



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Dark photons from charm mesons at LHCb

Philip Ilten, Jesse Thaler, Mike Williams, and Wei Xue

Phys. Rev. D **92**, 115017 — Published 21 December 2015

DOI: [10.1103/PhysRevD.92.115017](https://doi.org/10.1103/PhysRevD.92.115017)

Dark photons from charm mesons at LHCb

Philip Ilten,^{1,*} Jesse Thaler,^{2,†} Mike Williams,^{1,‡} and Wei Xue^{2,§}

¹Laboratory for Nuclear Science, Massachusetts Institute of Technology, Cambridge, MA 02139, U.S.A.

²Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, U.S.A.

We propose a search for dark photons A' at the LHCb experiment using the charm meson decay $D^{*0}(2007)^0 \rightarrow D^0 A'$. At nominal luminosity, $D^{*0} \rightarrow D^0 \gamma$ decays will be produced at about 700 kHz within the LHCb acceptance, yielding over 5 trillion such decays during Run 3 of the LHC. Replacing the photon with a kinetically-mixed dark photon, LHCb is then sensitive to dark photons that decay as $A' \rightarrow e^+ e^-$. We pursue two search strategies in this paper. The displaced strategy takes advantage of the large Lorentz boost of the dark photon and the excellent vertex resolution of LHCb, yielding a nearly background-free search when the A' decay vertex is significantly displaced from the proton-proton primary vertex. The resonant strategy takes advantage of the large event rate for $D^{*0} \rightarrow D^0 A'$ and the excellent invariant mass resolution of LHCb, yielding a background-limited search that nevertheless covers a significant portion of the A' parameter space. Both search strategies rely on the planned upgrade to a triggerless-readout system at LHCb in Run 3, which will permit identification of low-momentum electron-positron pairs online during data taking. For dark photon masses below about 100 MeV, LHCb can explore nearly all of the dark photon parameter space between existing prompt- A' and beam-dump limits.

I. INTRODUCTION

Rare decays of mesons are a powerful probe of physics beyond the standard model (SM). Precise measurements of branching fractions and decay kinematics indirectly constrain extensions of the SM by bounding symmetry-violating or higher-dimensional operators. More directly, non-SM particles could be produced in meson decays when kinematically allowed, and depending on their lifetimes, these particles could yield striking signals with displaced vertices. A well-motivated hypothetical particle is the dark photon A' which inherits a small coupling to the SM via kinetic mixing with the ordinary photon γ [1–6]. Indeed, some of the most stringent constraints on the properties of dark photons come from rare decays of mesons, including $\pi^0 \rightarrow \gamma A'$ [7–14], $\eta/\eta' \rightarrow \gamma A'$ [15, 16], and $\phi \rightarrow \eta A'$ [17, 18].

The minimal dark photon scenario involves a single broken $U(1)$ gauge symmetry, along with mixing between the A' and SM hypercharge fields via the operator $F'_{\mu\nu} B^{\mu\nu}$. After electroweak symmetry breaking and diagonalizing the gauge boson kinetic terms, the dark photon gains a suppressed coupling to the ordinary electromagnetic current J_{EM}^μ , where the relevant terms in the Lagrangian are

$$\mathcal{L} \supset -\frac{1}{4} F'_{\mu\nu} F'^{\mu\nu} + \frac{1}{2} m_{A'}^2 A'_\mu A'^\mu + \epsilon e A'_\mu J_{\text{EM}}^\mu. \quad (1)$$

This minimal scenario has two free parameters: the dark photon mass $m_{A'}$ and the kinetic-mixing parameter ϵ (often reported in terms of ϵ^2). The constraints placed on dark photons in the $m_{A'}-\epsilon^2$ plane are shown in Fig. 1

for $2m_e < m_{A'} < 5$ GeV, assuming that the A' dominantly decays into visible SM states (see Ref. [19] for a review).¹ For $\epsilon^2 \gtrsim 10^{-6}$, the most stringent bounds come from searches for prompt A' decays at collider and fixed-target experiments [14, 30–32]. As ϵ decreases, the A' lifetime increases, while as $m_{A'}$ decreases, the lifetime and Lorentz boost factor both increase. Therefore, the constraints obtained from beam-dump experiments exclude wedge-shaped regions in the $m_{A'}-\epsilon^2$ plane [6, 15, 16, 33–43]. Also shown in Fig. 1 are electron $g-2$ bounds [44–46]², the preferred region to explain the muon $g-2$ anomaly [47], and supernova bounds from cooling [48] and emissions [49]. Anticipated limits from other planned experiments are shown later in Fig. 9.

In this paper, we propose a search for dark photons through the rare charm meson decay

$$D^{*0} \rightarrow D^0 A', \quad A' \rightarrow e^+ e^-, \quad (2)$$

at the LHCb experiment during Run 3 of the LHC (scheduled for 2021–23).³ The goal of this search is to explore the region between the prompt- A' and beam-dump limits for the range $m_{A'} \in [10, 100]$ MeV, which roughly includes $\epsilon^2 \in [10^{-10}, 10^{-6}]$. Reaching such small values of ϵ^2 is only possible for decays where the yield of the corresponding SM process (i.e. replacing A' with γ) is at least $\mathcal{O}(10^{10})$. Within the LHCb acceptance, over five trillion $D^{*0} \rightarrow D^0 \gamma$ decays will be produced in proton-proton (pp) collisions at 14 TeV during Run 3, making this decay channel a suitable choice.

* philten@cern.ch

† jthaler@mit.edu

‡ mwill@mit.edu

§ weixue@mit.edu

¹ There are also interesting searches where the dark photon decays invisibly to dark matter [20–29].

² Since we follow the analysis in Ref. [19], we obtain more conservative bounds from $(g-2)_e$ than shown in Ref. [46].

³ Throughout this paper, $D^{*0} \equiv D^{*0}(2007)^0$ and the inclusion of charge-conjugate processes is implied.

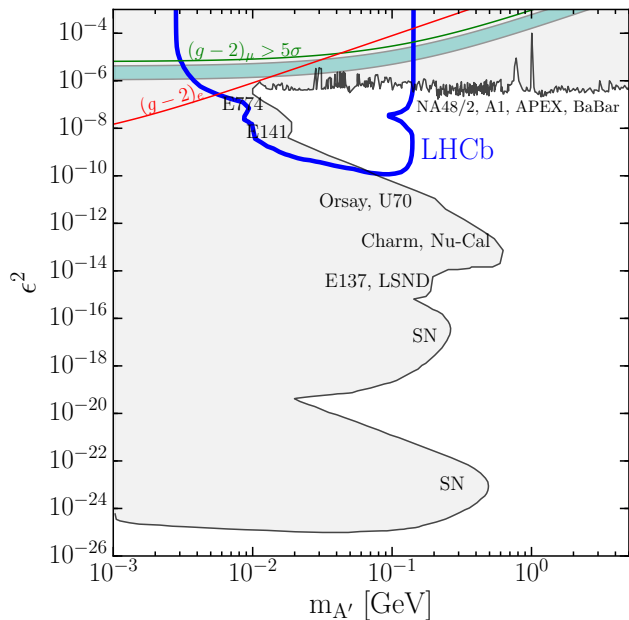


FIG. 1. Current bounds on dark photons with visible decays to SM states, adapted and updated from Ref. [19]. The upper bounds are from prompt- A' searches while the wedge-shaped bounds are from beam-dump searches and supernova considerations. The LHCb search region in Fig. 2 covers most of the gap between these bounds for $m_{A'} \lesssim 100$ MeV, with a reach extending to $m_{A'} \lesssim 140$ MeV. Anticipated limits from other planned experiments are shown in Fig. 9.

The range of $m_{A'}$ values that is in principle accessible in this search is $m_{A'} \in [2m_e, \Delta m_D]$, where [50]

$$\Delta m_D \equiv m_{D^{*0}} - m_{D^0} = 142.12 \pm 0.07 \text{ MeV}. \quad (3)$$

The proximity of Δm_D to m_{π^0} leads to phase-space suppression of the decay $D^{*0} \rightarrow D^0 \pi^0$, which results in a sizable branching fraction of about 38% for the decay $D^{*0} \rightarrow D^0 \gamma$.⁴ The small value of Δm_D , however, also leads to typical electron momenta of $\mathcal{O}(\text{GeV})$ within the LHCb acceptance. Therefore, the planned upgrade to a triggerless-readout system employing real-time calibration at LHCb in Run 3 [51]—which will permit identification of relatively low-momentum e^+e^- pairs online during data taking—will be crucial for carrying out this search.

