

Searching for long-lived particles: A compact detector for exotics at LHCb

Vladimir V. Gligorov,¹ Simon Knapen,^{2,3} Michele Papucci,^{2,3} and Dean J. Robinson⁴

¹*LPNHE, Université Pierre et Marie Curie, Université Paris Diderot, CNRS/IN2P3, Paris, France*

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

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

⁴*Physics Department, University of Cincinnati, Cincinnati, Ohio 45221, USA*



(Received 5 October 2017; published 31 January 2018)

We advocate for the construction of a new detector element at the LHCb experiment, designed to search for displaced decays of beyond Standard Model long-lived particles, taking advantage of a large shielded space in the LHCb cavern that is expected to soon become available. We discuss the general features and putative capabilities of such an experiment, as well as its various advantages and complementarities with respect to the existing LHC experiments and proposals such as SHiP and MATHUSLA. For two well-motivated beyond Standard Model benchmark scenarios—Higgs decay to dark photons and B meson decays via a Higgs mixing portal—the reach either complements or exceeds that predicted for other LHC experiments.

DOI: [10.1103/PhysRevD.97.015023](https://doi.org/10.1103/PhysRevD.97.015023)

I. INTRODUCTION

Deep and long-standing questions concerning the constitution of Nature remain unanswered. These include the (un)naturalness and structure of the electroweak vacuum, as well as the origins of dark matter, the baryon asymmetry, neutrino masses, flavor hierarchies and CP violation. Well-motivated theoretical frameworks that attempt to address these questions often predict the existence of metastable states, also known as long-lived particles (LLPs): LLPs with lifetimes up to the subsecond regime are broadly consistent with cosmological bounds, opening a large parameter space to be explored.

Many extensions of the Standard Model motivated by the hierarchy problem predict LLPs: LLPs have been studied in the context of gauge-mediated supersymmetry [1,2], R-parity violating supersymmetry [2–5], stealth supersymmetry [6], minisplit supersymmetry [7,8], neutral naturalness [9–11] and certain relaxion models [12,13]. An extensive literature contemplates LLPs in models of dark matter [14–21], baryogenesis [22–24], neutrino masses [25–31] and hidden valleys [32–34,34], predicting a wide range of LLP production and decay morphologies. Finally, the discovery of the Higgs boson [35,36] opens up the possibility of Higgs mixing portals that generically admit exotic Higgs decays into LLPs.

ATLAS, CMS and LHCb have developed numerous programs to search for LLPs. Because of their large geometric acceptance and integrated luminosity, ATLAS and CMS will typically provide the best reach for charged, colored or relatively heavy LLPs (see e.g. [37–40]). Softer and/or short-lived final states tend to be more problematic due to triggering challenges and the large irreducible backgrounds inherent to high luminosity hadron collisions. LHCb, by contrast, may search for soft $\mathcal{O}(\text{GeV})$ final states, so long as the decay occurs within the Vertex Locator system [41–44]. LHCb also adds sensitivity to states with a combination of high mass and short decay time [45–49]. Longer-lived, GeV-scale LLPs are difficult for all three experiments, but are theoretically well-motivated by various models of, e.g., asymmetric dark matter [15,17], strongly interacting dark matter [19,50] and neutral naturalness [51]. Some scenarios may be searched for with current and future beam dump experiments such as NA62 [52] or SHiP [53], but comprehensive coverage should include a large sample of Higgs bosons, which at the moment can only be supplied by the LHC.

In this paper we propose to take advantage of a large shielded space in the LHCb cavern that is expected to become available after the pre-Run 3 upgrade, to construct a Compact Detector for Exotics at LHCb (“CODEX-b”). Apart from generic LLP searches, by virtue of its location, such a detector is well-suited to probe for GeV-scale LLPs, for instance generated by Higgs mixing or dark photon portals. With a modest amount of additional shielding from the primary interaction point, CODEX-b can operate in a low background environment, eliminating the triggering challenges associated with ATLAS and CMS. The modest

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

size of CODEX-b is anticipated to translate to relatively low construction and maintenance costs and a relatively short construction timescale with proven, off-the-shelf components. It should also be emphasized that CODEX-b may provide complementary data, at relatively low cost, to potential discoveries in other proposed or existing experiments.

We are aware of several other recent proposals in the same spirit. The most directly comparable is the MATHUSLA proposal [54], which intends to operate at the surface above ATLAS or CMS, using an air-filled fiducial volume equipped with multiple horizontal tracking surfaces, and which may in some cases probe up to $\mathcal{O}(1\text{ s})$ LLP lifetimes. A recent analysis of the various key features and capabilities for this LLP search, using a template of tracking and partial reconstruction techniques, and an in-depth analysis of relevant backgrounds, can be found in Refs [54,55]. CODEX-b necessarily has various commonalities with this proposal, though the backgrounds and configuration we examine here both differ in several critical aspects, and the theoretical prospects may also differ depending on the ultimate configuration and technologies chosen. Other planned or proposed LLP detectors located near the LHC experiments, include: milliQan [56,57] which makes use of a drainage gallery above the CMS interaction point and is intended specifically to search for millicharged particles; and FASER [58], which intends to operate a few hundred meters downstream of the ATLAS or CMS interaction point, looking for forward-produced light weakly coupled particles.

In what follows we describe the general features and putative capabilities of CODEX-b and estimate the reach by focusing on two benchmark models that illustrate the complementarity and relative advantages of CODEX-b compared to other LHC experiments. These include: (i) a spin-1 massive gauge boson, γ_d , that is produced through the exotic Higgs decay $h \rightarrow \gamma_d \gamma_d$. This dark photon can subsequently decay to charged Standard Model fermions through mixing with the photon [59–62]. (ii) A light scalar field, ϕ , that mixes with the SM Higgs boson. Such a particle can be abundantly produced in inclusive $b \rightarrow s\phi$ decays [63–65], subsequently decaying back to SM fermions through the same Higgs portal. We further study a proof-of-concept implementation of tracking, in order to verify that reconstruction of the decay vertices can be done efficiently and explore the capabilities to reconstruct the exotic state mass using timing information. Details of the commonalities and complementarities between CODEX-b and other proposed LLP searches will also be discussed.

II. GENERAL STRATEGY

A. Layout and capabilities

The LHCb experiment is a single arm forward spectrometer, located at interaction point 8 (IP8) on the LHC beam line, previously occupied by the DELPHI experiment

at LEP. Apart from LHCb itself, key features for this discussion are the data acquisition (DAQ) space and DELPHI barrel exhibit, both located behind a 3 m thick concrete radiation shield in UXA cavern, approximately 25 m from IP8. The planned relocation of the DAQ to the surface during the upcoming pre-Run 3 upgrade [67] leaves available a ready-made underground cavity, which is large enough to house a detector with fiducial volume $10 \times 10 \times 10$ m in size, or $20 \times 10 \times 10$ m should DELPHI also be removed.¹ We shall, for the sake of brevity, refer to the fiducial volume as “the box”. For all sensitivity projections we use a $10 \times 10 \times 10$ m fiducial volume located as close as possible to the IP; specifically, the coordinates of the fiducial volume are from $x = 26$ to 36 m (transverse), $y = -7$ to 3 m (vertical) and $z = 5$ to 15 m (forward). The current layout of the LHCb cavern is shown in Fig. 1, with the potential CODEX-b location overlaid in red.

Our proposal, in short, is to instrument this space with tracking layers, in order to search for decays-in-flight of LLPs generated at IP8. (Additional shielding near the IP will be required to suppress backgrounds, discussed in detail below.) To our knowledge, no space of this size is available near the ATLAS and CMS interaction points; for example the drainage gallery that is to house MilliQan is only ~ 3 m in diameter [57].

At the very least, in this type of experimental setup (cf. e.g. the MATHUSLA proposal [54]) one may search for at least two tracks originating from the same reconstructed vertex: a signal of the LLP decay $\phi \rightarrow ff$ or $\phi \rightarrow ffX$, where ϕ is the LLP, f is any trackable SM final state, and X is anything else. Apart from 2-body decays to SM particles, this morphology can include decays with additional invisible particles, e.g. $\phi \rightarrow ff\phi'$ or semileptonic $\phi \rightarrow f\ell\nu$. The former are of interest for (super)symmetric theories with squeezed spectra, in which the LLP may be an excited state with a small mass splitting from a lighter stable state. The signature of such multibody decays is acoplanarity of the tracks with respect to the decay vertex displacement from the IP. Multibody signatures may be more difficult to reconstruct, for example if the tracking efficiency drops sharply at low momentum due to multiple scattering. On the other hand, if these lower momentum tracks can be reconstructed, we may be able to determine their momentum with reasonable accuracy using time-of-flight information (see also Sec. III C). A multi-track decay signature consistent in time and space with a mother particle originating from an LHC collision and

¹The old DAQ area also housed the low and high voltage power supplies of the different subdetectors, which would then have to move elsewhere in the cavern, for example to what is now the DELPHI public viewing gallery. Depending on the required infrastructure in UXA cavern, it may therefore actually be more feasible to remove DELPHI and house the detector volume in its place, but such optimization is beyond the scope of this paper.

