP searches is also determined by the *morphology* of LLP decays. The simplest scenario contemplates a large branching ratio for 2-body LLP decays to two charged SM particles – for instance $\ell^+\ell^-$, $\pi^+\pi^-$ or K^+K^- . In many well-motivated benchmark scenarios (see Sect. 2), however, the LLP may decay to various final states involving missing energy, photons, or high multiplicity, softer final states. In any experimental environment, these more complex decay morphologies can be much more challenging to detect or reconstruct: Reconstructing missing energy final states requires the ability to measure track momenta; detecting photons requires a calorimeter element or preshower component; identifying high mul-

tiplicity final states requires the suppression of soft hadronic backgrounds. The CODEX-b baseline concept, as described below in Sect. 1.3, is well-suited to reconstruct several of these morphologies, in addition to the simple 2-body decays. Extensions of the baseline design may permit some calorimetry or pre-shower capabilities, which would enable the reconstruction of photons and other neutral hadrons.

1.3 Baseline detector concept

The proposed CODEX-b location is in the UX85 cavern, roughly 25 meters from the interaction point 8 (IP8), with a nominal fiducial volume of $10\text{ m} \times 10\text{ m} \times 10\text{ m}$ (see Fig. 3a) [8]. Specifically, the fiducial volume is defined by $26\text{ m} < x < 36\text{ m}$, $-7\text{ m} < y < 3\text{ m}$ and $5\text{ m} < z < 15\text{ m}$, where the z direction is aligned along the beam line and the origin of the coordinate system is taken to be the interaction point. This location roughly corresponds to the pseudorapidity range $0.13 < \eta < 0.54$. Passive shielding is partially provided by the existing UXA wall, while the remainder is achieved by a combination of active vetos and passive shielding located nearer to the IP. A detailed description of the backgrounds and the required amount of shielding can be found in Sect. 3.

The actual reach of any LLP detector will be tempered by various efficiencies, including efficiencies for tracking and vertex reconstruction. In particular, no magnetic field will be available in the CODEX-b fiducial volume. To design an LLP *detection* program, rather than only an exclusionary one, it is therefore important to be able to confirm the presence of exotic physics and reject possibly mis-modeled backgrounds. This requires capabilities for particle identification, mass reconstruction and/or event reconstruction.

To address these considerations, several detector concepts are being considered. The baseline CODEX-b conceptual design makes use of Resistive Plate Chambers (RPC) tracking stations with $\mathcal{O}(100)$ ps timing resolution. A hermetic detector, with respect to the LLP decay vertex, is needed to achieve good signal efficiency and background rejection. In the baseline design, this is achieved by placing six RPC layers on each surface of the detector. To ensure good vertex resolution five additional triplets of RPC layers are placed equally spaced along the depth of the detector, as shown in Fig. 3b. Other, more ambitious options are being considered, that use both RPCs as well as large scale calorimeter technologies such as liquid [10] or plastic scintillators, used in accelerator neutrino experiments such as NO ν A [11], T2K upgrade [12] or DUNE [13]. If deemed feasible, implementing one of these options would permit measurement of decay modes involving neutral final states, improved particle identification and more efficient background rejection techniques.

Because the baseline CODEX-b concept makes use of proven and well-understood technologies for tracking and

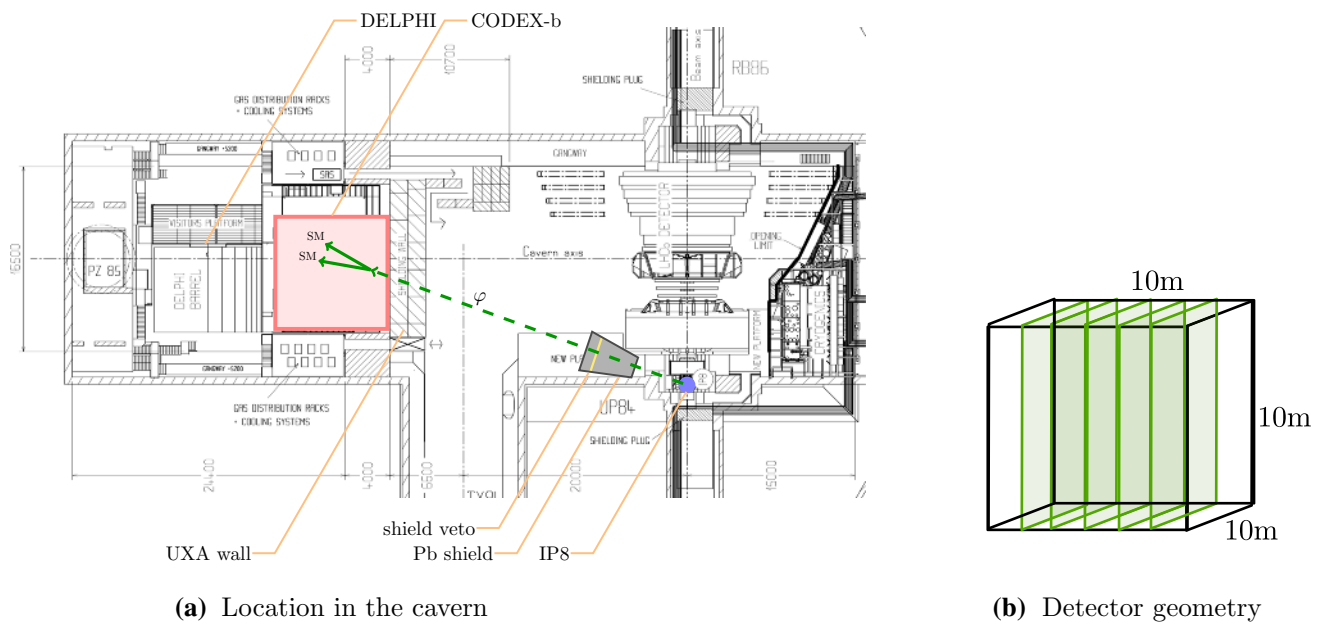


Fig. 3 Left: Layout of the LHCb experimental cavern UX85 at point 8 of the LHC [9], overlaid with the CODEX-b volume, as reproduced from Ref. [8]. Right: Schematic representation of the proposed detector geometry

precision timing resolution, any estimation or simulation of the net reconstruction efficiencies is expected to be reliable. These estimates must be ultimately validated by data-driven determinations from a demonstrator detector, which we call CODEX- β (see Sect. 4). Combined together, the baseline tracking and timing capabilities will permit mass reconstruction and particle identification for some benchmark scenarios.

The transverse location of the detector permits reliable background simulations based on well-measured SM transverse production cross-sections. The SM particle propagation through matter – necessary to simulate the response of the UXA radiation wall and the additional passive and active shielding – is also well understood for the typical particle energies generated in that pseudorapidity range. The proposed location behind the UXA radiation wall will also permit regular maintenance of the experiment, e.g. during technical or other stops. In addition to background simulations, the active veto and the ability to vary the amount of shielding over the detector acceptance permit LLP measurements or exclusions to be determined with respect to data-driven baseline measurements or calibrations of relevant backgrounds (see Sect. 3).

1.4 Search power, complementarities and unique features

Although ATLAS, CMS and LHCb were not explicitly designed with LLP searches in mind, they have been remarkably effective at probing a large region of the LLP parameter space (see [7, 14] for recent reviews). The main vari-

ables which provide the necessary discrimination for triggering and off-line background rejection are often the amount of energy deposited and/or the number of tracks connected to the displaced vertex. In most searches, the signal efficiency therefore drops dramatically for low mass LLPs, especially when they are not highly energetic (e.g. from Higgs decays.) For instance, the penetration of hadrons into the ATLAS or CMS muon systems, combined with a reduced trigger efficiency, attenuates the LHC reach for light LLPs, $m_{\text{LLP}} \lesssim 10 \text{ GeV}$, decaying in the muon systems.

Beam dump experiments such as SHiP [15–17], NA62 [18] in beam-dump mode, as well as forward experiments like FASER [19–21] evade this problem by employing passive and/or active shielding to fully attenuate the SM backgrounds. The LLPs are moreover boosted in a relatively narrow cone, and very high beam luminosities can be attained. This results in excellent reach for light LLPs that are predominantly produced at relatively low center-of-mass energy, such as a kinetically mixed dark photon. The main trade-off in this approach is, however, the limited partonic center-of-mass energy, which severely limits their sensitivity to heavier LLPs or LLPs primarily produced through heavy portals (e.g. Higgs decays).

Finally, proposals in pursuit of shielded, transverse, background-free detectors such as MATHUSLA [22, 23], CODEX-b [8] and AL3X [24] aim to operate at relatively low pseudorapidity η , but with far greater shielding compared to the ATLAS and CMS muon systems. This removes the background rejection and triggering challenges even for low mass LLPs, $m_{\text{LLP}} \lesssim 10 \text{ GeV}$, though at the expense of

a reduced geometric acceptance and/or reduced luminosity. Because of their location transverse from the beamline, they can access processes for which a high parton center-of-mass energy is needed, such as Higgs and Z production.

In this light, the regimes for which existing and proposed experiments have the most effective coverage can be roughly summarized as follows:

1. ATLAS and CMS: Heavy LLPs ($m_{\text{LLP}} \gtrsim 10 \text{ GeV}$) for all lifetimes ($c\tau \lesssim 10^7 \text{ m}$).
2. LHCb: Short to medium lifetimes ($c\tau \lesssim 1 \text{ m}$) for light LLPs ($0.1 \text{ GeV} \lesssim m_{\text{LLP}} \lesssim 10 \text{ GeV}$).
3. Forward/beam dump detectors (FASER, NA62, SHiP): Medium to long lifetime regime ($0.1 \lesssim c\tau \lesssim 10^7 \text{ m}$) for light LLPs ($m_{\text{LLP}} \lesssim \text{few GeV}$), for low \sqrt{s} production channels.
4. Shielded, transversely displaced detectors (MATHUSLA, CODEX-b, AL3X): Relatively light LLPs² ($m_{\text{LLP}} \lesssim 10\text{--}100 \text{ GeV}$) in the long lifetime regime ($1 \lesssim c\tau \lesssim 10^7 \text{ m}$), and high \sqrt{s} production channels.

In Fig. 4 we provide a visual schematic summarizing these LLP coverages, showing slices in the space of LLP mass, lifetime, and \sqrt{s} , that provide a sketch of the complementarity and unique features of various LLP search strategies and proposals. Relative to the existing LHC detectors, CODEX-b will be able to probe unique regimes of parameter space over a large range of well motivated models and portals, explored further in the Physics Case in Sect. 2. A more extensive discussion and evaluation of the landscape of LLP experimental proposals can be found in Refs. [7] and [25].

While the ambitiously sized ‘transverse’ detector proposals such as MATHUSLA and AL3X would explore even larger ranges of the parameter space, the more manageable and modest size of CODEX-b provides a substantially lower cost alternative with good LLP sensitivity. It also allows for the possibility of additional detector subsystems, such as precision tracking and calorimetry. Furthermore, the proximity of CODEX-b to the LHCb interaction point (IP8) and LHCb’s trigger-less readout (based on standardized and readily available technologies) makes it straightforward to integrate the detector into the LHCb readout for triggering and/or partial event reconstruction. This capability is not available to any other proposed LLP experiment at the LHC interaction points, and may prove crucial to authenticate any signals seen by CODEX-b. For a further discussion of the experimental design drivers and preliminary case studies of how different

detector capabilities can effect the sensitivity for different models, we refer to Sect. 5.

1.5 Timeline

The CODEX- β demonstrator detector is proposed for Run 3 and is therefore complementary in time to the other funded proposals such as FASER. In contrast, the full version of CODEX-b, as well as FASER2, SHiP, MATHUSLA, and AL3X are all proposed to operate in Runs 4 or 5 during the HL-LHC. We show the nominal timeline for CODEX- β and CODEX-b in Fig. 5. Results as well as design and construction lessons from CODEX- β are expected to inform the final design choices for the full detector, and may also inform the evolution of the schedule shown in Fig. 5. The modest size of CODEX-b, the accessibility of the DELPHI cavern, and the use of proven technologies in the baseline design, is expected to imply not only lower construction and maintenance costs but also a relatively short construction timescale. It should be emphasized that CODEX-b may provide complementary data both in reach and in time, at relatively low cost, to potential discoveries in other more ambitious proposals, should they be built, as well as to existing LHC experiments.

2 Physics case

2.1 Theory survey strategies

Long-lived particles occur generically in theories with a hierarchy of mass scales and/or couplings (see Sect. 1.1), such as the Standard Model and many of its possible extensions. This raises the question how best to survey the reach of any new or existing experiment in the theory landscape. Given the vast range of possibilities, injecting some amount of “theory prejudice” cannot be avoided. We therefore consider two complementary strategies to survey the theory space: (i) studying *minimal models* or “portals”, where one extends the Standard Model with a single new particle that is inert under all SM gauge interactions. The set of minimal modes satisfying this criteria is both predictive and relatively small – we restrict ourselves to the set of minimal models generating operators of dimension 4 or lower, as well as the well-motivated dimension 5 operators for axion-like particles. It is important to keep in mind, however, that minimal models are merely simplified models, meant to parametrize different classes of phenomenological features that may arise in more complete models. To mitigate this deficiency to some extent, we then also consider: (ii) studying a number of *complete models*, which are more elaborate but aim to address one or more of the outstanding problems of the Standard Model, such as the gauge hierarchy problem, the mechanism of baryogenesis, or the nature of dark matter. These complete models

² The degree to which each of these detectors can compete with ATLAS and CMS in the high mass regime, $m_{\text{LLP}} \gtrsim 10 \text{ GeV}$, depends on their angular acceptance and integrated luminosity. The larger volume MATHUSLA and AL3X proposals therefore typically remain more competitive with the main detectors for higher LLP masses than CODEX-b.

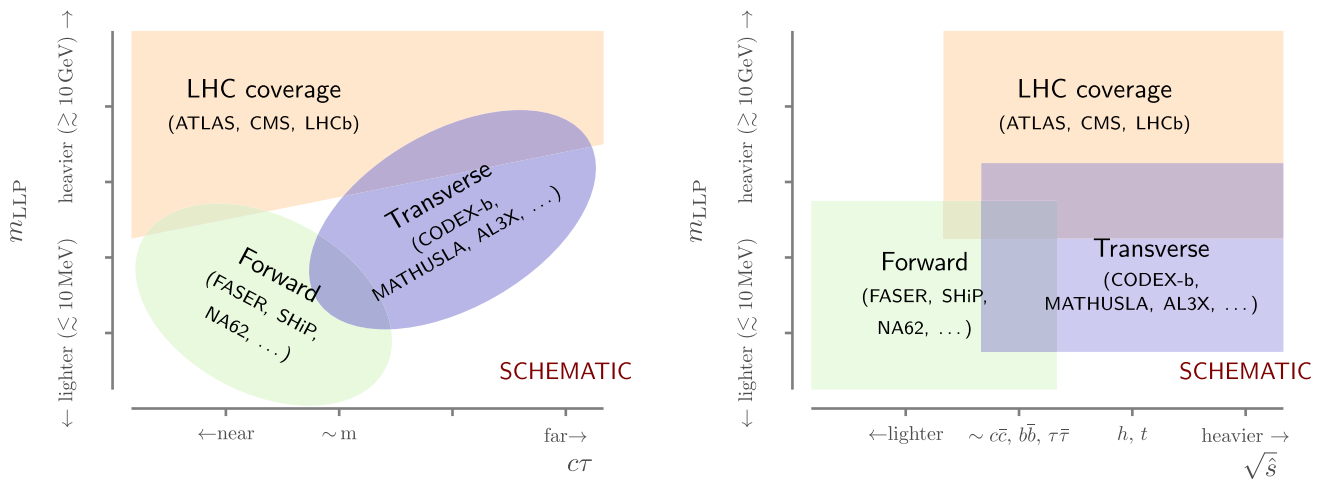


Fig. 4 Schematic summary of reach and coverage of various current, planned or proposed experiments in the LLP mass, lifetime, and \sqrt{s} space

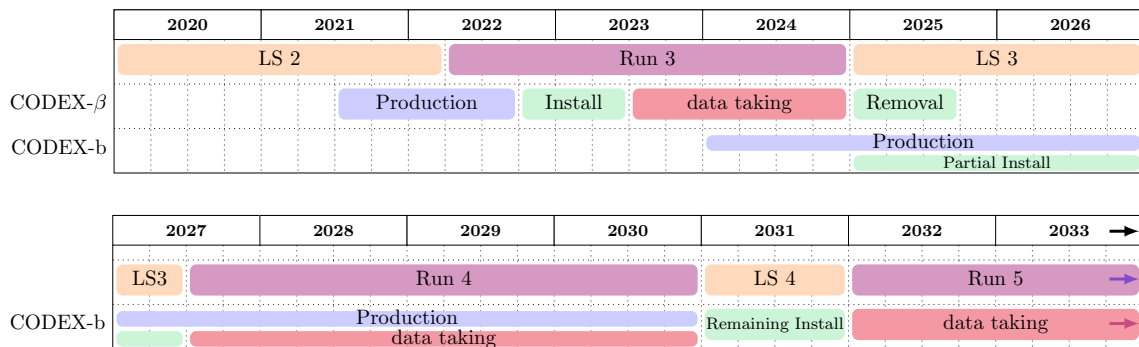


Fig. 5 Approximate production, installation and data-taking timelines for demonstrator (CODEX- β) and full (CODEX-b) detectors. (Updated according COVID to constraints.)

feature LLPs as a consequence of the proposed mechanisms introduced to solve these problems.

2.2 Novel studies

While many of the models surveyed below have been studied elsewhere and are recapitulated here, several of the studies in this section contain new and novel results, either correcting previous literature, recasting previous studies for the case of CODEX-b, or introducing new models not studied before. Specifically, we draw the reader’s attention to:

1. The axion-like particles (ALPs) minimal model (Sect. 2.3.3), which includes new contributions to ALP production from parton fragmentation. This can be very important in LHC collisions, significantly enhancing production estimates and consequent reaches, but was overlooked in previous literature.
2. The heavy neutral leptons (HNLs) minimal model (Sect. 2.3.4), which includes modest corrections to the HNL lifetime and τ branching ratios, compared to prior treatments.

3. The neutral naturalness complete model (Sect. 2.4.3), which is recast for CODEX-b from prior studies.
4. The coscattering dark matter complete model (Sect. 2.4.5), which contains LLPs produced through an exotic Z decay, and has not been studied previously.

2.3 Minimal models

The underlying philosophy of the minimal model approach is the fact that the symmetries of the SM already strongly restrict the portals through which a new, neutral state can interact with our sector. The minimal models can then be classified via whether the new particle is a scalar (S), pseudo-scalar (a), a fermion (N) or a vector (A'). In each case there are only a few operators of dimension 4 or lower (dimension 5 for the pseudo-scalar) which are allowed by gauge invariance. The most common nomenclature of the minimal models and their corresponding operators are

$$\text{Abelian hidden sector: } F_{\mu\nu}F'^{\mu\nu}, \quad hA'_\mu A'^\mu \quad (1a)$$

$$\text{Scalar-Higgs portal: } S^2H^\dagger H, \quad SH^\dagger H \quad (1b)$$

$$\text{Heavy neutral leptons: } \tilde{H}\tilde{L}N \quad (1c)$$

$$\begin{aligned} \text{Axion-like particles: } & \partial^\mu a \bar{\psi} \gamma_\mu \gamma_5 \psi, \quad a W_{\mu\nu} \tilde{W}^{\mu\nu}, \\ & a B_{\mu\nu} \tilde{B}^{\mu\nu}, \quad a G_{\mu\nu} \tilde{G}^{\mu\nu} \end{aligned} \quad (1d)$$

where $F'^{\mu\nu}$ is the field strength operator corresponding to a $U(1)$ gauge field A' , H is the SM Higgs doublet, and h the physical, SM Higgs boson.³ Where applicable, we consider cases in which a different operator is responsible for the production and decay of the LLP, as summarized in Fig. 6. Note that the $h A'_\mu A'^\mu$ and $S^2 H^\dagger H$ operators respect a \mathbb{Z}_2 symmetry for the new fields and will not induce a decay for the LLP on their own.

For the axion portal, the ALP can couple independently to the $SU(2)$ and $U(1)$ gauge bosons. In the infrared, only the linear combination corresponding $a F \tilde{F}$ survives, though the coupling to the massive electroweak bosons can contribute to certain production modes. Moreover, the gauge operators mix into the fermionic operators through renormalization group running. Classifying the models according to production and decay portals obscures this key point⁴, and we have therefore chosen to present the model space for the ALPs in Fig. 6 in terms of UV operators. Once the UV boundary condition at a scale Λ is given, such a choice fully specifies both the ALP production and the decay modes, which often proceed via a combination of the listed operators.

2.3.1 Abelian hidden sector

The Abelian hidden sector model [26–28] is a simple extension of the Standard Model, consisting of an additional, massive $U(1)$ gauge boson (A') and its corresponding Higgs boson (H') (see e.g. [29–35] for an incomplete list of References containing other models with similar phenomenology). The A' and the H' can mix with respectively the SM photon [36, 37] and Higgs boson, each of which provide a portal into this new sector. In the limit where the H' is heavier than the SM Higgs, it effectively decouples from the phenomenology, such that only the operators in (1a) remain in the low energy effective theory.

The mixing of the A' with the photon through the $F_{\mu\nu} F'^{\mu\nu}$ operator can be rewritten as a (millicharged) coupling of the A' to all charged SM fermions. In the limit that the $h A'_\mu A'^\mu$ coupling is negligible (along with higher dimension operators, such as $h F'_{\mu\nu} F'^{\mu\nu}$), the mixing with the photon alone

can induce both the production and decay of the A' in a correlated manner, which has been studied in great detail (see e.g. [25] and references therein). CODEX-b has no sensitivity to this scenario, because the large couplings required for sufficient production cross-sections imply an A' lifetime that is too short for any A' 's to reach the detector. However, the LHCb VELO and various forward detectors are already expected to greatly improve the reach for this scenario [19, 38–45].

The $h A'_\mu A'^\mu$ operator, by contrast, is controlled by the mixing of the H' with the SM Higgs. This can arise from the kinetic term $|D_\mu H'|^2$, with $\langle H' \rangle \neq 0$ and $H - H'$ mixing. This induces the exotic Higgs decay $h \rightarrow A' A'$. In the limit where the mixing with the photon is small, this becomes the dominant production mode for the A' , which then decays through the kinetic mixing portal to SM states. CODEX-b would have good sensitivity to this mixing due to its transverse location, with high \sqrt{s} . Importantly, the coupling to the Higgs and the mixing with the photon are independent parameters, so that the lifetime of the A' and the $h \rightarrow A' A'$ branching ratio are themselves independent, and therefore convenient variables to parameterize the model. Figure 7 shows the reach of CODEX-b for two different values of the A' mass, as done in Ref. [8] (see commentary therein), as well as the reach of AL3X [24] and MATHUSLA [22].

For ATLAS and CMS, the muon spectrometers have the largest fiducial acceptance as well as the most shielding, thanks to the hadronic calorimeters. The projected ATLAS reach for 3 ab^{-1} was taken from Ref. [46] for the low mass benchmark. In Ref. [47] searches for one displaced vertex (1DV) and two displaced vertices (2DV) were performed with 36.1 fb^{-1} of 13 TeV data. We use these results to extrapolate the reach of ATLAS for the high mass benchmark to the full HL-LHC dataset, where the widths of the bands corresponds to a range between a ‘conservative’ and ‘optimistic’ extrapolation for each of the 1DV and 2DV searches. Concretely, the 1DV search in Ref. [47] is currently background limited, with comparable systematic and statistical uncertainties. For our optimistic 1DV extrapolation we assume that the background scales linearly with the luminosity and that the systematic uncertainties can be made negligible with further analysis improvements. This corresponds to a rescaling of the current expected limit with $\sqrt{36.1 \text{ fb}^{-1}/3000 \text{ fb}^{-1}}$. For our conservative 1DV extrapolation we assume the systematic uncertainties remain the same, with negligible statistical uncertainties. This corresponds to an improvement of the current expected limit with roughly a factor of ~ 2 . The 2DV search in Ref. [47] currently has an expected background of 0.027 events, which implies ~ 3 expected background events, if the background is assumed to scale linearly with the luminosity. For our optimistic 2DV extrapolation we assume the search remains background free, which corresponds to a rescaling of the current expected limits with

³ The second operator in Eq. (1a) is strictly speaking not gauge invariant, but can be trivially generated by the kinetic term of a heavy dark scalar charged under the $U(1)$ that acquires a vacuum expectation value (VEV) and mixes with the SM Higgs (see e.g. [26–28]).

⁴ An example of when such identification is not as straightforward may be provided by the case of ALP coupled to photons, where the main production mechanism relevant for CODEX-b is via an effective ALP coupling to quarks.

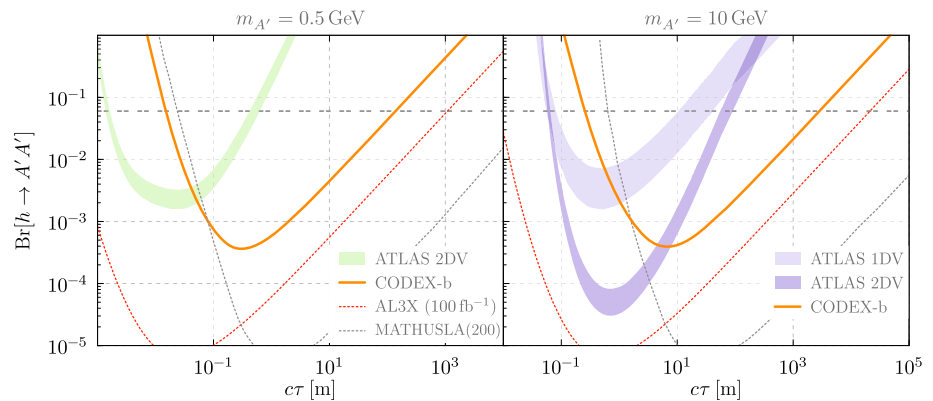
Vector (A')	$hA'A'$	$F'F$				
$F'F$	Fig. 7	no reach				
Scalar (S)	$SH^\dagger H$	$S^2H^\dagger H$				
$SH^\dagger H$	Fig. 9a	Fig. 9b				
HNL (N)	$\tilde{H}\tilde{L}N$					
$\tilde{H}\tilde{L}N$	Fig. 14					
ALP (a)	$\partial_\mu a \bar{q}\gamma^\mu \gamma^5 q$	$a\tilde{G}G$	$a\tilde{F}F$	$a(W\tilde{W} - B\tilde{B})$		
	Fig. 11	Fig. 12	pending	pending		

	Production portal
	Decay portal
	UV operator

Fig. 6 Minimal model tabular space formed from either: production (green) and decay (blue) portals, where well-defined by symmetries or suppressions; or, UV operators (orange), where either the production and decay portal may involve linear combinations of operators under RG evolution or field redefinitions. Each table cell corresponds to a

minimal model: cells for which the CODEX-b reach is known refer to the relevant figure in this document; cells denoted ‘pending’ indicate a model that may be probed by CODEX-b, but no reach projection is presently available

Fig. 7 Reach for $h \rightarrow A'A'$, as computed in Ref. [8]. Shaded bands refer to the optimistic and conservative estimates of the ATLAS sensitivity [46,47] for 3 ab^{-1} , as explained in the text. The horizontal dashed line represents the estimated HL-LHC limit on the invisible branching fraction of the Higgs [48]. The MATHUSLA reach is shown for its $200\text{ m} \times 200\text{ m}$ configuration with 3 ab^{-1} ; for AL3X 100 fb^{-1} of integrated luminosity was assumed



$36.1\text{ fb}^{-1}/3000\text{ fb}^{-1}$. For the conservative 2DV extrapolation we assume 10 expected and observed background events, leading to a slightly weaker limit than with the background free assumption.

Upon rescaling $c\tau$ to account for difference in boost distributions, the maximum CODEX-b reach is largely insensitive to the mass of the A' , modulo minor differences in reconstruction efficiency for highly boosted particles (see Sect. 5.1). This is not the case for ATLAS and CMS, where higher masses generate more activity in the muon spectrometer, which helps greatly with reducing the SM backgrounds.

2.3.2 Scalar-Higgs portal

The most minimal extension of the SM consists of adding a single, real scalar degree of freedom (S). Gauge invariance restricts the Lagrangian to

$$\mathcal{L} \supset A_S SH^\dagger H + \frac{\lambda}{2} S^2 H^\dagger H + \dots \quad (2)$$

where the ellipsis denotes higher dimensional operators, assumed to be suppressed. This minimal model is often referred to as simply the ‘‘Higgs portal’’ in the literature,

though the precise meaning of the latter can vary depending on the context. LHCb has already been shown to have sensitivity to this model [49,50], and CODEX-b would greatly extend its sensitivity into the small coupling/long lifetime regime.

The parameter A_S can be exchanged for the mixing angle, $\sin\theta$, of the S with the physical Higgs boson eigenstate. In the mass eigenbasis, the new light scalar therefore inherits all the couplings of the SM model Higgs: Mass hierarchical couplings with all the SM fermions, as well couplings to photons and gluons at one loop. All such couplings are suppressed by the small parameter $\sin\theta$. The couplings induced by Higgs mixing are responsible not only for the decay of S [51,52,52–55], but also contribute to its production cross-section. Concretely, for $m_K < m_S < m_B$, the dominant production mode is via the $b \rightarrow s$ penguin in Fig. 8a [56–58], because S couples most strongly to the virtual top quark in the loop. If the quartic coupling λ is non-zero, the rate is supplemented by a penguin with an off-shell Higgs boson, shown in Fig. 8b [59], as well as direct Higgs decays, shown in Fig. 8c.

In Fig. 9 we show the reach of CODEX-b taking two choices of λ , following [25]: (i) $\lambda = 0$, corresponding to

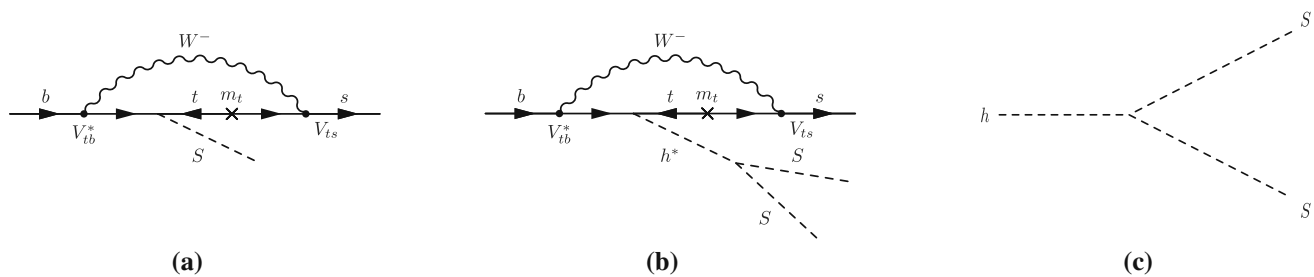


Fig. 8 Diagrams responsible for S production in a minimal extended Higgs sector. **a** is proportional to the mixing between S and Higgs, $\sin^2 \theta$, while **b** and **c** are proportional to the square of the quartic coupling, λ^2

the most conservative scenario, in which the production rate is smallest; (ii) $\lambda = 1.6 \times 10^{-3}$ was chosen such that the $\text{Br}[h \rightarrow SS] = 0.01$.⁵ The latter roughly corresponds to the future reach for the branching ratio of the Higgs to invisible states. In this sense it is the most optimistic scenario that would not be probed already by ATLAS and CMS. The reach for other choices of λ therefore interpolates between Fig. 8a and Fig. 8b. Also shown are the limits from LHCb [49, 50] and CHARM [60], and projections for MATHUSLA [61], FASER2 [62], SHiP [63], AL3X [24] and LHCb, where for the latter we extrapolated the limits from [49, 50], assuming (optimistically) that the large lifetime signal region remains background free with the HL-LHC dataset.

The scalar-Higgs portal is, by virtue of its minimality, very constraining as a model. When studying LLPs produced in B decays, it is therefore worthwhile to relax its assumptions, in particular relaxing the full correlation between the lifetime and the production rate – the $b \rightarrow sS$ branching ratio – as is the case in a number of non-Minimally Flavor Violating (MFV) models (see e.g. [64–67]). Fig. 10 shows the CODEX-b reach in the $b \rightarrow sS$ branching ratio for a number of benchmark LLP mass points, as done in Ref. [8]. The LHCb reach and exclusions are taken and extrapolated from Refs. [49, 50], assuming 30% (10%) branching ratio of $S \rightarrow \mu\mu$ for the 0.5 GeV (1 GeV) benchmark (see Ref. [8]). Also shown are the current and projected limits for $B \rightarrow K^{(*)}\nu\nu$ [68, 69]. A crucial difference compared to LHCb is that the CODEX-b reach depends only on the total branching ratio to charged tracks, rather than on the branching ratio to muons.

Interestingly, the CODEX- β detector proposed for Run 3 (see Sect. 4) may already have novel sensitivity to the $b \rightarrow sS$ branching ratio, as shown in Fig. 29. This reach is estimated under the requirement that the number of tracks in the final state is at least four, in order to control relevant backgrounds (see Sect. 3). A more detailed discussion of this reach is reserved for Sect. 4.