To cover the desired dark photon parameter space, we employ two different search strategies, shown in Fig. 2. The *displaced search*, relevant at smaller values of ϵ^2 , looks for an $A' \rightarrow e^+e^-$ decay vertex that is significantly displaced from the pp collision. This search benefits from the sizable Lorentz boost factor of the produced dark

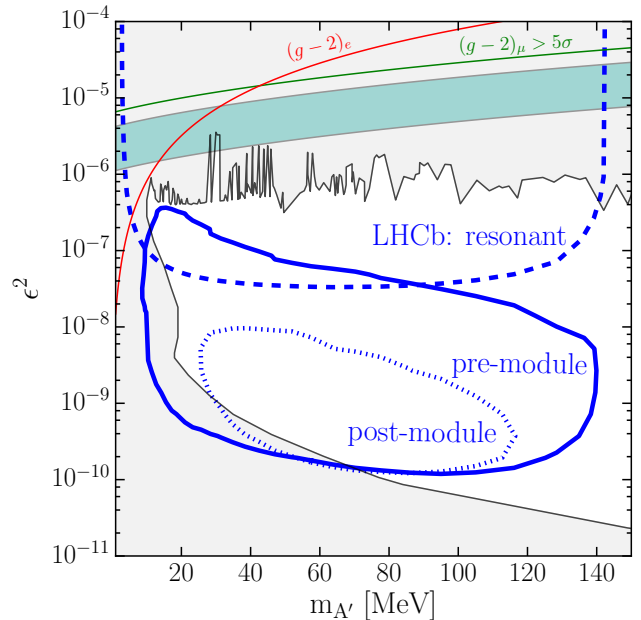


FIG. 2. Potential bounds from LHCb after Run 3, for both the displaced (pre-module, solid blue) and resonant (dashed blue) searches. Also shown is an alternative displaced search strategy (post-module, dotted blue) that looks for A' vertices downstream of the first tracking module.

photons and the excellent vertex resolution of LHCb. Our main displaced search looks for A' decays within the beam vacuum upstream of the first tracking module (i.e. pre-module), where the dominant background comes from misreconstructed prompt $D^{*0} \rightarrow D^0 e^+e^-$ events.⁵ Because the A' gains a transverse momentum kick from pp collisions, the A' flight trajectory intersects the LHCb detector, making it possible to identify displaced e^+e^- pairs with smaller opening angles than the HPS experiment [52]. We also present an alternative displaced search for A' decays downstream of the first tracking module (i.e. post-module), where the dominant background comes from $D^{*0} \rightarrow D^0 \gamma$ events with $\gamma \rightarrow e^+e^-$ conversion within the LHCb material.

The *resonant search*, relevant at larger values of ϵ^2 , looks for an $A' \rightarrow e^+e^-$ resonance peak over the continuum SM background. This search benefits from the large yield of $D^{*0} \rightarrow D^0 A'$ decays during LHC Run 3, which is larger than the A' yield in fixed-target experiments like MAMI/A1 [30, 31] and APEX [32]. Furthermore, the narrow width of the D^{*0} meson, which is less than the detector invariant-mass resolution, provides kinematical constraints that can be used to improve the resolution on $m_{e^+e^-}$. This resonant search can also be employed for non-minimal dark photon scenarios

⁴ This explains why we choose the decay $D^{*0} \rightarrow D^0 A'$ instead of $D^*(2010)^+ \rightarrow D^+ A'$, since the corresponding branching fraction $D^{*+} \rightarrow D^+ \gamma$ is only 1.6%.

⁵ We thank Natalia Toro for extensive discussions regarding this background.

where the A' might also decay invisibly into dark matter, shortening the A' lifetime. In that case, the anticipated limits in Fig. 2 would roughly apply to the combination $\epsilon^2 \times \mathcal{B}(A' \rightarrow e^+e^-)$.

The remainder of this paper is organized as follows. In Sec. II, we estimate the A' signal and SM background cross sections, extracting the D^{*0} production rate and D^0 decay modes from an event generator and estimating the D^{*0} decay rates using a simple operator analysis. In Sec. III, we describe the LHCb detector and charged-particle tracking, provide the selection requirements applied to D^0 and D^{*0} meson candidates, and derive the A' mass resolution. We present the pre-module displaced A' search in Sec. IV, a post-module variant in Sec. V, and the resonant A' search in Sec. VI. Possible improvements are outlined in Sec. VII and a comparison to other experiments (especially HPS) is given in Sec. VIII. We summarize in Sec. IX and discuss how the LHCb dark photon search strategy might be extended above the Δm_D threshold.

II. SIGNAL AND BACKGROUND RATES

Dark photon production in D^{*0} meson decays proceeds mainly via

$$D^{*0} \rightarrow D^0 A', \quad D^{*0} \rightarrow D^0 \pi^0 (\gamma A'), \quad (4)$$

though the low-energy photon in the latter decay is unlikely to be detected at LHCb. Here and throughout, we use the notation $X(YZ)$ to mean $X \rightarrow YZ$ in a subsequent decay. Because $\Delta m_D < 2m_\mu$, $A' \rightarrow e^+e^-$ is the only relevant visible decay channel.

The dominant backgrounds to the pre-module displaced A' search (Sec. IV) are $D^{*0} \rightarrow D^0 e^+e^-$ and $D^{*0} \rightarrow D^0 \pi^0 (\gamma e^+e^-)$, where the e^+e^- pair is misreconstructed as being displaced due to a hard electron scatter in material. These backgrounds can be highly suppressed by requiring that the e^+e^- kinematics are consistent with a displaced A' vertex occurring in the proper decay plane. The dominant background to the post-module displaced A' search (Sec. V) is $D^{*0} \rightarrow D^0 \gamma$, where the γ converts into an e^+e^- pair during interactions with the detector material. This background can be highly suppressed by requiring that the e^+e^- vertex position is not consistent with the location of any detector material. The dominant backgrounds to the resonant search (Sec. VI) are again $D^{*0} \rightarrow D^0 e^+e^-$ and $D^{*0} \rightarrow D^0 \pi^0 (\gamma e^+e^-)$, where the A' has been replaced by an off-shell γ^* . The first background is irreducible, making the resolution on $m_{e^+e^-}$ the driving factor in the resonant search reach.

A. D^{*0} Meson Production

We simulate D^{*0} production in pp collisions at a center-of-mass energy of 14 TeV using PYTHIA 8.201 [53]

with the default settings. Since a large fraction of charm quarks are produced from gluon splitting and since we need to model forward physics at small transverse momentum p_T , we run all soft QCD processes in PYTHIA (i.e. `SoftQCD:all = on`). While the D^{*0} production cross section is not yet known at 14 TeV, the result obtained using PYTHIA for the inclusive $pp \rightarrow D^{*+}$ production cross section at 7 TeV agrees with the measured value by LHCb [54] to within about 5%.⁶ Since PYTHIA does not record the spin of the D^{*0} mesons, they are treated as unpolarized in this analysis.

To define the fiducial region, we require the D^0 meson to satisfy the following transverse momentum and pseudorapidity requirements:

$$p_T(D^0) > 1 \text{ GeV}, \quad 2 < \eta(D^0) < 5. \quad (5)$$

Note that this requirement is placed on the D^0 meson, not on the D^{*0} , to suppress backgrounds to the D^0 component of the signal. The D^{*0} production cross section within this fiducial region is

$$\sigma(pp \rightarrow D^{*0} \rightarrow D_{\text{fid}}^0) = 0.95 \text{ mb}, \quad (6)$$

excluding secondary production of D^{*0} mesons from b -hadron decays. It may be possible to make use of some secondary decays; in this analysis, however, we require that the A' originates from the pp collision to suppress backgrounds (see Sec. IV A).

The nominal instantaneous luminosity expected at LHCb during Run 3 is 2 nb^{-1} per second [51], which will produce D^{*0} mesons at a rate of almost 2 MHz (equivalently, $D^{*0} \rightarrow D^0 \gamma$ at 0.7 MHz). Assuming an integrated luminosity of 15 fb^{-1} in Run 3,⁷ this results in an estimated yield of 14 trillion D^{*0} mesons produced within this fiducial region, or

$$N(D^{*0} \rightarrow D^0 \gamma) = 5.4 \times 10^{12}, \quad (7)$$

which we use as the baseline for our estimated reach.

B. D^{*0} Meson Decays

The D^{*0} meson is an $I(J^P) = \frac{1}{2}(1^-)$ state with a mass of $2006.96 \pm 0.10 \text{ MeV}$ and a width less than 2.1 MeV. It decays promptly mainly into two final states with branching fractions of

$$\mathcal{B}(D^{*0} \rightarrow D^0 \pi^0) = (61.9 \pm 2.9)\%, \quad (8)$$

$$\mathcal{B}(D^{*0} \rightarrow D^0 \gamma) = (38.1 \pm 2.9)\%, \quad (9)$$

⁶ During the final preparation of this article, LHCb presented the first prompt charm cross section measurement at 13 TeV [55]. Based on this result, we estimate that the relevant cross section for determining the dark photon reach should be about 20% higher than the one used in this paper.

⁷ The length of Run 3 is scheduled to be about the same as Run 1. LHCb collected a total of 3 fb^{-1} in Run 1. The instantaneous luminosity will be five times higher in Run 3. Therefore, assuming the LHC performance is the same (including the slow ramp up), this gives an estimate of 15 fb^{-1} in Run 3.

where the D^0 meson is a $\frac{1}{2}(0^-)$ state [50]. As mentioned above, $D^{*0} \rightarrow D^0 e^+ e^-$ is the dominant background to the pre-module displaced search as well as to the resonant search. To our knowledge, this branching fraction has not yet been measured; therefore, we will estimate the rate for this decay using an operator analysis. This same approach is used to determine the $D^{*0} \rightarrow D^0 A'$ rate.