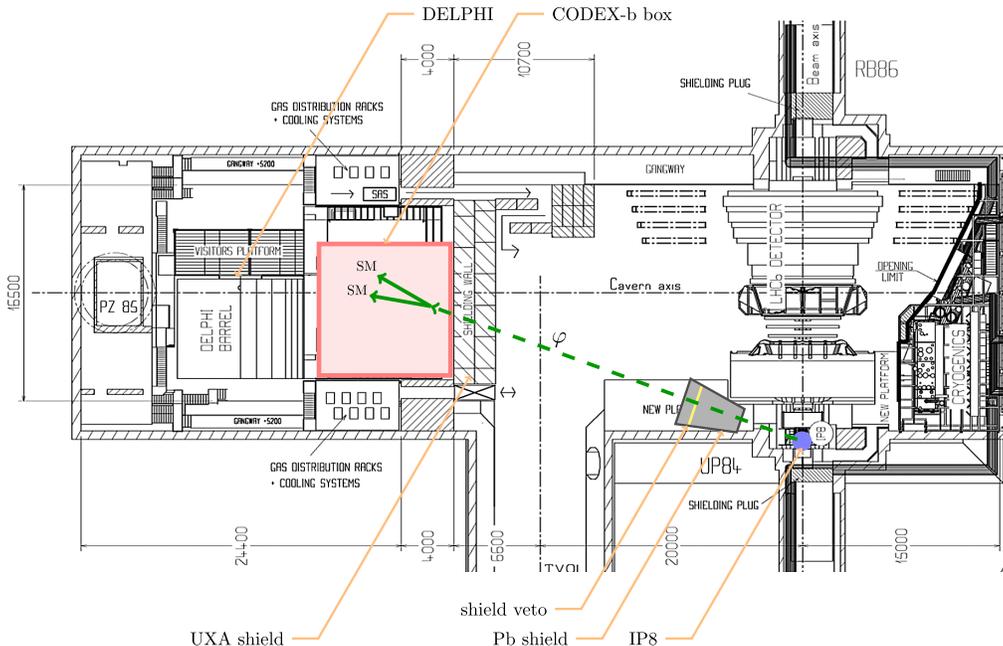


FIG. 1. Layout of the LHCb experimental cavern UX85 at point 8 of the LHC [66], overlaid with the CODEX-b apparatus.

inconsistent with the mass of a known SM particle would therefore be a very compelling sign of new physics (NP).

As we explore in some detail below, the relative proximity of the box to the interaction point, and the experimentally clean environment, produces a new physics reach that may complement or exceed those of other LHC experiments, despite the lower luminosity at the LHCb interaction point. Moreover, this setup may feature various other possibilities and capabilities:

- (i) A large amount of empty space between IP8 and the radiation shield can also be exploited if the LLP has significant branching fraction to muons, that typically punch through the UXA cavern radiation shield. Here the signal comprises two muons whose tracks reconstruct a decay vertex in the region before the box itself. We call this region the “muon shadow”. Vertex reconstruction and control of backgrounds inside the muon shadow may be challenging, so we do not focus on this option hereafter, except for benchmarks with relatively hard muon final states.
- (ii) The proximity of the box to LHCb—approximately only four bunch crossing times for relativistic objects—may permit it to interface with LHCb’s planned triggerless readout, allowing for identification and at least partial reconstruction of the LLP event. For the benchmarks we consider here, this may enable one to tag a vector boson fusion jet for Higgs decays, or an associated $K^{(*)}$ for B decays.
- (iii) The modest size of the fiducial volume may also permit, in principle, implementation of more ambitious detection technologies such as calorimetry or

time-of-flight, providing momentum reconstruction and particle identification that will aid in the confirmation of a discovery.

B. Reach intuition

The geometric acceptance of the CODEX-b box is $\sim 1\%$ (normalized to 4π). The LLP reach is attenuated further by the distribution of the LLP production and interplay between the LLP lifetime τ and the box depth. The number of LLP decay vertices expected in the box

$$N_{\text{box}} = \mathcal{L}_{\text{LHCb}} \times \sigma_{pp \rightarrow \phi X} \times \int_{\text{vol}} \frac{d\epsilon(r, \eta)}{dV} dV, \quad (1)$$

where the location of the box is specified by an azimuthal angle, the distance from the IP, r , and the pseudorapidity, η . In these coordinates, the differential fiducial efficiency is

$$\frac{d\epsilon(r, \eta)}{dV} = \frac{1}{2\pi r^2 c\tau} \int d\beta w(\beta, \eta) \times \frac{e^{-r/(c\tau\beta\gamma)}}{\beta\gamma}, \quad (2)$$

with γ and β the usual kinematic variables. The function $w(\beta, \eta)$ is the differential probability of producing the LLP with pseudorapidity η and velocity β , and is typically obtained from Monte Carlo.

To gain a rough sense of the achievable fiducial efficiency, let us assume w is factorizable into a δ -function in β at $\beta_0\gamma_0 \sim 3$ and a flat distribution in pseudorapidity for $|\eta| < \eta_0 \sim 5$. This is a reasonable approximation for, e.g., an exotic Higgs decay. That is, $w(\beta, \eta) \approx \delta(\beta - \beta_0)/(2\eta_0)$

on the box domain $\eta \in [0.2, 0.6]$. The fiducial efficiency is then approximately

$$\varepsilon_{\text{box}} \simeq \frac{0.4}{2\eta_0} \frac{|\phi_2 - \phi_1|}{2\pi} [e^{-r_1/r_0} - e^{-r_2/r_0}], \quad (3)$$

with $r_0 = c\tau\beta_0\gamma_0$. Using $|\phi_2 - \phi_1| \sim 10/25$, $r_1 \sim 25$ m, $r_2 \sim 35$ m, one estimates a maximum fiducial efficiency $\varepsilon_{\text{box}} \sim 10^{-3}$. In the long (short) lifetime regime $c\tau \gg (r_{1,2})$ ($c\tau \ll r_{1,2}$), this efficiency is linearly (exponentially) suppressed by $|r_2 - r_1|/r_0$ (e^{-r_1/r_0}). In the case of Higgs decay to dark photons, e.g., this translates to a maximal 2σ exclusion reach of $\text{Br}[h \rightarrow 2\gamma_d] \sim 10^{-4}$, for $\mathcal{L} = 300 \text{ fb}^{-1}$ expected after Run 5. We confirm this estimate with a more detailed simulation below.

C. Tracking

In order to demonstrate the feasibility of the proposed detector, we have studied a simple tracking layout based on resistive plate chamber (RPC) strip modules with 1 cm^2 effective granularity. Such modules typically also have 1 ns or better timing resolution, which may be useful for background rejection or improving the reconstruction of slow-moving signals.

As there is no magnetic field to provide a momentum estimate, it is crucial that we have the best possible spatial resolution on the decay vertex of our hypothetical long-lived particle. The most important parameter for vertex resolution is the distance between the vertex and the first measured point on the decay product trajectories. Further, softer LLPs may have large opening angles, so that their decay products can exit the box through any of its six faces with non-negligible probability. For these reasons, we instrument each of the six box faces with a sextet of RPC layers, and we also add five equally spaced triplets of RPC layers along the depth of the box (i.e. $y-z$ aligned planes with transverse spacing $\delta x = 1.67$ m). The layers in each RPC sextet or triplet are spaced at 4 cm intervals. For the purpose of our estimates of the efficiency, we consider an event to be reconstructed if (i) each track has at least 6 hits, (ii) each track has at least 600 MeV of momentum (cf. Ref. [68]) and (iii) none of the hits are shared between the tracks. A detailed study of tracking performance for different particle species, including the effects of multiple scattering, is left for a future study.