⁵ In the specific context of this minimal model, this size of quartic implies that m_S is rather severely fine-tuned for $m_S \lesssim 10$ GeV.

2.3.3 Axion-like particles

Axion-like particles (ALPs) are pseudoscalar particles coupled to the SM through dimension-5 operators. They arise in a variety of BSM models and when associated with the breaking of approximate Peccei–Quinn-like symmetries they tend to be light. Furthermore, their (highly) suppressed dimension-5 couplings naturally renders them excellent candidates for LLP searches. The Lagrangian for an ALP, a , can be parameterized as [70]

$$\begin{aligned} \mathcal{L} \supset & \frac{1}{2} (\partial_\mu a)^2 - \frac{1}{2} m_a^2 a^2 + \frac{c_q^{ij}}{2\Lambda} (\partial_\mu a) \bar{q}_i \gamma^\mu \gamma^5 q_j \\ & + \frac{c_\ell^{ij}}{2\Lambda} (\partial_\mu a) \bar{\ell}_i \gamma^\mu \gamma^5 \ell_j \\ & + \frac{4\pi\alpha_s c_G}{\Lambda} a G_{\mu\nu}^a \tilde{G}^{a,\mu\nu} + \frac{4\pi\alpha_2 c_W}{\Lambda} a W_{\mu\nu}^a \tilde{W}^{a,\mu\nu} \\ & + \frac{4\pi\alpha_1 c_B}{\Lambda} a B_{\mu\nu} \tilde{B}^{\mu\nu} + \dots \end{aligned} \quad (3)$$

where $\tilde{G}_{\mu\nu} = 1/2 \epsilon_{\mu\nu\rho\sigma} G^{\rho\sigma}$. The couplings to fermions do not have to be aligned in flavor space with the SM Yukawas, leading to interesting flavor violating effects. The gauge operators mix into the fermionic ones at 1-loop, and therefore in choosing a benchmark model one needs to specify the values of these couplings as a UV boundary condition at a scale Λ . In the following we will focus on the same benchmark models chosen in the Physics Beyond Colliders (PBC) community study [25] based on the ALP coupling to photons (“BC9”, defined as $c_W + c_B \neq 0$), universally to quarks and leptons (“BC10”, $c_q^{ij} = c \delta^{ij}$, $c_\ell^{ij} = c \delta^{ij}$, $c \neq 0$) and to gluons (“BC11”, $c_G \neq 0$). Another interesting benchmark to consider is the so-called photophobic ALP [71], in which the ALP only couples to the $SU(2) \times U(1)$ gauge bosons such that it is decoupled from the photons in the UV and has highly suppressed photon couplings in the IR.

CODEX-b is expected to have a potentially interesting reach for all these cases. This is true with the nominal design provided the ALP has a sizable branching fraction into visible final states, while for ALPs decaying to photons one would

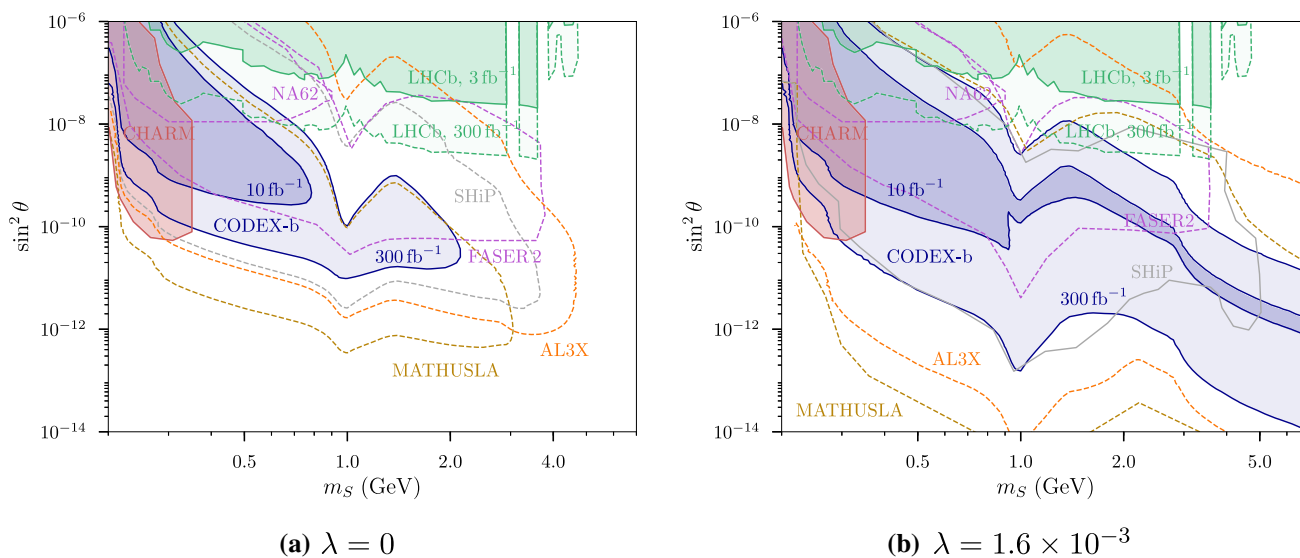


Fig. 9 Reach of CODEX-b as a function of $\sin^2 \theta$, for two representative values of λ . The MATHUSLA reach is shown for its 200 m \times 200 m configuration for 3 ab^{-1} ; for AL3X 100 fb^{-1} of integrated luminosity was assumed

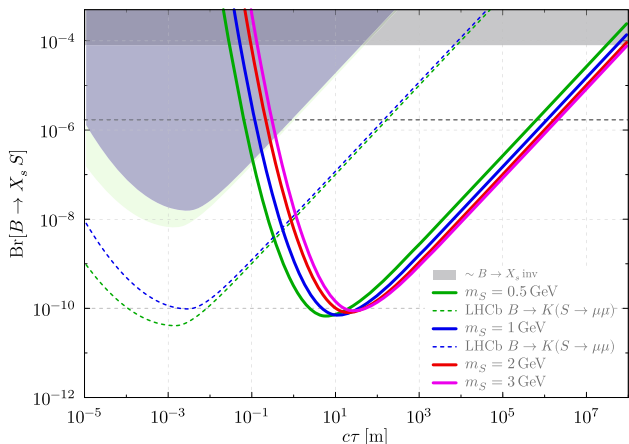


Fig. 10 Reach of CODEX-b for LLPs produced in B -meson decays, in a non-minimal model. Also shown is the current (shaded) and projected (dashed) reach for: LHCb via $B \rightarrow K^{(*)}(S \rightarrow \mu\mu)$, for $m_S = 0.5$ GeV (green) and $m_S = 1$ GeV (blue), assuming a muon branching ratio of 30% and 10%, respectively (see Ref. [8]); and for $B \rightarrow K^{(*)} + \text{inv.}$ (gray)

require a calorimeter element, as discussed below in Sect. 5.3. In this section, we will present the updated reach plots for BC10 and BC11 and leave the ALP with photon couplings (BC9) and the photophobic case for future study.

ALPs coupled to quark and gluons can be copiously produced at the LHC even though their couplings are suppressed enough to induce macroscopic decay lengths. They therefore provide an excellent target for LLP experiments such as CODEX-b. Based on the fragmentation of partons to hadrons in LHC collisions, we can divide the ALP production into four different mechanisms:

1. radiation during partonic shower evolution (using the direct ALP couplings to quarks and/or gluons),
2. production during hadronization of quarks and gluons via mixing with $(J^{PC} =)0^{-+} \bar{q}q$ operators (dominated at low ALP masses via mixing with π^0, η, η'),
3. production in hadron decays via mixing with neutral pseudoscalar mesons, and
4. production in flavor-changing neutral current bottom and strange hadron decays, via loop-induced flavor-violating penguins.

The last mechanism has been already considered extensively in the literature. The ALP production probability scales parametrically as $(m_t/\Lambda)^2$ and is proportional to the number of strange or b -hadrons produced. In general, the population of ALPs produced by this mechanism is not very boosted at low pseudorapidities. For the PBC study, it was the only production mechanism considered for BC10, and it was included in BC11.

The second and third mechanism are related as they both incorporate how the ALP couples to low energy QCD degrees of freedom. Conventionally the problem is rephrased into ALP mixing with neutral pseudoscalar mesons. This production is parametrically suppressed by $(f_\pi/\Lambda)^2$ and it quickly dies off for ALP masses much above 1 GeV. The population of ALPs produced by these mechanisms is not very boosted at low pseudorapidities, while the forward experiments will have access to very energetic ALPs. Compared to the PBC study, we treat separately the two cases of hadronization and hadron decays as they give rise to populations of ALPs with different energy distributions, and include them both in BC10 and BC11.

Finally, the first mechanism listed above has been so far overlooked in the literature. However, emission in the parton shower can be the most important production mechanism at transverse LHC experiments such as CODEX-b. Emission of (pseudo)scalars is expected to exhibit neither collinear nor soft enhancements, such that ALPs emitted in the shower may then carry an $\mathcal{O}(1)$ fraction of the parent energy and can be emitted at large angles.

For the case of quark-coupled ALPs (BC10), emission in the parton shower is suppressed by the quark mass – a consequence of the soft pion theorem – i.e. by m_q^2/Λ^2 (or by loop factors to the induced gluon coupling, as below). The shower contribution may nevertheless still dominate at high ALP masses, where the other production mechanisms are forbidden by phase space or kinematically suppressed. For gluon-coupled ALPs (BC11), however, no such suppression arises in the shower. In a parton shower approximation, the ALP emission is attributed to a single parton with a given probability: While interference terms between ALP emissions from adjacent legs – e.g. in $g \rightarrow gga$ – cannot be neglected, such an approximation still captures the bulk of the production, even when the ALP is not emitted in the soft limit.

The parton shower approximation greatly simplifies the description of ALP emission, allowing the implementation in existing Monte Carlo tools. In this approximation, the probability for a parton to fragment into an ALP scales parametrically as Q^2/Λ^2 , with Q of the order of the virtuality of the parent parton. For example, the $g \rightarrow ga$ splitting function

$$P_{g \rightarrow ag}(t, z) = \pi \alpha_s c_G^2 \frac{t}{\Lambda^2} \left(1 - \frac{m_a^2}{t}\right)^2 \quad (4)$$

where $t = Q^2$. While the population of partons with large energies is much smaller than the final number of hadrons, the production rate is enhanced by a large $\mathcal{O}(Q^2/f_\pi^2)$ factor, compared to the second and third mechanisms. In LHC collisions this is sufficient to produce a large population of energetic ALPs at low pseudorapidities, with boosts exceeding 10^3 for ALP masses in the 0.1–1 GeV range. The CODEX-b reach can therefore be extended to higher ALP masses and larger couplings compared to previous estimates if very collimated LLP decays can be detected.

We estimate these production mechanisms using `Pythia 8`, with the code modified to account for the production of ALPs during hadronization. We include ALP production in decays by extending its decay table in such a way that for each decay mode containing a π^0, η, η' meson in the final state, we add another entry with the meson substituted by the ALP. The branching ratio is rescaled by the ALP mixing factor and phase space differences.

The ALP production from the shower is computed by navigating through the generated QCD shower history and for

each applicable parton an ALP is generated by re-decaying that parton with a weight: The ratio of the ALP branching (integrated) probability over the total (SM+ALP) (integrated) probabilities. This is correct for time-like showers in the limit that the ALP branching probability is small, because in this limit the branching scale is still controlled by the SM Sudakov factor. This procedure is not applicable to space-like showers, without also incorporating information from parton distribution functions. Such space-like showers, however, provide only a sub-leading contribution to transverse production, i.e. at the low pseudorapidities for the CODEX-b acceptance, and we therefore neglect them. For forward experiments at the LHC, such as FASER2, we do not include any shower contribution in the reach estimates, since a more complete treatment is required in order to fully estimate ALP production at high pseudorapidities, and the effect is expected to be at most $\mathcal{O}(1)$. For the case of a fermion-coupled ALP we include both the emission from heavy quark lines, proportional to $(m_q c_q/\Lambda)^2$, and from loop-induced coupling to gluons, taking $c_G = N_f c_q/32\pi^2$ in Eq. (4) above [70], where N_f is the number of flavors. Further details will be given in upcoming work.

The reach predictions are shown in Fig. 11 for a fermion-coupled ALP (BC10) and in Fig. 12 for an ALP coupled to gluons (BC11), for the case of the nominal – i.e. tracker-only – CODEX-b design. In this case, an ALP decaying only to neutral particles such as photons is invisible, and highly boosted ALPs may decay to merged tracks, such that the signature resembles more closely a single appearing track inside the detector volume. For such a signature, the CODEX-b baseline design is not background-free; we use the background estimates presented in Table 3 (see Sect. 3 below), corresponding to 50 events of background in the entire detector in 300 fb^{-1} . The CODEX-b reach with a calorimeter option (shown here as a dashed line) is further discussed in Sect. 5.3.

The MATHUSLA estimates and the CHARM exclusion in Fig. 12 have been recomputed, while all the other curves have been taken from [25], after rescaling them to the appropriate lifetime and branching ratio expressions used in our plots. For MATHUSLA we used the $200 \text{ m} \times 200 \text{ m}$ configuration and assumed that a floor veto for upward going muons entering the decay volume is available with a rejection power of 10^5 . Based on the estimates of 10^7 upward going muons [23, 72], we therefore used 100 events as the background for unresolved highly boosted ALPs.

For the case of the fermion-coupled ALP we have further improved the lifetime and branching ratio calculations compared to those used in Refs. [25, 73], by including the partial widths of ALP into light QCD degrees of freedom, using the same procedure as in Ref. [74]. The result is shown in Fig. 13. In particular in the $1 \lesssim m_a \lesssim 3 \text{ GeV}$ range, for a given coupling the ALP lifetime is $\mathcal{O}(10)$ smaller than previ-

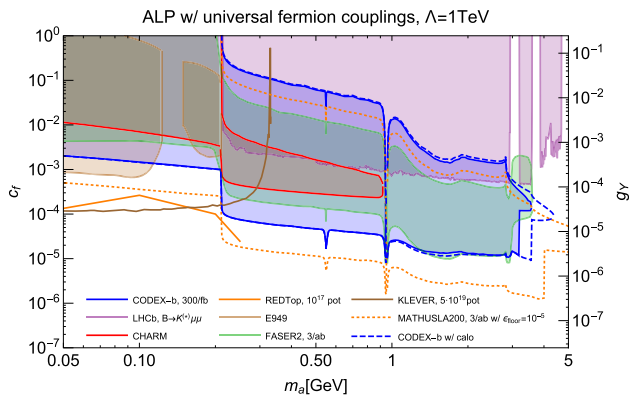


Fig. 11 Reach of CODEX-b for fermion-coupled ALPs. The vertical axis on the left corresponds to the couplings defined in Eq. (3), while the one on the right to the normalization used in the PBC study [25]. The baseline (tracker only) CODEX-b design is shown as solid, while the gain by a calorimeter option is shown as dashed. All the curves for the other experiments except MATHUSLA are taken from [25], after rescaling with to the different lifetime/branching ratio calculation used here. The MATHUSLA reach is based on our estimates, see text for details

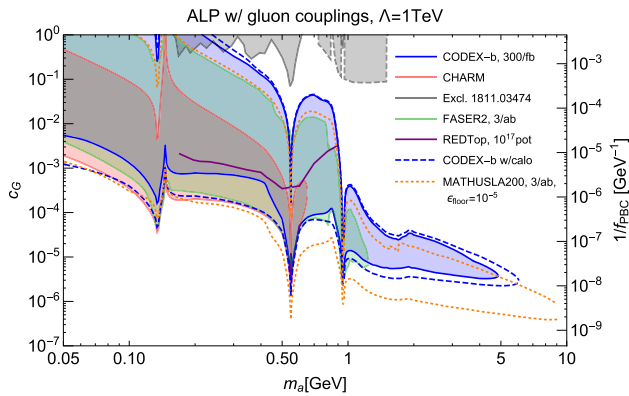


Fig. 12 Reach of CODEX-b for gluon-coupled ALPs. The vertical axis on the left corresponds to the couplings defined in Eq. (3), while the one on the right to the normalization used in the PBC study [25]. The baseline (tracker only) CODEX-b design is shown as solid, while the gain by a calorimeter option is shown as dashed. See Fig. 34 for further information about how the CODEX-b reach changes with different detector designs. FASER2 and REDTop curves are taken from [25], after rescaling to the different lifetime/branching ratio calculation used here. The CHARM curve has been recomputed with the same assumptions used for the CODEX-b curve. The MATHUSLA reach is based on our estimates, see text for details

ously assumed, and the decays are mostly to hadrons instead of muon pairs.

2.3.4 Heavy neutral leptons

Heavy neutral leptons (HNLs) may generically interact with the SM sector via the lepton Yukawa portal, mediated by the marginal operator $\bar{L}_i \tilde{H} N$, or may feature in a range of simplified NP models coupled to the SM via various higher-

dimensional operators. In the $m_N \sim 0.1\text{--}10\text{ GeV}$ regime, that we consider below, these models can be motivated e.g. by explanations for the neutrino masses [75], the νMSM [76, 77], dark matter theories [78], or by models designed to address various recent semileptonic anomalies [79–81].

UV completions of SM–HNL operators typically imply an active-sterile mixing $\nu_\ell = U_{\ell j} \nu_j + U_{\ell N} N$, where ν_j and N are mass eigenstates, and U is an extension of the PMNS neutrino mixing matrix to incorporate the active-sterile mixings $U_{\ell N}$. If $|U_{\ell N}|$ are the dominant couplings of N to the SM and N has negligible partial width to any hidden sector, then the N decay width is electroweak suppressed, scaling as $\Gamma \sim G_F^2 |U_{\ell N}|^2 m_N^5$. Because the mixing $|U_{\ell N}|$ can be very small, N can then become long-lived. We assume hereafter for the sake of simplicity that N couples predominantly to only a single active neutrino flavor, i.e.

$$U_{\ell_i N} \gg U_{\ell_j \neq \ell_i N}, \tag{5}$$

and refer to ℓ as the ‘valence’ lepton.

The width of the HNL can be expressed as

$$\frac{\Gamma_N}{s |U_{\ell_i N}|^2} = \sum_M (\Gamma_{\nu_i M} + \Gamma_{\ell_i M}) + \sum_j (\Gamma_{\ell_i \ell_j \nu_j} + \Gamma_{\nu_i \ell_j \ell_j}) + \sum_{q, q'} \Gamma_{\ell_i q q'} + \sum_q \Gamma_{\nu_i q q} + \Gamma_{\nu_i \nu \nu}, \tag{6}$$

where $s = 1$ ($s = 2$) for a Dirac (Majorana) HNL and the final state M corresponds to a single kinematically allowed (ground-state) meson. Specifically, M considers: charged pseudoscalars, π^\pm, K^\pm ; neutral pseudoscalars π^0, η, η' ; charged vectors, $\rho^\pm, K^{*\pm}$; and neutral vectors, ρ^0, ω, ϕ . For $m_N > 1.5\text{ GeV}$, we switch from the exclusive meson final states to the inclusive decays widths $\Gamma_{\ell_i q q'}$ and $\Gamma_{\nu_i q q}$, which are disabled below 1.5 GeV. Expressions for each of the partial widths may be found in Ref. [82]; each is mediated by either the W or Z , generating long lifetimes for N once one requires $U_{\ell N} \ll 1$. Apart from the 3ν , and some fraction of the νM and $\nu q q$ (e.g. $\nu \pi^0 \pi^0$) decay modes, all the N decays involve two or more tracks, so that the decay vertex will be reconstructible in CODEX-b, up to $\mathcal{O}(1)$ reconstruction efficiencies. We model the branching ratio to multiple tracks by considering the decay products of the particles produced. Below 1.5 GeV, we consider the decay modes of the meson M to determine the frequency of having 2 or more charged tracks; above 1.5 GeV where $\nu q q$ production is considered instead of exclusive single meson modes, we conservatively approximate the frequency of having two or more charged tracks as 2/3.

HNLs may be abundantly produced by leveraging the large $b\bar{b}$ and $c\bar{c}$ production cross-section times branching ratios into semileptonic final states. In particular, for $0.1\text{ GeV} \lesssim m_N \lesssim 3\text{ GeV}$, the dominant production modes are the typically fully inclusive $c \rightarrow s\ell N$ and $b \rightarrow c\ell N$. In order to

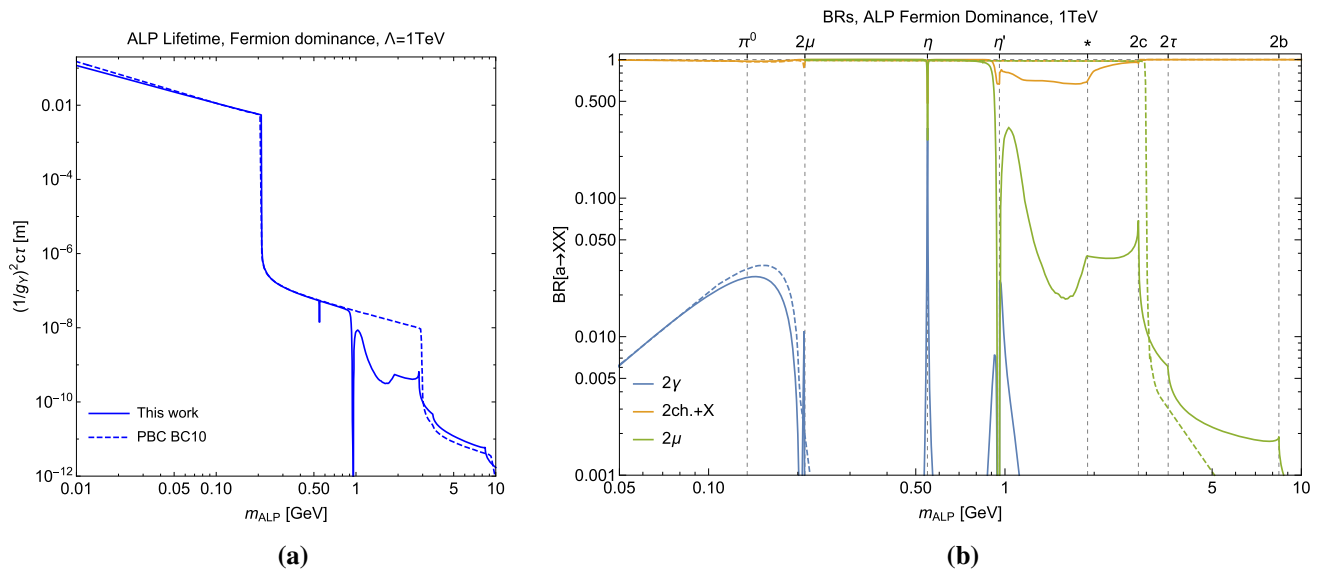


Fig. 13 Lifetime (left) and branching ratios (right) for an ALP coupled to fermions, used in Fig. 11. For comparison we plot as dashed lines the corresponding values used for BC10 in the PBC document

capture mass threshold effects, production from these heavy flavor semileptonic decays is estimated by considering a sum of exclusive modes. The hadronic form factors are treated as constants: An acceptable estimate for these purposes, as corrections are expected to be small, $\sim \Lambda_{\text{qcd}}/m_{c,b}$. In certain kinematic regimes, the on-shell (Drell–Yan) $W^{(*)} \rightarrow \ell N$ or $Z^{(*)} \rightarrow \nu N$ channels can become important, as can the two-body $D_s \rightarrow \ell N$ and $B_c \rightarrow \ell N$ decays (a prior study in Ref. [83] for HNLs at CODEX-b neglected the latter contributions).

In our reach projections, we assume the production cross-section $\sigma(b\bar{b}) \simeq 500 \mu\text{b}$ and $\sigma(c\bar{c})$ is taken to be 20 times larger, based on FONLL estimates [84,85]. The EW production cross-sections used are $\sigma(W \rightarrow \ell\nu) \simeq 20\text{nb}$ and $\sum_j \sigma(Z \rightarrow \nu_j \nu_j) \simeq 12\text{nb}$ [86]. The $\sigma(D_s)/\sigma(D)$ production fraction is taken to be 10% [87,88], and we assume a production fraction $\sigma(B_c)/\sigma(B) \simeq 2 \times 10^{-3}$ [89,90].

In the case of the τ valence lepton, with $m_N < m_\tau$, the HNL may be produced not only in association with the τ , but also as its daughter. For example, both $b \rightarrow c\tau N$ and $b \rightarrow c(\tau \rightarrow N e \nu_e)\nu$ are comparable production channels. When kinematically allowed, we approximate this effect by including for the valence τ case an additional factor of $1 + \text{BR}(\tau \rightarrow N + X)/|U_\tau|^2$, where $\text{BR}(\tau \rightarrow N + X)$ is the HNL mass dependent BR of the tau into a valence τ HNL plus anything [82]. HNL production from Drell–Yan τ 's is also included, but typically sub-leading: The relevant production cross-section is estimated with MadGraph [91] to be $\sigma(\tau_{\text{DY}}) \simeq 37\text{nb}$.

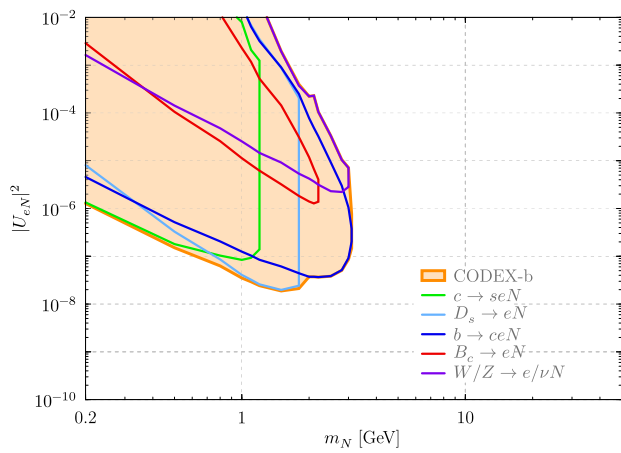
The projected sensitivity of CODEX-b to HNLs in the single flavor mixing regime is shown in Fig. 14. The breakdown in terms of the individual production modes is shown in the

left panels, while the right panels compare CODEX-b sensitivity versus constraints from prior experiments, including BEBC [92], PS191 [93], CHARM [94–96], JINR [97], and NuTeV [98], DELPHI [99], and ATLAS [100] (shown collectively by gray regions). Also included are projected reaches for other current or proposed experiments, including NA62 [101], DUNE [102], SHiP [17], FASER [103], and MATHUSLA [83]. We adopt the Dirac convention in all our reach projections; the corresponding reach for the Majorana case is typically almost identical, though relevant exclusions may change.⁶

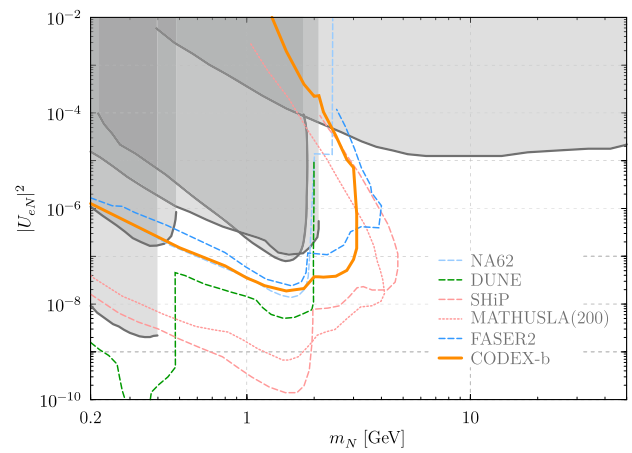
2.4 Complete models

The LLP search program at the LHC is extensive and rich. In the context of complete models, it has been driven so far primarily by searches for weak scale supersymmetry, along with searches for dark matter, mechanisms of baryogenesis, and hidden valley models. In this section, we review the part of the theory space relevant for CODEX-b, which is typically the most difficult to access with the existing experiments. A comprehensive overview of all known possible signatures is neither feasible nor necessary, the latter thanks to the inclu-

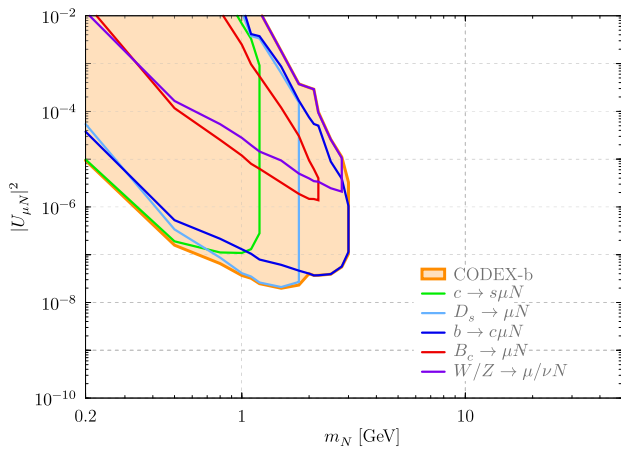
⁶ With regard to prior measurements, the PS191 [93] and CHARM [96] measurements are explicitly quoted for the Dirac HNL case, and the DELPHI measurements [99] appear also to be implicitly for a Dirac HNL. Recent ATLAS [100] and CMS [104] measurements are sensitive to prompt Majorana HNLs decays only via trilepton searches that reject opposite-sign same-flavor lepton pairs that would be produced in the Dirac case; the lepton number conserving Dirac case is, however, probed via a displaced decay search in Ref. [100]. The convention of prior CHARM measurements [94,95], as well as for BEBC, JINR and NuTeV [92,97,98] are unclear.



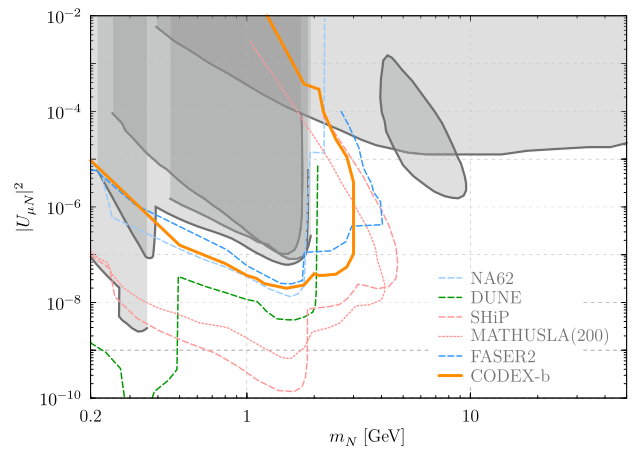
(a) $U_{eN} \gg U_{\mu N}, U_{\tau N}$ (Channels)



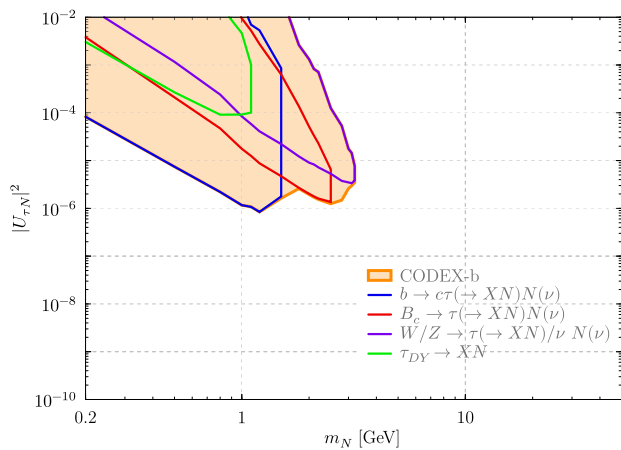
(b) $U_{eN} \gg U_{\mu N}, U_{\tau N}$ (Combined)



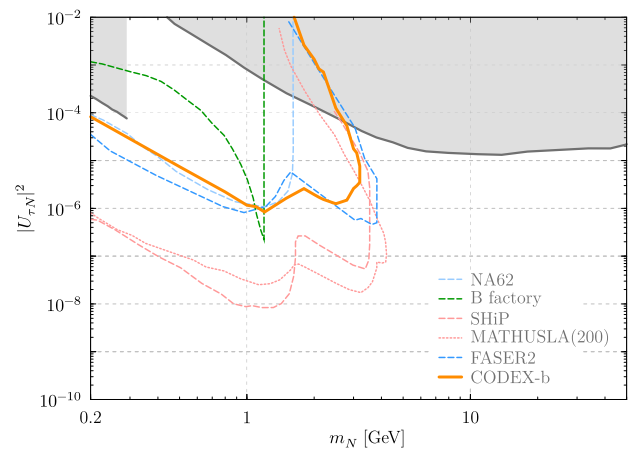
(c) $U_{\mu N} \gg U_{eN}, U_{\tau N}$ (Channels)



(d) $U_{\mu N} \gg U_{eN}, U_{\tau N}$ (Combined)



(e) $U_{\tau N} \gg U_{eN}, U_{\mu N}$ (Channels)



(f) $U_{\tau N} \gg U_{eN}, U_{\mu N}$ (Combined)

Fig. 14 Projected sensitivity of CODEX-b to Dirac heavy neutral leptons. Left: Contributions from the individual decay channels with the net result (orange). Right: Comparison with current constraints (gray) and other proposed experiments, including NA62 [101], DUNE [102], SHiP [17], FASER2 [103], and MATHUSLA [83], which is shown

for its 200 m \times 200 m configuration. ATLAS [100] and CMS [104] constraints on prompt Majorana HNL decays are not shown as the sensitivity is currently subdominant to the DELPHI exclusions, and they moreover use a lepton number violating final state only accessible to Majorana HNLs $-\mu^\pm\mu^\pm e^\mp$ or $e^\pm e^\pm\mu^\mp$ to place limits

sive setup of CODEX-b. Instead we restrict ourselves to a few recent and representative examples. For a more comprehensive overview of the theory space we refer to Ref. [105].