To calculate these $D^{*0} \rightarrow D^0$ transition amplitudes, we must first determine the $\langle D^{*0} | J_{\text{EM}}^\mu | D^0 \rangle$ matrix element. By parity, time reversal, and Lorentz invariance, this transition dipole matrix element can be written in the form

$$\langle D^{*0} | J_{\text{EM}}^\mu | D^0 \rangle = \mu_{\text{eff}}(k^2) \epsilon^{\mu\alpha\beta\lambda} v_\alpha k_\beta \epsilon_\lambda, \quad (10)$$

where v_α is the four velocity of the D^{*0} meson, k_β is the momentum flowing out of the current, and ϵ_λ is the polarization of the D^{*0} meson. Here, μ_{eff} is a k -dependent effective dipole moment, whose value could be determined using a simple quark model (see, e.g., Ref. [56]) or using a more sophisticated treatment with heavy meson chiral perturbation theory (see, e.g., Ref. [57]). For our purposes, we simply need to treat μ_{eff} as being roughly constant over the range $k^2 \in [0, \Delta m_D^2]$, which is a reasonable approximation given that $\Delta m_D^2 < \Lambda_{\text{QCD}}^2$. (Indeed, this relation is always satisfied in the heavy charm quark limit, where $\Delta m_D \propto \Lambda_{\text{QCD}}^2/m_c$.) The precise value of μ_{eff} is irrelevant for our analysis since it cancels out when taking ratios of partial widths.

Using Eq. (10), we estimate the decay rate for $D^{*0} \rightarrow D^0 \gamma$ within the SM and in the $\Delta m_D \ll m_D$ limit to be

$$\Gamma(D^{*0} \rightarrow D^0 \gamma) = \frac{\alpha_{\text{EM}}}{3} \mu_{\text{eff}}^2 \Delta m_D^3, \quad (11)$$

where $\alpha_{\text{EM}} = e^2/4\pi$. To calculate the $D^{*0} \rightarrow D^0 e^+ e^-$ decay rate, the off-shell photon propagator must be included. In the $m_e = 0$ limit, the amplitude for this process is

$$|\mathcal{M}_{D^{*0} \rightarrow D^0 e^+ e^-}|^2 = -\frac{2e^4 \mu_{\text{eff}}^2}{3} \left[1 - \frac{(k_1 \cdot v)^2 + (k_2 \cdot v)^2}{k_1 \cdot k_2} \right], \quad (12)$$

where k_1 and k_2 are the electron and positron momenta. The ratio of partial widths is determined numerically to be

$$\frac{\Gamma(D^{*0} \rightarrow D^0 e^+ e^-)}{\Gamma(D^{*0} \rightarrow D^0 \gamma)} = 6.4 \times 10^{-3}. \quad (13)$$

Since the dark photon also couples to J_{EM}^μ , we use Eq. (10) to calculate the $D^{*0} \rightarrow D^0 A'$ decay rate. The ratio of partial widths is

$$\frac{\Gamma(D^{*0} \rightarrow D^0 A')}{\Gamma(D^{*0} \rightarrow D^0 \gamma)} = \epsilon^2 \left(1 - \frac{m_{A'}^2}{\Delta m_D^2} \right)^{3/2}, \quad (14)$$

where we assume $m_{A'}, \Delta m_D \ll m_D$. This expression has the expected kinetic-mixing and phase-space suppressions. Since the D^{*0} meson is treated as unpolarized in PYTHIA, we ignore spin correlations in the subsequent $A' \rightarrow e^+ e^-$ decay.⁸

C. Rare π^0 Decays

To determine the $D^{*0} \rightarrow D^0 \pi^0(\gamma A')$ decay rate in Eq. (4), we start by estimating the rate of the decay $\pi^0 \rightarrow \gamma A'$ using the SM effective Lagrangian

$$\mathcal{L} = \frac{\alpha_{\text{EM}}}{2\pi f_\pi} \pi^0 \epsilon^{\mu\nu\rho\sigma} F_{\mu\nu} F_{\rho\sigma}, \quad (15)$$

where f_π is the pion decay constant and the pion form factor is ignored. The dark photon is accounted for by making the replacement

$$F_{\mu\nu} \rightarrow F_{\mu\nu} + \epsilon F'_{\mu\nu}, \quad (16)$$

which leads to the ratio of partial widths

$$\frac{\Gamma(\pi^0 \rightarrow \gamma A')}{\Gamma(\pi^0 \rightarrow \gamma \gamma)} = 2\epsilon^2 \left(\frac{m_\pi^2 - m_{A'}^2}{m_\pi^2} \right)^3. \quad (17)$$

The same effective Lagrangian can also be used for the SM decay $\pi^0 \rightarrow \gamma e^+ e^-$. The amplitude is

$$|\mathcal{M}_{\pi^0 \rightarrow \gamma e^+ e^-}|^2 = \frac{4\alpha_{\text{EM}}^3}{\pi f_\pi^2 m_\pi^2} \left(m_{\pi^0}^4 + 2m_{\gamma e^-}^4 + m_{e^+ e^-}^4 + 2m_{\gamma e^-}^2 m_{e^+ e^-}^2 - 2m_{\pi^0}^2 (m_{\gamma e^-}^2 + m_{e^+ e^-}^2) \right). \quad (18)$$

The ratio of partial widths is obtained numerically to be

$$\frac{\Gamma(\pi^0 \rightarrow \gamma e^+ e^-)}{\Gamma(\pi^0 \rightarrow \gamma \gamma)} = 0.012, \quad (19)$$

which agrees with the nominal value for this ratio [50].

D. Dark Photon Decays

Assuming the only allowed decay mode is $A' \rightarrow e^+ e^-$, the total width of the A' is

$$\Gamma_{A'} = \frac{\epsilon^2 \alpha_{\text{EM}}}{3} m_{A'} \left(1 + 2 \frac{m_e^2}{m_{A'}^2} \right) \sqrt{1 - 4 \frac{m_e^2}{m_{A'}^2}}. \quad (20)$$

In the lab frame, the mean flight distance of the dark photon is approximately

$$\ell_{A'} \simeq 16 \text{ mm} \left(\frac{\gamma_{\text{boost}}}{10^2} \right) \left(\frac{10^{-8}}{\epsilon^2} \right) \left(\frac{50 \text{ MeV}}{m_{A'}} \right), \quad (21)$$

where γ_{boost} is the Lorentz boost factor. In Fig. 3 we show some example spectra of A' boost factors from simulated $D^{*0} \rightarrow D^0 A'$ decays, where both electrons are

⁸ As a technical note, to generate $D^{*0} \rightarrow D^0 A'$ events, we reweight

a sample of $D^{*0} \rightarrow D^0 \gamma$ events from PYTHIA. In particular, we implement $D^{*0} \rightarrow D^0 A'$ in the D^{*0} meson rest frame, boost to match the D^{*0} kinematics from PYTHIA, and then boost the D^0 decay products to account for the altered D^0 momentum. A similar strategy is employed for generating all other decays in our study.

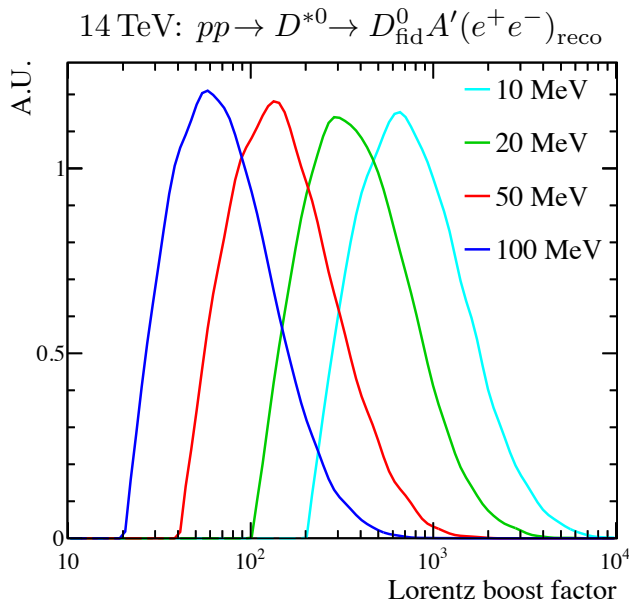


FIG. 3. Dark photon Lorentz boost factors for $m_{A'} = \{10, 20, 50, 100\}$ MeV. These factors are independent of ϵ^2 .

required to satisfy $|\vec{p}_e| \equiv p_e > 1$ GeV and $2 < \eta_e < 5$ so that they can be reconstructed by LHCb (see Sec. III A below). The A' inherits sizable momentum from the D^{*0} meson, leading to γ_{boost} factors that reach $\mathcal{O}(10^4)$. The corresponding spectra for the total (ℓ) and transverse (ℓ_T) flight distance of the A' are shown in Fig. 4. For $\epsilon^2 \lesssim 10^{-7}$ the displacement between the A' and pp -collision vertices is resolvable by LHCb.