We study the vertex reconstruction efficiency achieved by the above-specified layout for the $h \rightarrow \gamma_d\gamma_d$ and $B \rightarrow \varphi X_s$ benchmarks over a range of γ_d, φ masses and lifetimes. Details of these models and the reach analysis are provided below in Sec. III. For the purposes of this study we decay the long-lived particle into two charged tracks, not least because a realistic estimate of the performance for multibody decays will depend much more strongly on assumptions about the detector layout and the impact of multiple scattering. The reconstruction efficiencies for the

TABLE I. Efficiency of reconstructing at least two tracks for a LLP decaying in the fiducial volume, for both $B \rightarrow X_s\varphi$ and $h \rightarrow \gamma_d\gamma_d$ scenarios, for various lifetimes. Masses are in GeV.

$c\tau$ (m)	$m_\varphi[B \rightarrow X_s\varphi]$			$m_{\gamma_d}[h \rightarrow \gamma_d\gamma_d]$				
	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

benchmark scenarios are shown in Table I. In the case of $B \rightarrow \varphi X_s$ decays, the efficiency is dominated by the momenta of the decay products. By contrast, the main source of inefficiency for $h \rightarrow \gamma_d\gamma_d$ decays, particularly at low φ mass, is the small opening angle of the decay, which results in partially overlapping decay products. We currently reject all signal decays for which the first measured points of the decay products are within the nominal detector granularity (1 cm) of each other, however this inefficiency may be reducible by improving the detector granularity or by allowing fewer than six hits on a track.

D. Backgrounds and shielding

Backgrounds from cosmics are suppressed because of the shielded underground location, by the fact that most cosmics are out of time with respect to the LHC collisions, and because the signal track velocities should have a significant horizontal component, as the IP is horizontally displaced from the detector. The rate of cosmic muons incident on CMS was simulated to be $\sim 1 \text{ kHz}$ [69], the vast majority of which would traverse the box from the top to bottom. Once timing and directionality are accounted for, this background is therefore subdominant to muons from the primary vertex, which have a rate of the order 1–10 kHz [70].

Apart from the existing UXA cavern radiation shield, additional shielding is required to attenuate collision backgrounds to manageable levels. Prompt muons, or muons from decays-in-flight of charged pions and kaons, may typically generate tracks in the detector volume, but the likelihood of two such muon tracks themselves reconstructing a fake vertex inside the box is expected to be negligible once muons that point back to vicinity of the IP are vetoed. However, muons may scatter on air inside the detector volume, producing tracks either from ionized electrons, elastic pair production or muon-nucleon deep inelastic scattering (see Ref. [71] for a review). Other backgrounds arise from neutral kaon decays producing track pairs, and neutrons or neutrinos producing scattering events with

morphologies resembling a signal. On top of the primary fluxes of leptons and hadrons, also secondary production in the shield itself can be significant, such that the inclusion of active veto components in the lead shielding is required. A “shield veto” for charged particles (mainly μ 's and π^\pm) would remove events where, for example, a muon produces secondary neutrons or kaons in the last few interaction lengths of the concrete shield, that then enter the box and either scatter on nuclei or decay, respectively.

The minimum amount of shielding can be estimated by sufficiently reducing the prompt K_L and neutron components. From a simulated minimum bias sample generated with Pythia 8 [72,73], the fiducial efficiency for promptly produced K_L 's and n to enter the box is $\varepsilon_{\text{box}} \approx 4.8 \times 10^{-3}$ and 5.1×10^{-3} , respectively. For a QCD cross section $\sigma_{\text{qcd}} \approx 100$ mb, at 300 fb^{-1} of integrated luminosity one naively requires $\approx 32\lambda$ of shielding— λ is an interaction length—to suppress these backgrounds to $\lesssim 1$ event. The natural location for such a shield is as close to the IP as possible, as this reduces the volume of shielding material required. Based on the cavern geometry shown in Fig. 1, we focus on the case that a Pb shield is installed covering the detector geometric acceptance, beginning approximately 5m from the IP and several meters in depth. (One may also consider other shielding materials such as tungsten, but the qualitative conclusions remain the same.) As a proof-of-concept, based on the quoted Particle Data Group nuclear interaction lengths for energetic neutrons, respectively 18 cm and 42 cm [74], we further focus on the following shielding configuration: 3.6 m of Pb, corresponding to 20λ ; a shield veto comprising RPC or scintillator layers with efficiency $\varepsilon_{\text{veto}}$; an additional 0.9 m of Pb, corresponding to 5λ ; and finally the UXA cavern concrete shield, comprising approximately 7λ . We refer to this as a “(20 + 5) λ ” Pb shielding configuration and consider it as our benchmark in the following. Further optimization of the shield design is beyond the scope of this study.

We model the shielding response to backgrounds—propagation of primary muons, pions, kaons, neutrons and protons and secondary production of these and neutrinos and photons—with a preliminary Geant4 [75–77] simulation which includes propagation in lead, air and concrete. We simulate the initial muon flux for $E_\mu \lesssim 7$ GeV with Pythia 8; pion, kaon and heavy flavor decays in minimum bias events produce the dominant primary contribution in this regime. Based on the geometry of Fig. 1, we conservatively include secondary muons produced by pions or kaons that decay within 5 m transverse of the IP, and within 3 m forward of the IP. The resulting estimated muon flux for $E_\mu \gtrsim 3$ GeV is typically $\mathcal{O}(\text{few})$ larger than a naive estimation of the same from inclusive muon spectra measured at ATLAS [78]. To estimate the spectrum for $E_\mu > 7$ GeV, we use the shape of measured muon inclusive spectra at ATLAS [79] and rescale them to match onto our simulated results, assuming the muon spectrum is roughly

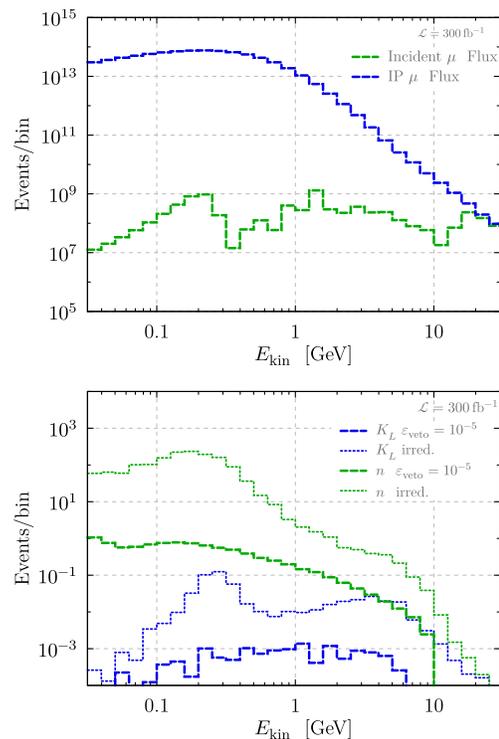


FIG. 2. (Top) Energy dependence of muon number flux before propagation through shielding (blue), and of muon number flux incident on the fiducial volume (green) after propagation through shielding. (Bottom) Energy dependence of neutral kaon (blue) and neutron (green) shield vetoable and irreducible fluxes. The vetoable fluxes are specified with a rejection factor $\varepsilon_{\text{veto}} = 10^{-5}$ already applied.

flat in the $|\eta| < 2$ regime and modified by only $\mathcal{O}(1)$ factors upon mapping from 7 to 14 TeV, the latter verified using FONLL [80,81]. Primary fluxes and spectra of kaons, pions, protons and neutrons in the box acceptance are generated from a minimum bias sample produced with Pythia 8.

The simulation of the $(20 + 5)\lambda$ shielding response proceeds as follows: all primary fluxes are propagated through 20λ of Pb up to the shield veto, then partitioned into intermediate charged and neutral fluxes, since only charged fluxes can be rejected by the active veto. We then propagate these charged and neutral fluxes through the remaining lead, air and concrete layers separately and refer to them as the reducible and irreducible fluxes respectively: The veto has power to reject secondaries produced from charged states downstream—in the matter between the veto and the detector—but not neutral secondaries produced upstream—in the matter between the IP and the veto—or produced by neutral fluxes.

In Fig. 2 we show the kinetic energy distributions for the flux of muons, neutral kaons and neutrons entering the box, after propagating through the shield. Both neutrons and kaons are quite soft, so that only a small fraction are kinematically capable of producing multiple tracks and are sufficiently boosted to do so within timing constraints. In

TABLE II. Particle yields in CODEX-b from the preliminary Geant4 background simulation. Yields are classified by particle species and calculated for 300 fb^{-1} of integrated luminosity using the $(20 + 5)\lambda$ Pb shield. Separate yields for (ir)reducibility by the active shield veto are indicated. The irreducible π^\pm, K^\pm yields can be vetoed by the box tracking layers themselves. The yield of ~ 7 irreducible neutrons corresponds to $\lesssim 0.4$ BG events, assuming the box volume is filled with air. The applied cuts conservatively select the kinematic thresholds necessary for producing charged tracks in the box.