2.4.1 R-parity violating supersymmetry

The LHC has placed strong limits on supersymmetric particles in a plethora of different scenarios. The limits are especially strong if the colored superpartners are within the kinematic range of the collider. If this is not the case, the limits on the lightest neutralino ($\tilde{\chi}_1^0$) are remarkably mild, especially if the lightest neutralino is mostly bino-like. In this case $\tilde{\chi}_1^0$ can still reside in the \sim GeV mass range, and be arbitrarily separated from the lightest chargino. Such a light neutralino must be unstable to prevent it from overclosing the universe, which will happen if R-parity is violated [106]. The $\tilde{\chi}_1^0$ then decays through an off-shell sfermion coupling to SM particles through a potentially small R-parity violating coupling. The combination of these effects typically provide a macroscopic $\tilde{\chi}_1^0$ proper lifetime.

The sensitivity of CODEX-b to this scenario was recently studied for $\tilde{\chi}_1^0$ production through exotic B and D decays [107], as well as from exotic Z^0 decays [83]. Dercks et al. [107] studied the interaction

$$W_{RPV} = \lambda'_{ijk} L_i Q_j D_k^c \quad (7)$$

and considered five benchmarks, corresponding to different choices for the matrix λ'_{ijk} , each with a different phenomenology. We reproduce here their results for their benchmarks 1 and 4, and refer the reader to Ref. [107] for the remainder. The parameter choices, production modes and main decay modes are summarized in Table 1. The reach of CODEX-b is shown in Fig. 15. In both benchmarks, CODEX-b would probe more than 2 orders of magnitude in the coupling constants. For benchmark 4 the reach would be substantially increased if the detector is capable of detecting neutral final states by means of some calorimetry.

The above results assume the wino and higgsino multiplets are heavy enough to be decoupled from the phenomenology. This need not be the case. For instance, the current LHC bounds allow for a higgsino as light as ~ 150 GeV [108], as long as the wino is kinematically inaccessible and the bino decays predominantly outside the detector. In this case, the mixing of the bino-like $\tilde{\chi}_1^0$ can be large enough to induce a substantial branching ratio for the $Z \rightarrow \tilde{\chi}_1^0 \tilde{\chi}_1^0$ process. Helo et al. [83] showed that the reach of CODEX-b would exceed the $Z \rightarrow$ invisible bound for $0.1 \text{ GeV} < m_{\tilde{\chi}_1^0} < m_Z/2$ and $10^{-1} \text{ m} < c\tau < 10^6 \text{ m}$, as shown in Fig. 16. The reach is independent of the flavor structure of the RPV coupling(s), so long as the branching ratio to final states with at least two charged tracks is unsuppressed. It should be noted that the ATLAS searches in the muon chamber [46,47] are expected

to have sensitivity to this scenario, although no recasted estimate is currently available. As with exotic Higgs decays in Sect. 2.3.1, the expectation is, however, that CODEX-b would substantially improve upon the ATLAS reach for low $m_{\tilde{\chi}_0}$.

2.4.2 Relaxion models

Relaxion models rely on the cosmological evolution of a scalar field – the relaxion – to dynamically drive the weak scale towards an otherwise unnaturally low value [109]. The relaxion sector therefore must be in contact with the SM electroweak sector, and the implications of relaxion-Higgs mixing have been studied extensively [109–113]. The phenomenological constraints were mapped out in detail in Refs. [114,115] (see [52] for similar phenomenology in a model where the light scalar is identified with the inflaton). Following the discussion in Ref. [105], the phenomenologically relevant physics of the relaxion, ϕ , is contained in the term

$$\mathcal{L} \supset 2C \frac{h^2}{\Lambda} \Lambda_N^3 \cos\left(\frac{\phi}{f} + \delta\right), \quad (8)$$

in which h is the real component of the SM Higgs field that obtains a vacuum expectation value v , Λ is the cut-off scale of the effective theory, Λ_N is the scale of a confining hidden sector, f is the scale at which a UV $U(1)$ symmetry is broken spontaneously, and finally, C and δ are real constants. After ϕ settles into its vacuum expectation value, ϕ_0 , Eq. (8) can be expanded in large ϕ_0/f , such that

$$\begin{aligned} \mathcal{L} \supset & 2\lambda' \sin\left(\frac{\phi_0}{f} + \delta\right) \frac{v^2}{f} h^2 \phi \\ & + \lambda' \cos\left(\frac{\phi_0}{f} + \delta\right) \frac{v^2}{f^2} h^2 \phi^2 + \dots, \end{aligned} \quad (9)$$

with $\lambda' = C \Lambda_N^3 / v^2 \Lambda$. The model in Eq. (9) now directly maps onto the scalar-Higgs portal in Eq. (2) of Sect. 2.3.2. CODEX-b and other intensity and/or lifetime frontier experiments can then probe the model in the regime $\lambda' \sim 1$ and $f \sim \text{TeV}$. The angle $\phi_0/f + \delta$ controls whether the mixing or quartic term is most important: On the one hand, if it is small, the lifetime of ϕ increases but the quartic in Eq. (9) can be sizable, enhancing the $h \rightarrow \phi\phi$ branching ratio (Fig. 9b). On the other hand, for $\phi_0/f + \delta \simeq \pi/2$ the quartic is negligible and the phenomenology is simply that of a scalar field mixing with the Higgs (Fig. 9a).

2.4.3 Neutral naturalness

The Abelian hidden sector model in Sect. 2.3.1 has enough free parameters to set the mass ($m_{A'}$), the Higgs branching ratio ($\text{Br}(h \rightarrow A'A')$) and the width ($\Gamma_{A'}$) independently. It therefore allows for a very general parametrization of the

Table 1 Summary of two of the five benchmark models considered in Ref. [107]

	Coupling	Production	Decay products
Benchmark 1	$\lambda'_{122}, \lambda'_{112}$	$D_s^\pm \rightarrow \tilde{\chi}_1^0 + e^\pm$	$\eta, \eta', \phi, K^{0,\pm} + \nu_e, e^\mp$
Benchmark 4	$\lambda'_{131}, \lambda'_{121}$	$B^{0,\pm} \rightarrow \tilde{\chi}_1^0 + X^{0,\pm}$	$D^\pm, D^{*\pm} + e^\mp$

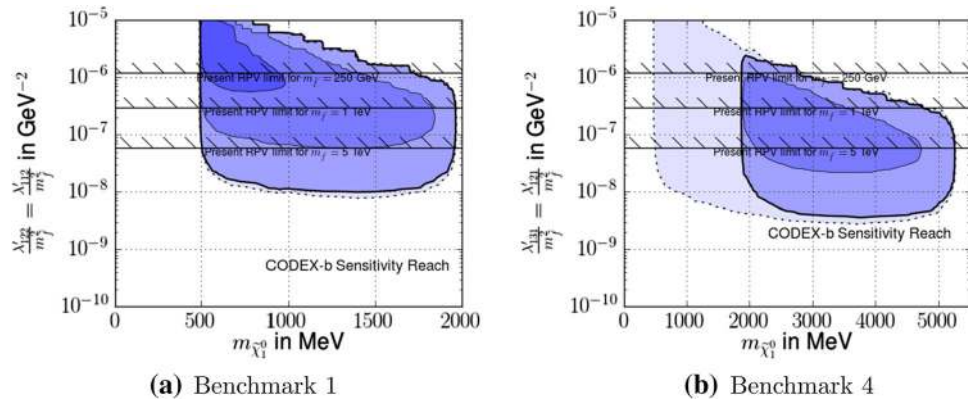


Fig. 15 Projected sensitivity of CODEX-b for light neutralinos with R-parity violating coupling produced in D and B meson decays, reproduced from [107] with permission of the authors. The light blue, blue and dark blue regions enclosed by the solid black lines correspond to $\gtrsim 3, 3 \times 10^3$ and 3×10^6 events respectively. The dashed curve represents the extended sensitivity if one assumes CODEX-b could also detect the neutral decays of the neutralino. The hashed solid lines indicate the single RPV coupling limit for different values of the sfermion masses. See Ref. [107] for details

resents the extended sensitivity if one assumes CODEX-b could also detect the neutral decays of the neutralino. The hashed solid lines indicate the single RPV coupling limit for different values of the sfermion masses. See Ref. [107] for details

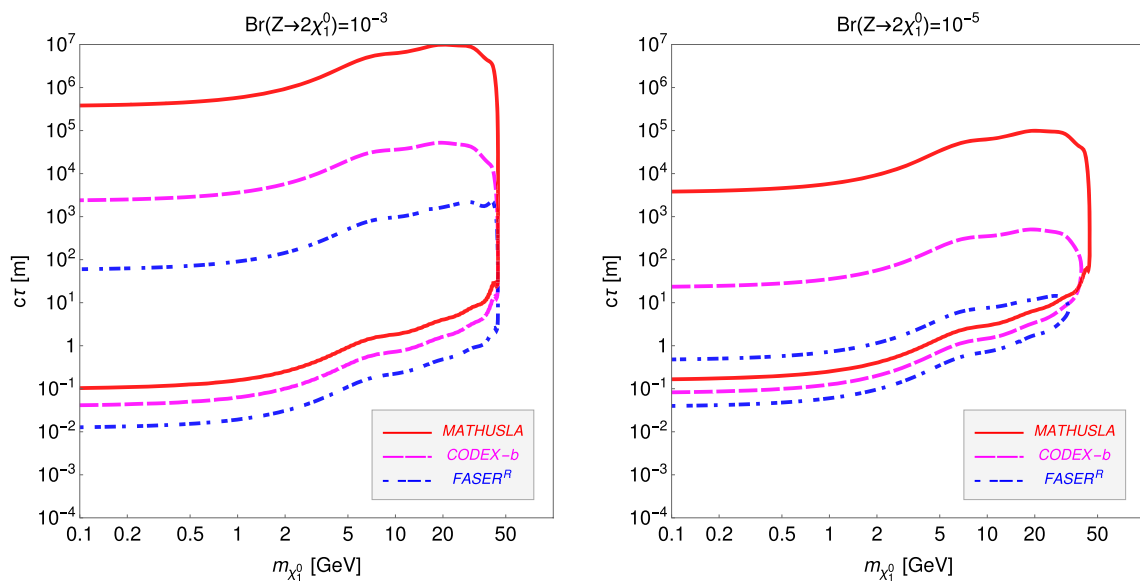


Fig. 16 Projected sensitivity of CODEX-b for light neutralinos with R-parity violating coupling, as produced in Z decays, reproduced from Ref. [83] with permission of the authors. Also shown are projections

for the $200\text{ m} \times 200\text{ m}$ MATHUSLA configuration and FASER^R, the 1 m radius configuration (referred to as FASER2 elsewhere in this document)

reach for exotic Higgs decays in terms of the lifetime, mass and production rate of the LLP. The downside of this generality is that the model has too many independent parameters to be very predictive. In many models, however, the lifetime has a very strong dependence on the mass, favoring long

lifetimes for low mass states. We therefore provide a second, more constrained example where the lifetime is not a free parameter.

The example we choose is the fraternal twin Higgs [116], which is a recent incarnation of the Twin Higgs paradigm

[117, 118], which is designed to address the little hierarchy problem. It is itself an example of a hidden valley [119, 120]. The model consists of a dark or “twin” sector containing an $SU(2) \times SU(3)$ gauge symmetry, that are counterparts of the SM weak and color gauge groups. It further contains a dark b -quark and a number of heavier states which are phenomenologically less relevant. The most relevant interactions are

$$\mathcal{L} \supset \lambda \left(\mathcal{H} \mathcal{H}^\dagger - f^2 \right)^2 + y'_t H' q'_L t'_R + y'_b H'^c q'_L b'_R, \quad \mathcal{H} \equiv \begin{pmatrix} H \\ H' \end{pmatrix}, \quad (10)$$

with H the SM Higgs doublet and H' the dark sector Higgs doublet. The “twin quarks” q'_L , b'_R and t'_R are dark sector copies of the 3rd generation quarks.

The Higgs potential of this model has an accidental $SU(4)$ symmetry, which protects the Higgs mass at one loop provided that $y'_t \approx y_t$, with y_t the SM top Yukawa coupling. The corresponding top partner – the “twin top” – carries color charge under the twin sector’s $SU(3)$ rather than SM color, and is therefore not subject to existing collider constraints from searches for colored top partners. The accidental symmetry exchanging $H \leftrightarrow H'$ may further be softly broken, such that $\langle H \rangle = v$ and $\langle H' \rangle \approx f$. The parameter f is typically expressed in terms of the mass of the twin top quark, m_T , through the relation $m_T = y_t f / \sqrt{2}$. The existing constraints on the branching ratio of the SM Higgs already demand $m_T/m_t \gtrsim 3$ [121].

We consider the scenario in which the b' mass is heavier than the dark $SU(3)$ confinement scale, Λ' , such that the lightest state in the hadronic spectrum is the 0^{++} glueball [122, 123], with a mass $m_0 \approx 6.8\Lambda'$. The 0^{++} glueball mixes with the SM Higgs boson through the operator

$$\mathcal{L} \supset -\frac{\alpha'_3 y_t}{6\pi\sqrt{2}} \frac{m_t}{m_T^2} h \text{Tr} [G'^{\mu\nu} G'_{\mu\nu}] \quad (11)$$

where h is the physical Higgs boson and α'_3 the twin QCD gauge coupling. After mapping the gluon operator to the low energy glueball field, this leads to a very suppressed decay width of the 0^{++} state, even for moderate values of m_t/m_T . In particular, the lifetime is a very strong function of the mass, and can be roughly parametrized as

$$c\tau_{0^{++}} \sim 18 m \times \left(\frac{10 \text{ GeV}}{m_0} \right)^7 \times \left(\frac{m_T/m_t}{3} \right)^4. \quad (12)$$

This is naturally in the range where displaced detectors like CODEX-b, AL3X and MATHUSLA are sensitive. The full lifetime curve is shown in the left hand panel of Fig. 17, where we have accounted for the running of α'_3 , as in Ref. [116, 124].

For simplicity we assume that the second Higgs is too heavy to be produced in large numbers at the LHC, as is typ-

ical in composite UV completions. However, even in this pessimistic scenario the SM Higgs has a substantial branching ratio to the twin sector. Specifically, this Higgs has a branching ratio of roughly $\sim m_t^2/m_T^2$ for the $h \rightarrow b'b'$ channel. The b' quarks subsequently form dark quarkonium states, which in turn can decay to lightest hadronic states in the hidden sector. While this branching ratio is large, the phenomenology of the dark quarkonium depends on the detailed spectrum of twin quarks (see e.g. Ref. [124]). There is however a smaller but more model-independent branching ratio of the SM Higgs directly to twin gluons, given by [125]

$$\text{Br}[h \rightarrow g'g'] \approx \text{Br}[h \rightarrow gg] \times \left(\frac{\alpha'_s(m_h)}{\alpha_s(m_h)} \frac{m_t^2}{m_T^2} \right)^2 \quad (13)$$

with $\text{Br}[h \rightarrow gg] = 0.086$. $\alpha_s(m_h)$ and $\alpha'_s(m_h)$ are the strong couplings, respectively in the SM and twin sectors, evaluated at m_h . The hidden glueball hadronization dynamics is not known from first principles, and we have assumed that the Higgs decays to the twin sector on average produces two 0^{++} glueballs. Especially at the rather low m_0 of interest for CODEX-b, this is likely a conservative approximation.

The projected reach of CODEX-b, MATHUSLA and ATLAS is shown in the right hand panel of Fig. 17. The projections for ATLAS were obtained as in Sect. 2.3.1. The high mass, short lifetime regime may be covered with new tagging algorithms for the identification of merged jets at LHCb [126, 127]. We find that CODEX-b would significantly extend the reach of ATLAS for models of neutral naturalness. For hidden glueballs, the factor of ~ 30 larger geometric acceptance times luminosity for MATHUSLA only results in roughly a factor of ~ 2 more reach in m_0 for a fixed m_T , because of the scaling in Eq. (12). For higher glueball masses, CODEX-b outperforms MATHUSLA due to its shorter baseline. However, this region will likely be covered by ATLAS.

In summary, this hidden glueball model serves to illustrate an important point: For light hidden sector states, the lifetime often grows as a strong power-law of its mass, as illustrated by Fig. 17. For ATLAS and CMS, this means that the standard background rejection strategy of requiring two vertices becomes extremely inefficient for such light hidden states. Instead, displaced detectors like CODEX-b, MATHUSLA and FASER are needed to cover the low mass part of the parameter space.

2.4.4 Inelastic dark matter

Berlin and Kling [128] have studied the reach for various (proposed) LLP experiments in the context of a simple model for inelastic dark matter [129, 130]. The ingredients are two Weyl spinors with opposite charges under a dark, higgsed $U(1)$ gauge interaction. In the low energy limit, the model

Fig. 17 Left: Lifetime of the 0^{++} glueball as a function of its mass. Right: Projected reach of CODEX-b, MATHUSLA (200 m \times 200 m) and ATLAS at the full luminosity of the HL-LHC. The solid (dashed) ATLAS contours refer to the optimistic (conservative) extrapolations of the ATLAS reach, as discussed in Sect. 2.3.1. The horizontal dashed line indicates the reach of precision Higgs coupling measurements at the LHC [48]

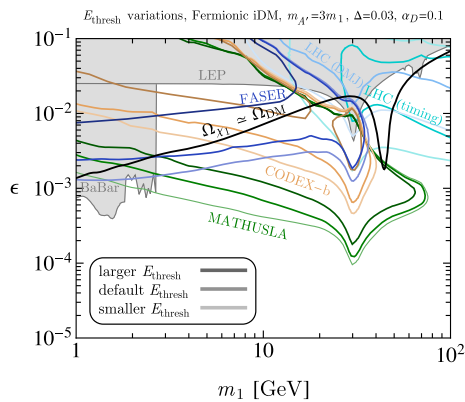
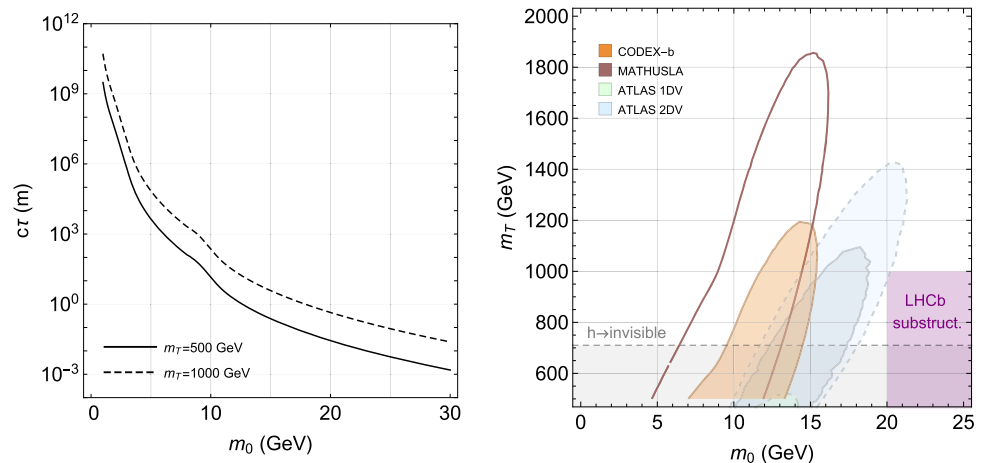


Fig. 18 Sensitivity estimates for the inelastic dark matter benchmark, reproduced from Ref. [128] with permission of the authors. The black line indicates the line on which the correct dark matter relic density is predicted by the model. Darker/lighter shades correspond to larger/smaller minimum energy thresholds for the decay products of χ_2 , with CODEX-b shown in orange shades. For CODEX-b and MATHUSLA (200 m \times 200 m), the minimum energy is taken to 1200, 600, or 300 MeV per track. For FASER2 (1 m radius), the total visible energy deposition is taken to be greater than 200, 100, or 50 GeV. For a displaced muon-jet search at ATLAS/CMS and a timing analysis at CMS with a conventional monojet trigger, the minimum required transverse lepton momentum is 10, 5, or 2.5 GeV and 6, 3, or 1.5 GeV, respectively

reduces to

$$\mathcal{L} \supset -ie_D \bar{\chi}_2 \not{A}' \chi_1 + \frac{\epsilon}{2} F'^{\mu\nu} F_{\mu\nu} + \dots, \quad (14)$$

where the second term indicates the mixing of the dark gauge boson with the SM photon. The ellipsis represents sub-leading terms which do not significantly contribute to the phenomenology. The pseudo-Dirac fermions χ_1 and χ_2 are naturally close in mass, which leads to a phase space suppression of the width of χ_2 . The fractional mass difference is parameterized by $\Delta \equiv (m_2 - m_1)/m_1 \ll 1$.

At the LHC, the production occurs through $q\bar{q} \rightarrow A' \rightarrow \chi_2 \chi_1$, which is controlled by the mixing parameter ϵ . The

decay width of χ_2 is given by

$$\Gamma(\chi_2 \rightarrow \chi_1 \ell^+ \ell^-) \simeq \frac{4\epsilon^2 \alpha_{em} \alpha_D \Delta^5 m_1^5}{15\pi m_{A'}^4}, \quad (15)$$

where $\alpha_D = e_D^2/4\pi$ is the dark gauge coupling. CODEX-b, MATHUSLA, FASER and the existing LHC experiments can search for the pair of soft, displaced fermions from the χ_2 decay. The expected sensitivity of the various experiments is shown in Fig. 18 for an example slice of the parameter space. In particular, CODEX-b will be able to probe a large fraction of the parameter space that produces the observed dark matter relic density, as indicated by the black line in Fig. 18. It is worth noting that for this model, the minimum energy threshold per track is an important parameter in determining the reach, which should inform the design of the detector. For more benchmark points and details regarding the cosmology, we refer to Ref. [128].

2.4.5 Dark matter coscattering

The process of coscattering [131, 132] has been studied as a way to generate the correct relic DM abundance. Coscattering has a similar framework to coannihilating dark matter models: Both models contain at least one dark matter particle χ , a second state charged under the \mathbf{Z}_2 of the dark sector, ψ , and a third particle X that allows the two particles to transition into one another via an interaction such as a Yukawa, $yX\chi\psi$. In many coannihilation scenarios $\psi\psi \leftrightarrow XX$ (or SM) is an efficient annihilation mechanism, while $\chi\chi, \chi\psi \leftrightarrow XX$ (or SM) is not. Throughout the coannihilation, the ‘‘coscattering’’ process $\psi X \leftrightarrow \chi X$ (or similar) remains efficient and allows the χ and ψ species to interchange, without changing the dark particle number. Eventually, $\psi\psi \leftrightarrow XX$ freezes out, and the total dark particle number is fixed.

By contrast, one may consider coscattering DM [131], in which the $\psi X \leftrightarrow \chi X$ coscattering process drops out of equilibrium before the $\psi\psi \leftrightarrow XX$ coannihilation process.

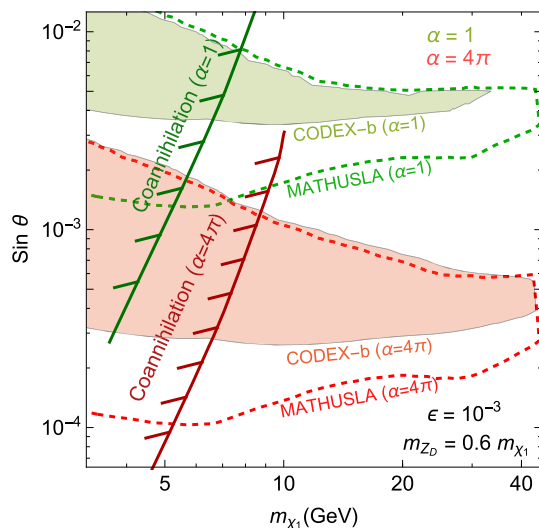


Fig. 19 Projected sensitivity to the coscattering dark matter benchmark for $\alpha_D = 1$ (green) and $\alpha_D = 4\pi$ (red). The shaded region represents the reach for CODEX-b with 300 fb^{-1} . The dashed line is the reach for MATHUSLA in the $200 \text{ m} \times 200 \text{ m}$ configuration with 3 ab^{-1} . To the left of the dark hatched line, the coannihilation process of $\chi_2 \bar{\chi}_1 \rightarrow Z_D Z_D$ remains active long enough to deplete the relic abundance of χ_1 below the observed amount so that the model is inconsistent to the left of these lines

This requires three ingredients: $m_X \sim m_\psi \sim m_\chi$; a large $\psi\psi \leftrightarrow XX$ cross-section; and a small $\psi X \leftrightarrow \chi X$ cross-section. As χ does not have any sizable interactions other than with ψ by assumption, there are no interactions beyond $\psi X \leftrightarrow \chi X$ that allow for χ to maintain a thermal distribution while it is in the process of decoupling from the thermal bath. This results in important non-thermal corrections that require tracking the full phase space density, rather than just the particle number n_χ , in order to correctly evaluate the relic abundance [131].

The vector portal model we consider throughout the rest of this subsection is similar to the one in the Sect. 2.4.4. Here we introduce a new $U(1)_D$ gauge group with fairly strong couplings, a scalar charged under the $U(1)_D$ that obtains a VEV, a Dirac spinor χ_2 charged under the gauge group, and a second Dirac spinor χ_1 that is not. The Lagrangian for the model is

$$\mathcal{L} \supset -i g_D \bar{\chi}_2 \not{Z}_D \chi_2 + m_2 \bar{\chi}_2 \chi_2 + m_1 \bar{\chi}_1 \chi_1 + y_{21} \phi \bar{\chi}_2 \chi_1 + y_{12} \phi^* \bar{\chi}_1 \chi_2 + \frac{\epsilon}{2 \cos \theta_W} Z_D^{\mu\nu} B_{\mu\nu} + \dots \quad (16)$$

The scalar VEV $\langle \phi \rangle$ gives a mass to the dark vector and generates a small mixing between the $U(1)_D$ active χ_2 and sterile χ_1 . For simplicity, we set $y \equiv y_{12} = y_{21}$. When $\Delta m \equiv m_2 - m_1 \gg y \langle \phi \rangle$ a small mixing angle $\theta \approx y \langle \phi \rangle / \Delta m$ is generated. We assume that $m_\phi \gtrsim m_{Z_D}$, so that then when $y \ll g_D$, the phenomenology is insensitive to the presence of the scalar.

The mixing of Z_D with the Z boson allows for $Z \rightarrow \chi_2 \bar{\chi}_2$ with a branching ratio of

$$\text{BR}(Z \rightarrow \chi_2 \bar{\chi}_2) = \frac{\alpha_D \epsilon^2 \tan^2 \theta_W m_Z^5}{12(m_Z^2 - m_{Z_D}^2)^2 \Gamma_Z} \times \left(1 + 2 \frac{m_{\chi_2}^2}{m_Z^2} \right) \sqrt{1 - \frac{m_{\chi_2}^2}{m_Z^2}}. \quad (17)$$

The daughter χ_2 particles from the Z decay can propagate several meters before decaying to χ_1 through an off-shell dark photon, i.e. $\chi_2 \rightarrow \chi_1 f f$. The ‘ ff ’ indicates a pair of SM fermions, which CODEX-b can detect. The decay rate is dictated by the splitting between the two states. For example, the partial width to electrons, neglecting the electron mass

$$\Gamma_{ee} = \frac{\alpha_D \alpha_{em} \epsilon^2 \sin^2 \theta}{24 m_{Z_D}^4 \pi m_{\chi_2}^3} \left[m_{\chi_2}^8 - 2 m_{\chi_1} m_{\chi_2}^7 - 8 m_{\chi_1}^2 m_{\chi_2}^6 - 18 m_{\chi_1}^3 m_{\chi_2}^5 + 18 m_{\chi_1}^5 m_{\chi_2}^3 + 8 m_{\chi_1}^6 m_{\chi_2}^2 + 2 m_{\chi_1}^7 m_{\chi_2} - m_{\chi_1}^8 - 24 m_{\chi_1}^3 m_{\chi_2}^3 \right] \times (m_{\chi_1}^2 + m_{\chi_1} m_{\chi_2} + m_{\chi_2}^2) \ln \left(\frac{m_{\chi_1}}{m_{\chi_2}} \right). \quad (18)$$

From this expression we can approximate the lifetime as

$$c\tau_{\chi_2} = \frac{\text{BR}(Z_D \rightarrow ee; \Delta m)}{\Gamma_{ee}}, \quad (19)$$

where $\text{BR}(Z_D \rightarrow ee; \Delta m)$ is the branching ratio for a kinetically mixed dark vector of mass Δm into ee . This is done to approximate the inclusion of additional accessible final states, as splittings in this model are commonly $\mathcal{O}(\text{GeV})$. While a more thorough treatment would integrate over phase space for each massive channel separately, this approximation captures the leading effect to well within the precision desired here. Additionally, χ_2 pairs can be directly produced through an off-shell Z_D . Because the Z_D is off-shell, this does not generate a large contribution unless $m_{\chi_2} \lesssim 10 \text{ GeV}$. This model provides a scenario containing an exotic Z decay into long-lived particles.

In Fig. 19 we show the projected sensitivity for CODEX-b (shaded) and MATHUSLA (dashed) [105] to the model setting $\epsilon = 10^{-3}$, $m_{Z_D} = 0.6 m_{\chi_1}$, and for two choices of $\alpha_D = 1$ and 4π (green and red, respectively). With these parameters fixed, the choice of $\sin \theta$ fixes the mass splitting from the DM relic abundance criteria. At small masses, the $\bar{\chi}_2 \chi_1 \leftrightarrow Z_D Z_D$ coannihilation process remains in equilibrium long enough to deplete the χ_1 number density below the relic abundance today. This region is illustrated by the dark hatched lines.

2.4.6 Dark matter from sterile coannihilation

D’Agnolo et. al. [133] have explored the mechanism of sterile coannihilation, for which the number density in the dark sector is set by the annihilation of states that are heavier than the dark matter. In this scenario the dark matter remains in chemical equilibrium with these heavy states until after their annihilation process freezes out, which naturally allows for much lighter dark matter than in standard thermal freeze-out models.

Concretely, the example model that is considered in Ref. [133] is given by

$$-i\mathcal{L} \supset \frac{1}{2}m_\phi^2\phi^2 + \frac{1}{2}m_\psi\psi^2 + \delta m\chi\psi + \frac{1}{2}m_\chi\chi^2 + \frac{y}{2}\phi\psi^2 \tag{20}$$

where the parameter $\delta m \ll m_\psi, m_\chi$ generates a small mixing between ψ and χ . For the choice $m_\psi \gtrsim m_\chi > m_\phi$, the relic density of χ is effectively set by $\psi\psi \rightarrow \phi\phi$ annihilations. Finally, ϕ is assumed to mix with the SM Higgs, and it is this coupling which keeps the dark sector in thermal equilibrium with the SM sector. For a summary of the direct detection and cosmological probes of this model, we refer to Ref. [133]. From a collider point of view, the most promising way to probe the model is to search for the scalar ϕ through its mixing with the Higgs. This scenario is identical to the scalar-Higgs portal model with $\lambda = 0$, which is discussed in Sect. 2.3.2. Fig. 20 shows the projected reach for CODEX-b, overlaid with the relevant constraints and projections from dark matter direct detection and CMB measurements.

2.4.7 Asymmetric dark matter

In many asymmetric dark matter models, the DM abundance mass is directly tied to the matter anti-matter asymmetry in the SM sector [134–136]. Therefore the generic expectation for the DM is to carry $B - L$ quantum numbers and have a mass \simeq GeV. For this mechanism to operate, the DM sector interactions with the SM should be suppressed and both sectors communicate in the early universe through operators of the form

$$\frac{1}{M^{\Delta_{SM}+\Delta_X-4}}\mathcal{O}_X\mathcal{O}_{SM}, \tag{21}$$

where \mathcal{O}_X and \mathcal{O}_{SM} are operators consisting of dark sector and SM fields respectively, with Δ_X and Δ_{SM} their respective operator dimensions. In supersymmetric models of asymmetric dark matter [134], the simplest operators in the superpotential are of the form

$$W = XLH, \quad \frac{1}{M}Xu^c d^c d^c, \quad \frac{1}{M}XQLd^c, \quad \frac{1}{M}XLLe^c, \tag{22}$$

with $X \equiv \tilde{x} + \theta x$ the chiral superfield containing the DM, denoted by x . The phenomenology of this scenario is very similar to that of RPV supersymmetry, with decay chains such as $\tilde{\chi}^0 \rightarrow \tilde{x}u^c d^c d^c$ (see Sect. 2.4.1). To accommodate the correct cosmology, macroscopic lifetimes $c\tau \sim 10$ m are typically required [135, 137]. Moreover, \tilde{x} itself may or may not be stable, depending on the model.