III. BASELINE LHCb SELECTION

The LHCb detector is a single-arm spectrometer covering the forward region of $2 < \eta < 5$ [58, 59]. The detector, which was built to study the decays of hadrons containing b and c quarks, includes a high-precision tracking system capable of measuring charged-particle momenta with a resolution of about 0.5% in the region of interest for this search.⁹ The silicon-strip vertex locator (VELO) that surrounds the pp interaction region measures heavy-flavor hadron lifetimes with an uncertainty of about 50 fs [60]. Different types of particles are distinguished using information from two ring-imaging Cherenkov (RICH) detectors [61], a calorimeter system, and a system of muon chambers [62]. Both the momentum resolution and reconstruction efficiency are $\mathcal{O}(10)$ times worse for neutral particles than for charged ones. For this reason, the analysis strategy outlined below is based entirely on charged-particle information.

⁹ The precision of electron momentum measurements is limited by bremsstrahlung radiation; see Sec. III D.

A. Track Types

After exiting the VELO a distance of $\mathcal{O}(1$ m) from the pp collision, charged particles next traverse the first RICH detector (RICH1) before reaching a large-area silicon-strip detector located just upstream of a dipole magnet with a bending power of about 4 Tm [63]. Downstream of the magnet, there are three stations of silicon-strip detectors and straw drift tubes. All tracking systems will be upgraded for Run 3, though only the changes to the tracking systems upstream of the magnet are relevant here. The VELO has been redesigned to use pixels and is expected to have slightly better lifetime resolution and a lower material budget in Run 3 [64]. The tracking station just upstream of the magnet will also be replaced by a pixel detector and provide better coverage in η than the current detector [65]. This tracking station is known as the upstream tracker (UT).

In LHCb jargon, there are two types of tracks relevant for this search:

- LONG tracks that have hits in the VELO, the UT, and the stations downstream of the magnet. These tracks have excellent momentum resolution in both magnitude and direction.
- UP tracks that have hits in the VELO and the UT, but not in the stations downstream of the magnet. These tracks have excellent directional resolution obtained from the VELO. Since the curvature measurement is based only on the fringe field in which the UT operates, however, the uncertainty on the magnitude of the momentum is about 12% [65].

We also note that LHCb defines DOWN tracks which have hits in the UT and downstream of the magnet but no hits in the VELO. While DOWN tracks are not used in this search, they could be useful for other searches involving long-lived particles.

Charged particles may end up being reconstructed as UP tracks if they are swept out of the LHCb acceptance by the dipole magnetic field. This may occur if a particle is produced near the edge of the detector or if it is produced with low momentum. For simplicity, we take any charged particle with $2 < \eta < 5$ and $p > 3$ GeV to have 100% efficiency of being reconstructed as a LONG track. Any track that is not LONG, but satisfies $2 < \eta < 5$ and $p > 1$ GeV is assigned as an UP track. In reality, the reconstruction efficiency is not a step function—particles with $p < 3$ GeV may be reconstructed as LONG tracks, while particles with $p > 3$ GeV may produce UP tracks or not be reconstructed at all—but this simple choice reproduces well the overall tracking performance. The momentum resolution for each track type is derived in App. A.

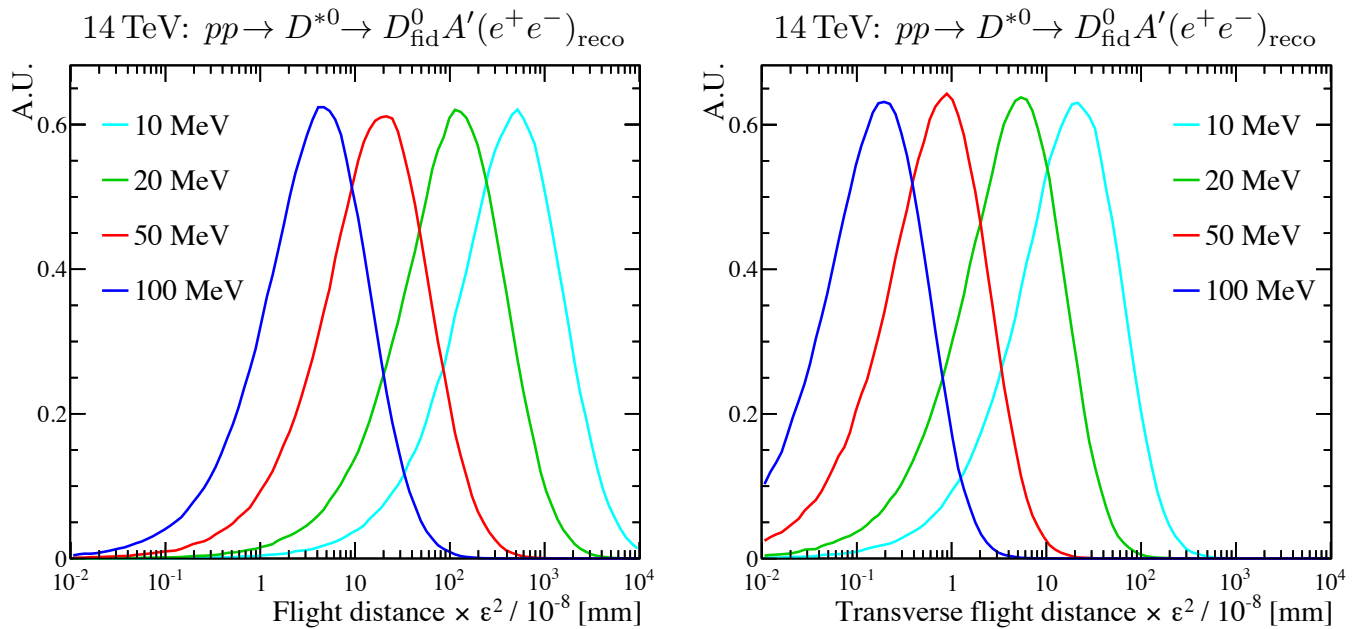


FIG. 4. Flight distance distributions for $m_{A'} = \{10, 20, 50, 100\}$ MeV showing (left) $\ell \times (\epsilon^2/10^{-8})$ and (right) $\ell_T \times (\epsilon^2/10^{-8})$.

B. D^0 Reconstruction

The D^0 meson momentum must be reconstructed for this search, since the kinematic constraints imposed by the D^{*0} mass will be used to suppress backgrounds and to improve the resolution on $m_{e^+e^-}$. We consider two categories of D^0 reconstruction.

- F-type: All of the D^0 children are charged particles so that the D^0 can be *fully reconstructed*. At least two of the decay products must be reconstructed as LONG tracks. This suppresses combinatorial backgrounds and provides excellent resolution on the location of the D^0 decay vertex and on the D^0 momentum \vec{p}_D . The remaining decay products are permitted to be reconstructed as either LONG or UP tracks.
- P-type: At least two of the D^0 children are reconstructed as LONG tracks so that there is excellent resolution on the location of the D^0 decay vertex (there may be UP tracks as well). Requiring significant D^0 flight distance then permits reconstructing with good precision the direction of the D^0 momentum \hat{p}_D using the vector from the pp collision to the D^0 decay vertex. For the case where the invariant mass of the missing particle(s) is known, $|\vec{p}_D|$ can be solved for as discussed below. In this way, the D^0 is *pseudo-fully reconstructed*.

The F-type decays considered in this search are given in Table I. Each is of the form $D^0 \rightarrow hh$ or $D^0 \rightarrow hhhh$, where $h = K^\pm$ or π^\pm . We do not consider doubly Cabibbo-suppressed decays (e.g. $D^0 \rightarrow K^+\pi^-$) since

Decay	\mathcal{B}	$\mathcal{B} \times \text{eff}_D^F$	$\mathcal{B} \times \text{eff}_D^P$
$D^0 \rightarrow \{K^-\pi^+, KK, \pi\pi\}$	4.4%	2.5%	–
$D^0 \rightarrow \{K^-\pi^+, 2K^0\pi^+, 4\pi^+\}$	9.1%	4.5%	1.0%
$D^0 \rightarrow K\ell(\nu)$	6.8%	–	2.0%
$D^0 \rightarrow K\pi(\pi^0)_{[0, m_{K^0}]}$	22.0%	–	6.6%
$D^0 \rightarrow KK(K^0)_{[\text{all}]}$	1.5%	–	0.5%
$D^0 \rightarrow K3\pi(\pi^0)_{[0, m_{K^0}]}$	8.5%	–	1.4%
Total		7.0%	11.5%

TABLE I. Decays of D^0 mesons used in this search. The branching fraction \mathcal{B} and efficiency-corrected branching fraction are given for each decay, for both the F-type (fully reconstructed) and P-type (pseudo-fully reconstructed) selections. The notation (x) denotes that x is not reconstructed. Entries with an $[a, b]$ subscript count any decay where the invariant mass of the non-reconstructed system satisfies $a \leq m_{\text{mis}} \leq b$ as signal.

they have small branching fractions and can be difficult experimentally to separate from the related Cabibbo-favored decays. LHCb has already published results using most of the F-type decays listed here (see, e.g., Refs. [66, 67]), and each decay is expected to have minimal combinatorial background contamination even with only loose selection criteria applied. Here we assume a baseline F-type D^0 selection efficiency of 90%. The total efficiency is then $\text{eff}_D^F \approx 50\%$, which is dominated by the requirement that all decay products are reconstructed by LHCb. As shown in App. A, the resolution on \vec{p}_D for F-type decays is excellent.