BG species	Particle yields		Baseline cuts
	Irreducible by shield veto	Reducible by shield veto	
$n + \bar{n}$	7	5×10^4	$E_{\text{kin}} > 1 \text{ GeV}$
K_L^0	0.2	9×10^2	$E_{\text{kin}} > 0.5 \text{ GeV}$
$\pi^\pm + K^\pm$	0.5	3×10^4	$E_{\text{kin}} > 0.5 \text{ GeV}$
$\nu + \bar{\nu}$	0.5	2×10^6	$E > 0.5 \text{ GeV}$

Table II we show the results of this preliminary analysis, including approximate energy thresholds to produce charged tracks and/or satisfy timing constraints. The dominant contributions in both shield reducible and irreducible populations arise from neutrons. [The neutrinos have a sub-pb cross section, resulting in $\ll 1$ scattering event for a population of $\mathcal{O}(10^6)$ size.] The neutron incoherent cross section on nitrogen ~ 1 b, so that the probability of a neutron scattering on air somewhere in the box is $\sim 5\%$. From Table II, the number of neutron scattering background events is then ~ 0.4 . Moreover, if needed, any surviving unvetoable neutrons can be further suppressed with a marginal increase in the Pb shield, or based on the event morphology given their softness. The vetoable neutron population suggests $\epsilon_{\text{veto}} \sim 10^{-5}$ is required, which is used in Fig. 2 (bottom) for the neutron and K_L fluxes. These results broadly agree with a simplified propagation model making use of the muon average energy loss $\langle -dE_\mu/dx \rangle$ for Pb and standard concrete, taken from Particle Data Group data [74], the quoted neutron interaction lengths [74], the measured neutral kaon absorption cross section in Pb at low energies [82,83] and the inclusive muoproduction cross section for neutral kaons [84–88].

The muon flux entering the box is shown in Fig. 2 (top). To conservatively account for dangerous muon-air interactions, we (over)estimate the cross section via the total muon-proton cross section with energy exchange $\delta E_\mu > 0.1 \text{ GeV}$. This cross section includes contributions from: elastic pair-production, $\mu p \rightarrow \mu e^+ e^- p$; inelastic photonuclear processes, $\mu n/p \rightarrow \mu X$, with X a hadronic final state containing at least one baryon; bremsstrahlung, for which the radiated photon has low, but nonzero conversion probability; and ionization. Between 10 and 100 GeV, this cross section is roughly constant, being $\mathcal{O}(\text{mb})$ [71,89]. Keeping only muons above 0.1 GeV and including a 10^{-5} rejection factor by the shield veto, the

number of air scattering muons in the box is $\mathcal{O}(70)$ for $\mathcal{L} = 300 \text{ fb}^{-1}$. These events can be rejected by the RPC layers on the detector faces.

Neutrino air-scattering inside the detector volume, arises from either the atmospheric ν flux or from the IP. For neutrinos with $E \sim 1 \text{ GeV}$, the total neutrino-nucleon cross section is roughly $0.01 \text{ pb} \times (E_\nu/\text{GeV})$ [90], which implies that the probability for a $\sim 1 \text{ GeV}$ neutrino to scatter in the fiducial volume is $\mathcal{O}(10^{-14})$. For an atmospheric neutrino flux of $\sim 1 \text{ cm}^{-2} E_\nu^{-1} s^{-1}$ [91], we expect the roughly 0.01 collisions per year, which is revised to ~ 0.05 if the energy dependence of the cross section and neutrino flux are accounted for. If needed, this background can be reduced further by imposing timing and directionality constraints.

The neutrino background from the IP requires a proper simulation of the material budget in the LHCb cavern, which is beyond this proof-of-concept study. Once an energy threshold cut $E_\nu > 0.5 \text{ GeV}$ is included, the contribution to this flux from the decays of hadrons (mostly π^\pm) propagating through the shield is negligible, as indicated by Table II. Similarly, pions decaying in any dense material inside the cavern or in its walls overwhelmingly produce only soft neutrinos. Instead, the main contribution to the neutrino flux comes from pion decays in the empty spaces within the cavern. We conservatively estimate this flux by simulating the truth-level neutrino energy spectrum from pion and kaon decays with Pythia 8. We impose a $|\eta| < 5$ cut and assume that the remaining neutrinos radiate out spherically from the IP. We further require that the decay lengths of the pion to be shorter than 20 m in the forward direction, $z > 0$, and require the neutrino energy to be larger than 0.5 GeV as in Table II. [Harder neutrinos from heavy flavor decays have a flux at least $\mathcal{O}(10^{-2})$ smaller than in this estimate.] Folding this flux with the total neutrino-nucleon cross section [90] to estimate an upper limit on the number of neutrinos scattering against air in the fiducial volume, we find roughly $\mathcal{O}(3)$ potential events, but we expect that this estimate is still very conservative. In particular, the neutrino cross section falls quadratically with decreasing energy below 1 GeV, no pion energy losses have been included, and the size of the empty cavern spaces have been overestimated. This conclusion is also supported by a more realistic simulation for the MATHUSLA proposal [54], which estimates at most $\mathcal{O}(3)$ IP generated neutrino events at the HL-LHC, but for ~ 100 times larger luminosity times geometric acceptance compared to CODEX-b.

We finally note that the location of CODEX-b and the fact that all the relevant backgrounds originate in the LHC collisions provides a simple way of calibrating them: all that is required is to install a telescope consisting of a few RPC layers in the LHCb cavern. Using this one can measure background fluxes with and without the additional shielding proposed for CODEX-b, and obtain a data-driven measurement of the background rates from which to

extrapolate the expected numbers of background events in CODEX-b itself.

III. REACH ANALYSIS

A. Exotic B decays

A Higgs mixing portal admits exotic inclusive $B \rightarrow X_s \varphi$ decays, in which φ is a light CP -even scalar that mixes with the Higgs, with mixing angle $\theta \ll 1$. The inclusive branching ratio [63–65]

$$\frac{\text{Br}[B \rightarrow X_s \varphi]}{\text{Br}[B \rightarrow X_c e \nu]} = \frac{27g^2 s_\theta^2}{256\pi^2} \frac{m_t^4}{m_b^2 m_W^2} \frac{(1 - m_\varphi^2/m_b^2)^2}{f(m_c/m_b)} \left| \frac{V_{ts} V_{tb}}{V_{cb}} \right|^2, \quad (4)$$

with $s_\theta \equiv \sin\theta$. This yields $\text{Br}[B \rightarrow X_s \varphi] \approx 6 \cdot s_\theta^2 (1 - m_\varphi^2/m_b^2)^2$ from the measured inclusive $B \rightarrow X_c e \nu$ semileptonic rate [74]. This ratio is expected to break down for m_φ approaching m_b . Approximating the inclusive rate instead by adding the various exclusive contributions, we have checked that we do not recover sensitivity in this high mass region. In our simple benchmark model, the scalar φ decays exclusively to SM states, such that its lifetime is also controlled by s_θ^2 , though complicated by effects of hadronic resonances when $m_\varphi \sim 1$ GeV, as well as threshold effects. The theory uncertainties in this region are substantial, and we use the data-driven result from [92,93]. Another common choice is the perturbative spectator model [94]. The difference in reach predicted by both models can be rather large, especially in the region around the masses of the f_0 mesons. As an example, for $m_\varphi = 1$ GeV the branching ratio is related to the lifetime via $\text{Br}[B \rightarrow X_s \varphi](c\tau_\varphi/1 \text{ m}) \sim 6 \times 10^{-10}$ in the data driven model, while the analogous number in the spectator model is 2×10^{-8} .

We generate a B meson sample with Pythia 8, enforcing the exclusive decay $B \rightarrow K\varphi$ as a proxy to estimate the box fiducial efficiency for $B \rightarrow X_s \varphi$. We do not include muon shadow contributions, as the muon energy in these LLP decays is typically low, $E_\mu \lesssim \text{few GeV}$, so that muon penetration (scattering) through the concrete shield might be unacceptably low (high) for a decay vertex reconstruction. The peak box fiducial efficiency is $\sim 10^{-4}$ at $c\tau_\varphi \sim 10$ m: the LLPs captured by the transversely located box are typically only mildly boosted.