More generally, if the dark sector has additional symmetries and multiple states in the GeV mass range, as occurs naturally in hidden valley models with asymmetric dark matter (see e.g. Refs. [119, 138]), these excited states often must decay to the DM plus some SM states. Such decays must necessarily occur through higher dimensional operators, and macroscopic lifetimes are therefore generic. As for previous portals, LLP searches in the GeV mass range are best suited to displaced, background-free detectors such as CODEX-b.

2.4.8 Other dark matter models

There are many other dark matter models that could provide signals observable with CODEX-b. Presenting projections for all possibilities is beyond the scope of this work, but here we briefly summarize many of the existing scenarios that can provide long-lived particles. Below we detail: SIMPs, ELDERs, co-decaying DM, dynamical DM, and freeze-in DM.

Strongly interacting massive particles (SIMPs) [139–142] obtain their relic density through a $3 \rightarrow 2$ annihilation process mediated by a strong, hidden sector force. The preferred mass scale in this scenario is the GeV scale, and the strong nature of the hidden sector implies the presence a whole spectrum of dark pions, dark vector mesons etc. The $3 \rightarrow 2$ annihilation process, however, heats up the dark sector in the early universe, which would drive the dark matter to be exponentially hotter than the SM if it were completely decoupled. Since the dark matter is known to be cold, this means there must exist a sufficiently strong portal keeping both sectors in thermal equilibrium. These interactions predict a variety of signatures, in conventional dark matter detection and at colliders. In this sense SIMP models provide more motivation for exploring the Hidden Valley framework: At the LHC, it is possible to produce the hidden quarks through the aforementioned portal (e.g. a kinetically mixed dark photon), which would subsequently shower and fragment to hidden mesons with masses around the GeV scale. Some of these mesons will be stable and invisible, such as the dark matter, while other will decay back to the standard model, often with macroscopic displacements. This phenomenology is studied in detail in Refs. [143, 144].

ELastically DEcoupling Relics (ELDERs) [145] share many of the features of the SIMP models, including the strong $3 \rightarrow 2$ annihilation process in the hidden sector and the mandatory portal with the SM to prevent the dark sector from

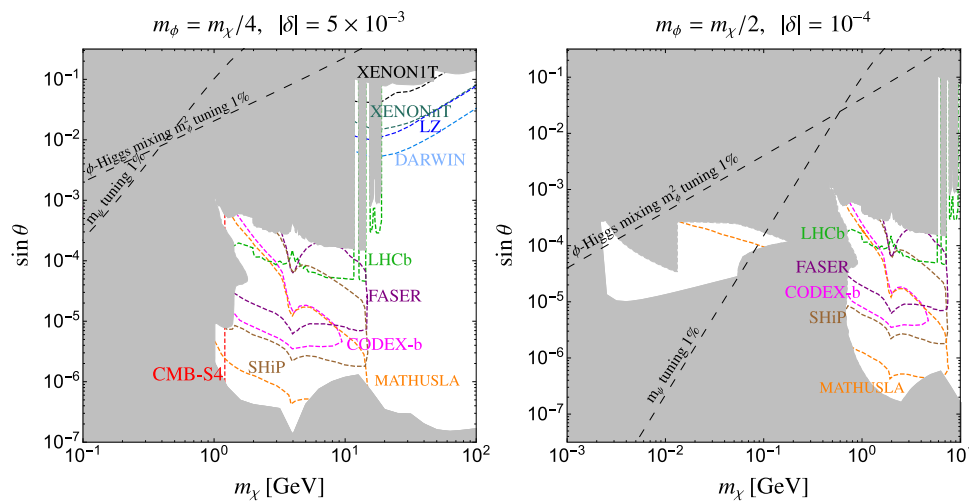


Fig. 20 Projected sensitivity to the coannihilation dark matter benchmark, reproduced from Ref. [133] with permission of the authors. The projections are effectively the same as those in Fig. 9a. In every point of the plot, Δ is fixed to reproduce the observed relic density. Left: The remaining parameters are set as $y = e^{i\pi/4}$, $m_\phi = m_\chi/4$ and

$\delta = 5 \times 10^{-3} e^{i\pi/4}$, with $\delta \equiv \delta m/m_\chi$. Right: The remaining parameters are set instead as $y = e^{i\pi/4}$, $m_\phi = m_\chi/2$ and $\delta = 10^{-4} e^{i\pi/4}$. MATHUSLA sensitivity is shown for the $200 \text{ m} \times 200 \text{ m}$ configuration, while FASER sensitivity is shown for the 1 m radius configuration, now referred to as FASER2

overheating. In contrast to SIMP models, the elastic scattering processes between the dark sector and the SM freeze out *before* the end of the $3 \rightarrow 2$ annihilations, such that the dark matter cannibalizes itself for some time during the evolution of the universe. ELDER models are also examples of hidden valleys, and the collider phenomenology is therefore qualitatively similar to that of SIMP models.

In **co-decaying dark matter models** [146–148], the dark matter state, χ_1 , is kept in equilibrium with a slightly heavier dark state, χ_2 , though efficient $\chi_1\chi_1 \leftrightarrow \chi_2\chi_2$ processes in the early universe, but the dark sector does not maintain thermal equilibrium with the SM. The χ_2 state is, however, unstable and decays back to the SM. Because both states remain in equilibrium, this also depletes the χ_1 number density once the temperature of the dark sector drops below the mass of χ_2 . For this mechanism to operate, χ_2 should have a macroscopic lifetime. On the one hand, the heavier χ_2 could very well be produced at the LHC through a heavy portal, however this is not strictly required for the co-decaying dark matter framework to operate. On the other hand, if implemented in the context of e.g. neutral naturalness [149, 150], a production mechanism at the LHC is typically a prediction and the phenomenology is once again that of a hidden valley.

In **dynamical dark matter models** [151–153], the dark sector contains a large ensemble of decaying dark states with a wide range of lifetimes. Their collective abundance makes up the DM abundance we see today, by balancing their share in the universe’s energy budget against their lifetime. Some of the states in the ensemble are expected to have lifetimes that can be resolved on collider length scales. Just as for co-decaying dark matter, an observable cross-section at the

LHC is possible but not required. If the dynamical dark sector can be accessed, however, the collider phenomenology is rich [154–156] and auxiliary, displaced, background-free LLP detectors can play an important role [157].

Finally, in **Freeze-in models** [158], the dark matter is never in equilibrium with the SM sector, but instead the dark sector is slowly populated through either scattering or the decay of a heavy state. This mechanism demands very weak couplings, which in the case of freeze-in through decay predicts a long-lived state decaying to DM plus a number of SM states. In the models considered so far, the preferred mass range for the decaying state tends to be in the 100 GeV to 1 TeV regime [159–162], such that ATLAS and CMS ought to be well equipped to find these decays. Should the final states however prove to be difficult to resolve at ATLAS and CMS, or should the parent particle be lighter than currently predicted, CODEX-b could provide the means to probe these models.

2.4.9 Baryogenesis

There exists a wide range of models explaining the baryon asymmetry in the Universe, some of which reside in the deep UV, while others are tied to the weak scale, such as electroweak baryogenesis (see e.g. Ref. [163]) and WIMP baryogenesis [164, 165]. The latter in particular predicts long-lived particles at LHC, with a phenomenology that is qualitatively similar to displaced decays for RPV supersymmetry (see Sect. 2.4.1). We refer to Ref. [105] for a discussion of the discovery potential of WIMP baryogenesis at the lifetime frontier.

Instead we focus here in more depth on a recent idea which generates the baryon asymmetry through the CP-violating oscillations of heavy flavor baryons [166, 167] (see Refs. [66, 67] for similar ideas involving heavy flavor mesons and Ref. [168] for a supersymmetric realization). This enables very low reheating temperatures and is moreover directly testable by experiments such as CODEX-b, as well as Belle II.

The model relies on the presence of the light Majorana fermions χ_1 and χ_2 which carry baryon number. (Two generations are needed to allow for CP-violation.) They couple to the SM quarks through

$$\mathcal{L} \supset \frac{g_{ijkl}}{\Lambda^2} \chi_i u_j d_k d_\ell + \text{h.c.} \quad (23)$$

The out-of-equilibrium condition necessary for baryogenesis can be satisfied, for instance, by a late decay of a third dark fermion to the SM heavy flavor quarks. The operator in Eq. (23) generates a dimension-9 operator with $\Delta B = 2$, of the form $(u_j d_k d_\ell)^2$, that is responsible for the baryon oscillations. For these oscillations to be sufficiently large to generate the observed baryon asymmetry, one needs $m_{\chi_{1,2}} \lesssim m_B$, which has intriguing phenomenological consequences. Moreover, stringent constraints from the dinucleon decay of O^{16} imply that

$$c\tau_\chi \gtrsim 100 \text{ m} \left(\frac{5 \text{ GeV}}{m_{\chi_{1,2}}} \right)^5 \left(\frac{\Lambda/\sqrt{g_{uss}}}{20 \text{ TeV}} \right)^4. \quad (24)$$

These low masses and long lifetimes are precisely where CODEX-b would have a substantial advantage over ATLAS, CMS and LHCb. The branching ratios for B baryons and mesons into $\chi_{1,2}$ can even be as large as 10^{-3} , which means that the rate of $\chi_{1,2}$ production at IP8 could be very large. Part of the parameter space of this model might therefore be probed already by the CODEX- β during Run 3 (see Sect. 4.3).

2.4.10 Hidden valleys

Hidden Valley models [119] are hidden sectors with non-trivial dynamics, which can lead to a relatively large multiplicity of final states in decays of hidden particles. Confining hidden sectors provide a canonical example, because of their non-trivial spectrum of hidden sector hadrons and the “dark shower” that may arise when energy is injected in the hidden sector through a high energy portal. (Fully perturbative examples can also be constructed easily.) Some hidden valleys naturally arise in models which aim to address various shortcomings of the SM: Examples discussed in the preceding sections are neutral naturalness (Sect. 2.4.3), asymmetric dark matter (Sect. 2.4.7) and SIMP dark matter (Sect. 2.4.8).

The phenomenology of hidden valleys can vary widely [120, 138, 169–173], both in terms of the energy and angular distributions of the final states, as well as the lifetime of

the dark sector particles. A handful of initial searches have already been performed at ATLAS, CMS, and LHCb (see e.g. Refs. [174–177]). The various opportunities afforded by the experiments, as well as the challenges involved in constructing a comprehensive search plan were recently summarized in Ref. [7]. In particular,

- While the LLPs are typically in the \sim GeV range, their lifetimes can easily take phenomenologically relevant values spanning many orders of magnitude. In the short lifetime regime, backgrounds can be suppressed by demanding multiple displaced vertices in the same event, provided that a suitable trigger can be found, but this strategy is much less effective in the long-lifetime regime.
- There are generically multiple species of LLPs, with vastly different lifetimes, and some may decay (quasi-)promptly or to a high multiplicity of soft final states. In practice, this means that a displaced decay to SM final states from a dark shower is likely to fail traditional isolation criteria or p_T thresholds, further complicating searches at ATLAS, CMS and LHCb.
- The energy flow in the event may be non-standard, and is poorly understood theoretically. This means that standard jet-clustering algorithms are expected to fail for a large subclass of models.

Because of both its ability to search inclusively for LLP decays and its background-free setup, CODEX-b is not limited by many of these challenges, and would be sensitive to any hidden valley model which has at least one LLP species in the spectrum with both a sizable branching fraction to charged final states and a moderately large lifetime, i.e. $c\tau \gtrsim 1$ m. The latter requirement in particular makes CODEX-b highly complementary to ATLAS, CMS and LHCb in the context of these models, as the short lifetime regime can be probed with a multi-vertex strategy in the main detectors, provided that the putative trigger challenges can be addressed.

3 Backgrounds

Crucial to the CODEX-b programme is the creation and maintenance of a background-free environment. An in-depth discussion of relevant primary and secondary backgrounds may be found in Ref. [8] as well as Ref. [24].

In this section we re-examine the core features of the relevant backgrounds, and the required active and passive shielding required to ensure a background-free environment in the detector. This study includes an updated and more realistic Geant4 simulation of the shielding response that incorporates uncertainties, charged-neutral particle correlations, as well as an updated simulation of the high energy tails of the

primary backgrounds, and simulation of multitrack production in the detector volume. Further, the details and results of a measurement campaign, conducted in the LHCb cavern in 2018 as a preliminary data-driven validation of these simulations, are presented.

3.1 Overview

An LHC interaction point produces a large flux of primary hadrons and leptons. Many of these may be fatal to a background-free environment either because they are themselves neutral long-lived particles, e.g. (anti)neutrons and K_L^0 's, that can enter the detector and then decay or scatter into tracks, or because they may generate such neutral LLP secondaries by scattering in material, e.g. muons, pions or even neutrinos. In the baseline CODEX-b design, LLP-like events are comprised of tracks originating within the detector volume, with the track momentum as low as 400 MeV. This threshold is conservative with respect to likely minimum tracking requirements for a signal (cf. Ref. [178]).

Suppression of primary hadron fluxes can be achieved with a sufficient amount of shielding material: roughly 10^{14} neutrons and K_L^0 's are produced per 300 fb^{-1} at IP8, requiring $\log(10^{14}) \simeq 32\lambda$ of shield for full attenuation, where λ is a nuclear interaction length. In the nominal CODEX-b design, the 3 m of concrete in the UXA radiation wall, corresponding to 7λ ,⁷ is supplemented with an additional 25 λ Pb shield, corresponding to about 4.5 m, as shown in Fig. 21. (We focus here on a shield comprised of lead, though composite shielding making use of e.g. tungsten might also be considered, with similar performance [24].) However, this large amount of shielding material in turn may act as a *source* of neutral LLP secondaries, produced by muons (or neutrinos, see Sect. 3.2.3) that stream through the shielding material. The most concerning neutral secondaries are those produced in the last few λ by high energy muons that themselves slow down and stop before reaching the detector veto layers. Such parent muons are not visible to the detector, while the daughter neutral secondaries, because they pass through only a few λ of shield, may themselves escape the leeward side of the shield and enter the detector volume: We call these ‘stopped-parent secondaries’; a typical topology is shown in Fig. 21.

As a rough example, a 10 GeV muon has a ‘CSDA’ (Continuous Slowing-Down Approximation) range in lead of approximately 6 m [68], corresponding to 32λ ($\lambda_{\text{Pb}} \simeq 0.18 \text{ m}$). Approximately 10^9 such muons are produced per 300 fb^{-1} in the CODEX-b acceptance (see Sect. 3.2.1), and by the last few λ they have slowed to $\lesssim \text{GeV}$ kinetic energy.

⁷ The UXA wall is 3.2 m in depth, corresponding to $\simeq 7.5\lambda$ for ‘standard’ concrete [68]. Since the precise composition of the concrete used in the wall is not available we treat the wall as 3 m of standard concrete as a conservative estimation.

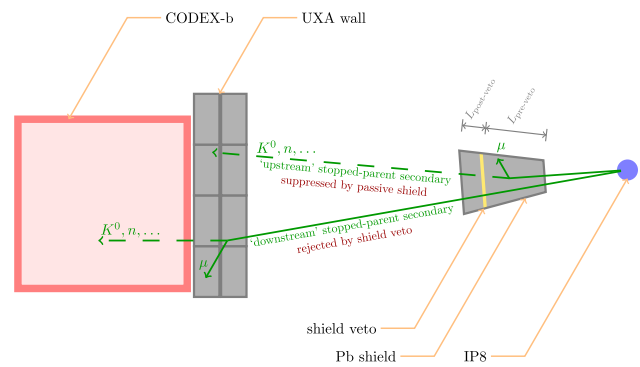


Fig. 21 Cross-section of the shielding configuration of the Pb shield (gray), active shield veto (gold), and concrete UXA wall with respect to IP8 and the detector volume. Also shown are typical topologies for production of upstream and downstream stopped-parent secondaries, which are suppressed by passive shielding or rejected by the active shield veto, respectively

The strange muoproduction cross-section for a GeV muon is $\sim 0.01 \mu\text{b}$ per nucleon, so that in the last λ approximately few $\times 10^3$ K_L^0 's are produced by these muons. The kaon absorption cross-section on a Pb atom is $\sim 2\text{b}$, so the reabsorption probability in the last λ is $\sim 30\%$, with the result that $\sim 10^3$ stopped-parent secondary K_L^0 's can still escape into the detector. This behavior is more properly modelled by a system of linear differential equations, that capture the interplay of the muon dE/dx with the energy-dependence of the secondary muoproduction cross-section and their (re)absorption cross-sections; in practice we simulate this with *Geant4*, as described below in Sect. 3.2.

The CODEX-b proposal resolves this secondary background problem by the addition of a veto layer placed deep inside the shield itself: an active shield element, shown in gold in Fig. 21. This veto layer may then trigger on the parent muons before they produce neutral secondaries and stop. The veto layer must be placed deep enough in the shield – shielded sufficiently from the IP – so that the efficiency required to veto the stopped-parent secondaries produced downstream is not too high. At the same time there must be sufficient shielding downstream from the veto to attenuate the stopped-parent secondaries with respect to the shield veto itself: That is, neutrals produced upstream before the veto layer that could still reach the detector (see Fig. 21). An additional consideration is that the veto rejection rate itself should be much smaller than the overall event rate, in order not to degrade the LLP signal detection efficiency. The nominal shield we consider has a so-called ‘(20 + 5) λ ’ configuration, with 20λ of Pb before the shield veto and 5λ afterwards, plus the additional 3 m of concrete (7λ) from the UXA wall.

3.2 Simulation

3.2.1 Primary fluxes

Generation of the primary IP fluxes is achieved via simulation of the production of pions, neutral and charged kaons, (anti)muons, (anti)neutrons, (anti)protons and neutrino fluxes with `Pythia 8` [179, 180]. Included production channels span minimum bias (QCD), heavy flavor decays (HF), as well as Drell–Yan production (DY). Leptons produced from pion decay vertices inside a cylindrical radius $r < 5$ m and $z < 2$ m are included. We simulate weighted `Pythia 8` events biasing the primary collisions in \hat{p}_T in order to achieve approximately flat statistical errors in $\log(\hat{s})$ up to $\sqrt{\hat{s}}$ of a few TeV. Under the same procedure we also combine soft and hard QCD processes with a \hat{p}_T cut of 20 GeV. A similar cut is used to define the HF sample. For the DY case we both standard $2 \rightarrow 2$ Drell–Yan processes and $V + j$, suitably combined to avoid double counting. In Fig. 22 we show all the relevant generated fluxes, broken down by production channel. In most cases, QCD production dominates, however, the HF and DY production can be important for high energy muon tails.

3.2.2 Simulated shield propagation

Particle propagation and production of secondary backgrounds inside the shield is simulated with `Geant4 10.3` using the `Shielding 2.1` physics list. The `FTFP_BERT` physics list is used to model high-energy interactions, based on the Fritiof [181–184] and Bertini intra-nuclear cascade [185–187] models and the standard electromagnetic physics package [188].

Propagating $\sim 10^{14}$ – 10^{17} particles through the full shield is obviously computationally prohibitive. Instead, as in Refs. [8] and [24] we use a “particle-gun” on a shield subelement, typically either 5 or 2λ deep for Pb, and 7λ for concrete. The subelement geometry is chosen to be a conical section with the same opening angle as the CODEX-b geometric acceptance – approximately 23° – in order to conservatively capture forward-propagating backgrounds after mild angular rescattering. The particle-gun input and output is binned logarithmically in kinetic energy, in 20 bins from $E_{\text{kin}} = 10^{-1.6}$ GeV to 100 GeV and by particle species, including: γ , e^\pm , p^\pm , n^\pm , $\pi^{\pm,0}$, K^\pm , $K_{S,L}^0$, μ^\pm , ν . Propagation of charged and neutral particles and anti-particles are treated separately for kaons, pions, neutrons and muons. For each particle-gun energy bin and species, 10^5 events are simulated; 10^7 events are simulated for muons and anti-muons to properly capture strange muoproduction of secondary K_L^0 s. To also properly capture the ‘CSDA’ or slowing-down behavior of high energy muons when transiting through a large number of shield subelements, the particle-gun energy for

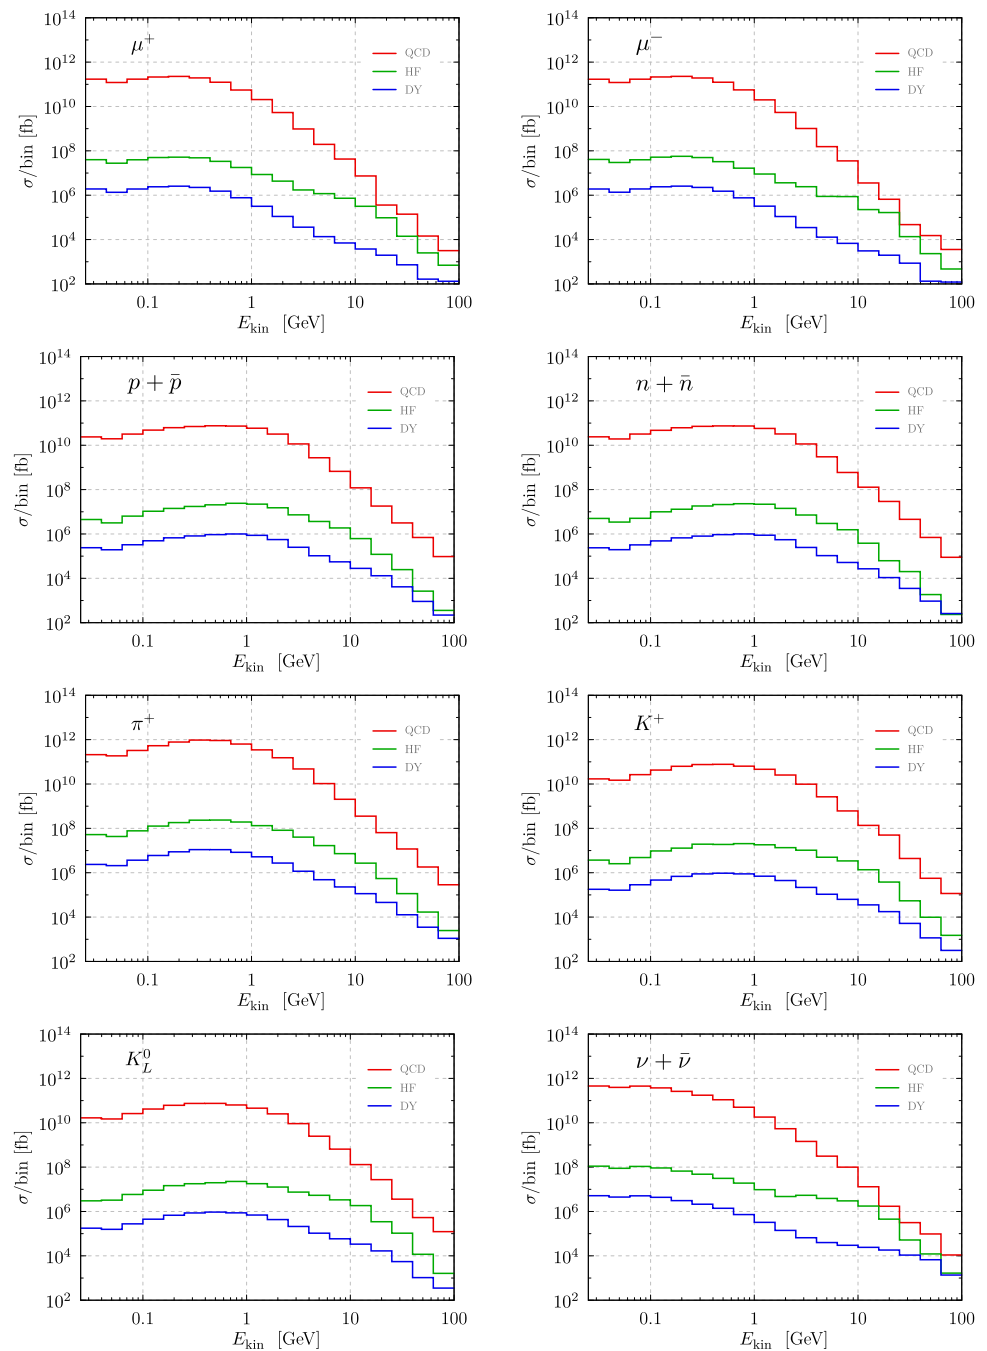
muons was distributed uniformly in kinetic energy, within each bin.

Combining these results together one generates a “transfer matrix” between all incoming and outgoing background fluxes in the shield subelement for each chosen depth and material type. These transfer matrices may then be further composed together with the primary IP fluxes to obtain the attenuation and response of the full shield. Neutrino production of neutral hadrons occurs at a prohibitively small rate and is included separately. As muons may often generate problematic secondaries via forward scattering $\mu \rightarrow X\mu$, an additional handle on the capability to veto neutral secondaries is obtained by keeping track of the presence of any associated charged particles in the particle-gun event, that may trigger relevant veto layers: a charged-neutral correlation. This ‘correlation veto’ is implemented by an additional binning in outgoing particle kinetic energy and the kinetic energy of the hardest charged particle in the event. This information is then used to generate an additional transfer matrix in which the outgoing particles are produced in association with a charged particle above a chosen kinetic energy threshold. We conservatively set the correlation veto threshold to be $E_{\text{kin}} > 0.6$ GeV. At both the shield veto and detector we apply an additional suppression of neutral secondaries according to their charged-neutral correlation.

In order to incorporate statistical uncertainties in the `Geant4` simulation, an array of 50 pseudo-datasets are generated by Poisson-distributing the statistics of each simulated particle-gun event. Thus in practice one obtains 50 separate transfer matrix compositions for the shield simulation, from which the statistics of overall shield performance and uncertainties may be extracted.

In Table 2 we show the results of this simulation for $(20 + 5)\lambda$ shield configuration made of Pb, plus the 3 m concrete UXA wall, with a shield veto efficiency of $1 - \varepsilon_{\text{veto}} = 10^{-4}$. For outgoing neutral particle fluxes in Table 2, a kinetic energy cut $E_{\text{kin}} > 0.4$ GeV was applied, as required by minimum tracking requirements to produce at least one track. Table 2 includes the background fluxes rejected by the shield veto, both with and without application of the charged-neutral correlation veto in the detector. The ~ 60 neutrons may each produce a single track scattering event along the 10 m depth of detector (see Sect. 3.2.5). The neutron incoherent scattering cross-section on air is ~ 1 b, so that the probability of a neutron scattering on air into two tracks along the 10 m depth of detector is at most $\sim 5\%$. Requiring neutrons with $E_{\text{kin}} > 0.8$ GeV for at least two tracks, results in a total neutron flux of ~ 3 per 300 fb^{-1} , so that the net scattering rate to two or more tracks is < 1 . The shield veto rejection rate for the $(20 + 5)\lambda$ configuration is $\simeq 2.2$ kHz, assuming the projected instantaneous luminosity of $10^{-34} \text{ cm}^{-2} \text{ s}^{-1}$ at IP8. (By comparison, for only 15λ of Pb before the veto, this would increase to $\simeq 6.0$ kHz.) This is dominated mainly

Fig. 22 IP production cross-section per kinetic energy bin, for minimum bias (QCD), Heavy Flavor (HF), and Drell–Yan (DY) production channels



by the incoming muon flux, and is far smaller than the total event rate, and therefore has a negligible effect on the detector efficiency.

One sees in Table 2 that the shield veto is crucial to achieve a zero background environment. Moreover, for a low, rather than zero, background environment, the shield veto data provides a data-driven means to calibrate the background simulation, from which any residual backgrounds in the detector can then be more reliably estimated and characterized.

In Fig. 23 we show the net background flux distributions in kinetic energy (blue) for a variety of neutral and charged

species, including uncertainties, and without any E_{kin} cuts. These may be compared to the background flux distributions of particles reaching the detector that are rejected by the shield veto (red, with 10^{-4} scaling) and the IP fluxes (green, with 10^{-12} scaling).

3.2.3 Neutrinos

An additional background may arise through production of neutral secondaries from neutrinos, that stream through the shield unimpeded. In particular, with a sufficiently high

Table 2 Results from the Geant4 background simulation for (20 + 5)λ Pb shield, i.e. with an active shield veto at 20λ, applying a veto efficiency of $1 - \epsilon_{\text{veto}} = 10^{-4}$. For outgoing neutral particles a kinetic energy cut $E_{\text{kin}} > 0.4\text{ GeV}$ was applied as required by minimum tracking requirements, except for anti-neutrons in order to exclude $\bar{n} + N$ annihilation processes. We also show the rate for neutrons with $E_{\text{kin}} > 0.8\text{ GeV}$, required for production of at least two tracks via scattering. For total luminosity $\mathcal{L} = 300\text{ fb}^{-1}$, the column “Net ($E_{\text{kin}}^{\text{neutral}} > 0.4\text{ GeV}$)” shows the net background particle yields after traversing the shield plus veto rejection, including veto correlations

(denoted ‘±/0’) between charged particles with $E_{\text{kin}} > 0.6\text{ GeV}$ and neutral particles. The column “Shield veto rejection (total)” shows the corresponding background particle yields entering the detector subject to the shield veto rejection alone, without applying the charged-neutral correlation veto. The column “Shield veto rejection (±/0 correlation)” shows the corresponding background particle yields entering the detector subject to the shield veto rejection, after application of the charged-neutral correlation veto on the detector front face. The final column lists the net background yield including detector rejection, scattering or decay probabilities

BG species	Particle yields			Net yield
	Net ($E_{\text{kin}}^{\text{neutral}} > 0.4\text{ GeV}$)	Shield veto rejection (total)	Shield veto rejection (±/0 correlation)	
γ	0.54 ± 0.12	$(8.06 \pm 0.60) \times 10^4$	$(2.62 \pm 1.03) \times 10^3$	–
n	58.10 ± 4.63	$(4.59 \pm 0.15) \times 10^5$	$(3.44 \pm 0.51) \times 10^4$	–
$n (> 0.8\text{ GeV})$	2.78 ± 0.25	$(1.03 \pm 0.06) \times 10^5$	$(7.45 \pm 1.92) \times 10^3$	$\lesssim 1$
\bar{n} (no cut)	$(3.24 \pm 0.72) \times 10^{-3}$	34.40 ± 25.80	$(7.12 \pm 2.19) \times 10^{-2}$	$\ll 1$
K_L^0	0.49 ± 0.05	$(1.94 \pm 0.74) \times 10^3$	54.40 ± 19.20	$\lesssim 0.1$
K_S^0	$(6.33 \pm 1.39) \times 10^{-3}$	93.90 ± 45.80	0.74 ± 0.19	$\ll 1$
$\nu + \bar{\nu}$	$(5.69 \pm 0.00) \times 10^{13}$	$(7.35 \pm 0.12) \times 10^6$	$(7.31 \pm 0.11) \times 10^6$	–
p^\pm	$(2.07 \pm 0.26) \times 10^2$	$(9.24 \pm 0.36) \times 10^5$	$(9.24 \pm 0.36) \times 10^5$	–
e^\pm	$(4.53 \pm 0.02) \times 10^3$	$(4.38 \pm 0.02) \times 10^7$	$(4.38 \pm 0.02) \times 10^7$	–
π^+	34.70 ± 2.27	$(2.96 \pm 0.20) \times 10^5$	$(2.96 \pm 0.20) \times 10^5$	–
π^-	31.40 ± 2.12	$(2.68 \pm 0.19) \times 10^5$	$(2.68 \pm 0.19) \times 10^5$	–
K^+	0.83 ± 0.30	$(3.08 \pm 1.24) \times 10^3$	$(3.08 \pm 1.24) \times 10^3$	–
K^-	0.23 ± 0.12	$(1.12 \pm 0.63) \times 10^3$	$(1.12 \pm 0.63) \times 10^3$	–
μ^+	$(1.04 \pm 0.00) \times 10^6$	$(1.04 \pm 0.00) \times 10^{10}$	$(1.04 \pm 0.00) \times 10^{10}$	–
μ^-	$(8.07 \pm 0.01) \times 10^5$	$(8.07 \pm 0.01) \times 10^9$	$(8.07 \pm 0.01) \times 10^9$	–

neutrino flux, $\bar{\nu}p \rightarrow \ell n$ quasi-elastic scattering may produce a non-negligible amount of neutrons in the last few λ that reach the detector volume, while the charged lepton is too soft or misses the acceptance. (The cross-section for the neutral current scattering $\nu n \rightarrow \nu n$ or $\bar{\nu}n \rightarrow \bar{\nu}n$ is approximately 10 times smaller than for the charged current process [189].) From Table 2 and Fig. 23, approximately 5×10^{13} neutrinos are produced per 300 fb^{-1} at IP8 in the CODEX-b acceptance, with $E_\nu > 0.4\text{ GeV}$. The neutrino flux is approximately power-law suppressed by a quartic above $E_\nu \sim 1\text{ GeV}$. Hence, although the charged current cross-section for $\bar{\nu}p \rightarrow \ell n$ is only $\sim 0.01(E_\nu/\text{GeV})\text{ pb}$ [189], the large flux of $\mathcal{O}(\text{GeV})$ neutrinos streaming through the shield implies as many as ~ 10 neutrons might be generated per λ of concrete in the UX A wall with $E_{\text{kin}} > 0.4\text{ GeV}$. Neutral kaon production, such as $\nu n \rightarrow \nu K_L^0 \Lambda$, has a cross-section $\sim 0.1\text{ fb}$ for $E_\nu \sim 3.5\text{ GeV}$ [190], scaling approximately linearly with neutrino energy, and therefore may be safely neglected.