In P-type decays, we can use the measured flight direction to pseudo-fully reconstruct \vec{p}_D . The direction \hat{p}_D is a unit-normalized vector from the pp collision to the D^0 decay vertex. The magnitude is $|\vec{p}_D| = (\vec{p}_{\text{vis}} \cdot \hat{p}_D + \vec{p}_{\text{mis}} \cdot \hat{p}_D)$, where \vec{p}_{vis} is the reconstructed (visible) momentum and \vec{p}_{mis} is the non-reconstructed (missing) momentum. Balancing the momentum transverse to the direction of flight, requires $p_{\text{mis}}^\perp \equiv |\vec{p}_{\text{mis}} \times \hat{p}_D| = |\vec{p}_{\text{vis}} \times \hat{p}_D|$. Assuming that the invariant mass of the missing decay products is known, $|\vec{p}_D|$ can be solved for in the D^0 rest frame using conservation of energy and the known D^0 meson mass. Since p_{mis}^\perp is invariant under boosts along \hat{p}_D , $|\vec{p}_{\text{mis}} \cdot \hat{p}_D|$ in the D^0 rest frame is easily obtained. Finally, $\vec{p}_{\text{mis}} \cdot \hat{p}_D$ can be determined in the lab frame up to a two-fold ambiguity that arises because the sign of $\vec{p}_{\text{mis}} \cdot \hat{p}_D$ in the D^0 rest frame is not known. However, once the D^0 is combined with an $A' \rightarrow e^+e^-$ candidate to form a D^{*0} candidate, the vast majority of the time only the correct solution produces an invariant mass consistent with that of the D^{*0} meson. As described in App. A, we take the baseline selection efficiency for P-type decays to be 50%, since the D^0 flight distance must be large relative to the vertex resolution to obtain good resolution on \vec{p}_D .

The P-type decays considered in this search are given in Table I. We again do not consider doubly Cabibbo-suppressed decays. Other decays that are ignored include those where the missing mass cannot be reliably predicted, such as $D^0 \rightarrow \pi^- \ell^+ (X)$, which dominantly has $X = K^0 \nu$ as the missing system. Note that solving for $|\vec{p}_D|$ in P-type decays requires using the known missing mass as a constraint. That said, the resolution is only degraded slightly if the true missing mass differs from that used in the reconstruction by up to about $0.2m_{D^0}$. For example, when the visible part of the decay is $K^- \pi^+$, the most likely missing system is a single π^0 ; if the missing mass is taken to be m_{π^0} , but the actual decay is $D^0 \rightarrow K^- \pi^+ K^0$, the resolution obtained on $m_{e^+e^-}$ by applying the “wrong” kinematical constraints to the D^{*0} candidate is only worse by about 10%. In Table I, we list the missing mass ranges considered as signal for each P-type decay. Candidates where the missing mass falls outside of these windows are ignored in this analysis, since they have worse resolution and anyways make up a small fraction of the P-type decays. A derivation of the P-type D^0 resolution is given in App. A. The resolution on \vec{p}_D is about an order of magnitude worse in P-type than F-type decays; however, the $m_{A'}$ resolution after performing a mass-constrained fit is similar (as shown in Fig. 5 below).

C. D^{*0} Reconstruction

To reduce the background from unassociated $D^0 e^+e^-$ combinations, we require that the reconstructed mass difference

$$\Delta m_D^{\text{reco}} = m_{\text{reco}}(D^0 e^+e^-) - m_{\text{reco}}(D^0) \quad (22)$$

satisfies

$$-50 \text{ MeV} < \Delta m_D^{\text{reco}} - \Delta m_D < 20 \text{ MeV}. \quad (23)$$

The looser requirement is placed on the lower edge due to bremsstrahlung by the electrons. This mass requirement highly suppresses the decay $D^{*0} \rightarrow D^0 \pi^0 (\gamma e^+e^-)$ and its A' counterpart, except when $m_{e^+e^-}$ is large (see Sec. IV D below). The efficiency of this requirement is about $\text{eff}_{\Delta m_D} \approx 85\%$. Note that this cut can be tightened at the expense of signal efficiency if combinatorial backgrounds turn out to be problematic (see Sec. IV C below).

D. A' Reconstruction

The reconstructed electrons produced in $A' \rightarrow e^+e^-$ decays are a mixture of UP and LONG tracks. Only a few percent of the electrons have momenta large enough that equivalent-momenta non-electrons would be able to emit Cherenkov light in RICH1. Therefore, identification of the e^+ and e^- should be highly efficient with a low hadron-misidentification rate. Furthermore, the signature of a maximum-Cherenkov-angle ring in coincidence with a track should suppress the fake-track background which can be sizable at low momenta.

Bremsstrahlung radiation and multiple scattering of the electrons significantly affect the $m_{e^+e^-}$ resolution. We implement this numerically in our simulation following Refs. [50, 68] and using the Run 3 LHCb VELO [64], RICH1 [69], and UT [65] material budgets. Bremsstrahlung downstream of the magnet does not affect the momentum measurement and is ignored.

In Fig. 5, we show the resolution on $m_{e^+e^-}$ for several values of $m_{A'}$, where the A' candidates are constrained to originate from the pp collision. Bremsstrahlung creates large low-mass tails resulting in poor resolution on $m_{e^+e^-}$. Since the D^{*0} mass is known and its width is less than the detector resolution, though, we can correct the $m_{e^+e^-}$ distribution once we identify the D^{*0} candidate and apply the Δm_D^{reco} cut. As a heuristic, one can rescale the $m_{e^+e^-}$ value by a simple correction factor

$$m_{e^+e^-}^{\text{corr}} = m_{e^+e^-}^{\text{reco}} \left(2 - \frac{\Delta m_D^{\text{reco}}}{\Delta m_D} \right). \quad (24)$$

A more sophisticated approach involves performing a mass-constrained fit to enforce energy-momentum conservation and the known D^{*0} mass using the covariance matrices of all reconstructed particles. Using this fit, we find 10–20% improvement in $\sigma(m_{e^+e^-})$ relative to the simple correction given in Eq. (24). As shown in Fig. 5, the resolution on $m_{e^+e^-}$ after the applying the kinematic fit is 2–3 MeV using F-type D^0 candidates, and 2–5 MeV using P-type D^0 candidates.

The key difference between the pre-module displaced, post-module displaced, and resonant searches are the requirements placed on the A' flight distance. These are described in more detail in the subsequent sections.

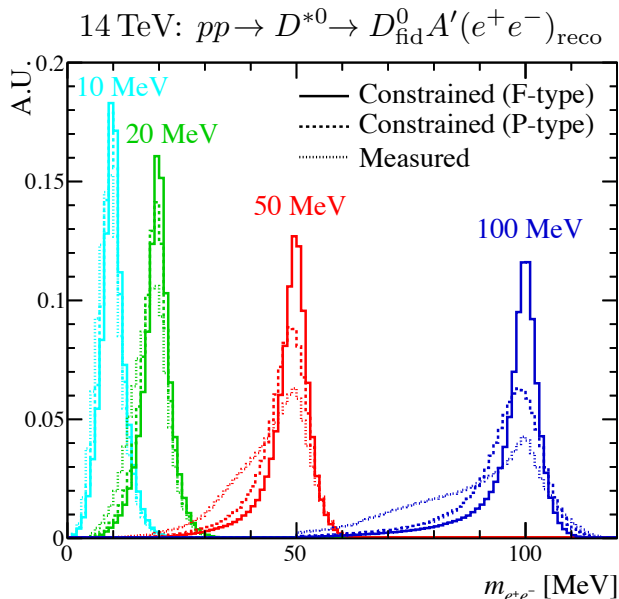


FIG. 5. Distribution of $m_{e^+e^-}$ with (solid, dashed) and without (dotted) incorporating the D^{*0} mass constraint for $m_{A'} = \{10, 20, 50, 100\}$ MeV. The solid curve shows better performance than the dashed one because F-type D^0 candidates have better momentum resolution than P-type ones.