In Fig. 3 we show the CODEX-b reach (exclusion at greater than 95% CL) on the Higgs mixing portal s_θ^2 - m_φ parameter space, compared with existing bounds from CHARM [95] and LHCb [41], as well as projected reaches for LHCb, MATHUSLA and SHiP. We assume a $b\bar{b}$ production cross section of $500 \mu\text{b}$. For the projected LHCb reach we rescaled the existing $B \rightarrow K(\varphi \rightarrow \mu\mu)$ limit [41] under the (optimistic) assumption of zero background, implying that the limit on the fiducial rate scales

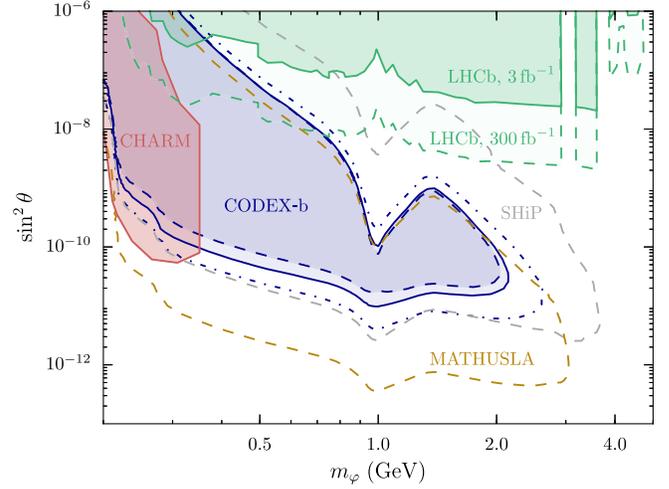


FIG. 3. CODEX-b reach for $B \rightarrow X_s \varphi$ in the s_θ^2 - m_φ plane. Solid (dashed) blue line assumes 100% (Table I) tracking efficiency. Dot-dashed line indicates the reach for $\mathcal{L} = 1 \text{ ab}^{-1}$.

linearly with the integrated luminosity. (A similar limit from $B \rightarrow K^*(\varphi \rightarrow \mu\mu)$ is slightly weaker [42].) The sensitivity of MATHUSLA to this signature has concurrently been pointed out in [96]. The curve here is our own recast of the MATHUSLA reach and agrees with the results in [96], up to small differences which can be attributed to slightly different assumptions regarding the width of φ . The original SHiP projection [97] was computed using a perturbative spectator model for the width of φ . To properly compare all experiments, the curve shown in Fig. 3 is a recast to the data-driven model in [92,93], where we use the efficiency maps provided in [97].

The lower extent of the reach in s_θ^2 is determined by the total number of beauty hadrons and the CODEX-b fiducial efficiency, while the upper extent of the s_θ^2 reach is controlled by the φ lifetime: a larger s_θ^2 implies a larger rate of φ production along with a shorter φ lifetime, such that most φ 's decay before they reach the detector. One finds that CODEX-b would significantly extend the reach of LHCb, and complement part of the projected parameter reach for SHiP as well as for MATHUSLA.

One may also consider more general portals that do not feature the fixed branching ratio-lifetime relations predicted by the simplest Higgs portal models. In Fig. 4 we show the branching ratio reach for such theories, for various φ mass benchmarks. Compared to LHCb, which searches for $B \rightarrow K(\varphi \rightarrow \mu\mu)$, a key advantage is that the reach is not sensitive to the model-dependent muonic branching ratio, only requiring instead that the final states are trackable. [Decays into neutral hadron pairs, such as $\pi^0\pi^0$, cannot be seen without calorimetry, however such final states comprise at most $\mathcal{O}(30\%)$ of final states for $2m_\pi < m_\varphi \lesssim 1$ GeV and typically otherwise comprise a much smaller contribution.] While the muon branching ratio is typically $\mathcal{O}(1)$ for $m_\varphi < 2m_K$ from kinematic

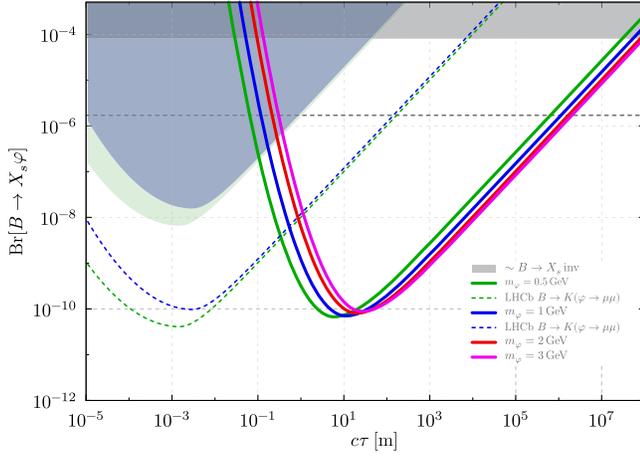


FIG. 4. Inclusive CODEX-b $B \rightarrow X_s \varphi$ reach (solid lines). The shaded regions (dashed lines) indicate current LHCb limits (300 fb^{-1} projection) from $B \rightarrow K(\varphi \rightarrow \mu\mu)$, rescaled to the inclusive process using the ratio of Eq. (4) and the theory predictions for the exclusive branching ratio [98,99], and assuming $\text{Br}[\varphi \rightarrow \mu\mu] \approx 30\%$ and 10% for $m_\varphi = 0.5 \text{ GeV}$ and 1 GeV , respectively. Approximate current [74] and Belle II projected [100] limits from $B \rightarrow K^{(*)} \nu \bar{\nu}$ precision measurements are also shown (gray shading and dashed line).

considerations, at higher masses this branching ratio may drop precipitously to the subpercent level. As an example, we show the projected LHCb reach in Fig. 4 for $m_\varphi = 0.5 \text{ GeV}$ compared to $m_\varphi = 1 \text{ GeV}$.

B. Exotic Higgs decays

Exotic Higgs decays to two dark photons may be generated by a kinetic mixing portal (e.g. [59–62]). In the short lifetime limit, dark photons can be searched for with the main LHCb detector, in D^* decays [101] or with an inclusive search [102]. To estimate the CODEX-b fiducial efficiency, we simulate gluon fusion Higgs production at IP8 with Pythia 8, with subsequent $h \rightarrow \gamma_d \gamma_d$ decay. The dark photon branching ratios to various SM final states are approximated from existing e^+e^- data [103], which is relevant if one exploits the muon shadow. In Fig. 5 we show the expected reach in $\text{Br}[h \rightarrow \gamma_d \gamma_d]$ for $m_{\gamma_d} = 0.5$ and 10 GeV benchmarks as a function of dark photon lifetime, for both the CODEX-b fiducial volume and for the case that the muon shadow can be used. For the 0.5 GeV benchmark, the larger $\gamma_d \rightarrow \mu\mu$ branching ratio enhances the reach of the muon shadow, compared to the 10 GeV case.

A displaced vertex search at ATLAS/CMS has geometric acceptance ~ 1 (normalized to 4π), and approximately 10 times higher luminosity. Other than the trigger challenges associated with LLPs, a second crucial distinction is that the calorimeters comprise only $\sim 10\lambda$ of shielding compared to the 32λ shield in the CODEX-b setup. Searches for light displaced objects in the ATLAS/CMS muon system are therefore expected to suffer from significant backgrounds

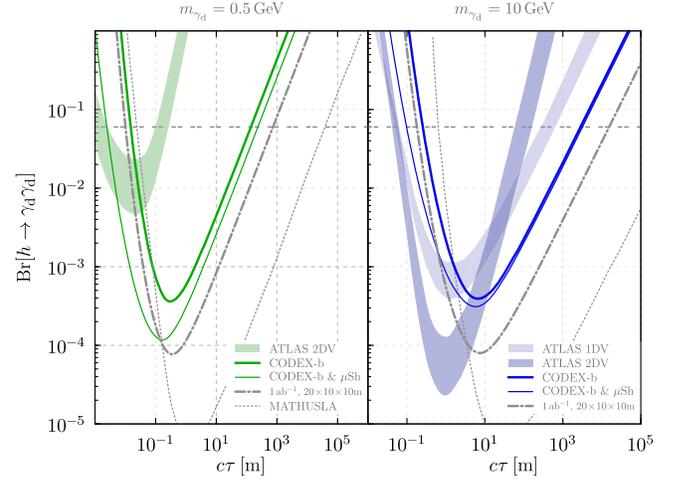


FIG. 5. Higgs decay to dark photon reach, using the CODEX-b fiducial volume alone and with the muon shadow “ μSh ”. The $\gamma_d \rightarrow \mu\mu$ branching ratio is taken from e^+e^- data [103]. Also shown is the CODEX-b reach with $\mathcal{L} = 1 \text{ ab}^{-1}$ and a larger box, should DELPHI be removed. The approximate reach for MATHUSLA (gray dotted), rescaled from [54], and $h \rightarrow$ invisibles is also shown (horizontal gray dashed) [104].

from punch-through jets. To heavily reduce these backgrounds, it is often necessary to require two displaced objects, which is a significant penalty in reach for the long lifetime regime. We compare our $m_{\gamma_d} = 10 \text{ GeV}$ benchmark point with the projected sensitivity of searches for one and two displaced jets [105]. The latter estimate is based on the existing ATLAS displaced dijet search [106]. Its sensitivity deteriorates when the displaced vertices generate a low number of tracks, which occurs both for $m_{\gamma_d} < 10 \text{ GeV}$ and for models with small hadronic branching ratios. Neither difficulty applies to CODEX-b.