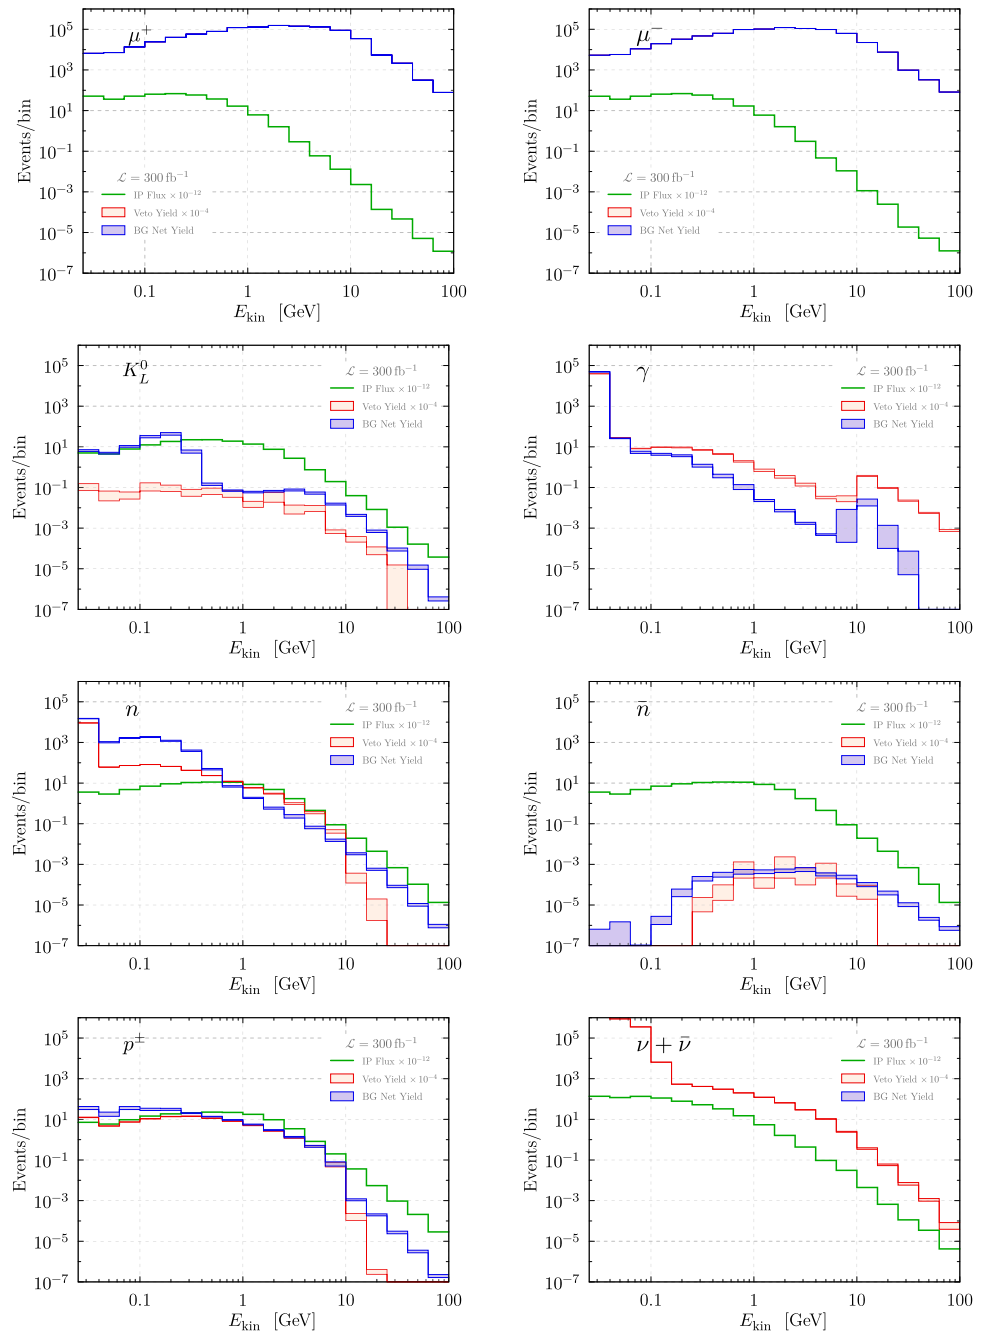
Composing the IP flux of anti-neutrinos in Fig. 22 with $E_\nu > 0.4\text{ GeV}$ against the measured energy-dependent $\bar{\nu}p \rightarrow \ell n$ cross-section [189], and including subsequent

attenuation as characterized by the nuclear interaction length λ, one may estimate a conservative upper bound on the number of neutrons that might reach the detector with $E_{\text{kin}} > 0.4\text{ GeV}$. From this procedure one finds that there are at most approximately 5 neutrino-produced neutrons of this type. Since this estimate is extremely conservative, neutron production from neutrinos is expected to be negligible compared to secondary neutrons from other primary fluxes.

3.2.4 Shielding marginal performance

As the detector tolerance of backgrounds may vary depending on the ultimately implemented detector technologies, it is instructive to assess the performance of the shield under variation of the shielding configuration, including variation of the total shield depth, L_{shield} , and the placement and efficiency of the shield veto. We illustrate the marginal changes in shielding performance in terms of the total neutron and K_L^0 fluxes for kinetic energy $E_{\text{kin}} > 0.4\text{ GeV}$, by varying the shield configuration as combinations of the 5 or 2λ transfer matrices permit, and by permitting the veto efficiency to range from $1 - \epsilon_{\text{veto}} = 10^{-5}$ up to 10^{-2} . In the top panel of Fig. 24, the

Fig. 23 Background fluxes per kinetic energy bin, comparing primary IP fluxes $\times 10^{-12}$ (green) with the irreducible background flux (blue) after the $(20 + 5)\lambda$ passive and active shield. Also shown are background fluxes $\times 10^{-4}$ entering the detector that are rejected by the shield veto (red)



corresponding variation of the neutron (black-blue palette) and K_L^0 (red-yellow palette) background fluxes are shown, taking all combinations of $L_{\text{pre-veto}} \in \{15, 17, 19, 20\}\lambda$ and $L_{\text{post-veto}} \in \{4, 5, 6\}\lambda$, as defined in Fig. 21, and $1 - \varepsilon_{\text{veto}} \in \{10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}\}$.

The simulated background fluxes are generally insensitive to marginal variation in the location of the shield veto, e.g. one sees that $(19 + 4)\lambda$ performs similarly to the nominal $(20 + 5)\lambda$ configuration. At very high efficiencies, i.e. $1 - \varepsilon_{\text{veto}} < 10^{-4}$, both background fluxes are roughly expo-

entially distributed in shield depth. In this case the backgrounds are either unsuppressed primaries or stopped-parent secondaries produced upstream from the shield.

As the shield veto efficiency is reduced, however, one sees a departure from the exponential suppression: Contributions from stopped-parent secondaries produced downstream from the shield veto begin to dominate. For the K_L^0 flux, this departure occurs only at $1 - \varepsilon_{\text{veto}} > 10^{-2}$ and at a larger L_{shield} , compared to the neutrons. This arises because of a somewhat larger charged-neutral correlation for production of K_L^0 's:

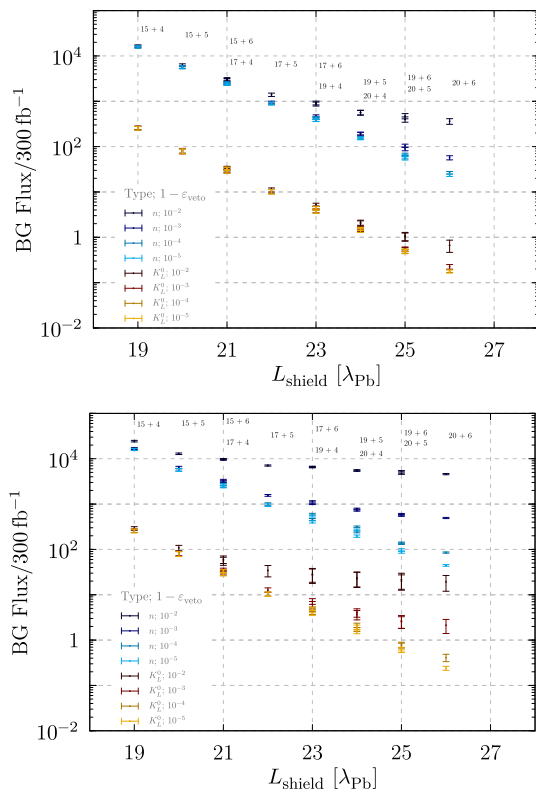


Fig. 24 Neutron (black-blue palette) and K_L^0 (red-yellow palette) background fluxes with kinetic energy $E_{kin} > 0.4\text{ GeV}$ versus total Pb shield depth in λ_{Pb} , under variation of the shield configuration and veto efficiency, including the charged-neutral correlation veto (top) and without the charged-neutral correlation veto (bottom). For each value of the total shield depth, the (possibly multiple) corresponding configurations “ $L_{pre-veto} + L_{post-veto}$ ” are shown by the adjacent labels in units of λ_{Pb}

Their parent muons are typically somewhat hard and may reach the veto or detector.

One may assess the degree of this effect – effectively, the amount of non-stopped parent secondaries – by considering the case that the charged-neutral correlation veto is not applied (in practice, it may happen that the associated charged particles do not always trigger the veto). We show the corresponding shielding performance in the bottom panel of Fig. 24. For $1 - \epsilon_{veto} > 10^{-3}$, the background fluxes become substantially larger. One deduces that, especially for K_L^0 s, charged parents that produce secondaries downstream of the shield may typically reach the detector. In Fig. 25 we show the neutron and K_L^0 fluxes in the absence of a shield veto ($1 - \epsilon_{veto} = 1$), both with (light palette) and without (dark palette) the charged-neutral correlation veto in the detector (denoted ‘ $\pm/0$ ’). One sees that while the charged-neutral correlation veto can significantly reduce the total fluxes, substantial net backgrounds for K_L^0 s and neutrons remain.

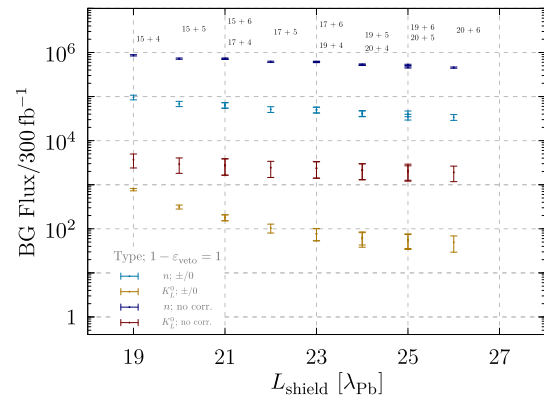


Fig. 25 Variation of neutron and K_L^0 background fluxes with kinetic energy $E_{kin} > 0.4\text{ GeV}$ versus total Pb shield depth in λ_{Pb} under variation of the shield configuration with no shield veto ($1 - \epsilon_{veto} = 1$), including the charged-neutral correlation veto in the detector (light palette) and without the charged-neutral correlation veto in the detector (dark palette). For each value of the total shield depth, the (possibly multiple) corresponding configurations “ $L_{pre-veto} + L_{post-veto}$ ” are shown by the adjacent labels in units of λ_{Pb}

3.2.5 Simulated track production

The simulated neutral background fluxes entering the detector may be folded against the probability of scattering into one or more tracks on material inside the detector, or against the probability of decays into one or more tracks in the detector interior: Nominally the number of two or more tracks should be < 1 to ensure a background-free environment. In Table 3 we show the rates of multitrack production for 1–9 tracks from scattering of the neutral fluxes on air in the $10 \times 10 \times 10\text{ m}^3$ detector volume, for the nominal $(20 + 5)\lambda$ Pb shield configuration with $1 - \epsilon_{veto} = 10^{-4}$. This production is simulated with Geant4 as in Sect. 3.2.2, requiring each track to have kinetic energy $E_{kin} > 0.4\text{ GeV}$.

For total luminosity $\mathcal{L} = 300\text{ fb}^{-1}$, one sees that the total number of scatterings or decays into two or more tracks is $\simeq 0.22 \pm 0.03$. This comports with our simulation and estimation of the background effective yields in Table 2.

3.3 Measurement campaign

To verify the background simulation, a data-driven calibration is needed, using data taken during collisions at IP8. This section serves as a brief summary of an initial measurement campaign that was undertaken during Run 2 operations at various locations in the UXA cavern, shielded only by the UXA wall, in August 2018. For a detailed description we refer to Ref. [191], which also features additional detailed background simulations, including the effects of infrastructure in the LHCb cavern.

The detector setup used scintillators, light-guides and photomultiplier tubes (PMT) from the HeRSChel experiment

Table 3 Multitrack production on air from the Geant4 background simulation for the $(20 + 5)\lambda$ Pb shield, in the $10 \times 10 \times 10 \text{ m}^3$ detector volume for total luminosity $\mathcal{L} = 300 \text{ fb}^{-1}$, requiring $E_{\text{kin}} > 0.4 \text{ GeV}$

Tracks	$(20 + 5)\lambda$ Pb shield $1 - \varepsilon_{\text{veto}} = 10^{-4}$	Run 3 (CODEX- β)	Run 3 (CODEX- β) K_L^0 contribution
1	53.90 ± 5.51	$(3.87 \pm 0.11) \times 10^8$	$(2.94 \pm 0.07) \times 10^8$
2	0.21 ± 0.02	$(4.09 \pm 0.13) \times 10^7$	$(3.74 \pm 0.13) \times 10^7$
3	$(1.36 \pm 0.34) \times 10^{-2}$	$(5.96 \pm 1.01) \times 10^5$	$(2.92 \pm 0.45) \times 10^5$
4	$(1.51 \pm 0.30) \times 10^{-3}$	$(6.78 \pm 1.22) \times 10^4$	$(5.12 \pm 1.19) \times 10^4$
5	$(3.80 \pm 0.87) \times 10^{-4}$	$(1.69 \pm 0.50) \times 10^4$	$(1.42 \pm 0.50) \times 10^4$
6	$(1.09 \pm 0.27) \times 10^{-4}$	$(3.23 \pm 0.79) \times 10^3$	$(2.21 \pm 0.79) \times 10^3$
7	$(1.84 \pm 1.41) \times 10^{-4}$	$(4.23 \pm 2.30) \times 10^3$	$(1.75 \pm 0.77) \times 10^3$
8	$(2.98 \pm 1.31) \times 10^{-5}$	$(1.04 \pm 0.63) \times 10^3$	$(8.45 \pm 6.11) \times 10^2$
9	$(1.07 \pm 0.33) \times 10^{-5}$	$(2.41 \pm 0.43) \times 10^2$	$(1.37 \pm 0.35) \times 10^2$

per track. Also shown are corresponding rates for total neutral and K_L^0 multitrack production during Run 3 in the CODEX- β volume for total luminosity $\mathcal{L} = 15 \text{ fb}^{-1}$ (see Sect. 4)

[192] at LHCb. The detector itself consisted of two parallel scintillator plates with surface area $300 \times 300 \text{ mm}^2$: Details can be found in Ref. [191]. Before transporting the setup to IP8 it was tested with cosmic rays, indicating an efficiency $> 95\%$ for minimum ionizing particles (MIPs), i.e. for those particles with kinetic energy $\gtrsim 100 \text{ MeV}$. In the simulation of Ref. [191], it was further verified that no collision event produced more than one hit on the scintillator acceptance, such that pile up of hits within the trigger window can be ruled out.

A two-fold coincidence between both scintillator plates was required to trigger the detector, and the trigger was not synchronized with the collisions. However, the coincidence time-window was 5 ns , much shorter than the 25 ns collision frequency at IP8, so that spill-over effects can be neglected. Two waveforms were recorded from each scintillator and the timestamp for all MIP hits. This timestamp was used to correlate the events with the beam status during data-taking.

The measurements were taken on the ‘‘D3 platform’’ level in the UXA hall, behind the concrete UXA wall. Figure 26 shows this platform, and the different locations and configurations used for the data-taking. The detector was deployed at three different positions on the passerelle between Data Acquisition (DAQ) server racks and the UXA wall, as well as at one location between the DELPHI barrel exhibit and the DAQ racks. The scintillator stand was oriented either parallel (‘‘||’’), rotated 45° or perpendicular (‘‘ \perp ’’) to the beam line.

The measurement period spanned 17 days between 25th July and 10th August, 2018 with 52,036 recorded triggers observed during the run. The instantaneous luminosity at IP8 was stable during the measurement period. There was no beam until July 30th due to machine development and an inadvertent power loss. Figure 27 show the main results from the measurement campaign. The red data points represent the instantaneous luminosity measured by LHCb in

Hz/nb. The green and blue data points indicate the hit rate in Hz, where the setup was alternated between the six different configurations/positions. The plots are shown in both linear and logarithmic scales.

Table 4 contains the hit rate from ambient background without beam, with an average hit rate at each position and configuration of 2 mHz . The ambient background can therefore be considered negligible for this measurement. Table 5 shows the rate during stable beam. This rate is non-negligible, even for the small $300 \times 300 \text{ mm}^2$ area of the scintillators. The rate increases from location P1 to P2 to P4, which, from Fig. 26, implies that there is more activity in the downstream region. This dependence on the η arises from additional concrete near IP8, which screens part of the CODEX-b acceptance, see Ref. [191]. Moreover, by comparing the rate at P2 with P5, behind the DAQ server racks, one can see that the racks are adding some amount of shielding material. As expected, the flux also depends on the orientation with respect to the beam direction, as indicated by the difference in rate between P5 and P6. In absolute numbers, the rate just behind the concrete wall is roughly 0.5 Hz over the 900 cm^2 scintillator area.

The predicted charged particle flux from the Geant4 simulation of Sect. 3.2 with just the concrete UXA wall acting as a shield, predicts a hit rate $\sim 10 \text{ Hz}$ at position P2, assuming an instantaneous luminosity $\sim 0.4 \text{ Hz nb}^{-1}$, as in Fig. 27. This reduces to $\sim 5 \text{ Hz}$ treating the full width of UXA wall as 7.5λ of standard concrete. This prediction will likely further receive $\mathcal{O}(1)$ reductions from: relaxing our conservative treatment of forward-propagating backgrounds under angular rescattering; accounting for the longer propagation path length through the wall at higher angles of incidence; variations or uncertainties in the simulation of the primary muon fluxes; and accounting for possibly additional material in the line-of-sight, such as concrete nearby the IP and the platform

Fig. 26 The four measurement locations on the D3 level in the LHCb cavern, shown by red, orange, green and blue dots. The configurations are labelled P1–P6. Figure reproduced and modified from Ref. [191], as adapted from Ref. [8]

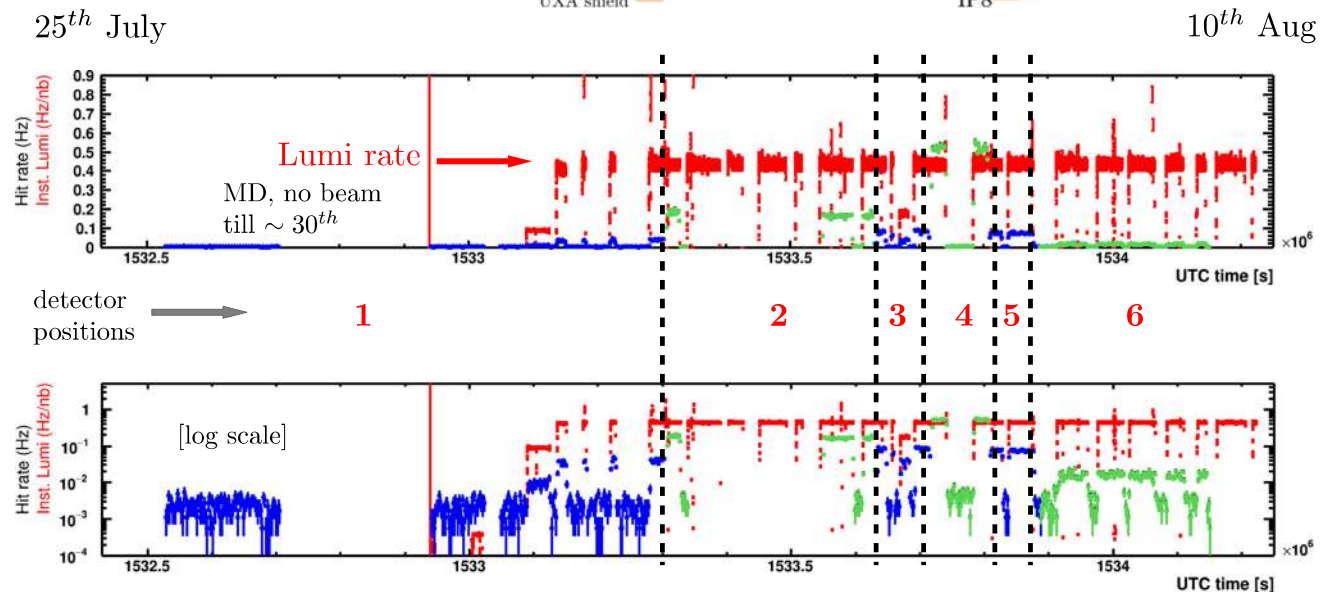
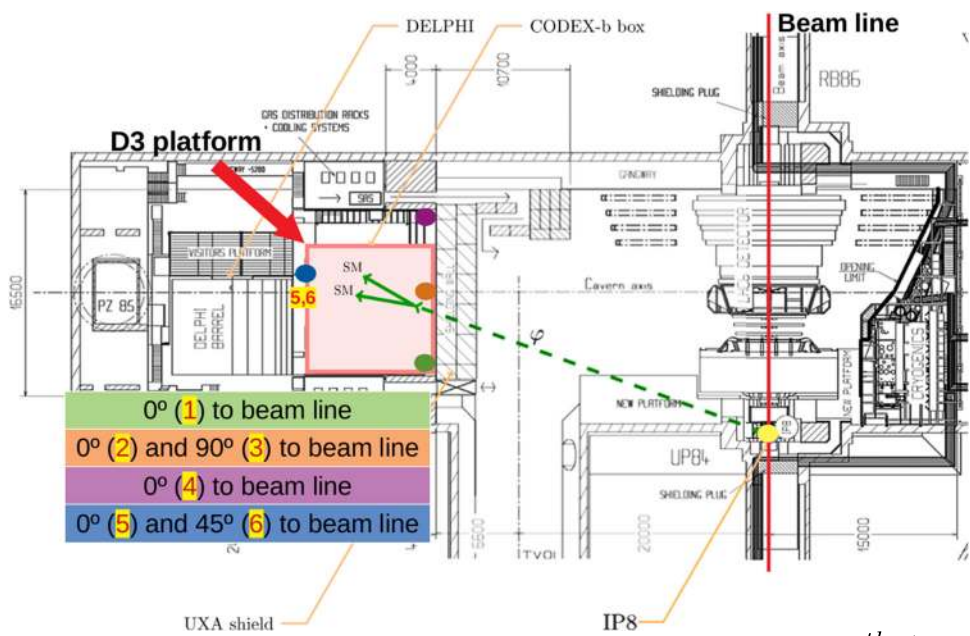


Fig. 27 Hit rates during the run based on the six P1–P6 positions/configurations on a linear (top) and log (bottom) scale. Red data points denote the luminosity rate of LHCb, blue and green data points denote hit rates. Figure reproduced from Ref. [191]

Table 4 Background hit rates based on each position/configuration with no beam. Table reproduced from Ref. [191]

Position	Description	Hit rate (mHz)
P1	UXA wall, right corner, to beam	1.99 ± 0.07
P2	UXA wall, center, to beam	2.76 ± 0.03
P3	UXA wall, center, \perp to beam	2.26 ± 0.03
P4	UXA wall, left corner, to beam	3.11 ± 0.03
P5	UXA wall + DAQ racks, center, to beam	1.95 ± 0.03
P6	UXA wall + DAQ racks, center, 45° to beam	2.22 ± 0.02

Table 5 Average hit rates measured during stable beam, at various configurations. Table reproduced from Ref. [191]

Position	Description	Hit rate (mHz)
P1	UXA wall, right corner, to beam	38.99 ± 0.99
P2	UXA wall, center, to beam	167.10 ± 1.43
P3	UXA wall, center, \perp to beam	82.81 ± 1.55
P4	UXA wall, left corner, to beam	517.45 ± 3.52
P5	UXA wall + DAQ racks, center, to beam	73.58 ± 1.18
P6	UXA wall + DAQ racks, center, 45° to beam	15.71 ± 0.33

adjacent to the LHCb magnet. A more complete estimate requires a full simulation of the LHCb cavern itself, such as described in Ref. [191]. Nonetheless comparison to the measured 0.2 Hz rate demonstrates that the Geant4 simulation of Sect. 3.2 provides conservative estimates of the expected backgrounds.

4 CODEX- β

To validate the CODEX-b concept, a proposal has been developed for a small, $2 \times 2 \times 2 \text{ m}^3$ demonstrator detector – “CODEX- β ” – which will be operational during Run 3. This detector will be placed in the proposed location for CODEX-b (UXA hall, sometimes referred to as the ‘DELPHI cavern’) shielded only by the existing, concrete UXA radiation wall.

4.1 Motivation

The main goals of the CODEX- β setup are enumerated as follows:

- (a) *Demonstrate the ability to detect and reconstruct charged particles which penetrate into the UXA hall, as well as the decay products of neutral particles decaying within the UXA hall.*

This is desirable to provide an accurate and fully data-driven estimate of the backgrounds, so that the design of the eventual shield (both passive and active, instrumented) needed by the full experiment can be optimized to be as small as possible. We have already made preliminary background measurements in the UXA hall using a pair of scintillators during Run 2 (see Sect. 3.3). However these measurements were simply of hit counts; we could not reconstruct particle trajectories. CODEX- β will allow us to track particles within a volume similar to the CODEX-b fiducial volume, and in particular separate charged particles produced outside the decay volume from backgrounds induced by particle scattering inside the decay volume itself.

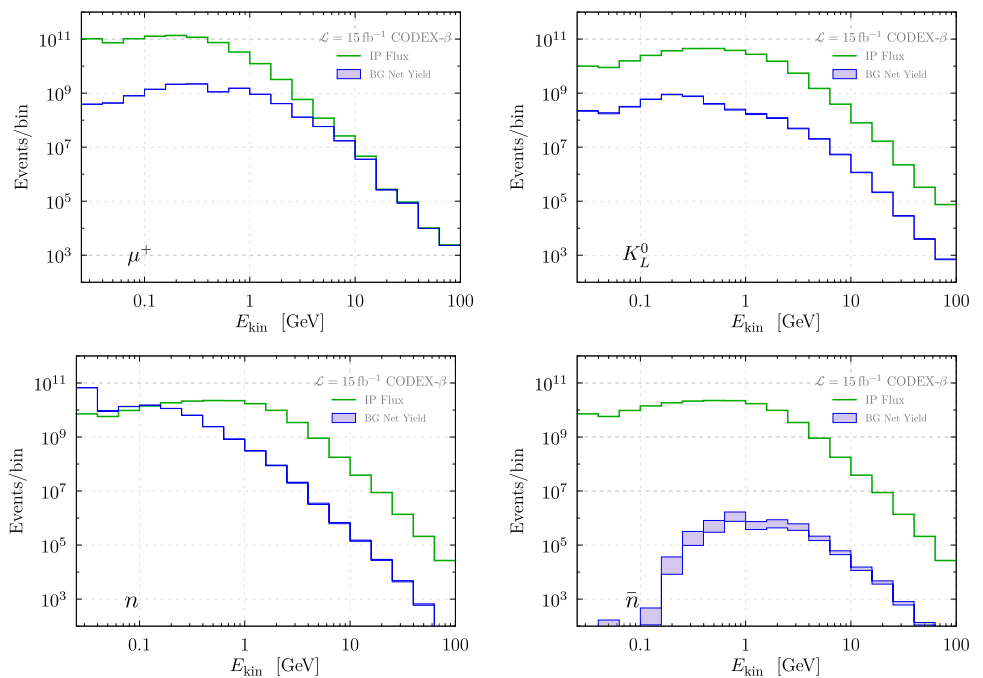
As shown in Sect. 3, the residual neutron scattering inside the decay volume is one of the most important backgrounds identified in the original CODEX-b proposal [8]. The tracking capability of CODEX- β will also allow us to measure the origin of charged particle backgrounds, and in particular potential soft charged particles which could be swept towards CODEX-b by LHCb’s magnet “focusing” and thereby evade the Pb shield.

- (b) *Detect and reconstruct a significant sample of neutral particles decaying inside the hermetic detector volume.*

This will allow us to observe e.g. K_L^0 decays and use this data to calibrate our detector simulation. Aside from measuring background levels, observing long-lived SM particles decaying inside the detector acceptance will allow us to calibrate the detector reconstruction and the RPC timing resolution. The most natural candidates are K_L^0 mesons: In Fig. 28 we show the expected differential fluxes of neutrons, antineutrons, K_L^0 s and (anti)muons, with respect to their kinetic energy, for an integrated luminosity of 15 fb^{-1} after propagation through the UXA wall. Also shown are the primary fluxes of the same species. In Table 3 we show the expected multitrack production from decay or scattering on air by neutral fluxes entering CODEX- β , requiring $E_{\text{kin}} > 0.4 \text{ GeV}$ per track. We also show the multitrack contribution just from K_L^0 s entering CODEX- β .

One sees that approximately a few $\times 10^7$ K_L^0 decays to two tracks are expected in the CODEX- β volume per nominal year of data taking in Run 3. The results of the background simulation show that we will be able to reconstruct a variety of K_L^0 decays in the CODEX-b demonstrator volume. The decay vertex and decay product trajectories moreover allow the boost to be reconstructed independently of the time-of-flight information. Comparing the boost distribution of K_L^0 mesons observed in CODEX-b, as well as the K_L^0 mean decay time which can be inferred from this distribution, will allow us to calibrate and validate our detector simulation and reconstruction.

Fig. 28 Background fluxes per kinetic energy bin, comparing primary IP fluxes (green) with the background flux (blue) in the CODEX- β acceptance after passage through the 3 m UXA concrete wall



(c) Show that CODEX-b can be integrated into the LHCb DAQ and demonstrate an ability to give a trigger to LHCb to retain an event that looks interesting in CODEX-b.

The RPC readout is compatible with LHCb’s data acquisition hardware. Some relatively straightforward firmware development will be required to enable LHCb’s usual FPGA backend readout boards to receive the CODEX-b data. Based on expected data rates, we estimate that a single FPGA backend readout board will be comfortably able to read out the full CODEX- β detector. From an LHCb point of view the simplest solution would be that this board also clusters the RPCs and performs a basic track reconstruction, so that events which look interesting for CODEX- β can be kept for further inspection by LHCb’s High Level Trigger simply by reading the CODEX-b data raw bank. Given that CODEX-b is about the same distance from the interaction point as LHCb’s muon system, latency should not be an issue. Our background measurements in the cavern during Run 2 indicated hit (not track) rates of maximum 500 mHz across a scintillator area of order 10^{-1} m^2 . Therefore, even a simple track reconstruction should allow all interesting events in CODEX-b to be kept for offline inspection. It will be desirable to have a possibility to read CODEX-b out during beam-off periods, for cosmic-ray data taking and calibration. CODEX-b will therefore ideally appear as a sub-detector within LHCb, though one whose presence/readiness is not required for nominal LHCb data taking.

(d) Integrate the detector into the gas and electricity services in the UXA hall, and demonstrate stable operational behavior.

(e) Demonstrate a viable and modular support structure for the RPC layers, to form a basis for the eventual support structure for the full detector.

(f) Search for multi-track signatures from BSM physics: Despite its limited acceptance and large backgrounds, CODEX- β is expected to have some new reach for LLPs produced in exotic B-meson decays. This is discussed below in Sect. 4.3.

4.2 Technical description and timetable

The high-level requirements listed above drive the design of CODEX- β to be a $2 \times 2 \times 2 \text{ m}^3$ cube. Each side of the cube will consist of 2 RPC panels, each of which is $2 \times 1 \text{ m}^2$ in area. Each such panel block will contain a triplet of RPC layers. In addition there will be two panels of the same $2 \times 1 \text{ m}^2$ area placed in the middle of the cube, for a total of $(6+1) \times 2 \times 3 = 42$ such $2 \times 1 \text{ m}^2$ RPC layers. CODEX- β is proposed for installation in the barrack which housed LHCb’s Run 1 and Run 2 High Level Trigger farm, and which will be empty in Run 3 as the High Level Trigger will be housed in a dedicated data processing center on the surface. As a result CODEX- β will have ample space and straightforward access to all required detector services. The proposed detector technology for CODEX- β is that of the ATLAS RPCs for phase I upgrade, while the full CODEX-b detector would follow the phase II design (see [193] for technical descriptions).