IV. DISPLACED A' SEARCH (PRE-MODULE)

The A' typically has a large Lorentz boost factor, resulting in the A' decay vertex being significantly displaced from the pp collision for $\epsilon^2 \lesssim 10^{-7}$. The combined signature of a displaced D^0 decay vertex, a displaced $A' \rightarrow e^+e^-$ vertex, $m(D^0A')$ consistent with $m(D^{*0})$, and a consistent decay topology will result in a nearly background-free search. This pre-module displaced search is aimed at A' decay vertices that occur within the beam vacuum upstream of the first VELO module intersected by the A' trajectory.

A. Conversion and Misreconstruction Backgrounds

At LHCb, the first layer of material is the foil that separates the beam vacuum from the VELO vacuum. This foil is corrugated to accommodate the VELO modules, such that if the A' decays prior to the foil, it still effectively decays within the VELO tracking volume. The average transverse distance that the A' will travel before hitting a VELO module is 6 mm [64], which, because of the corrugated foil geometry, is roughly the average transverse flight distance to the foil as well.

To effectively eliminate backgrounds from $\gamma \rightarrow e^+e^-$ conversions in the foil, we require the A' decay vertex to be reconstructed upstream of the foil. Furthermore, each reconstructed electron must have an associated hit in the first relevant VELO module given the location of

the reconstructed A' decay vertex. These hits are required to have at least one vacant VELO pixel between them to avoid any charge-sharing issues, imposing an effective buffer distance between the A' decay vertex and the foil:

$$D \approx \frac{0.123 \text{ mm}}{\alpha_{e^+e^-}}, \quad (25)$$

where $\alpha_{e^+e^-}$ is the electron-positron opening angle. In reality, the VELO pixels in Run 3 will be $55 \times 55 \mu\text{m}^2$ squares; the definition of D is based on treating the pixels as circles with 0.123 mm being twice the effective diameter (the precise value used here has no impact on our search). The pre-module A' requirement can then be approximated by requiring the A' transverse flight distance to satisfy

$$\ell_T < 6 \text{ mm} - D_T, \quad D_T = D \sin \theta, \quad (26)$$

where θ gives the A' flight direction. To remove A' trajectories that first intersect the foil far from a module, we require $\eta_{A'} > 2.6$. We also impose $\eta_{A'} < 5$ to avoid possible contamination due to pp collisions that are not properly reconstructed.¹⁰

Having suppressed conversion backgrounds, the dominant background comes from prompt $D^{*0} \rightarrow D^0 e^+e^-$ events where the e^+e^- vertex is misreconstructed as being displaced because of multiple scattering of the electrons in the detector material. We estimate this background in a toy simulation of the Run 3 VELO, taking scattering angle distributions from a GEANT simulation which includes non-Gaussian Molière scattering tails.¹¹ Many of these fake A' vertices can be eliminated by requiring a consistent decay topology, in particular that the angle between $\vec{p}_{A'}$ and the vector formed from the pp collision to the A' decay vertex is consistent with zero, and the electrons travel within a consistent decay plane.

The remaining misreconstructed background events have a consistent topology, so a cut on transverse flight distance ℓ_T is required to ensure a significant displaced A' vertex. To avoid fake displaced vertices from one electron experiencing a large-angle scattering, we also require both the electron and positron to have a non-trivial impact parameter (IP) with respect to the pp collision. These requirements are summarized by

$$\ell_T > n \sigma_{\ell_T}, \quad \text{IP}_{e^\pm} > \frac{n}{2} \sigma_{\text{IP}}, \quad (27)$$

¹⁰ An A' candidate may be accidentally formed from a prompt e^+e^- pair produced in a pp collision if the event is not properly reconstructed. In particular, if a D^0 meson is produced in another pp collision upstream of that interaction point, the “displaced” A' would produce a consistent decay topology, albeit with $\eta_{A'} \rightarrow \infty$.

¹¹ It is likely that GEANT overestimates the probability for large-angle scatterings (see Ref. [70]). If so, our results are conservative, since these scattering tails effectively define the reach for the pre-module A' search.

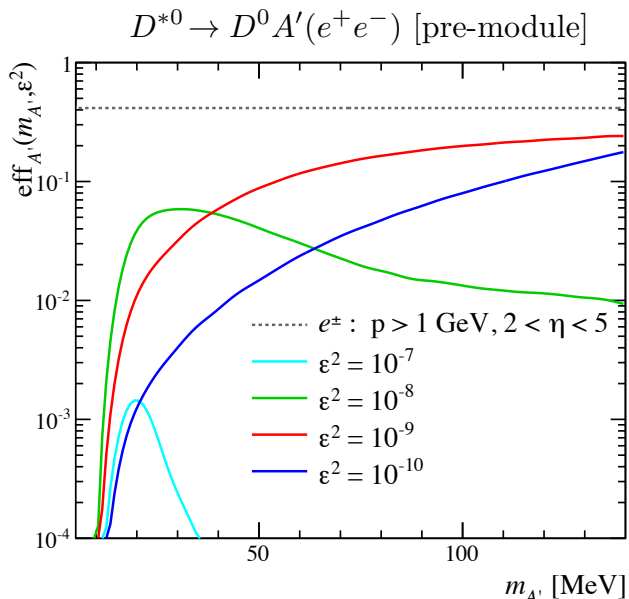


FIG. 6. Pre-module displaced search A' signal efficiency $\text{eff}_{A'}(m_{A'}, \epsilon^2)$ for several ϵ^2 values. The dashed line shows the efficiency of the track-reconstruction requirements placed on the electrons alone, while the solid lines show the total A' efficiency. This efficiency does not include the contribution from the D^0 or D^{*0} selection (i.e. $\text{eff}_D^{\text{F,P}}$ or $\text{eff}_{\Delta m_D}$).

where the value of n is adjusted to yield ≈ 1 background event in each A' mass window, with n ranging from 3 to 5 as a function of $m_{A'}$. The selection in Eq. (27) is meant to be simple and robust, and could certainly be optimized in a full analysis. See App. B for details on the ℓ_T and IP resolution.

B. Event Selection

Summarizing, the event selection for the pre-module displaced A' search is:

- F-type or P-type D^0 candidate;
- e^+ and e^- from LONG or UP tracks with hits in the first VELO module they intersect;
- reconstructed D^{*0} candidate from the D^0 , e^+ , and e^- ;
- reconstructed $A' \rightarrow e^+e^-$ satisfying the conversion veto ($\ell_T < 6 \text{ mm} - D_T$, $\eta_{A'} \in [2.6, 5]$);
- reconstructed $A' \rightarrow e^+e^-$ with significant displacement ($\ell_T > n \sigma_{\ell_T}$, $\text{IP}_{e^\pm} > \frac{n}{2} \sigma_{\text{IP}}$).

In Fig. 6, we show the resulting A' signal efficiency $\text{eff}_{A'}(m_{A'}, \epsilon^2)$. For $\epsilon^2 = 10^{-10}$, which is near the low end of the reach, the efficiency is limited by the efficiency of the conversion veto. As ϵ increases, the requirement of

a significant displacement ultimately limits the reach of the displaced search.

C. Additional Backgrounds

Beyond the misreconstructed $D^{*0} \rightarrow D^0 e^+ e^-$ background, a full accounting of the potential backgrounds for the displaced search is difficult since all SM processes with large rates are highly suppressed. Therefore, any additional backgrounds will be dominated by extremely rare processes or highly unlikely coincidences.

One possible source of backgrounds would be $B \rightarrow D^{*0}(D^0 e^+ e^-)X$ decays, since the resulting e^+e^- vertex is truly displaced. Such decays can be suppressed by making the following requirements: the D^0 and A' momenta must intersect the pp -collision point when traced upstream from their respective decay vertices; the D^{*0} decay vertex is consistent with the pp -collision vertex; and there are no additional tracks consistent with originating from the A' decay vertex. Furthermore, one could require that the A' decay vertex is downstream of the D^0 decay vertex, which would be efficient for the smaller ϵ values probed in this search. Therefore, we do not expect a significant amount of background coming from $B \rightarrow D^{*0}(D^0 e^+ e^-)X$ decays.

The decays of other long-lived mesons could also be sources of displaced e^+e^- vertices. Decays of charged pions and kaons that produce an e^+e^- pair are rare, though, and the probability for these particles to decay in the VELO is small. A more likely source is the decay $\pi^0 \rightarrow e^+e^-\gamma$, where the π^0 is produced in the decay of a long-lived meson. All of the other meson-decay products must be neutral, of course, otherwise the presence of additional charged particles consistent with originating from the e^+e^- vertex could be used as a veto. For example, the decays $K_S \rightarrow \pi^0 \pi^0$ and $D^0 \rightarrow K_S \pi^0$ occur with huge rates within the LHCb VELO. Such decays, however, are unlikely to result in the A' candidate momentum intersecting the pp -collision point or to occur in coincidence with a D^0 meson such that $m(D^0 A')$ is consistent with $m(D^{*0})$.