For our $m_{\gamma_d} = 0.5 \text{ GeV}$ benchmark, we compare with the ATLAS search for a pair of displaced lepton jets [107]. This search is currently systematic limited, so that the range for our estimate of the HL-LHC ATLAS reach in the left panel of Fig. 5 is bounded above by the current expected limit [107] and bounded below by the current expected limit, rescaled under the assumption that the systematic uncertainties can be reduced with a factor of five. We expect backgrounds for a single displaced dilepton search would be prohibitively large.

C. Mass measurement

Aside from the discovery potential outlined above, CODEX-b should also be capable of measuring the velocity of the LLP. For a given assumption on the production mechanism, this then allows for a mass measurement of the new state on a statistical basis. As is well known (see e.g. Ref. [55]), the geometry of two-body γ_d/φ decays to massless final states can provide information about their velocity and the ability to discriminate between different

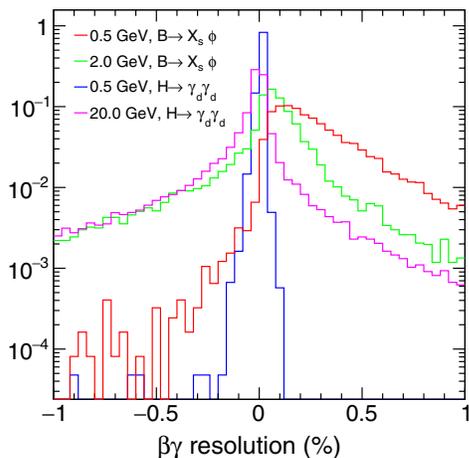


FIG. 6. Boost ($\beta\gamma$) resolution for different signals.

γ_d/φ masses. More complex final states may be possible as well, but in this proof-of-concept study we restrict ourselves to two body decays only.

The β resolution which can be achieved with our simple tracking layout, using geometric information only and assuming massless final states, is shown in Fig. 6. It is Gaussian for the $h \rightarrow \gamma_d\gamma_d$ benchmarks, while in the $B \rightarrow X_s\varphi$ case it is non-Gaussian and biased, because the φ decay products are so slow that the approximation of $\beta = 1$ begins to break down for them. Nevertheless we can still reconstruct the γ_d/φ velocity to better than 1% in all cases, which in practice means that the ability to discriminate between different γ_d/φ masses is largely dominated by the actual distribution of γ_d/φ velocities for a given mass, and not by the detector resolution.

The corresponding distribution of reconstructed γ_d/φ boosts is shown in Fig. 7 for different $B \rightarrow X_s\varphi$ and $h \rightarrow \gamma_d\gamma_d$ masses and lifetimes. We achieve good discrimination across a wide range of masses in the $h \rightarrow \gamma_d\gamma_d$ case, but perhaps more surprisingly we also have some discriminating power between different φ masses for the $B \rightarrow X_s\varphi$ benchmark.

This approach to a mass measurement uses only spatial information, but complementary information may be provided by using timing information from the RPC stations to measure the velocity of the decay products. Such a mass reconstruction can be useful to discriminate between slow-moving new states, or to veto unexpected slow-moving SM backgrounds which decay or scatter in the box. We show the performance in the $B \rightarrow X_s\varphi$ case in Fig. 8, assuming timing resolutions of 100 ps or 50 ps for the RPCs, which may be affordably achievable.² in the next decade. The reconstructed mass resolution depends very strongly on the mass, because heavier φ particles produce decay products

²The calibration and time alignment of the box can be performed with cosmic rays, just as the initial time alignment of LHCb was.

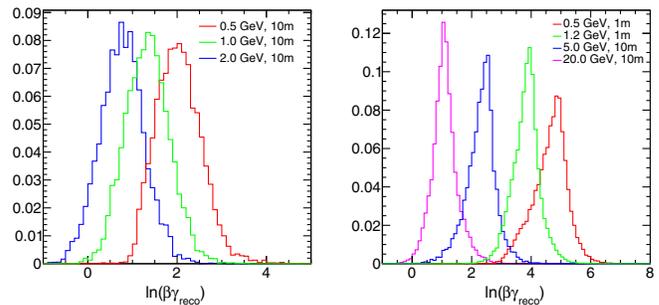


FIG. 7. Reconstructed LLP boost for different $B \rightarrow X_s\varphi$ (left) and $h \rightarrow \gamma_d\gamma_d$ (right) mass benchmarks and $c\tau = 1$ or 10 m.

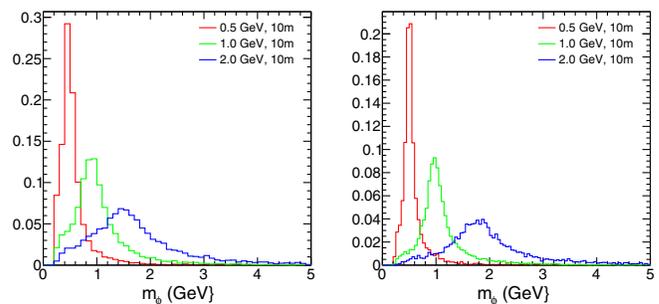


FIG. 8. Reconstructed LLP mass for different $B \rightarrow X_s\varphi$ mass benchmarks and $c\tau = 10$ m, for 100 ps (left) and 50 ps (right) time resolution.

which are significantly faster on the scales relevant here. Nevertheless even in the 100 ps case it is largely possible to separate the 0.5 GeV and 2.0 GeV signals, which also means that we should be able to reconstruct, e.g., any residual background $K_S^0 \rightarrow \pi\pi$ decays that might occur in the box. In the Higgs case the two-body decay products are always too fast to allow this kind of mass reconstruction, though conversely this also allows them to be discriminated from SM backgrounds and from lighter exotic long-lived states produced, for instance, in heavy flavour decays.

IV. DISCUSSION AND CONCLUSIONS

The results for the $h \rightarrow \gamma_d\gamma_d$ and $b \rightarrow s\varphi$ benchmarks provide a proof-of-concept that CODEX-b may significantly extend the reach of LLP searches beyond what the current LHC experiments can achieve, while also covering significant regions of the parameter spaces that other proposed LLP searches seek to explore. However, both for these generic benchmarks and for more specific theories, there is little theoretical guidance as to where the focus of LLP searches should fall within the wide range of unconstrained LLP masses, branching ratios and lifetimes, and it is difficult to physically construct a single detector capable of exploring this entire parameter space. Thus a picture of complementary detectors with decorrelated backgrounds emerges as the natural choice for a comprehensive LLP search program.

In this picture, ATLAS and CMS emerge as the most suitable tools to search for heavy and/or charged LLPs for a wide range of lifetimes, due to their $\sim 4\pi$ coverage and rather large size. LHCb adds complementary sensitivity for lower mass LLPs with relatively short lifetimes. Beam dump experiments such as NA62 and SHiP, and ultra-forward LHC detectors like FASER, benefit from reduced backgrounds and a focused LLP source, and can therefore be instrumented with a magnet while also covering a wide range of potential lifetimes. However, for LLPs above a few GeV in mass or LLPs which are produced through a heavy portal, they have limited reach, and in particular they can only explore the Higgs sector indirectly. By contrast, collider-based and shielded detectors such as CODEX-b or MATHUSLA can explore further in mass, but require larger geometric sizes to ensure adequate fiducial acceptances. This and other practical considerations make it more difficult to instrument this type of detector with a magnet, so that instead a combination of time-of-flight information, possible calorimetry, and kinematic constraints must be used in order to infer the mass and/or boost of any observed signal candidates.