The timetable for installation is driven by the primary consideration to not interfere with the building or commissioning of the LHCb upgrade. For this reason, we originally proposed installation in winter 2021/2022, integration in the LHCb DAQ during spring 2022 and first data taking in summer

2022. Given the COVID pandemic and the subsequent delay of LHC Run 3 datataking to spring 2022, we have chosen to push back installation and all subsequent steps by one year. Although this will mean missing out on 2022 datataking, that is not so crucial for a demonstrator and avoids having to simultaneously commission CODEX- β and the upgraded LHCb detector. Given the modest size of CODEX- β and the use of well-understood detector components, we estimate around six months are needed to produce and qualify the RPCs. Therefore, if approved, it is realistic to complete the bulk of the construction during the first half of 2021. The mechanical support structure will build on existing structures used in ATLAS but modified to be modular. It will provide the required stability for a cubic arrangement of the detector layers; the design and construction of this structure is expected to take place in 2020–2021. The total cost of the detector components is expected to be roughly 150k €.

4.3 New physics reach

The acceptance of CODEX- β is roughly only 8×10^{-3} times that of the full CODEX-b detector, and no shielding beyond the existing concrete wall will be in place. Its reach for BSM physics is therefore limited due to its reduced acceptance and high background environment. However, roughly 10^{13} b -hadrons will be produced at IP8 during Run 3. This enables CODEX- β to probe some new regions of parameter space for those cases in which the LLP production branching ratio from e.g. B decays is independent from its lifetime (cf. Fig. 10).

This scenario can arise e.g. in models that address the baryogenesis puzzle [66,67,166–168], as described in Sect. 2.4.9. We take this model as a representative example: In our simplified phenomenological setup we consider a new particle χ , with a coupling

$$\mathcal{L} \supset \lambda_{ijk} \chi u_i d_j d_k, \quad (25)$$

where we assume for simplicity that the λ_{bsu} and λ_{udd} couplings are independent. The former is responsible for the production via e.g. $B \rightarrow X_s \chi$ decays, where X_s here is a SM (multi)hadronic state with baryon number ± 1 ; the latter induces the decay of χ to an (anti-)baryon plus a number of light mesons. $\text{Br}[B \rightarrow X_s \chi]$ and $c\tau$ are then independent parameters. The λ_{udd} coupling moreover must be parametrically small, to avoid exotic dinucleon decays [194,195], implying that χ must be long-lived.

In Run 3, as shown in Fig. 28 we expect roughly 10^9 K_L and 4×10^9 neutrons to enter the CODEX- β fiducial volume for an integrated luminosity of 15 fb^{-1} , requiring $E_{\text{kin}} > 0.4 \text{ GeV}$. This is desirable for calibration purposes, as explained above. An additional 4×10^6 antineutrons also enter, with no kinetic energy cut. We show the corresponding multitrack production rate for CODEX- β in Table 3. The multitrack background falls relatively fast with the number

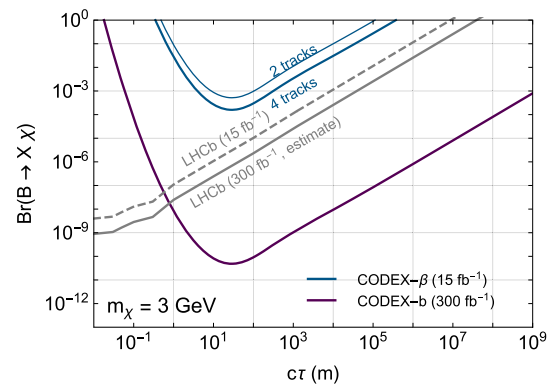


Fig. 29 CODEX- β reach for a long-lived χ decaying hadronically

of tracks, partly because of the relative softness of the fluxes emanating from the shield and partly because the K_L^0 s mainly decay to no more than two tracks. We therefore define the LLP signal region as those events with 4 or more reconstructed tracks, also requiring $E_{\text{kin}} > 0.4 \text{ GeV}$ from expected minimum tracking requirements. The expected number of background events in the signal region is then roughly 8.5×10^4 per 15 fb^{-1} . In the actual experimental setup this number can be calibrated from a control sample with less than 4 tracks, if the ratio of both regions is taken from Monte Carlo.

For a signal benchmark with $m_\chi = 3 \text{ GeV}$, the probability of decaying to 4 tracks with $E_{\text{kin}} > 0.4 \text{ GeV}$ is roughly 15%, as estimated with Pythia 8. In Fig. 29 we show the estimated 2σ limit reach under these assumptions. For comparison, we also show the reach of the full CODEX-b detector and the reach using a ≥ 2 tracks selection, for which the background is roughly three orders of magnitude higher. (This is partially compensated for by a higher signal efficiency.) At this stage, no attempt has been made to discriminate signal from background by making use of angular variables, in particular pointing to the interaction point. In this sense the estimated reach is therefore conservative.

For completeness, we also include a preliminary estimate of the reach of LHCb itself for this signature with 15 fb^{-1} of data, analyzing all decay products that can be reconstructed as a track at LHCb: e, μ, p, K^\pm and π^\pm . For reconstructing a χ vertex, we first require all pairs of tracks to be vertexed not more than 1 mm away from each other. We build the position of these vertices by finding the point that minimizes the distance to each pair of tracks. We then average all the resulting vertices to generate the χ decay vertex. In order to build the B^+ decays, the χ vertex is required not to be more than 1 mm away from a K^+ . For the background, SM B^+ decays are considered, subject to the same reconstruction criteria. All other backgrounds are neglected.

The analysis cuts are included in Table 6. For the rest of the experimental efficiencies, we estimate a 97% efficiency

Table 6 List of analysis cuts for the LHCb reach estimate for a hadronically decaying, long-lived particle (χ). SV_z and SV_R respectively stand for the longitudinal and transverse position of the secondary vertex, and θ is the angle of the track with respect to the beam axis

Cut	Value
p_T of χ daughters	> 500 MeV/ c
θ of χ daughters	< 400 mrad
SV_z	< 400 mm
SV_R	$\in [14, 25]$ mm
p_T of K^+	> 500 MeV/ c
θ of K^+	< 400 mrad

per track [196], and take the remaining efficiencies to be 100%. Signal and background were binned according to 4 or more tracks and 6 or more tracks in the secondary vertex. We estimate the mass resolution, σ , for 4 and 6 body-decays of the χ particle to be ~ 12 and 21 MeV, respectively. This estimate is based on a study of 3 and 4-body B and D meson decays at LHCb [197–200], interpolating or extrapolating to the appropriate track multiplicity. Following a similar procedure, we estimate the B^+ meson mass resolution to be ~ 24 and ~ 36 MeV for 5 and 7 body decays, respectively. To determine the background yields, we cut in $\pm 2\sigma$ windows around the B^+ and χ invariant masses. To determine the limits, we take σ_{bb} at $\sqrt{s} = 14$ TeV to be $500 \mu b$ [201], and the fraction of b quarks hadronising to a B^+ is taken to be 40% [202]. Combining this, we compute the limits on the branching ratio of the B^+ decay for both the 4+ and 6+ track bins. The projected limit shown in Fig. 29 is the strongest of both limits, for each $c\tau$ point. In addition, we show a rough estimate of the reach for the HL LHC, by rescaling the limits with the square root of the ratio of the luminosities.

One sees that CODEX- β and the main LHCb detector will have complementary sensitivity to this benchmark scenario, with likely better sensitivity from the LHCb search. However, it is conceivable that CODEX- β may set an earlier limit than an LHCb analysis on Run 3 data, especially given the comparatively simpler analysis required for the former. In both estimates no attempt was made to further reduce the backgrounds by means of kinematic cuts, so both projections are conservative.

5 Detector case studies

In this section we discuss various detector studies, as well as possible extensions of the baseline CODEX-b detector configuration. As CODEX- β is based on the same underlying technology as the full detector, most results apply directly to it as well. One caveat, however, is that CODEX- β will use a simplified front-end readout based on FPGA cards and there-

fore have a significantly poorer timing resolution of around $800/\sqrt{12}$ ps per gas gap. This resolution will nevertheless be comfortably sufficient to integrate CODEX- β into the LHCb readout and to validate the detector concept.

5.1 Tracking

5.1.1 Design drivers

The geometry and required capabilities of the tracking stations are informed by the signal benchmarks in Sect. 2. The main design drivers are:

(a) Hermeticity:

As discussed in Sect. 1.2, the primary motivation for CODEX-b is to cover relatively low energy signals, as compared to e.g. SUSY signatures. In many benchmark models (see in particular Sects. 2.3.2 and 2.3.4) the LLPs are therefore only moderately boosted and large opening angles are common. To achieve optimal signal efficiency, it is therefore desirable to place tracking stations on the back-end, top, bottom and sides of the fiducial volume. This is illustrated in Fig. 30 which shows the distribution of hits on the various faces of the detector for an example $h \rightarrow A'A'$ model (see Sect. 2.3.1) with $m_{A'} = 5$ GeV and $A' \rightarrow \tau\tau$, and the τ 's decaying in the 3-prong mode. The majority of the hits land on the back-face of the fiducial volume, but the sides, top and bottom cannot be neglected. This is despite the relatively high boost of this benchmark model, as compared to models in which, e.g., the LLP is produced in a heavy flavor decay. It may be feasible to instrument those faces that see typically fewer hits more sparsely than the nominal design outlined in Sect. 1.3. Studies to this effect are under way for a wider range of models and will inform the final design.

Finally, some tracking stations on the front face are needed to reject backgrounds from charged particles emanating from the shield, primarily muons (see Sect. 3). For these stations, resolution is less important than efficiency, and alternative technologies (e.g. scintillator planes) may be considered.

(b) Vertex resolution:

Assuming that no magnetic field will be available in the CODEX-b decay volume, a good vertex resolution is essential to convincingly demonstrate a signal. The baseline design calls for 6 RPC layers in each tracking station covering the walls, while the five internal tracking stations would have 3 RPC layers each. This is however subject to further optimization, as less RPC layers may be needed on some of the less important wall of the cube, as indicated by Fig. 30. The most important parameter is the distance to the first tracker plane, which motivates five additional tracking

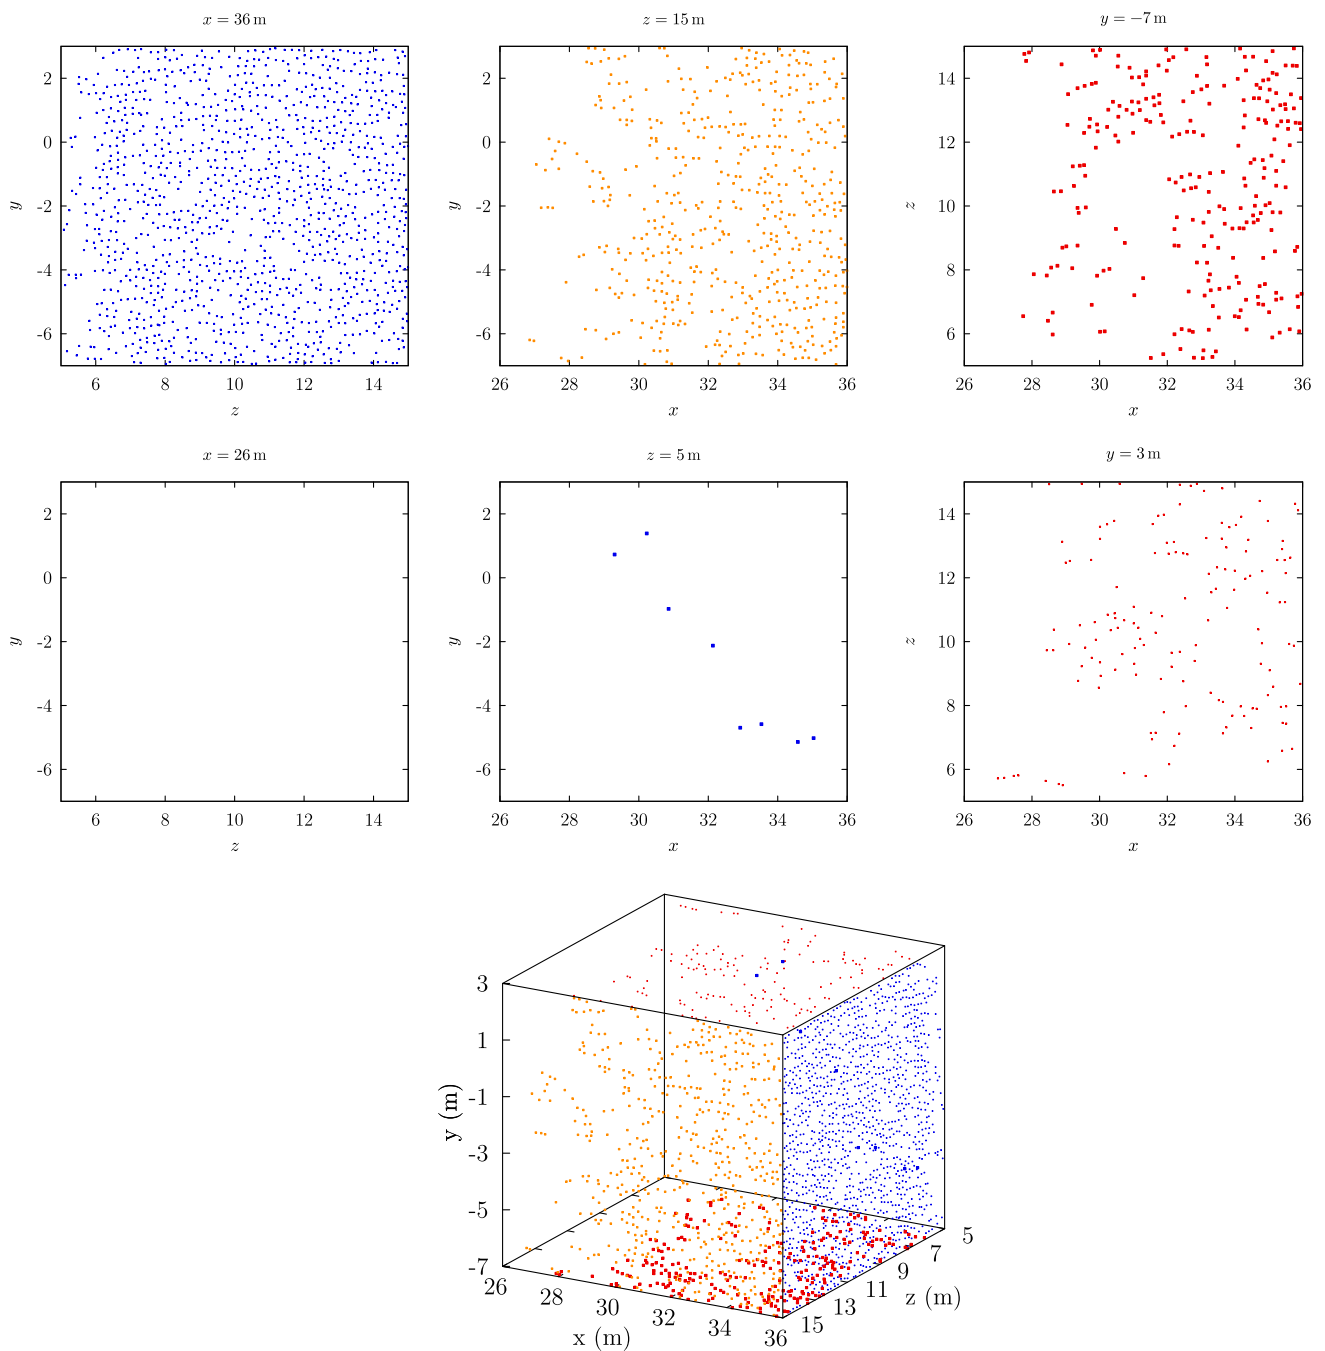


Fig. 30 Above: Scatter plot of hits on the faces of the box for 496 decays inside the box, from a $h \rightarrow A'A' \rightarrow 4\tau$ benchmark. (From left to right) Top row: $x = 36\text{ m}$, $z = 15\text{ m}$, $y = -7\text{ m}$. Bottom row:

$x = 26\text{ m}$, $z = 5\text{ m}$, $y = 3\text{ m}$. Below: The corresponding scatter plots shown in projection on the surfaces of the detector, with the same color scheme

stations spread throughout the fiducial volume, in order to achieve vertex resolutions on the order of millimeters rather than centimeters. For signals characterized by a high boost (e.g. Higgs decays, Sect. 2.3.1) the vertex resolution also impacts the signal reconstruction efficiency, as tracks tend to merge. The reconstruction efficiency corresponding to the nominal design for tracks from such an exotic Higgs decay,

$h \rightarrow A'A'$, is shown in the right-hand panel of Table 7, under the requirements:

- the track momentum $> 600\text{ MeV}$ (trivially satisfied for this benchmark);
- each track has at least 6 hits; and,
- the first hit of each track is unique.

It is the last requirement which can fail for a highly boosted LLP.

Once the LLP is required to decay in the fiducial volume, its proper lifetime, $c\tau$, is inversely correlated with its boost, so that a larger $c\tau$ typically implies a better reconstruction efficiency, as shown in the right-hand panel of Table 7. For particles with a boost factor of $\mathcal{O}(100)$ or more – roughly, $c\tau < 0.1$ m – the nominal design nonetheless achieves $\mathcal{O}(1)$ reconstruction efficiencies.

(c) Track momentum threshold:

The momentum threshold that can be achieved is especially relevant for two types of scenarios:

- LLPs produced in hadron decays are typically relatively soft, and in order to maintain an $\mathcal{O}(1)$ reconstruction efficiency, the track momentum threshold should be kept roughly around 600 MeV or lower. This is illustrated by the efficiency numbers in the left-hand panel of Table 7 for the $B \rightarrow X_s S$ portal of Sect. 2.3.2. Here, all losses in efficiency are because of the 600 MeV threshold that was assumed in the simulation.
- Inelastic dark matter models (see Sect. 2.4.4) are characterized by an LLP decaying to a nearby invisible state – the dark matter – and a number of soft SM tracks. Given the low amount of phase space available to the SM decay products, the reach of CODEX-b for this class of models is very sensitive to the threshold that can be achieved, as shown in Fig. 18.

5.1.2 Studies performed

A number of initial tracking studies have been performed to validate the design requirement outlined above. They further explore a number of CODEX-b design configurations, a variety of signals, and novel methods for particle boost reconstruction [203, 204]. Work is ongoing to integrate the CODEX-b detector into the LHCb simulation framework. This will facilitate the optimization of different reconstruction algorithms, as well as the study of how the information from both detectors could be integrated.

For these initial studies, a simplified Geant4 [205] description of CODEX-b was implemented, following the nominal design specifications [8], but without RPC faces on the top or bottom of the detector. The active detector material was modeled using silicon planes with 2 cm^2 granularity and the same radiation length as the RPCs proposed in the nominal design. Signal events of di-electron and di-muon candidates were then passed through this simulation to determine the detector response. Within this preliminary study, no attempt was made at modeling detector noise.

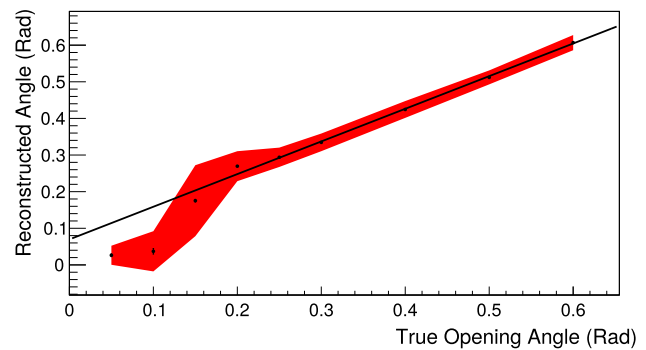


Fig. 31 Example opening angle reconstruction and resolution for an initial cluster and track building algorithm using 1 GeV electrons produced from a two-body decay at the front of the detector [203]

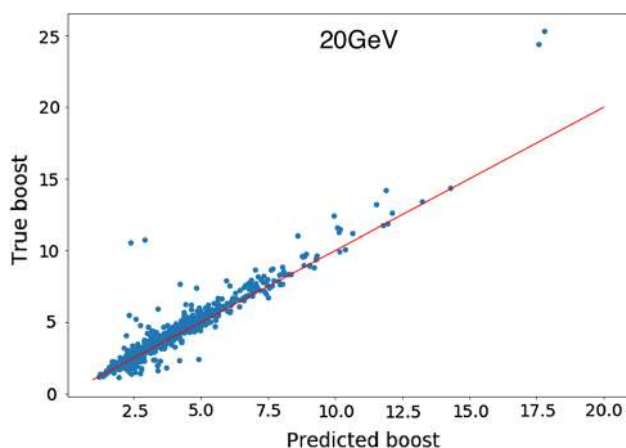
An initial clustering algorithm, based on a CALICE hadronic shower clustering algorithm [206], was designed to combine nearest neighbor energy deposits within the RPC layers, passing a minimum threshold, into clusters. As expected, this clustering was found to be necessary for electrons, but had little impact on reconstructing charged pion and muon signals. After clustering, an iterative linear track-finding algorithm was run. Both back-to-forward and forward-to-back algorithms were implemented, as well as various iterative approaches. For the expected signals within CODEX-b the back-to-forward tracking algorithm was found to provide the best performance. The tracking also performed well for more complex n -body signals without a common decay vertex, e.g. emerging jets.

The opening angle reconstruction as a function of the true opening angle for electrons with momenta of 1 GeV, produced from a two-body decay at the front of the detector, is shown in Fig. 31. For opening angles above 0.2 rad the algorithm provides a flat resolution of $\simeq 20\%$ and a ratio close to unity between the true and reconstructed opening angle. The tracking efficiency for single electrons with momentum greater than 0.5 GeV is $\simeq 0.95$. This efficiency is also dependent upon the local detector occupancy: For two-body decays at the front of the detector with an opening angle less than 0.05 radians the efficiency rapidly drops off, as the individual tracks of the two electron candidates can no longer be resolved. For muons these efficiencies are closer to $\simeq 1.0$, even down to momenta of 0.5 GeV.

While momentum information is not available for individual tracks, it is still possible to estimate the boost of an n -body signal decay. For a two-body decay with small masses, the parent signal boost can be analytically approximated by assuming relativistic decay products [207]. A study was performed looking at a six-body decay of the form $X \rightarrow \tau\tau$ with the simplified final state decay $\tau^- \rightarrow \pi^- \pi^- \pi^+ \nu$, where the missing energy of the neutrino and the resonance structure of decay was ignored [204]. A neural net was trained on eight decay topology features: the six opening angles of the pions

Table 7 Efficiency of reconstructing at least two tracks with the nominal design for both $B \rightarrow X_s S$ (Sect. 2.3.2) and $h \rightarrow A' A'$ (Sect. 2.3.1) scenarios, for various lifetimes [8]

$c\tau$ (m)	m_S (GeV) [$B \rightarrow X_s S$]			$m_{A'}$ (GeV) [$h \rightarrow A' A'$]				
	0.5	1.0	2.0	0.5	1.2	5.0	10.0	20.0
0.05	–	–	–	0.39	0.48	0.50	–	–
0.1	–	–	–	0.48	0.63	0.73	0.14	–
1.0	0.71	0.74	0.83	0.59	0.75	0.82	0.84	0.86
5.0	0.55	0.64	0.75	0.60	0.76	0.83	0.86	0.88
10.0	0.49	0.58	0.74	0.59	0.75	0.84	0.86	0.88
50.0	0.38	0.48	0.74	0.57	0.75	0.82	0.87	0.88
100.0	0.39	0.45	0.73	0.62	0.77	0.83	0.87	0.89
500.0	0.33	0.40	0.75	–	–	–	–	–

**Fig. 32** Six-body boost reconstruction for a narrow 20 GeV resonance decay into $\tau\tau$ with $\tau^- \rightarrow \pi^- \pi^- \pi^+$ [204]

and the angles of the two three-pion combinations most consistent with a τ -decay topology. The true boost versus the reconstructed boost for a narrow 20 GeV resonance is shown in Fig. 32. The boost resolution approaches $\simeq 4\%$ for resonance masses greater than 30 GeV. While this is a preliminary study, it demonstrates promise for reconstructing complex final states.

5.2 Timing

The baseline design employs RPC tracking detectors, which are expected to have a timing resolution of 350 ps per single gas gap of 1 mm [208], which corresponds to $350 \text{ ps}/\sqrt{6} \simeq 142 \text{ ps}$ resolution per station of 6 layers. The primary function of the timing capability is to synchronize the detector with the main LHCb detector. This enables one to match LHCb events with CODEX-b events, and to characterize and reject possible backgrounds. In particular, backgrounds induced by cosmic muons will be out of time with the collisions. As explained in Sect. 3, there is a sizable flux of relatively soft,

neutral hadrons emanating from the shield. These hadrons can scatter or decay in the detector, respectively in the case of neutrons and K_L 's, leading to a number of slow moving tracks. For example, over a distance of 2 m between two stations, a timing resolution of $\sim 150 \text{ ps}$ would allow one to reliably identify particles traveling at $\beta \lesssim 0.975$. For the example of a π^\pm , this corresponds to a momentum $\lesssim 0.6 \text{ GeV}$.

The timing capabilities of the RPCs is driven by the fluctuations of the primary ionization in the gas gap, and not to the readout electronics which can be designed to achieve resolutions of the order of 10 ps. It is possible that further development of the RPC technology may allow us to push this intrinsic timing resolution below 150 ps. Figure 33 shows the degree of separation which could be achieved for the $B \rightarrow X_s S$ portal (see Sect. 2.3.2) with some more optimistic assumptions for the timing resolution.

5.3 Calorimetry

Calorimetry would provide several important capabilities, notably particle identification (PID) via energy measurement and mass reconstruction, and the ability to expand the visible final states to include neutral hadrons and photons.

PID itself permits determination of the LLP decay modes, which could be crucial to identifying the quantum numbers and physics of the LLP itself. For instance, reconstructing a $\mu^+ \pi^- \pi^0$ final state might suggest a leptonic LLP coupling to a charged current, while even measuring the relative $e^+ e^-$ versus $\mu^+ \mu^-$ branching ratios could distinguish a vector from a scalar state. Moreover, the ability to reliably reconstruct the LLP mass would provide an additional crucial property of the new particle, while also providing an additional handle to reject SM LLP backgrounds. Detection and reconstruction of neutral hadrons, especially the π^0 , may permit rejection of K_L^0 backgrounds from the $\pi^+ \pi^- \pi^0$ final state.

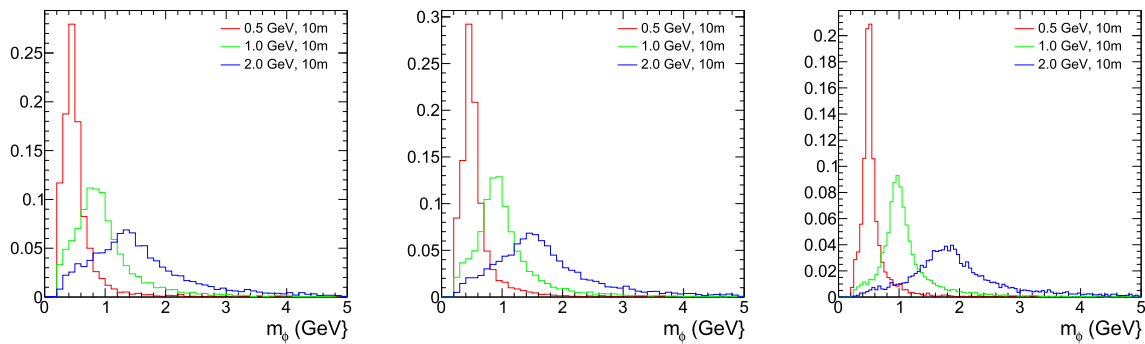


Fig. 33 Reconstructed LLP mass for different $B \rightarrow X_s \phi$ benchmarks with $c\tau = 10$ m, for 150 ps (left), 100 ps (middle), and 50 ps (right) timing resolution [8]

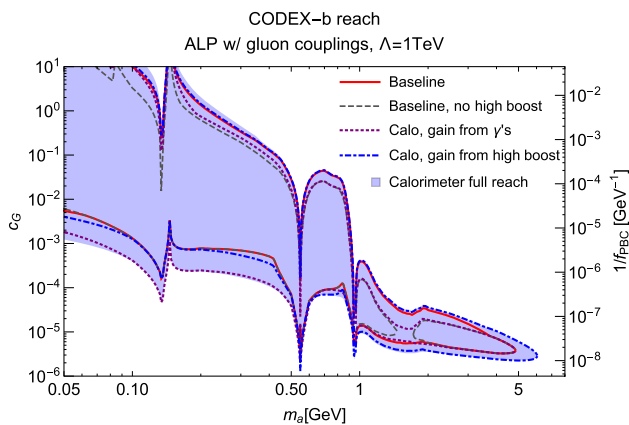


Fig. 34 Projected sensitivity of CODEX-b to gluon-coupled ALPs with a calorimeter element (blue, shaded) compared to the baseline, i.e. tracking only, detector (solid red). Also shown are the gains in reach coming from the ability to detect highly boosted LLPs (blue, dot-dashed) and from the ability to reconstruct photon final states (purple, dotted) separately. The gray dashed line corresponds to the baseline detector reach in which highly boosted ALPs are discarded instead of being considered with 50 events of background

Calorimeter elements may also improve characterization of highly-boosted LLPs, especially when their decay products start to become so collimated that it becomes difficult to separate them from single tracks given the finite track-to-track separation capabilities of the detector. Concretely, the track-to-track separation equals $2 \times \text{pitch} / \sqrt{12}$, for which we take 1 cm as a benchmark, similar to the expected performance ATLAS phase II RPCs. This can however be lowered if needed. With a tracker-only option, merged tracks will reconstruct as a single ‘appearing’ track in the tracking volume. However, for highly boosted LLPs, such as the ALPs of Sect. 2.3.3, that may have hadronic final states, such hadrons would develop energetic showers inside the calorimeter. This renders a signature strikingly different compared to e.g. low-energy neutron scattering. (Assuming each of the $\sim 10^6$ background muons passes through at least six tracking layers that are each 95% efficient, then the expected number

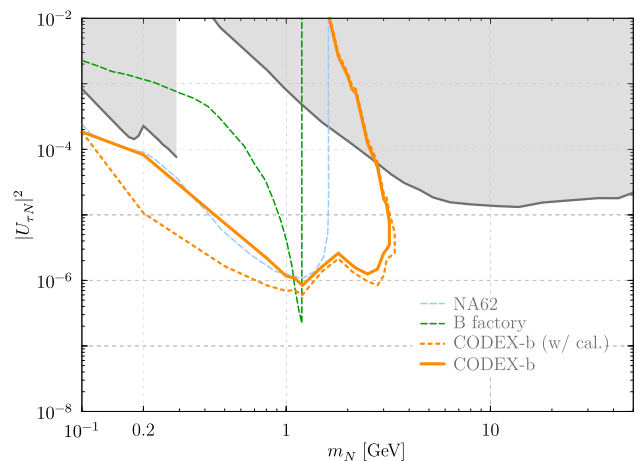


Fig. 35 Projected sensitivity of CODEX-b to Dirac heavy neutral leptons for $U_{\tau N} \gg U_{eN}, U_{\mu N}$ for a tracking detector only (solid) compared to an optimistic case (dashed) with a calorimeter capable of reconstructing the $N \rightarrow \nu \pi^0$ final state, assuming the background differential flux of single π^0 's is negligible or reducible

of appearing tracks induced from the muon background is $\sim 10^{-2}$ per 300 fb^{-1} .)

Energy measurement and PID may also help in the rejection of backgrounds, because they permit comparison of signal and background differential rates (in kinetic energy), rather than just the overall fluxes. Further, a calorimeter element placed on the front face of the detector, i.e. closest to the IP, may detect and absorb the flux of incoming neutrons (see Sect. 3), that might otherwise scatter and produce signal-like tracks: As seen in Table 3, single track production from neutron scattering inside the detector is non-negligible, with ~ 50 such events expected.