To see whether we could estimate displaced combinatorial backgrounds in Monte Carlo, we generated a sample of 30 million PYTHIA pp collisions at 14 TeV. We found that no combination of a true D^0 with two displaced tracks (not necessarily electrons, but assigned the electron mass) had an invariant mass within the D^{*0} mass window. In this simulated sample, there are only three candidates with $m(D^0 e^+ e^-) - m(D^0) < 500 \text{ MeV}$ using true electrons and none within 150 MeV of our D^{*0} mass window. However, we are anticipating 10^8 times more D^{*0} meson decays in the full LHCb data sample, so it appears that it is not feasible to use Monte Carlo to precisely estimate the displaced combinatorial background. This type of background will therefore need to be examined in data using the Δm_D^{reco} sidebands. If specific sources of combinatorial background are identified as

problematic, then the selection will need to be adjusted to remove them.

D. Contribution from Pion Decays

Thus far, we have ignored the channel $D^{*0} \rightarrow D^0\pi^0(\gamma A')$, which is another potential source of signal events. Decays of this type are highly suppressed by the Δm_D^{reco} requirement in Eq. (23), though, unless $m_{A'}$ is large. For most allowed $m_{A'}$ values, one can choose whether or not to include such decays in the analysis by adjusting the Δm_D^{reco} requirement. After removing the Δm_D^{reco} requirement, the expected yields of $D^{*0} \rightarrow D^0\pi^0(\gamma A')$ and $D^{*0} \rightarrow D^0 A'$ decays are comparable, but so is the expected background contamination from misreconstructed $D^{*0} \rightarrow D^0\pi^0(\gamma e^+e^-)$ and $D^{*0} \rightarrow D^0 e^+e^-$. We choose not to include this channel when estimating the reach below, but note that such decays may prove useful in a complete analysis.

If one does try to use the $D^{*0} \rightarrow D^0\pi^0(\gamma A')$ channel, then one should be aware of an important subtlety when incorporating $m_{e^+e^-}$ information. As described in Sec. III D, a kinematic fit can be used to improve the $m_{e^+e^-}$ resolution. For $D^{*0} \rightarrow D^0\pi^0(A'\gamma)$ decays, however, the missing γ is not accounted for when enforcing energy-momentum conservation. We find that this results in the $m_{A'}$ peak being shifted up in mass by about 20 MeV, with the resolution on $m_{e^+e^-}$ degraded by about a factor of two. This results in two peaks in the reconstructed $m_{A'}$ spectrum, coming from the $D^{*0} \rightarrow D^0 A'$ and $D^{*0} \rightarrow D^0\pi^0(\gamma A')$ channels. If the background level is low, then this second peak could be used to boost the significance of an A' signal. Indeed, one narrow peak with a second, wider peak shifted in mass by a fixed amount would be a striking signature. If the background level is high, then this wider peak would largely be absorbed into the background and have no impact on the signal significance.

Finally, photon conversions arising from $D^{*0} \rightarrow D^0\pi^0(\gamma\gamma)$ decays are also highly suppressed by the Δm_D^{reco} requirement, and in the absence of misreconstruction, can be eliminated by the pre-module requirement in Eq. (26). We expect such conversions to contribute less than those from $D^{*0} \rightarrow D^0\gamma$ decays.

E. Reach

The expected signal yield for the displaced search as a function of $m_{A'}$ and ϵ^2 is given by

$$\begin{aligned} S(m_{A'}, \epsilon^2) &= N(D^{*0} \rightarrow D^0\gamma) \frac{\Gamma(D^{*0} \rightarrow D^0 A')}{\Gamma(D^{*0} \rightarrow D^0\gamma)} \text{eff}_{\Delta m_D} \\ &\quad \times \left(\text{eff}_D^{\text{F}} + \text{eff}_D^{\text{P}} \right) \text{eff}_{A'}(m_{A'}, \epsilon^2) \\ &\simeq 85 \left(\frac{\epsilon^2}{10^{-10}} \right) \left(1 - \frac{m_{A'}^2}{\Delta m_D^2} \right)^{3/2} \text{eff}_{A'}(m_{A'}, \epsilon^2). \end{aligned} \quad (28)$$

As discussed above, we adjusted the requirement in Eq. (27) to ensure ≈ 1 background event in any given $m_{A'}$ window. Assuming that all relevant backgrounds have been accounted for, the reach would be set at 95% confidence level by requiring $S \gtrsim 3$. However, to allow for the possibility of extremely rare background sources, the reach is set by requiring $S \geq 5$. In this way, we account for either additional background candidates in the final data sample or for a lower A' selection efficiency due to the criteria required to suppress these additional backgrounds. The reach is shown in Fig. 2 assuming 15 fb^{-1} of data collected by LHCb in Run 3, which covers a significant part of the allowed parameter space for $m_{A'} \lesssim 100 \text{ MeV}$.

V. DISPLACED A' SEARCH (POST-MODULE)

In order to capture more A' signal events, one can effectively reverse the pre-module requirements in Eq. (26) and search for post-module A' decays. Here, the dominant background is $D^{*0} \rightarrow D^0\gamma$, where the on-shell γ converts into e^+e^- via interactions with the detector material. As we will see, this post-module search does not cover much additional A' parameter space compared to the pre-module search, but is important as a cross check of a possible A' discovery.

A. Misreconstruction and Conversion Backgrounds

The background considerations in the post-module case are reversed compared to the pre-module case in Sec. IV A. Here, the background from misreconstructed $D^{*0} \rightarrow D^0 e^+e^-$ events can be effectively eliminated by requiring *no* hits in the first VELO module intersected by the reconstructed electron trajectories.

The dominant background in the post-module search comes from $D^{*0} \rightarrow D^0\gamma$ with photon conversions. We simulate this background using the Run 3 LHCb VELO material as described in Ref. [64] with the Bethe-Heitler $m_{e^+e^-}$ spectrum as given in Ref. [71].¹² We start with electron tracks that each have at least three hits in the VELO. This imposes an effective A' transverse flight distance requirement of

$$\ell_T \in [6 \text{ mm}, 22 \text{ mm}]. \quad (29)$$

In reality, the electron hit requirement does not result in a step function for the A' efficiency. However, in the long-lifetime limit, this simple approximation produces

¹² It is vital that all A' searches use the Bethe-Heitler spectrum, rather than the one produced by GEANT. GEANT vastly underestimates the fraction of conversions that produce large $m_{e^+e^-}$ due to the usage of a less-CPU-intensive approximation of the Bethe-Heitler equation.

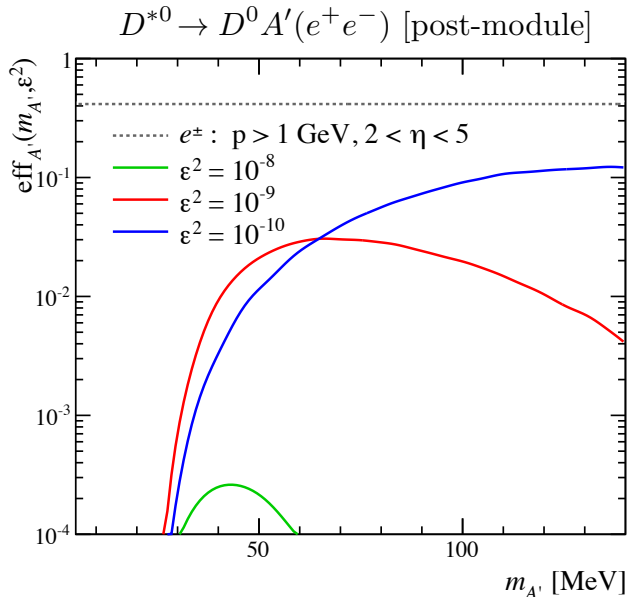


FIG. 7. Post-module displaced search A' signal efficiency $\text{eff}_{A'}(m_{A'}, \epsilon^2)$ for several ϵ^2 values. The dashed line shows the efficiency of the track-reconstruction requirements placed on the electrons alone, while the solid lines show the total A' efficiency. This efficiency does not include the contribution from the D^0 or D^{*0} selection (i.e. $\text{eff}_D^{\text{F,P}}$ or $\text{eff}_{\Delta m_D}$).

the same integrated efficiency. We then require the reconstructed A' vertex to be significantly displaced from the VELO material. This can be well-approximated by treating the VELO as a stack of modules located at longitudinal distances $z_i = i \cdot 25$ mm, where z is measured from the point where the A' has $\ell_T = 6$ mm (i.e. the average position where the A' trajectory crosses a VELO module). From a given A' decay vertex at a location z between modules i and $i + 1$, one requires

$$z - z_i > n\sigma_z, \quad z_{i+1} - z > D, \quad \text{IP}_{i,e^\pm} > \frac{n}{2}\sigma_{\text{IP}_i}, \quad (30)$$

where D is the same buffer distance in Eq. (25) and IP_i is defined with respect to the location where the A' trajectory intersects the i -th module (see App. B for the corresponding resolutions). We also impose the same $\eta_{A'} \in [2.6, 5]$ requirement as in Sec. IV A. Using our simulation, we adjust n such that ≈ 1 $D^{*0} \rightarrow D^0\gamma, \gamma \rightarrow e^+e^-$ event will survive these criteria in each A' mass window.