The location, size and dominant backgrounds for the latter two proposals are different—the MATHUSLA detector will nominally achieve both a larger solid angle coverage and up to ten times the integrated luminosity compared to CODEX-b, but is approximately ten times further from the IP. Consequently, the absolute reach, the LLP lifetime at their peak sensitivities, as well as their expected sensitivity to different LLP decay morphologies differs. Further significant differences are expected for the optimal tracker layouts; the potential to deploy more complex detector technologies such as calorimetry; alignment strategies; and reconstruction strategies.³ These differences enhance the complementarity between both experimental setups and imply that compatible LLP signals seen in both experiments would represent compelling evidence of physics beyond the Standard Model.

³An example of this is that optimization of the tracker layout for CODEX-b in Sec. II C requires minimization of the distance between the LLP decay vertex and the first measured point on the trajectories of the decay products. This parameter is particularly critical for CODEX-b because of the relatively short lever arm between the LLP production and decay vertices.

In summary, we have proposed the construction of a new detector element for the LHCb experiment, taking advantage of a soon-to-be available shielded space currently occupied by the data acquisition. The Compact Detector for Exotics at LHCb (“CODEX-b”) has the potential to significantly enhance the new physics reach and capabilities of the LHCb program, in ways complementary to, and in some cases exceeding, the reach of ATLAS and CMS and the LHCb main detector. We have verified in this proof-of-concept study that efficient vertex reconstruction can be achieved with an RPC tracking implementation, and that reconstruction of LLP mass and boost is feasible. In terms of reach, we project that for Higgs decays to sub-GeV dark photons and for exotic B -decays through a Higgs portal, CODEX-b would extend the reach of the currently planned LHC program by over one and two orders of magnitude, respectively.

ACKNOWLEDGMENTS

We thank Yasmine Amhis, Bob Cahn, Xabier Cid Vidal, David Curtin, Jared Evans, Tim Gershon, Heather Gray, Gaia Lanfranchi, Jacques Lefrancois, Alex Marks-Bluth, Simone Pagan-Griso, Chris Parkes, Giovanni Passaleva, Michael Sokoloff, Frederic Teubert and Mike Williams for helpful discussions. We especially thank Tim Gershon and Mike Williams for comments on the manuscript and Rolf Lindner for useful discussions about the layout of the LHCb cavern and for providing us with high resolution drawings of it. V. V. G. acknowledges funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program under Grant Agreement No. 724777 “RECEPT”. S. K. and M. P. are supported in part by the LDRD program of LBNL under Contract No. DE-AC02-05CH11231, and by the National Science Foundation (NSF) under Grants No. PHY-1002399 and No. PHY-1316783. D. R. acknowledges support from the University of Cincinnati. S. K., M. P. and D. R. also thank the Aspen Center of Physics, supported by the NSF Grant No. PHY-1066293, where this paper was completed. This research used resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the DoE under Contract No. DE-AC02-05CH11231.

- [1] S. Dimopoulos, M. Dine, S. Raby, and S. D. Thomas, *Phys. Rev. Lett.* **76**, 3494 (1996).
 [2] Z. Liu and B. Tweedie, *J. High Energy Phys.* **06** (2015) 042 (2015).

- [3] L. J. Hall and M. Suzuki, *Nucl. Phys.* **B231**, 419 (1984).
 [4] C. S. Aulakh and R. N. Mohapatra, *Phys. Lett.* **119B**, 136 (1982).

- [5] C. Csaki, E. Kuflik, S. Lombardo, O. Slone, and T. Volansky, *J. High Energy Phys.* **08** (2015) 016.
- [6] J. Fan, M. Reece, and J. T. Ruderman, *J. High Energy Phys.* **011** (2011) 012.
- [7] A. Arvanitaki, N. Craig, S. Dimopoulos, and G. Villadoro, *J. High Energy Phys.* **02** (2013) 126.
- [8] N. Arkani-Hamed, A. Gupta, D. E. Kaplan, N. Weiner, and T. Zorawski, [arXiv:1212.6971](https://arxiv.org/abs/1212.6971).
- [9] Z. Chacko, H.-S. Goh, and R. Harnik, *Phys. Rev. Lett.* **96**, 231802 (2006).
- [10] G. Burdman, Z. Chacko, H.-S. Goh, and R. Harnik, *J. High Energy Phys.* **02** (2007) 009.
- [11] H. Cai, H.-C. Cheng, and J. Terning, *J. High Energy Phys.* **05** (2009) 045.
- [12] P. W. Graham, D. E. Kaplan, and S. Rajendran, *Phys. Rev. Lett.* **115**, 221801 (2015).
- [13] T. Flacke, C. Frugiuele, E. Fuchs, R. S. Gupta, and G. Perez, *J. High Energy Phys.* **06** (2017) 050.
- [14] M. Baumgart, C. Cheung, J. T. Ruderman, L.-T. Wang, and I. Yavin, *J. High Energy Phys.* **04** (2009) 014.
- [15] D. E. Kaplan, M. A. Luty, and K. M. Zurek, *Phys. Rev. D* **79**, 115016 (2009).
- [16] K. R. Dienes, S. Su, and B. Thomas, *Phys. Rev. D* **86**, 054008 (2012).
- [17] I.-W. Kim and K. M. Zurek, *Phys. Rev. D* **89**, 035008 (2014).
- [18] A. Falkowski, Y. Hochberg, and J. T. Ruderman, *J. High Energy Phys.* **11** (2014) 140.
- [19] Y. Hochberg, E. Kuflik, and H. Murayama, *J. High Energy Phys.* **05** (2016) 090.
- [20] R. T. Co, F. D’Eramo, L. J. Hall, and D. Pappadopulo, *J. Cosmol. Astropart. Phys.* **12** (2015) 024.
- [21] A. Davoli, A. De Simone, T. Jacques, and V. Sanz, *J. High Energy Phys.* **11** (2017) 025.
- [22] Y. Cui and R. Sundrum, *Phys. Rev. D* **87**, 116013 (2013).
- [23] Y. Cui and B. Shuve, *J. High Energy Phys.* **02** (2015) 049.
- [24] K. Barry, P. W. Graham, and S. Rajendran, *Phys. Rev. D* **89**, 054003 (2014).
- [25] M. L. Graesser, *Phys. Rev. D* **76**, 075006 (2007).
- [26] M. L. Graesser, [arXiv:0705.2190](https://arxiv.org/abs/0705.2190).
- [27] J. C. Helo, M. Hirsch, and S. Kovalenko, *Phys. Rev. D* **89**, 073005 (2014); **93**, 099902E (2016).
- [28] A. Maiezza, M. Nemevšek, and F. Nesti, *Phys. Rev. Lett.* **115**, 081802 (2015).
- [29] B. Batell, M. Pospelov, and B. Shuve, *J. High Energy Phys.* **08** (2016) 052.
- [30] S. Antusch, E. Cazzato, and O. Fischer, *J. High Energy Phys.* **12** (2016) 007.
- [31] P. S. B. Dev, R. N. Mohapatra, and Y. Zhang, *Nucl. Phys.* **B923**, 179 (2017).
- [32] M. J. Strassler and K. M. Zurek, *Phys. Lett. B* **651**, 374 (2007).
- [33] T. Han, Z. Si, K. M. Zurek, and M. J. Strassler, *J. High Energy Phys.* **07** (2008) 008.
- [34] M. J. Strassler and K. M. Zurek, *Phys. Lett. B* **661**, 263 (2008).
- [35] S. Chatrchyan *et al.* (CMS), *Phys. Lett. B* **716**, 30 (2012).
- [36] G. Aad *et al.* (ATLAS), *Phys. Lett. B* **716**, 1 (2012).
- [37] CMS Collaboration, CERN Report No. CMS-PAS-EXO-16-036, 2016.
- [38] M. Aaboud *et al.* (ATLAS), *Phys. Rev. D* **93**, 112015 (2016).
- [39] CMS Collaboration, CERN Report No. CMS-PAS-EXO-16-003, 2017.
- [40] ATLAS Collaboration, CERN Report No. ATLAS-CONF-2016-103, 2016.
- [41] R. Aaij *et al.* (LHCb), *Phys. Rev. D* **95**, 071101 (2017).
- [42] R. Aaij *et al.* (LHCb), *Phys. Rev. Lett.* **115**, 161802 (2015).
- [43] A. Pierce, B. Shakya, Y. Tsai, and Y. Zhao, [arXiv:1708.05389](https://arxiv.org/abs/1708.05389).
- [44] S. Antusch, E. Cazzato, and O. Fischer, *Phys. Lett. B* **774**, 114 (2017).
- [45] R. Aaij *et al.* (LHCb), *Eur. Phys. J. C* **77**, 812 (2017).
- [46] R. Aaij *et al.* (LHCb), *Eur. Phys. J. C* **77**, 224 (2017).
- [47] R. Aaij *et al.* (LHCb), *Eur. Phys. J. C* **76**, 664 (2016).
- [48] R. Aaij *et al.* (LHCb), *Eur. Phys. J. C* **75**, 595 (2015).
- [49] R. Aaij *et al.* (LHCb), *Eur. Phys. J. C* **75**, 152 (2015).
- [50] Y. Hochberg, E. Kuflik, T. Volansky, and J. G. Wacker, *Phys. Rev. Lett.* **113**, 171301 (2014).
- [51] N. Craig, A. Katz, M. Strassler, and R. Sundrum, *J. High Energy Phys.* **07** (2015) 105.
- [52] E. Cortina Gil *et al.* (NA62), *J. Instrum.* **12**, P05025 (2017).
- [53] S. Alekhin *et al.*, *Rep. Prog. Phys.* **79**, 124201 (2016).
- [54] J. P. Chou, D. Curtin, and H. J. Lubatti, *Phys. Lett. B* **767**, 29 (2017).
- [55] D. Curtin and M. E. Peskin, *Phys. Rev. D* **97**, 015006 (2018).
- [56] A. Haas, C. S. Hill, E. Izaguirre, and I. Yavin, *Phys. Lett. B* **746**, 117 (2015).
- [57] A. Ball *et al.*, [arXiv:1607.04669](https://arxiv.org/abs/1607.04669).
- [58] J. Feng, I. Galon, F. Kling, and S. Trojanowski, *Phys. Rev. D*, [arXiv:1708.09389](https://arxiv.org/abs/1708.09389).
- [59] R. M. Schabinger and J. D. Wells, *Phys. Rev. D* **72**, 093007 (2005).
- [60] S. Gopalakrishna, S. Jung, and J. D. Wells, *Phys. Rev. D* **78**, 055002 (2008).
- [61] D. Curtin, R. Essig, S. Gori, and J. Shelton, *J. High Energy Phys.* **02** (2015) 157.
- [62] M. J. Strassler, [arXiv:0801.0629](https://arxiv.org/abs/0801.0629).
- [63] R. S. Willey and H. L. Yu, *Phys. Rev. D* **26**, 3287 (1982).
- [64] R. Chivukula and A. V. Manohar, *Phys. Lett. B* **207**, 86 (1988).
- [65] B. Grinstein, L. J. Hall, and L. Randall, *Phys. Lett. B* **211**, 363 (1988).
- [66] R. Lindner (private communication).
- [67] I. Bediaga *et al.* (LHCb Collaboration), Reports No. CERN-LHCC-2012-007, No. LHCb-TDR-12, 2017.
- [68] LHCb Collaboration, “Tracking and Alignment Plots for Conferences,” https://twiki.cern.ch/twiki/bin/view/LHCb/ConferencePlots#Tracking_efficiency.
- [69] V. Drollinger, Report No. CMS-NOTE-2005-012, 2005.
- [70] G. Aad *et al.*, Report No. ATL-DAQ-PUB-2016-001, 2016.
- [71] D. E. Groom, N. V. Mokhov, and S. I. Striganov, *At. Data Nucl. Data Tables* **78**, 183 (2001).
- [72] T. Sjostrand, S. Mrenna, and P. Z. Skands, *J. High Energy Phys.* **05** (2006) 026.
- [73] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, and P. Z. Skands, *Comput. Phys. Commun.* **191**, 159 (2015).