Diphoton final states may dominate the branching ratio of (pseudo)scalar LLPs, in particular the ALPs in the sub-GeV mass regime (see Sect. 2.3.3), such that reaches may be greatly improved with the ability to detect photons. Deposition of merged photons into the calorimeter will appear as a single highly energetic photon, to be compared with the relevant background photon fluxes shown in Fig. 23. Above

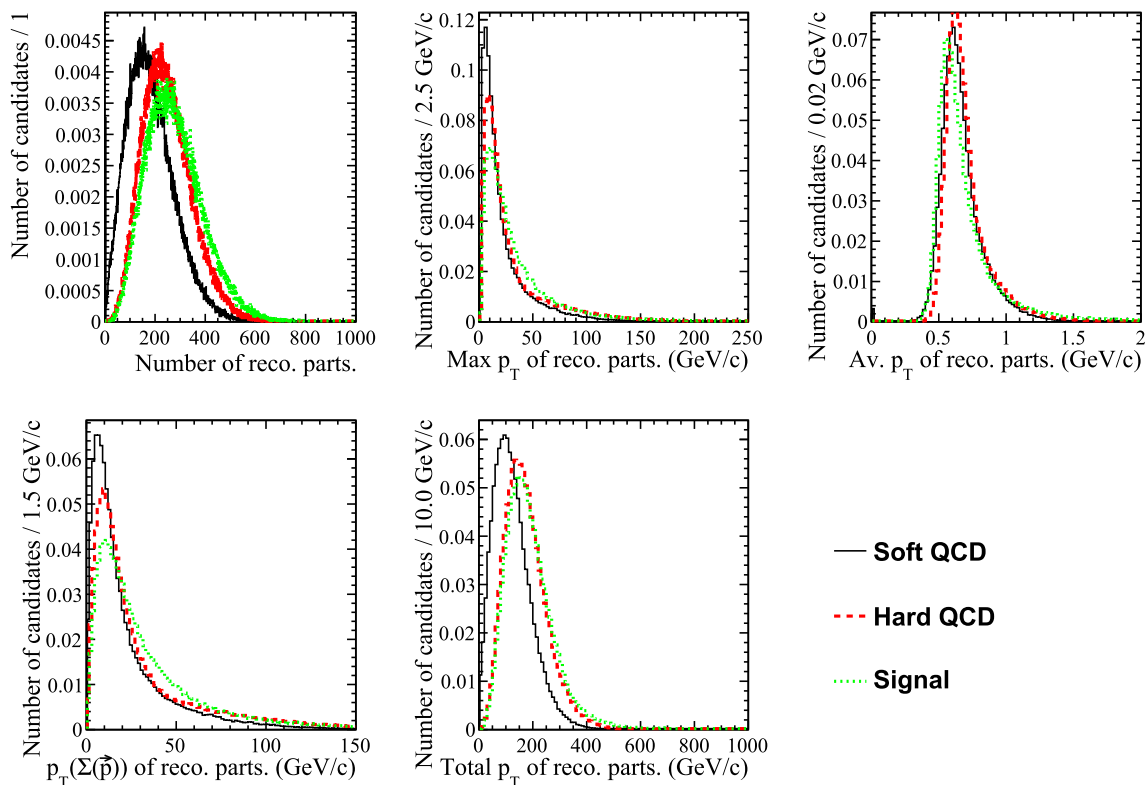


Fig. 36 Distribution of different general event variables at LHCb in events containing a $h \rightarrow A'(\mu^+\mu^-)A'(\mu^+\mu^-)$ decay (labelled as signal), soft and hard QCD processes. The processes were all generated with Pythia. The variables displayed are all generated using particles

reconstructible at LHCb. They correspond, from left to right and from top to bottom, to: the total number of particles in the event, the maximum p_T of all these particles, their average p_T , the p_T of the vectorial sum of all particles and the sum of their p_T

1 GeV, these fluxes are $\lesssim 10^{-1}$ per 300 fb^{-1} . In Fig. 34 we show the ALP coupled to gluon reach for a CODEX-b setup, assuming the that the baseline design can be extended to detect diphoton final states efficiently (shaded area). For comparison we also show the tracker-only baseline design (solid red line), and the gains obtained by detecting the highly-boosted ALPs with zero background (blue dot-dashed line) or detect the purely photonic decay modes (purple dotted line). For illustration we also include the baseline reach if one completely discards the highly-boosted ALPs (gray dashed line). The CODEX-b reach attainable with a calorimeter addition is striking, both at high and low ALP masses.

Even more striking improvements are attainable in models where the ALP decays to photon pairs most of the time (as in the Physics Beyond Colliders benchmark BC9), as the first detectable final state with the baseline detector, the Dalitz mode $a \rightarrow e^+e^-\gamma$, has a branching ratio of $\mathcal{O}(10^{-2})$.

Similarly, detection and reconstruction of neutral hadrons such as the π^0 may be important in capturing dominant branching ratios in certain heavy neutral lepton mass regimes. For example, for that case that the HNL is predominantly coupled to the τ , with $m_N < m_\tau$, the dominant decay mode is $N \rightarrow \nu_\tau \pi^0$. In Fig. 35 we show the improvement in

reach assuming this final state is reconstructible, compared to requiring at least two tracks. In practice, measurement of this final state requires an understanding of the background differential flux of single π^0 's. While the nominal flux of π^0 's is vanishing small, some might be produced from e.g. neutron scattering on air.

5.4 Tagging of events at LHCb

The LLPs detected at CODEX-b will be produced in events arising from pp collisions at Interaction Point 8. Therefore, these events could have information detectable at LHCb that is relevant to further help CODEX-b distinguish interesting phenomena from background. In this section we briefly review how – and how well – information from LHCb could be used to tag events at CODEX-b.

To study this tagging, we use as a benchmark a Higgs boson decaying to a pair of long-lived dark photons (see Sect. 2.3.1), which in turn decay to a pair of muons: $h \rightarrow A'(\mu^+\mu^-)A'(\mu^+\mu^-)$. The A' were assumed to have a mass of 1 GeV and a proper lifetime of $c\tau = 1 \text{ m}$. The decay was generated using Pythia [180] at a center-of-mass energy of $\sqrt{s} = 14 \text{ TeV}$.

The first aspect studied was the probability to detect an LLP decay *both* at CODEX-b and LHCb. For events in which one A' falls in the CODEX-b angular acceptance, $\sim 18\%$ have the other one in the LHCb acceptance. However, for the lifetimes of greatest interest with respect to the CODEX-b reach, hardly any of these produce any detectable decay object at LHCb. In particular, for the $c\tau = 1$ m benchmark, the probability for such a decay is only $\sim 10^{-5}$. In more complicated hidden sectors however, a high multiplicity of LLPs may be produced in the same event, so that one could be detected at LHCb and the other in CODEX-b.

The second possibility under study was how the LLP production mechanism can affect the underlying event seen at LHCb. This should be specifically relevant whenever the LLP is produced through the Higgs portal, such as in our benchmark example. We performed a general comparison of how events *look* at LHCb at truth level, with no reconstruction involved. To compare to the signal, we generated softQCD (minimum bias) and hardQCD ($b\bar{b}$) samples with `Pythia`, under the same conditions as the signal. For this comparison, we defined reconstructible particles at LHCb as those stable, charged particles that are produced in the LHCb acceptance. In Fig. 36 we show the distribution of different global variables of interest. While the figure shows a certain degree of discrimination between the different processes, more detailed studies will be needed. In particular, for this study gg fusion was chosen as Higgs production mode. Production via vector boson fusion, though having a smaller cross section, might provide more power to tag CODEX-b events at LHCb, by searching for a hard jet in the LHCb acceptance.

6 Outlook

The immediate priority for CODEX-b is the finalization of the design for the CODEX- β demonstrator and approval for its installation. A Letter of Intent for the full CODEX-b detector will follow this Expression of Interest in the near-term, including further developments of the detector design concept, although results from the CODEX- β demonstrator are expected to inform the final design choices for the detector. In particular, we intend to investigate in detail a realistic option for incorporating calorimetry in the CODEX-b design.

Acknowledgements We thank Wolfgang Altmannshofer, Brian Batell, James Beacham, Florian Bernlochner, David Curtin, Francesco D'Eramo, Jeff Dror, Gilly Elor, Iftah Galon, Stefania Gori, Felix Kling, Gaia Lanfranchi, Zoltan Ligeti, Simone Pagan Griso, Hiren Patel, Gilad Perez, Maxim Pospelov, Josh Ruderman, Bibhushan Shakya, Jessie Shelton, Brian Shuve, Yotam Soreq, Yuhsin Tsai, Jure Zupan for many helpful discussions. We similarly thank all the members of the BSM working group of the Physics Beyond Colliders report. In addition, we thank Asher Berlin, Raffaele Tito D'Agnolo, Daniel Dercks, Jordy De Vries, Herbi K. Dreiner, Juan Carlos Helo, Martin Hirsch, Felix Kling, Dave McKeen and Zeren Simon Wang for generously providing plots for

reproduction, with permission, as well as for comments on parts of the manuscript.

We thank the computing and simulation teams of the LHCb collaboration for their generous help in performing the simulation studies presented here. We thank the LHCb technical coordination for making the background studies in the DELPHI cavern possible. We thank Giovanni Passaleva for his support and encouragement, and for helpful advice regarding the baseline RPC technology and the relationship between CODEX-b and LHCb.

The work of XCV is supported by MINECO (Spain) through the Ramon y Cajal program RYC-2016-20073 and by XuntaGal under the ED431F 2018/01 project. JAE is supported by U.S. Department of Energy (DOE) grant DE-SC0011784. VVG is partially supported by ERC CoG "RECEPT" GA number 724777 within the H2020 framework programme. SK is supported by U.S. DOE grant DE-SC0009988 and the Paul Dirac fund at the Institute for Advanced Study. MP, DR and BN are supported by the U.S. DOE under contract DE-AC02-05CH11231. HR is supported in part by the U.S. DOE under contract DE-AC02-05CH11231. MW is supported by NSF grant PHY-1912836.

Significant parts of this work were performed at: the Aspen Center for Physics, supported by National Science Foundation grant PHY-1607611; the Munich Institute for Astro- and Particle Physics (MIAPP); and the Galileo Galilei Institute for Theoretical Physics. We thank all these institutions for their support and hospitality.

This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility operated under Contract No. DE-AC02-05CH11231. This work has received financial support from Xunta de Galicia (Centro singular de investigación de Galicia accreditation 2019–2022), by European Union ERDF, and by the "Mara de Maeztu" Units of Excellence program MDM-2016-0692 and the Spanish Research State Agency.

Data Availability Statement This manuscript has no associated data or the data will not be deposited. [Authors' comment: Uncalibrated data from the measurement campaign is available upon request.]

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>. Funded by SCOAP³.

References

1. ATLAS collaboration, G. Aad et al., Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Phys. Lett. B* **716**, 1–29 (2012). <https://doi.org/10.1016/j.physletb.2012.08.020>. arXiv:1207.7214
2. CMS collaboration, S. Chatrchyan et al., Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC, *Phys. Lett. B* **716**, 30–61 (2012). <https://doi.org/10.1016/j.physletb.2012.08.021>. arXiv:1207.7235
3. ATLAS, CMS collaboration, G. Aad et al., Measurements of the Higgs boson production and decay rates and constraints on its

- couplings from a combined ATLAS and CMS analysis of the LHC pp collision data at $\sqrt{s} = 7$ and 8 TeV. *JHEP* **08**, 045 (2016). [https://doi.org/10.1007/JHEP08\(2016\)045](https://doi.org/10.1007/JHEP08(2016)045). arXiv:1606.02266
4. M. H. Reno, D. Seckel, Primordial nucleosynthesis: the effects of injecting hadrons. *Phys. Rev. D* **37**, 3441–3462 (1988). <https://doi.org/10.1103/PhysRevD.37.3441>
 5. S. Dimopoulos, R. Esmailzadeh, L.J. Hall, G. Starkman, Is the universe closed by baryons? Nucleosynthesis with a late decaying massive particle. *Astrophys. J.* **330**, 545 (1988). <https://doi.org/10.1086/166493>
 6. S. Dimopoulos, R. Esmailzadeh, L.J. Hall, G. Starkman, Limits on late decaying particles from nucleosynthesis. *Nucl. Phys. B* **311**, 699–718 (1989). [https://doi.org/10.1016/0550-3213\(89\)90173-9](https://doi.org/10.1016/0550-3213(89)90173-9)
 7. J. Alimena et al., Searching for long-lived particles beyond the Standard Model at the Large Hadron Collider. *J. Phys. G: Nucl. Part. Phys.* **47**, 090501 (2020). <https://doi.org/10.1088/1361-6471/ab4574>. arXiv:1903.04497
 8. V.V. Gligorov, S. Knapen, M. Papucci, D.J. Robinson, Searching for long-lived particles: a compact detector for exotics at LHCb. *Phys. Rev. D* **97**, 015023 (2018). <https://doi.org/10.1103/PhysRevD.97.015023>. arXiv:1708.09395
 9. R. Lindner, Private communication
 10. S. Wagner, M. Grassi, A. Cabrera, A hybrid organic/inorganic scintillator for high performance measurements. arXiv:1807.00628
 11. NOvA collaboration, D.S. Ayres et al., The NOvA Technical Design Report, Tech. Rep. FERMILAB-DESIGN-2007-01 (2007). <https://doi.org/10.2172/935497>
 12. ND280 Upgrade work group collaboration, K. Iwamoto, The upgrade project of the T2K near detector. *PoS ICHEP2018*, 299 (2019). <https://doi.org/10.22323/1.340.0299>
 13. DUNE collaboration, G. Yang, 3D projection scintillator tracker in the DUNE near detector. *PoS ICHEP2018*, 868 (2019). <https://doi.org/10.22323/1.340.0868>
 14. L. Lee, C. Ohm, A. Soffer, T.-T. Yu, Collider searches for long-lived particles beyond the standard model. *Prog. Part. Nucl. Phys.* **106**, 210–255 (2019). <https://doi.org/10.1016/j.pnpnp.2019.02.006>. arXiv:1810.12602
 15. W. Bonivento et al., Proposal to search for heavy neutral leptons at the SPS. arXiv:1310.1762
 16. S. Alekhin et al., A facility to search for hidden particles at the CERN SPS: the SHiP physics case. *Rep. Prog. Phys.* **79**, 124201 (2016). <https://doi.org/10.1088/0034-4885/79/12/124201>. arXiv:1504.04855
 17. SHiP collaboration, M. Anelli et al., A facility to search for hidden particles (SHiP) at the CERN SPS. arXiv:1504.04956
 18. NA62 collaboration, E. Cortina Gil et al., The Beam and detector of the NA62 experiment at CERN. *JINST* **12**, P05025 (2017). <https://doi.org/10.1088/1748-0221/12/05/P05025>. arXiv:1703.08501
 19. J.L. Feng, I. Galon, F. Kling, S. Trojanowski, ForwArd Search Experiment at the LHC. *Phys. Rev. D* **97**, 035001 (2018). <https://doi.org/10.1103/PhysRevD.97.035001>. arXiv:1708.09389
 20. FASER collaboration, A. Ariga et al., Letter of intent for FASER: ForwArd Search Experiment at the LHC. arXiv:1811.10243
 21. FASER collaboration, A. Ariga et al., Technical proposal for FASER: ForwArd Search Experiment at the LHC. arXiv:1812.09139
 22. J.P. Chou, D. Curtin, H.J. Lubatti, New detectors to explore the lifetime frontier. *Phys. Lett. B* **767**, 29–36 (2017). <https://doi.org/10.1016/j.physletb.2017.01.043>. arXiv:1606.06298
 23. MATHUSLA collaboration, C. Alpigiani et al., A letter of intent for MATHUSLA: a dedicated displaced vertex detector above ATLAS or CMS. arXiv:1811.00927
 24. V.V. Gligorov, S. Knapen, B. Nachman, M. Papucci, D.J. Robinson, Leveraging the ALICE/L3 cavern for long-lived particle searches. *Phys. Rev. D* **99**, 015023 (2019). <https://doi.org/10.1103/PhysRevD.99.015023>. arXiv:1810.03636
 25. J. Beacham et al., Physics beyond colliders at CERN: beyond the standard model working group report. *J. Phys. G* **47**, 010501 (2020). <https://doi.org/10.1088/1361-6471/ab4cd2>. arXiv:1901.09966
 26. R.M. Schabinger, J.D. Wells, A minimal spontaneously broken hidden sector and its impact on Higgs boson physics at the large hadron collider. *Phys. Rev. D* **72**, 093007 (2005). <https://doi.org/10.1103/PhysRevD.72.093007>. arXiv:hep-ph/0509209
 27. S. Gopalakrishna, S. Jung, J.D. Wells, Higgs boson decays to four fermions through an abelian hidden sector. *Phys. Rev. D* **78**, 055002 (2008). <https://doi.org/10.1103/PhysRevD.78.055002>. arXiv:0801.3456
 28. D. Curtin, R. Essig, S. Gori, J. Shelton, Illuminating dark photons with high-energy colliders. *JHEP* **02**, 157 (2015). [https://doi.org/10.1007/JHEP02\(2015\)157](https://doi.org/10.1007/JHEP02(2015)157). arXiv:1412.0018
 29. M.J. Strassler, K.M. Zurek, Discovering the Higgs through highly-displaced vertices. *Phys. Lett. B* **661**, 263–267 (2008). <https://doi.org/10.1016/j.physletb.2008.02.008>. arXiv:hep-ph/0605193
 30. R. Barbieri, T. Gregoire, L.J. Hall, Mirror world at the large hadron collider. arXiv:hep-ph/0509242
 31. V. Barger, P. Langacker, M. McCaskey, M.J. Ramsey-Musolf, G. Shaughnessy, LHC phenomenology of an extended standard model with a real scalar singlet. *Phys. Rev. D* **77**, 035005 (2008). <https://doi.org/10.1103/PhysRevD.77.035005>. arXiv:0706.4311
 32. D.E. Morrissey, D. Poland, K.M. Zurek, Abelian hidden sectors at a GeV. *JHEP* **07**, 050 (2009). <https://doi.org/10.1088/1126-6708/2009/07/050>. arXiv:0904.2567
 33. D. Curtin et al., Exotic decays of the 125 GeV Higgs boson. *Phys. Rev. D* **90**, 075004 (2014). <https://doi.org/10.1103/PhysRevD.90.075004>. arXiv:1312.4992
 34. M. Bauer, P. Foldenauer, J. Jaeckel, Hunting all the hidden photons. *JHEP* **07**, 094 (2018). [https://doi.org/10.1007/JHEP07\(2018\)094](https://doi.org/10.1007/JHEP07(2018)094). arXiv:1803.05466
 35. F.F. Deppisch, S. Kulkarni, W. Liu, Searching for a light Z through Higgs production at the LHC. *Phys. Rev. D* **100**, 115023 (2019). <https://doi.org/10.1103/PhysRevD.100.115023>. arXiv:1908.11741
 36. B. Holdom, Two $U(1)$'s and epsilon charge shifts. *Phys. Lett. B* **166**, 196–198 (1986). [https://doi.org/10.1016/0370-2693\(86\)91377-8](https://doi.org/10.1016/0370-2693(86)91377-8)
 37. P. Galison, A. Manohar, TWO Z 's or not two Z 's? *Phys. Lett. B* **136**, 279–283 (1984). [https://doi.org/10.1016/0370-2693\(84\)91161-4](https://doi.org/10.1016/0370-2693(84)91161-4)
 38. LHCb collaboration, R. Aaij et al., Search for dark photons produced in 13 TeV pp collisions. *Phys. Rev. Lett.* **120**, 061801 (2018). <https://doi.org/10.1103/PhysRevLett.120.061801>. arXiv:1710.02867
 39. LHCb collaboration, R. Aaij et al., Search for $A' \rightarrow \mu^+ \mu^-$ decays. *Phys. Rev. Lett.* **124**, 041801 (2020). <https://doi.org/10.1103/PhysRevLett.124.041801>. arXiv:1910.06926
 40. P. Ilten, Y. Soreq, J. Thaler, M. Williams, W. Xue, Proposed inclusive dark photon search at LHCb. *Phys. Rev. Lett.* **116**, 251803 (2016). <https://doi.org/10.1103/PhysRevLett.116.251803>. arXiv:1603.08926
 41. P. Ilten, J. Thaler, M. Williams, W. Xue, Dark photons from charm mesons at LHCb. *Phys. Rev. D* **92**, 115017 (2015). <https://doi.org/10.1103/PhysRevD.92.115017>. arXiv:1509.06765
 42. S. Gardner, R.J. Holt, A.S. Tadepalli, New prospects in fixed target searches for dark forces with the SeaQuest experiment at Fermilab. *Phys. Rev. D* **93**, 115015 (2016). <https://doi.org/10.1103/PhysRevD.93.115015>. arXiv:1509.00050
 43. A. Berlin, S. Gori, P. Schuster, N. Toro, Dark sectors at the Fermilab SeaQuest experiment. *Phys. Rev. D* **98**, 035011 (2018). <https://doi.org/10.1103/PhysRevD.98.035011>. arXiv:1804.00661

44. SHiP collaboration, E. Graverini, SHiP sensitivity to Dark Photons, Tech. Rep. CERN-SHiP-NOTE-2016-004 (2016)
45. O. Moreno, The heavy photon search experiment at Jefferson Lab, in meeting of the APS division of particles and fields, 10 (2013). [arXiv:1310.2060](https://arxiv.org/abs/1310.2060)
46. ATLAS collaboration, Search prospects for dark-photons decaying to displaced collimated jets of muons at HL-LHC, Tech. Rep. ATL-PHYS-PUB-2019-002, CERN, Geneva (2019)
47. ATLAS collaboration, M. Aaboud et al., Search for long-lived particles produced in pp collisions at $\sqrt{s} = 13$ TeV that decay into displaced hadronic jets in the ATLAS muon spectrometer. *Phys. Rev. D* **99**, 052005 (2019). <https://doi.org/10.1103/PhysRevD.99.052005>. [arXiv:1811.07370](https://arxiv.org/abs/1811.07370)
48. M. Cepeda et al., Report from Working Group 2: Higgs physics at the HL-LHC and HE-LHC, vol. 7, pp. 221–584 (2019) <https://doi.org/10.23731/CYRM-2019-007.221> [arXiv:1902.00134](https://arxiv.org/abs/1902.00134)
49. LHCb collaboration, R. Aaij et al., Search for long-lived scalar particles in $B^+ \rightarrow K^+ \chi (\mu^+ \mu^-)$ decays. *Phys. Rev. D* **95**, 071101 (2017). <https://doi.org/10.1103/PhysRevD.95.071101>. [arXiv:1612.07818](https://arxiv.org/abs/1612.07818)
50. LHCb collaboration, R. Aaij et al., Search for hidden-sector bosons in $B^0 \rightarrow K^{*0} \mu^+ \mu^-$ decays., *Phys. Rev. Lett.* **115**, 161802 (2015). <https://doi.org/10.1103/PhysRevLett.115.161802>. [arXiv:1508.04094](https://arxiv.org/abs/1508.04094)
51. J.F. Donoghue, J. Gasser, H. Leutwyler, The decay of a light Higgs boson. *Nucl. Phys. B* **343**, 341–368 (1990). [https://doi.org/10.1016/0550-3213\(90\)90474-R](https://doi.org/10.1016/0550-3213(90)90474-R)
52. F. Bezrukov, D. Gorbunov, Light inflaton Hunter’s Guide. *JHEP* **05**, 010 (2010). [https://doi.org/10.1007/JHEP05\(2010\)010](https://doi.org/10.1007/JHEP05(2010)010). [arXiv:0912.0390](https://arxiv.org/abs/0912.0390)
53. A. Fradette, M. Pospelov, BBN for the LHC: constraints on lifetimes of the Higgs portal scalars. *Phys. Rev. D* **96**, 075033 (2017). <https://doi.org/10.1103/PhysRevD.96.075033>. [arXiv:1706.01920](https://arxiv.org/abs/1706.01920)
54. M.W. Winkler, Decay and detection of a light scalar boson mixing with the Higgs boson. *Phys. Rev. D* **99**, 015018 (2019). <https://doi.org/10.1103/PhysRevD.99.015018>. [arXiv:1809.01876](https://arxiv.org/abs/1809.01876)
55. F. Bezrukov, D. Gorbunov, I. Timiryasov, Uncertainties of hadronic scalar decay calculations. [arXiv:1812.08088](https://arxiv.org/abs/1812.08088)
56. R.S. Willey, H.L. Yu, Decays $K^\pm \rightarrow \pi^\pm l^+ l^-$ and limits on the mass of the neutral higgs boson. *Phys. Rev. D* **26**, 3287–3289 (1982). <https://doi.org/10.1103/PhysRevD.26.3287>
57. R. Chivukula, A.V. Manohar, Limits on a light Higgs boson. *Phys. Lett. B* **207**, 86–90 (1988). [https://doi.org/10.1016/0370-2693\(88\)90891-X](https://doi.org/10.1016/0370-2693(88)90891-X)
58. B. Grinstein, L.J. Hall, L. Randall, Do B meson decays exclude a light Higgs? *Phys. Lett. B* **211**, 363–369 (1988). [https://doi.org/10.1016/0370-2693\(88\)90916-1](https://doi.org/10.1016/0370-2693(88)90916-1)
59. B. Batell, M. Pospelov, A. Ritz, Multi-lepton signatures of a hidden sector in rare B decays. *Phys. Rev. D* **83**, 054005 (2011). <https://doi.org/10.1103/PhysRevD.83.054005>. [arXiv:0911.4938](https://arxiv.org/abs/0911.4938)
60. F. Bergsma et al., Search for axion-like particle production in 400 GeV proton–copper interactions. *Phys. Lett. B* **157**, 458–462 (1985). [https://doi.org/10.1016/0370-2693\(85\)90400-9](https://doi.org/10.1016/0370-2693(85)90400-9)
61. J.A. Evans, Detecting hidden particles with MATHUSLA. *Phys. Rev. D* **97**, 055046 (2018). <https://doi.org/10.1103/PhysRevD.97.055046>. [arXiv:1708.08503](https://arxiv.org/abs/1708.08503)
62. J.L. Feng, I. Galon, F. Kling, S. Trojanowski, Dark Higgs bosons at the ForwArd Search ExpeRiment. *Phys. Rev. D* **97**, 055034 (2018). <https://doi.org/10.1103/PhysRevD.97.055034>. [arXiv:1710.09387](https://arxiv.org/abs/1710.09387)
63. SHiP collaboration, G. Lanfranchi, Sensitivity of the SHiP experiment to a light scalar particle mixing with the Higgs, Tech. Rep. CERN-SHiP-NOTE-2017-001 (2017)
64. B. Batell, A. Freitas, A. Ismail, D. Mckeen, Flavor-specific scalar mediators. *Phys. Rev. D* **98**, 055026 (2018). <https://doi.org/10.1103/PhysRevD.98.055026>. [arXiv:1712.10022](https://arxiv.org/abs/1712.10022)
65. B. Batell, A. Freitas, A. Ismail, D. Mckeen, Probing light dark matter with a hadrophilic scalar mediator. *Phys. Rev. D* **100**, 095020 (2019). <https://doi.org/10.1103/PhysRevD.100.095020>. [arXiv:1812.05103](https://arxiv.org/abs/1812.05103)
66. G. Elor, M. Escudero, A. Nelson, Baryogenesis and dark matter from B mesons. *Phys. Rev. D* **99**, 035031 (2019). <https://doi.org/10.1103/PhysRevD.99.035031>. [arXiv:1810.00880](https://arxiv.org/abs/1810.00880)
67. A.E. Nelson, H. Xiao, Baryogenesis from B meson oscillations. *Phys. Rev. D* **100**, 075002 (2019). <https://doi.org/10.1103/PhysRevD.100.075002>. [arXiv:1901.08141](https://arxiv.org/abs/1901.08141)
68. Particle Data Group collaboration, M. Tanabashi et al., Review of particle physics. *Phys. Rev. D* **98**, 030001 (2018). <https://doi.org/10.1103/PhysRevD.98.030001>
69. Belle-II collaboration, W. Altmannshofer et al., The Belle II Physics Book. *PTEP* **2019**, 123C01 (2019). <https://doi.org/10.1093/ptep/ptz106>. [arXiv:1808.10567](https://arxiv.org/abs/1808.10567)
70. M. Bauer, M. Neubert, A. Thamm, Collider probes of axion-like particles. *JHEP* **12**, 044 (2017). [https://doi.org/10.1007/JHEP12\(2017\)044](https://doi.org/10.1007/JHEP12(2017)044). [arXiv:1708.00443](https://arxiv.org/abs/1708.00443)
71. N. Craig, A. Hook, S. Kasko, The photophobic ALP. *JHEP* **09**, 028 (2018). [https://doi.org/10.1007/JHEP09\(2018\)028](https://doi.org/10.1007/JHEP09(2018)028). [arXiv:1805.06538](https://arxiv.org/abs/1805.06538)
72. D. Curtin, Private communication
73. FASER collaboration, A. Ariga et al., FASER’s physics reach for long-lived particles. *Phys. Rev. D* **99**, 095011 (2019). <https://doi.org/10.1103/PhysRevD.99.095011>. [arXiv:1811.12522](https://arxiv.org/abs/1811.12522)
74. D. Aloni, Y. Soreq, M. Williams, Coupling QCD-scale axionlike particles to gluons. *Phys. Rev. Lett.* **123**, 031803 (2019). <https://doi.org/10.1103/PhysRevLett.123.031803>. [arXiv:1811.03474](https://arxiv.org/abs/1811.03474)
75. R.N. Mohapatra, J.W.F. Valle, Neutrino mass and baryon number nonconservation in superstring models. *Phys. Rev. D* **34**, 1642 (1986). <https://doi.org/10.1103/PhysRevD.34.1642>
76. T. Asaka, S. Blanchet, M. Shaposhnikov, The nuMSM, dark matter and neutrino masses. *Phys. Lett. B* **631**, 151–156 (2005). <https://doi.org/10.1016/j.physletb.2005.09.070>. [arXiv:hep-ph/0503065](https://arxiv.org/abs/hep-ph/0503065)
77. T. Asaka, M. Shaposhnikov, The ν MSM, dark matter and baryon asymmetry of the universe. *Phys. Lett. B* **620**, 17–26 (2005). <https://doi.org/10.1016/j.physletb.2005.06.020>. [arXiv:hep-ph/0505013](https://arxiv.org/abs/hep-ph/0505013)
78. B. Batell, T. Han, D. McKeen, B. Shams, Es Haggi, Thermal dark matter through the Dirac neutrino portal. *Phys. Rev. D* **97**, 075016 (2018). <https://doi.org/10.1103/PhysRevD.97.075016>. [arXiv:1709.07001](https://arxiv.org/abs/1709.07001)
79. A. Greljo, D.J. Robinson, B. Shakya, J. Zupan, $R(D^{(*)})$ from W' and right-handed neutrinos. *JHEP* **09**, 169 (2018). [https://doi.org/10.1007/JHEP09\(2018\)169](https://doi.org/10.1007/JHEP09(2018)169). [arXiv:1804.04642](https://arxiv.org/abs/1804.04642)
80. P. Asadi, M.R. Buckley, D. Shih, It’s all right(-handed neutrinos): a new W' model for the $R_{D^{(*)}}$ anomaly. *JHEP* **09**, 010 (2018). [https://doi.org/10.1007/JHEP09\(2018\)010](https://doi.org/10.1007/JHEP09(2018)010). [arXiv:1804.04135](https://arxiv.org/abs/1804.04135)
81. D.J. Robinson, B. Shakya, J. Zupan, Right-handed neutrinos and $R(D^{(*)})$. *JHEP* **02**, 119 (2019). [https://doi.org/10.1007/JHEP02\(2019\)119](https://doi.org/10.1007/JHEP02(2019)119). [arXiv:1807.04753](https://arxiv.org/abs/1807.04753)
82. K. Bondarenko, A. Boyarsky, D. Gorbunov, O. Ruchayskiy, Phenomenology of GeV-scale heavy neutral leptons. *JHEP* **11**, 032 (2018). [https://doi.org/10.1007/JHEP11\(2018\)032](https://doi.org/10.1007/JHEP11(2018)032). [arXiv:1805.08567](https://arxiv.org/abs/1805.08567)
83. J.C. Helo, M. Hirsch, Z.S. Wang, Heavy neutral fermions at the high-luminosity LHC. *JHEP* **07**, 056 (2018). [https://doi.org/10.1007/JHEP07\(2018\)056](https://doi.org/10.1007/JHEP07(2018)056). [arXiv:1803.02212](https://arxiv.org/abs/1803.02212)
84. M. Cacciari, M. Greco, P. Nason, The P(T) spectrum in heavy flavor hadroproduction. *JHEP* **05**, 007 (1998). <https://doi.org/10.1088/1126-6708/1998/05/007>. [arXiv:hep-ph/9803400](https://arxiv.org/abs/hep-ph/9803400)