- [74] C. Patrignani *et al.* (Particle Data Group), *Chin. Phys. C* **40**, 100001 (2016).
- [75] S. Agostinelli *et al.* (GEANT4), *Nucl. Instrum. Methods Phys. Res., Sect. A* **506**, 250 (2003).
- [76] J. Allison *et al.*, *IEEE Trans. Nucl. Sci.* **53**, 270 (2006).
- [77] M. Asai, A. Dotti, M. Verderi, and D. H. Wright (Geant4), *Ann. Nucl. Energy* **82**, 19 (2015).
- [78] ATLAS Collaboration, CERN Report No. ATLAS-CONF-2010-035, 2010.
- [79] G. Aad *et al.* (ATLAS), *Phys. Lett. B* **707**, 438 (2012).
- [80] M. Cacciari, M. Greco, and P. Nason, *J. High Energy Phys.* **05** (1998) 007.
- [81] M. Cacciari, S. Frixione, and P. Nason, *J. High Energy Phys.* **03** (2001) 006.
- [82] W. Lakin, E. Hughes, L. O'Neill, J. Otis, and L. Mandansky, *Phys. Lett.* **31B**, 677 (1970).
- [83] A. Gsponer, J. Hoffnagle, W. R. Molzon, J. Roehrig, V. L. Telegdi, B. Winstein, S. H. Aronson, G. J. Bock, D. Hedin, and G. B. Thomson, *Phys. Rev. Lett.* **42**, 9 (1979).
- [84] I. Cohen *et al.*, *Phys. Rev. Lett.* **40**, 1614 (1978).
- [85] R. G. Hicks *et al.*, *Phys. Rev. Lett.* **45**, 765 (1980).
- [86] M. Arneodo *et al.*, *Phys. Lett. B* **145**, 156 (1984).
- [87] R. Ent, *Nucl. Phys. A* **631**, 385 (1998).
- [88] M. Anelli *et al.* (SHiP), [arXiv:1504.04956](https://arxiv.org/abs/1504.04956).
- [89] A. Bulmahn and M. H. Reno, *Phys. Rev. D* **79**, 053008 (2009).
- [90] J. A. Formaggio and G. P. Zeller, *Rev. Mod. Phys.* **84**, 1307 (2012).
- [91] K. Daum, W. Rhode, P. Bareyre, R. Barloutaud, G. Chardin, B. Degrange, J. Ernwein, B. Kuznik, H. Meyer, L. Mosca, L. Moscoso, O. Perdereau, M. Schubnell, S. Tisserant, and Y. Wei, *Z. Phys. C* **66**, 417 (1995).
- [92] A. Fradette and M. Pospelov, *Phys. Rev. D* **96**, 075033 (2017).
- [93] F. Bezrukov and D. Gorbunov, *J. High Energy Phys.* **05** (2010) 010.
- [94] J. F. Gunion, H. E. Haber, G. L. Kane, and S. Dawson, *Front. Phys.* **80**, 1 (2000).
- [95] F. Bergsma *et al.*, *Phys. Lett.* **157B**, 458 (1985).
- [96] J. A. Evans, [arXiv:1708.08503](https://arxiv.org/abs/1708.08503).
- [97] G. Lanfranchi (SHiP Collaboration), <https://cds.cern.ch/record/2243034>.
- [98] B. Batell, M. Pospelov, and A. Ritz, *Phys. Rev. D* **83**, 054005 (2011).
- [99] G. Krnjaic, *Phys. Rev. D* **94**, 073009 (2016).
- [100] B2TiP (Belle II), The Belle II Physics Book, Report (to be published).
- [101] P. Ilten, J. Thaler, M. Williams, and W. Xue, *Phys. Rev. D* **92**, 115017 (2015).
- [102] P. Ilten, Y. Soreq, J. Thaler, M. Williams, and W. Xue, *Phys. Rev. Lett.* **116**, 251803 (2016).
- [103] P. Meade, M. Papucci, and T. Volansky, *J. High Energy Phys.* **12** (2009) 052.
- [104] CMS Collaboration, CERN Report No. CMS-NOTE-13-002, 2013.
- [105] A. Coccaro, D. Curtin, H. J. Lubatti, H. Russell, and J. Shelton, *Phys. Rev. D* **94**, 113003 (2016).
- [106] G. Aad *et al.* (ATLAS), *Phys. Rev. D* **92**, 012010 (2015).
- [107] ATLAS Collaboration, CERN Report No. ATLAS-CONF-2016-042, 2016.