85. M. Cacciari, S. Frixione, P. Nason, The $p(T)$ spectrum in heavy flavor photoproduction. *JHEP* **03**, 006 (2001). <https://doi.org/10.1088/1126-6708/2001/03/006>. arXiv:hep-ph/0102134
86. ATLAS collaboration, G. Aad et al., Measurement of W^\pm and Z -boson production cross sections in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Phys. Lett. B* **759**, 601–621 (2016). <https://doi.org/10.1016/j.physletb.2016.06.023>. arXiv:1603.09222
87. ALICE collaboration, B. Abelev et al., D_s^+ meson production at central rapidity in proton–proton collisions at $\sqrt{s} = 7$ TeV. *Phys. Lett. B* **718**, 279–294 (2012). <https://doi.org/10.1016/j.physletb.2012.10.049>. arXiv:1208.1948
88. ATLAS collaboration, G. Aad et al., Measurement of $D^{*\pm}$, D^\pm and D_s^\pm meson production cross sections in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector. *Nucl. Phys. B* **907**, 717–763 (2016). <https://doi.org/10.1016/j.nuclphysb.2016.04.032>. arXiv:1512.02913
89. LHCb collaboration, R. Aaij et al., Measurement of B meson production cross-sections in proton-proton collisions at $\sqrt{s} = 7$ TeV. *JHEP* **08**, 117 (2013). [https://doi.org/10.1007/JHEP08\(2013\)117](https://doi.org/10.1007/JHEP08(2013)117). arXiv:1306.3663
90. LHCb collaboration, R. Aaij et al., Observation of the decay $B_c^+ \rightarrow B_s^0 \pi^+$. *Phys. Rev. Lett.* **111**, 181801 (2013). <https://doi.org/10.1103/PhysRevLett.111.181801>. arXiv:1308.4544
91. J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer et al., The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations. *JHEP* **07**, 079 (2014). [https://doi.org/10.1007/JHEP07\(2014\)079](https://doi.org/10.1007/JHEP07(2014)079). arXiv:1405.0301
92. A. Cooper-Sarkar, S. Haywood, M. Parker, S. Sarkar, K. Barnham, P. Bostock et al., Search for heavy neutrino decays in the BEBC beam dump experiment. *Phys. Lett. B* **160**, 207–211 (1985). [https://doi.org/10.1016/0370-2693\(85\)91493-5](https://doi.org/10.1016/0370-2693(85)91493-5)
93. G. Bernardi, G. Carugno, J. Chauveau, F. Dicarolo, M. Dris, J. Dumarchez et al., Further limits on heavy neutrino couplings. *Phys. Lett. B* **203**, 332–334 (1988). [https://doi.org/10.1016/0370-2693\(88\)90563-1](https://doi.org/10.1016/0370-2693(88)90563-1)
94. P. Vilain, G. Wilquet, S. Petrak, R. Beyer, W. Flegel, H. Grote et al., Search for heavy isosinglet neutrinos. *Phys. Lett. B* **343**, 453–458 (1995). [https://doi.org/10.1016/0370-2693\(94\)01422-9](https://doi.org/10.1016/0370-2693(94)01422-9)
95. J. Dorenbosch, J. Allaby, U. Amaldi, G. Barbiellini, C. Berger, F. Bergsma et al., A search for decays of heavy neutrinos in the mass range 0.5–2.8 GeV. *Phys. Lett. B* **166**, 473–478 (1986). [https://doi.org/10.1016/0370-2693\(86\)91601-1](https://doi.org/10.1016/0370-2693(86)91601-1)
96. J. Orloff, A.N. Rozanov, C. Santoni, Limits on the mixing of tau neutrino to heavy neutrinos. *Phys. Lett. B* **550**, 8–15 (2002). [https://doi.org/10.1016/S0370-2693\(02\)02769-7](https://doi.org/10.1016/S0370-2693(02)02769-7). arXiv:hep-ph/0208075
97. S. Baranov, Y. Batusov, S. Bunyatov, O. Klimov, V. Lyukov, Y. Nefedov et al., Search for heavy neutrinos at the IHEP-JINR neutrino detector. *Phys. Lett. B* **302**, 336–340 (1993). [https://doi.org/10.1016/0370-2693\(93\)90405-7](https://doi.org/10.1016/0370-2693(93)90405-7)
98. NuTeV, E815 collaboration, A. Vaitaitis et al., Search for neutral heavy leptons in a high-energy neutrino beam. *Phys. Rev. Lett.* **83**, 4943–4946 (1999). <https://doi.org/10.1103/PhysRevLett.83.4943>. arXiv:hep-ex/9908011
99. Search for neutral heavy leptons produced in z decays. *Z. Phys. C Part. Fields* **74**, 57–71 (1997). <https://doi.org/10.1007/s002880050370>
100. ATLAS collaboration, G. Aad et al., Search for heavy neutral leptons in decays of W bosons produced in 13 TeV pp collisions using prompt and displaced signatures with the ATLAS detector. *JHEP* **10**, 265 (2019). [https://doi.org/10.1007/JHEP10\(2019\)265](https://doi.org/10.1007/JHEP10(2019)265). arXiv:1905.09787
101. M. Drewes, J. Hajer, J. Klaric, G. Lanfranchi, NA62 sensitivity to heavy neutral leptons in the low scale seesaw model. *JHEP* **07**, 105 (2018). [https://doi.org/10.1007/JHEP07\(2018\)105](https://doi.org/10.1007/JHEP07(2018)105). arXiv:1801.04207
102. LBNE collaboration, C. Adams et al., The long-baseline neutrino experiment: exploring fundamental symmetries of the universe. arXiv:1307.7335
103. F. Kling, S. Trojanowski, Heavy neutral leptons at FASER. *Phys. Rev. D* **97**, 095016 (2018). <https://doi.org/10.1103/PhysRevD.97.095016>. arXiv:1801.08947
104. CMS collaboration, A.M. Sirunyan et al., Search for heavy neutral leptons in events with three charged leptons in proton–proton collisions at $\sqrt{s} = 13$ TeV. *Phys. Rev. Lett.* **120**, 221801 (2018). <https://doi.org/10.1103/PhysRevLett.120.221801>. arXiv:1802.02965
105. D. Curtin et al., Long-lived particles at the energy frontier: the MATHUSLA physics case. *Rep. Prog. Phys.* **82**, 116201 (2019). <https://doi.org/10.1088/1361-6633/ab28d6>. arXiv:1806.07396
106. P. Bechtle et al., Killing the cMSSM softly. *Eur. Phys. J. C* **76**, 96 (2016). <https://doi.org/10.1140/epjc/s10052-015-3864-0>. arXiv:1508.05951
107. D. Dercks, J. De Vries, H.K. Dreiner, Z.S. Wang, R-parity violation and light neutralinos at CODEX-b, FASER, and MATHUSLA. *Phys. Rev. D* **99**, 055039 (2019). <https://doi.org/10.1103/PhysRevD.99.055039>. arXiv:1810.03617
108. ATLAS collaboration, M. Aaboud et al., Search for electroweak production of supersymmetric states in scenarios with compressed mass spectra at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Phys. Rev. D* **97**, 052010 (2018). <https://doi.org/10.1103/PhysRevD.97.052010>. arXiv:1712.08119
109. P.W. Graham, D.E. Kaplan, S. Rajendran, Cosmological relaxation of the electroweak scale. *Phys. Rev. Lett.* **115**, 221801 (2015). <https://doi.org/10.1103/PhysRevLett.115.221801>. arXiv:1504.07551
110. J.R. Espinosa, C. Grojean, G. Panico, A. Pomarol, O. Pujolàs, G. Servant, Cosmological Higgs-axion interplay for a naturally small electroweak scale. *Phys. Rev. Lett.* **115**, 251803 (2015). <https://doi.org/10.1103/PhysRevLett.115.251803>. arXiv:1506.09217
111. J.L. Evans, T. Gherghetta, N. Nagata, Z. Thomas, Naturalizing supersymmetry with a two-field relaxion mechanism. *JHEP* **09**, 150 (2016). [https://doi.org/10.1007/JHEP09\(2016\)150](https://doi.org/10.1007/JHEP09(2016)150). arXiv:1602.04812
112. J.L. Evans, T. Gherghetta, N. Nagata, M. Peloso, Low-scale D-term inflation and the relaxion mechanism. *Phys. Rev. D* **95**, 115027 (2017). <https://doi.org/10.1103/PhysRevD.95.115027>. arXiv:1704.03695
113. B. Batell, G.F. Giudice, M. McCullough, Natural heavy supersymmetry. *JHEP* **12**, 162 (2015). [https://doi.org/10.1007/JHEP12\(2015\)162](https://doi.org/10.1007/JHEP12(2015)162). arXiv:1509.00834
114. T. Flacke, C. Frugiuele, E. Fuchs, R.S. Gupta, G. Perez, Phenomenology of relaxion-Higgs mixing. *JHEP* **06**, 050 (2017). [https://doi.org/10.1007/JHEP06\(2017\)050](https://doi.org/10.1007/JHEP06(2017)050). arXiv:1610.02025
115. K. Choi, S.H. Im, Constraints on relaxion windows. *JHEP* **12**, 093 (2016). [https://doi.org/10.1007/JHEP12\(2016\)093](https://doi.org/10.1007/JHEP12(2016)093). arXiv:1610.00680
116. N. Craig, A. Katz, M. Strassler, R. Sundrum, Naturalness in the dark at the LHC. *JHEP* **07**, 105 (2015). [https://doi.org/10.1007/JHEP07\(2015\)105](https://doi.org/10.1007/JHEP07(2015)105). arXiv:1501.05310
117. Z. Chacko, H.-S. Goh, R. Harnik, The twin Higgs: natural electroweak breaking from mirror symmetry. *Phys. Rev. Lett.* **96**, 231802 (2006). <https://doi.org/10.1103/PhysRevLett.96.231802>. arXiv:hep-ph/0506256
118. Z. Chacko, H.-S. Goh, R. Harnik, A twin Higgs model from left–right symmetry. *JHEP* **01**, 108 (2006). <https://doi.org/10.1088/1126-6708/2006/01/108>. arXiv:hep-ph/0512088
119. M.J. Strassler, K.M. Zurek, Echoes of a hidden valley at hadron colliders. *Phys. Lett. B* **651**, 374–379 (2007). <https://doi.org/10.1016/j.physletb.2007.06.055>. arXiv:hep-ph/0604261

120. T. Han, Z. Si, K.M. Zurek, M.J. Strassler, Phenomenology of hidden valleys at hadron colliders. *JHEP* **07**, 008 (2008). <https://doi.org/10.1088/1126-6708/2008/07/008>. arXiv:0712.2041
121. G. Burdman, Z. Chacko, R. Harnik, L. de Lima, C.B. Verhaaren, Colorless top partners, a 125 GeV Higgs, and the limits on naturalness. *Phys. Rev. D* **91**, 055007 (2015). <https://doi.org/10.1103/PhysRevD.91.055007>. arXiv:1411.3310
122. C.J. Morningstar, M.J. Peardon, The Glueball spectrum from an anisotropic lattice study. *Phys. Rev. D* **60**, 034509 (1999). <https://doi.org/10.1103/PhysRevD.60.034509>. arXiv:hep-lat/9901004
123. Y. Chen et al., Glueball spectrum and matrix elements on anisotropic lattices. *Phys. Rev. D* **73**, 014516 (2006). <https://doi.org/10.1103/PhysRevD.73.014516>. arXiv:hep-lat/0510074
124. N. Craig, S. Knapen, P. Longhi, M. Strassler, The vector-like twin Higgs. *JHEP* **07**, 002 (2016). [https://doi.org/10.1007/JHEP07\(2016\)002](https://doi.org/10.1007/JHEP07(2016)002). arXiv:1601.07181
125. D. Curtin, C.B. Verhaaren, Discovering uncolored naturalness in exotic Higgs decays. *JHEP* **12**, 072 (2015). [https://doi.org/10.1007/JHEP12\(2015\)072](https://doi.org/10.1007/JHEP12(2015)072). arXiv:1506.06141
126. LHCb collaboration, R. Aaij et al., Physics case for an LHCb Upgrade II—opportunities in flavour physics, and beyond, in the HL-LHC era. arXiv:1808.08865
127. P. Bokan, Reconstructing low mass boosted $A' \rightarrow b\bar{b}$ at LHCb. A scalar particle in low mass region, Tech. Rep. CERN-STUDENTS-Note-2014-165 (2014)
128. A. Berlin, F. Kling, Inelastic dark matter at the LHC lifetime frontier: ATLAS, CMS, LHCb, CODEX-b, FASER, and MATHUSLA. *Phys. Rev. D* **99**, 015021 (2019). <https://doi.org/10.1103/PhysRevD.99.015021>. arXiv:1810.01879
129. D. Tucker-Smith, N. Weiner, Inelastic dark matter. *Phys. Rev. D* **64**, 043502 (2001). <https://doi.org/10.1103/PhysRevD.64.043502>. arXiv:hep-ph/0101138
130. E. Izaguirre, G. Krnjaic, B. Shuve, Discovering inelastic thermal-relic dark matter at colliders. *Phys. Rev. D* **93**, 063523 (2016). <https://doi.org/10.1103/PhysRevD.93.063523>. arXiv:1508.03050
131. R.T. D'Agnolo, D. Pappadopulo, J.T. Ruderman, Fourth exception in the calculation of relic abundances. *Phys. Rev. Lett.* **119**, 061102 (2017). <https://doi.org/10.1103/PhysRevLett.119.061102>. arXiv:1705.08450
132. R.T. D'Agnolo, D. Pappadopulo, J.T. Ruderman, P.-J. Wang, Thermal relic targets with exponentially small couplings. *Phys. Rev. Lett.* **124**, 151801 (2020). <https://doi.org/10.1103/PhysRevLett.124.151801>. arXiv:1906.09269
133. R.T. D'Agnolo, C. Mondino, J.T. Ruderman, P.-J. Wang, Exponentially light dark matter from coannihilation. *JHEP* **08**, 079 (2018). [https://doi.org/10.1007/JHEP08\(2018\)079](https://doi.org/10.1007/JHEP08(2018)079). arXiv:1803.02901
134. D.E. Kaplan, M.A. Luty, K.M. Zurek, Asymmetric dark matter. *Phys. Rev. D* **79**, 115016 (2009). <https://doi.org/10.1103/PhysRevD.79.115016>. arXiv:0901.4117
135. I.-W. Kim, K.M. Zurek, Flavor and collider signatures of asymmetric dark matter. *Phys. Rev. D* **89**, 035008 (2014). <https://doi.org/10.1103/PhysRevD.89.035008>. arXiv:1310.2617
136. K.M. Zurek, Asymmetric dark matter: theories, signatures, and constraints. *Phys. Rep.* **537**, 91–121 (2014). <https://doi.org/10.1016/j.physrep.2013.12.001>. arXiv:1308.0338
137. S. Chang, M.A. Luty, Displaced dark matter at colliders. arXiv:0906.5013
138. P. Schwaller, D. Stolarski, A. Weiler, Emerging jets. *JHEP* **05**, 059 (2015). [https://doi.org/10.1007/JHEP05\(2015\)059](https://doi.org/10.1007/JHEP05(2015)059). arXiv:1502.05409
139. Y. Hochberg, E. Kuflik, T. Volansky, J.G. Wacker, Mechanism for thermal relic dark matter of strongly interacting massive particles. *Phys. Rev. Lett.* **113**, 171301 (2014). <https://doi.org/10.1103/PhysRevLett.113.171301>. arXiv:1402.5143
140. Y. Hochberg, E. Kuflik, H. Murayama, T. Volansky, J.G. Wacker, Model for thermal relic dark matter of strongly interacting massive particles. *Phys. Rev. Lett.* **115**, 021301 (2015). <https://doi.org/10.1103/PhysRevLett.115.021301>. arXiv:1411.3727
141. S.-M. Choi, Y. Hochberg, E. Kuflik, H.M. Lee, Y. Mambrini, H. Murayama et al., Vector SIMP dark matter. *JHEP* **10**, 162 (2017). [https://doi.org/10.1007/JHEP10\(2017\)162](https://doi.org/10.1007/JHEP10(2017)162). arXiv:1707.01434
142. S.-M. Choi, H.M. Lee, P. Ko, A. Natale, Resolving phenomenological problems with strongly-interacting-massive-particle models with dark vector resonances. *Phys. Rev. D* **98**, 015034 (2018). <https://doi.org/10.1103/PhysRevD.98.015034>. arXiv:1801.07726
143. Y. Hochberg, E. Kuflik, H. Murayama, S.I.M.P. Spectroscopy. *JHEP* **05**, 090 (2016). [https://doi.org/10.1007/JHEP05\(2016\)090](https://doi.org/10.1007/JHEP05(2016)090). arXiv:1512.07917
144. A. Berlin, N. Blinov, S. Gori, P. Schuster, N. Toro, Cosmology and accelerator tests of strongly interacting dark matter. *Phys. Rev. D* **97**, 055033 (2018). <https://doi.org/10.1103/PhysRevD.97.055033>. arXiv:1801.05805
145. E. Kuflik, M. Perelstein, N.R.-L. Lorier, Y.-D. Tsai, Elastically decoupling dark matter. *Phys. Rev. Lett.* **116**, 221302 (2016). <https://doi.org/10.1103/PhysRevLett.116.221302>. arXiv:1512.04545
146. J.A. Dror, E. Kuflik, W.H. Ng, Codecaying dark matter. *Phys. Rev. Lett.* **117**, 211801 (2016). <https://doi.org/10.1103/PhysRevLett.117.211801>. arXiv:1607.03110
147. S. Okawa, M. Tanabashi, M. Yamanaka, Relic abundance in a secluded dark matter scenario with a massive mediator. *Phys. Rev. D* **95**, 023006 (2017). <https://doi.org/10.1103/PhysRevD.95.023006>. arXiv:1607.08520
148. M. Farina, D. Pappadopulo, J.T. Ruderman, G. Trevisan, Phases of cannibal dark matter. *JHEP* **12**, 039 (2016). [https://doi.org/10.1007/JHEP12\(2016\)039](https://doi.org/10.1007/JHEP12(2016)039). arXiv:1607.03108
149. M. Farina, Asymmetric twin dark matter. *JCAP* **1511**, 017 (2015). <https://doi.org/10.1088/1475-7516/2015/11/017>. arXiv:1506.03520
150. M. Freytsis, S. Knapen, D.J. Robinson, Y. Tsai, Gamma-rays from dark showers with twin Higgs models. *JHEP* **05**, 018 (2016). [https://doi.org/10.1007/JHEP05\(2016\)018](https://doi.org/10.1007/JHEP05(2016)018). arXiv:1601.07556
151. K.R. Dienes, B. Thomas, Dynamical dark matter: I. Theoretical overview. *Phys. Rev. D* **85**, 083523 (2012). <https://doi.org/10.1103/PhysRevD.85.083523>. arXiv:1106.4546
152. K.R. Dienes, B. Thomas, Dynamical dark matter: II. An explicit model. *Phys. Rev. D* **85**, 083524 (2012). <https://doi.org/10.1103/PhysRevD.85.083524>. arXiv:1107.0721
153. K.R. Dienes, B. Thomas, Phenomenological constraints on axion models of dynamical dark matter. *Phys. Rev. D* **86**, 055013 (2012). <https://doi.org/10.1103/PhysRevD.86.055013>. arXiv:1203.1923
154. K.R. Dienes, S. Su, B. Thomas, Distinguishing dynamical dark matter at the LHC. *Phys. Rev. D* **86**, 054008 (2012). <https://doi.org/10.1103/PhysRevD.86.054008>. arXiv:1204.4183
155. K.R. Dienes, S. Su, B. Thomas, Strategies for probing nonminimal dark sectors at colliders: the interplay between cuts and kinematic distributions. *Phys. Rev. D* **91**, 054002 (2015). <https://doi.org/10.1103/PhysRevD.91.054002>. arXiv:1407.2606
156. K.R. Dienes, S. Su, B. Thomas, Beyond the bump-hunt: a game plan for discovering dynamical dark matter at the LHC. *AIP Conf. Proc.* **1743**, 020013 (2016). <https://doi.org/10.1063/1.4953280>. <https://aip.scitation.org/doi/pdf/10.1063/1.4953280>
157. D. Curtin, K.R. Dienes, B. Thomas, Dynamical dark matter, MATHUSLA, and the lifetime frontier. *Phys. Rev. D* **98**, 115005 (2018). <https://doi.org/10.1103/PhysRevD.98.115005>. arXiv:1809.11021
158. L.J. Hall, K. Jedamzik, J. March-Russell, S.M. West, Freeze-in production of FIMP dark matter. *JHEP* **03**, 080 (2010). [https://doi.org/10.1007/JHEP03\(2010\)080](https://doi.org/10.1007/JHEP03(2010)080). arXiv:0911.1120

159. R.T. Co, F. D’Eramo, L.J. Hall, D. Pappadopulo, Freeze-in dark matter with displaced signatures at colliders. *JCAP* **1512**, 024 (2015). <https://doi.org/10.1088/1475-7516/2015/12/024>. arXiv:1506.07532
160. J.A. Evans, J. Shelton, Long-lived staus and displaced leptons at the LHC. *JHEP* **04**, 056 (2016). [https://doi.org/10.1007/JHEP04\(2016\)056](https://doi.org/10.1007/JHEP04(2016)056). arXiv:1601.01326
161. F. D’Eramo, N. Fernandez, S. Profumo, Dark matter freeze-in production in fast-expanding universes. *JCAP* **1802**, 046 (2018). <https://doi.org/10.1088/1475-7516/2018/02/046>. arXiv:1712.07453
162. J.M. No, P. Tunney, B. Zaldivar, Probing dark matter freeze-in with long-lived particle signatures: MATHUSLA. HL-LHC and FCC-hh. *JHEP* **03**, 022 (2020). [https://doi.org/10.1007/JHEP03\(2020\)022](https://doi.org/10.1007/JHEP03(2020)022). arXiv:1908.11387
163. D.E. Morrissey, M.J. Ramsey-Musolf, Electroweak baryogenesis. *New J. Phys.* **14**, 125003 (2012). <https://doi.org/10.1088/1367-2630/14/12/125003>. arXiv:1206.2942
164. Y. Cui, R. Sundrum, Baryogenesis for weakly interacting massive particles. *Phys. Rev. D* **87**, 116013 (2013). <https://doi.org/10.1103/PhysRevD.87.116013>. arXiv:1212.2973
165. Y. Cui, B. Shuve, Probing baryogenesis with displaced vertices at the LHC. *JHEP* **02**, 049 (2015). [https://doi.org/10.1007/JHEP02\(2015\)049](https://doi.org/10.1007/JHEP02(2015)049). arXiv:1409.6729
166. D. McKeen, A.E. Nelson, CP violating baryon oscillations. *Phys. Rev. D* **94**, 076002 (2016). <https://doi.org/10.1103/PhysRevD.94.076002>. arXiv:1512.05359
167. K. Aitken, D. McKeen, T. Neder, A.E. Nelson, Baryogenesis from oscillations of charmed or beautiful baryons. *Phys. Rev. D* **96**, 075009 (2017). <https://doi.org/10.1103/PhysRevD.96.075009>. arXiv:1708.01259
168. G. Alonso-Álvarez, G. Elor, A.E. Nelson, H. Xiao, A supersymmetric theory of baryogenesis and sterile sneutrino dark matter from B mesons. *JHEP* **03**, 046 (2020). [https://doi.org/10.1007/JHEP03\(2020\)046](https://doi.org/10.1007/JHEP03(2020)046). arXiv:1907.10612
169. M.J. Strassler, Possible effects of a hidden valley on supersymmetric phenomenology. arXiv:hep-ph/0607160
170. M.J. Strassler, Why unparticle models with mass gaps are examples of hidden valleys. arXiv:0801.0629
171. M.J. Strassler, On the phenomenology of hidden valleys with heavy flavor. arXiv:0806.2385
172. S. Knapen, S. Pagan Griso, M. Papucci, D. J. Robinson, Triggering soft bombs at the LHC. *JHEP* **08**, 076 (2017). [https://doi.org/10.1007/JHEP08\(2017\)076](https://doi.org/10.1007/JHEP08(2017)076). arXiv:1612.00850
173. A. Pierce, B. Shakya, Y. Tsai, Y. Zhao, Searching for confining hidden valleys at LHCb, ATLAS, and CMS. *Phys. Rev. D* **97**, 095033 (2018). <https://doi.org/10.1103/PhysRevD.97.095033>. arXiv:1708.05389
174. CMS collaboration, A.M. Sirunyan et al., Search for new particles decaying to a jet and an emerging jet. *JHEP* **02**, 179 (2019). [https://doi.org/10.1007/JHEP02\(2019\)179](https://doi.org/10.1007/JHEP02(2019)179). arXiv:1810.10069
175. LHCb collaboration, R. Aaij et al., Updated search for long-lived particles decaying to jet pairs. *Eur. Phys. J. C* **77**, 812 (2017). <https://doi.org/10.1140/epjc/s10052-017-5178-x>. arXiv:1705.07332
176. LHCb collaboration, R. Aaij et al., Searches for low-mass dimuon resonances. arXiv:2007.03923
177. ATLAS collaboration, G. Aad et al., Search for light long-lived neutral particles produced in pp collisions at $\sqrt{s} = 13$ TeV and decaying into collimated leptons or light hadrons with the ATLAS detector. *Eur. Phys. J. C* **80**, 450 (2020). <https://doi.org/10.1140/epjc/s10052-020-7997-4>. arXiv:1909.01246
178. LHCb collaboration, Tracking and alignment plots for conferences. https://twiki.cern.ch/twiki/bin/view/LHCb/ConferencePlots#Tracking_efficiency
179. T. Sjostrand, S. Mrenna, P.Z. Skands, PYTHIA 6.4 physics and manual. *JHEP* **05**, 026 (2006). <https://doi.org/10.1088/1126-6708/2006/05/026>. arXiv:hep-ph/0603175
180. T. Sjostrand et al., An introduction to PYTHIA 8.2. *Comput. Phys. Commun.* **191**, 159–177 (2015). <https://doi.org/10.1016/j.cpc.2015.01.024>. arXiv:1410.3012
181. B. Andersson, G. Gustafson, B. Nilsson-Almqvist, A model for low-PT hadronic reactions with generalizations to hadron–nucleus and nucleus–nucleus collisions. *Nucl. Phys. B* **281**, 289–309 (1987). [https://doi.org/10.1016/0550-3213\(87\)90257-4](https://doi.org/10.1016/0550-3213(87)90257-4)
182. B. Andersson, A. Tai, B.-H. Sa, Final state interactions in the (nuclear) FRITIOF string interaction scenario. *Z. Phys. C Part. Fields* **70**, 499–506 (1996). <https://doi.org/10.1007/s002880050127>
183. B. Nilsson-Almqvist, E. Stenlund, Interactions between hadrons and nuclei: the Lund Monte Carlo, Fritiof Version 1.6. *Comput. Phys. Commun.* **43**, 387 (1987). [https://doi.org/10.1016/0010-4655\(87\)90056-7](https://doi.org/10.1016/0010-4655(87)90056-7)
184. B. Ganhuayag, V. Uzhinsky, Modified FRITIOF code: negative charged particle production in high energy nucleus nucleus interactions. *Czech. J. Phys.* **47**, 913–918 (1997). <https://doi.org/10.1023/A:1021296114786>
185. M.P. Guthrie, R.G. Alsmiller, H.W. Bertini, Calculation of the capture of negative pions in light elements and comparison with experiments pertaining to cancer radiotherapy. *Nucl. Instrum. Method* **66**, 29–36 (1968). [https://doi.org/10.1016/0029-554X\(68\)90054-2](https://doi.org/10.1016/0029-554X(68)90054-2)
186. H.W. Bertini, M.P. Guthrie, News item results from medium-energy intranuclear-cascade calculation. *Nucl. Phys. A* **169**, 670–672 (1971). [https://doi.org/10.1016/0375-9474\(71\)90710-X](https://doi.org/10.1016/0375-9474(71)90710-X)
187. V.A. Karmanov, Light front wave function of relativistic composite system in explicitly solvable model. *Nucl. Phys. B* **166**, 378–398 (1980). [https://doi.org/10.1016/0550-3213\(80\)90204-7](https://doi.org/10.1016/0550-3213(80)90204-7)
188. H. Burkhardt, V.M. Grichine, P. Gumplinger, V.N. Ivanchenko, R.P. Kokoulin, M. Maire et al., Geant4 standard electromagnetic package for hep applications. In: *IEEE Symposium Conference Record Nuclear Science 2004*, vol. 3, pp. 1907–1910 (2004). <https://doi.org/10.1109/NSSMIC.2004.1462617>
189. J.A. Formaggio, G.P. Zeller, From eV to EeV: neutrino cross sections across energy scales. *Rev. Mod. Phys.* **84**, 1307–1341 (2012). <https://doi.org/10.1103/RevModPhys.84.1307>. arXiv:1305.7513
190. MINERvA collaboration, C.M. Marshall et al., Measurement of neutral-current K^+ production by neutrinos using MINERvA. *Phys. Rev. Lett.* **119**, 011802 (2017). <https://doi.org/10.1103/PhysRevLett.119.011802>. arXiv:1611.02224
191. B. Dey, J. Lee, V. Coco, C.-S. Moon, Background studies for the CODEX-b experiment: measurements and simulation. arXiv:1912.03846
192. K.C. Akiba et al., The HeRSChel detector: high-rapidity shower counters for LHCb. *JINST* **13**, P04017 (2018). <https://doi.org/10.1088/1748-0221/13/04/P04017>. arXiv:1801.04281
193. ATLAS collaboration, Technical Design Report for the Phase-II Upgrade of the ATLAS Muon Spectrometer, Tech. Rep. CERN-LHCC-2017-017. ATLAS-TDR-026, CERN, Geneva (2017)
194. Super-Kamiokande Collaboration collaboration, K. Abe, Y. Hayato, T. Iida, K. Ishihara, J. Kameda, Y. Koshio et al., Search for $n - \bar{n}$ oscillation in super-kamiokande. *Phys. Rev. D* **91**, 072006 (2015). <https://doi.org/10.1103/PhysRevD.91.072006>
195. E. Friedman, A. Gal, Realistic calculations of nuclear disappearance lifetimes induced by n anti- n oscillations. *Phys. Rev. D* **78**, 016002 (2008). <https://doi.org/10.1103/PhysRevD.78.016002>. arXiv:0803.3696
196. LHCb collaboration, R. Aaij et al., LHCb detector performance. *Int. J. Mod. Phys. A* **30**, 1530022 (2015). <https://doi.org/10.1142/S0217751X15300227>. arXiv:1412.6352

197. LHCb collaboration, R. Aaij et al., Measurement of angular and CP asymmetries in $D^0 \rightarrow \pi^+\pi^-\mu^+\mu^-$ and $D^0 \rightarrow K^+K^-\mu^+\mu^-$ decays. Phys. Rev. Lett. **121**, 091801 (2018). LHCb-PAPER-2018-020 CERN-EP-2018-162. <https://doi.org/10.1103/PhysRevLett.121.091801>. arXiv:1806.10793
198. LHCb collaboration, R. Aaij et al., Search for the suppressed decays $B^+ \rightarrow K^+K^+\pi^-$ and $B^+ \rightarrow \pi^+\pi^+K^-$. Phys. Lett. B **765**, 307 (2017). LHCb-PAPER-2016-023, CERN-EP-2016-177. <https://doi.org/10.1016/j.physletb.2016.11.053>. arXiv:1608.01478
199. LHCb collaboration, R. Aaij et al., Search for rare $B_{(s)}^0 \rightarrow \mu^+\mu^-\mu^+\mu^-$ decays. Phys. Rev. Lett. **110**, 211801 (2013). <https://doi.org/10.1103/PhysRevLett.110.211801>. arXiv:1303.1092
200. LHCb collaboration, R. Aaij et al., Measurement of the branching fractions of the decays $D^+ \rightarrow K^-K^+K^+$, $D^+ \rightarrow \pi^-\pi^+K^+$ and $D_s^+ \rightarrow \pi^-K^+K^+$. JHEP **03**, 176 (2019). [https://doi.org/10.1007/JHEP03\(2019\)176](https://doi.org/10.1007/JHEP03(2019)176). arXiv:1810.03138
201. LHCb collaboration, R. Aaij et al., Measurement of the b -quark production cross-section in 7 and 13 TeV pp collisions. Phys. Rev. Lett. **118**, 052002 (2017). LHCb-PAPER-2016-031, CERN-EP-2016-201. <https://doi.org/10.1103/PhysRevLett.118.052002>. arXiv:1612.05140
202. LHCb collaboration, R. Aaij et al., Measurement of the fragmentation fraction ratio f_s/f_d and its dependence on B meson kinematics. JHEP **04**, 001 (2013). CERN-PH-EP-2013-006, LHCb-PAPER-2012-037. [https://doi.org/10.1007/JHEP04\(2013\)001](https://doi.org/10.1007/JHEP04(2013)001). arXiv:1301.5286
203. G. Gibbons, Build a new framework to simulate the physics prospects of the proposed CODEX-b detector. Master thesis, University of Birmingham
204. R. Quessard, Searching for long-lived particles with CODEXb. Master thesis, LPNHE Paris (2019)
205. GEANT4 collaboration, S. Agostinelli et al., GEANT4: a simulation toolkit. Nucl. Instrum. Methods A **506**, 250–303 (2003). [https://doi.org/10.1016/S0168-9002\(03\)01368-8](https://doi.org/10.1016/S0168-9002(03)01368-8)
206. CALICE collaboration, G. Eigen et al., Characterisation of different stages of hadronic showers using the CALICE Si-W ECAL physics prototype. Nucl. Instrum. Methods A **937**, 41–52 (2019). <https://doi.org/10.1016/j.nima.2019.04.111>. arXiv:1902.06161
207. D. Curtin, M.E. Peskin, Analysis of long lived particle decays with the MATHUSLA detector. Phys. Rev. D **97**, 015006 (2018). <https://doi.org/10.1103/PhysRevD.97.015006>. arXiv:1705.06327
208. P. Pizzimento, Upgrade of the resistive plate chamber detector for the high-background environments, 15th Topical Seminar on Innovative Particle and Radiation Detectors (IPRD19), 14–17 October 2019 Siena