B. Event Selection and Reach

Summarizing, the event selection for the post-module displaced A' search is:

- F-type or P-type D^0 candidate;
- e^+ and e^- from LONG or UP tracks with no hits in the first VELO module they intersect;

- reconstructed D^{*0} candidate from the D^0, e^+ , and e^- ;
- reconstructed $A' \rightarrow e^+e^-$ satisfying the prompt veto ($\ell_T \in [6 \text{ mm}, 22 \text{ mm}]$, $\eta_{A'} \in [2.6, 5]$);
- reconstructed $A' \rightarrow e^+e^-$ with significant displacement from VELO modules ($\Delta z_i > n\sigma_z$, $\text{IP}_{i,e^\pm} > \frac{n}{2}\sigma_{\text{IP}_i}$, $\Delta z_{i+1} > D$).

In Fig. 7, we show the resulting A' signal efficiency $\text{eff}_{A'}(m_{A'}, \epsilon^2)$.

The formula for the reach for the post-module displaced search is identical to Eq. (28), albeit with a modified signal efficiency. We again set the reach using $S \geq 5$ to allow for rare, unaccounted for background sources. As shown in Fig. 2, the reach in the post-module search is not any better than the pre-module search. The reason is that for small enough ϵ^2 , the decay probability as a function of ℓ_T is constant, and there is comparable efficiency for the A' decay vertex to be upstream or downstream of the first VELO module. As ϵ^2 increases, the post-module $\ell_T > 6$ mm requirement becomes inefficient before the pre-module vertex requirements do, giving the pre-module search a better reach. Of course, a slightly better limit could be set by combining the pre-module and post-module requirements, but we do not do this here since the dominant background sources are different. We also note that if a discovery is made, the ability to confirm the presence of a signal in both displaced regions will provide a powerful systematic check.

C. Additional Backgrounds

The same additional backgrounds from Sec. IV C might affect the post-module search, with the exception of π^0 mesons produced in charm-meson decays. In addition, there is a potential background from improper reconstruction of photon conversion events.

One way to misreconstruct a γ conversion as an A' decay is if the γ converts in or just in front of a VELO module. Since the separation between the e^+ and e^- when they traverse the VELO module would be less than the hit resolution, only a single hit would be recorded. The positron track may be formed using this hit. The electron track would be missing a hit in this module and fail the selection requirements, unless an unassociated hit happens to occur close by. The hit occupancy expected in the VELO during Run 3 is about 0.08% in the innermost pixels, and less than 0.01% for pixels more than 10 mm from the beam line [64]. Since we require a full vacant pixel between the first e^+ and e^- hits, there must be a large angle between the electron momentum and the vector \vec{h} formed from the unassociated hit to the e^- hit in the second module. Therefore, a consistent track will only be formed if the electron undergoes an unlikely scatter in the second module so that the hit in the third module is consistent with \vec{h} . We cannot reliably estimate

the probability of this coincidence in our toy simulation, but expect it to be at worst comparable to the remaining conversion backgrounds left after the requirements in Sec. V A, so it should not affect the predicted reach.

VI. RESONANT SEARCH

When ϵ^2 is large enough, the A' decay vertex is no longer significantly displaced from the pp collision. In this case, we have to rely on reconstructing an A' mass peak. This search is also relevant for non-minimal scenarios with a larger dark photon width, such as when the dark photon has an invisible decay to dark matter particles (see, e.g., Refs. [20–29]).

For simplicity, we perform our analysis still assuming the minimal $\Gamma_{A'}$ value in Eq. (20). This means that conversion backgrounds are important, so we apply the conversion veto from Sec. IV A. The dominant background becomes prompt $D^{*0} \rightarrow D^0 e^+ e^-$, which is irreducible, and the resonant reach is limited by the $m_{e^+e^-}$ resolution shown in Fig. 5. To ensure good $m_{e^+e^-}$ resolution, we further impose an A' opening angle cut $\alpha_{e^+e^-} > 3$ mrad.

In summary, the event selection criteria for the resonant search are

- F-type or P-type D^0 candidate;
- e^+ and e^- from LONG or UP tracks with hits in the first VELO module they intersect;
- reconstructed D^{*0} candidate from the D^0 , e^+ , and e^- ;
- reconstructed $A' \rightarrow e^+e^-$ satisfying conversion veto ($\ell_T < 6$ mm $- D_T$, $\eta_{A'} \in [2.6, 5]$);
- A' decay opening angle > 3 mrad to ensure good $m_{e^+e^-}$ resolution.

In Fig. 8, we show the A' signal efficiency $\text{eff}_{A'}(m_{A'}, \epsilon^2)$ for the resonant search. Apart from the electron track-reconstruction, the dominant source of inefficiency is the opening angle requirement.

Due to the large irreducible background level in this search, we assume that $D^{*0} \rightarrow D^0 \pi^0 (\gamma A')$ decays provide negligible additional sensitivity and therefore ignore them. Background contamination from $D^{*0} \rightarrow D^0 \pi^0 (\gamma e^+ e^-)$ decays is still included, but it is highly suppressed by the Δm_D^{reco} requirement except when $m_{e^+e^-}$ is large.

To determine the reach in the resonant search, we require $S/\sqrt{B} > 2$ to obtain a 95% confidence limit. The signal significance is given by

$$\frac{S}{\sqrt{B}} = \Gamma(D^{*0} \rightarrow D^0 A') \sqrt{\frac{N(D^{*0} \rightarrow D^0 \gamma)}{\Gamma(D^{*0} \rightarrow D^0 \gamma) \Delta \Gamma}} \times \sqrt{\text{eff}_{\Delta m_D} \left(\text{eff}_D^{\text{F}} + \text{eff}_D^{\text{P}} \right) \text{eff}_{A'}(m_{A'}, \epsilon^2)}, \quad (31)$$

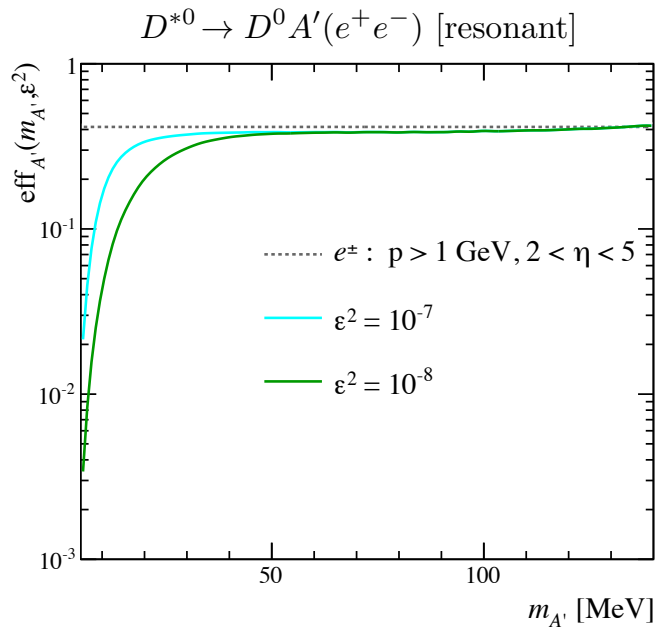


FIG. 8. Resonant search A' signal efficiency $\text{eff}_{A'}(m_{A'}, \epsilon^2)$. The dashed line shows the efficiency to reconstruct both electrons, while the solid line shows the total A' efficiency. The efficiency is nearly independent of ϵ for $\epsilon^2 > 10^{-7}$. This efficiency does not include the contribution from the D^0 or D^{*0} selection (i.e. $\text{eff}_D^{\text{F,P}}$ or $\text{eff}_{\Delta m_D}$).

where

$$\Delta \Gamma \equiv \int_{m_{A'} - \Delta m_{A'}}^{m_{A'} + \Delta m_{A'}} dm_{e^+e^-} \frac{d\Gamma(D^{*0} \rightarrow D^0 e^+ e^-)}{dm_{e^+e^-}} \quad (32)$$

and $\Delta m_{A'} = 2\sigma(m_{e^+e^-})$ is evaluated at $m_{A'}$.¹³ Here, we are assuming that $\text{eff}_{A'}(m_{A'}, \epsilon^2)$ is constant over $[m_{A'} - \Delta m_{A'}, m_{A'} + \Delta m_{A'}]$. The reach for the resonant search is shown in Fig. 2, and combined with the pre-module displaced search, it closes most of the available A' parameter space for $m_{A'} \lesssim 100$ MeV.

The triggerless-readout system, along with real-time data calibration, will make it possible to identify $D^{*0} \rightarrow D^0 A'$ candidates online during data taking. For the displaced searches, the number of candidates recorded for further analysis can be made as small as required since the dominant backgrounds are reducible. The resonant search, on the other hand, relies on looking for an $m_{e^+e^-}$ peak in $\mathcal{O}(10^9)$ $D^{*0} \rightarrow D^0 e^+ e^-$ decays. Recording billions of full events for such a search is simply not feasible. That said, LHCb is already commissioning in Run 2 partial event storage for use in studying high-rate processes [55]. This involves storing only the information

¹³ Because F-type and P-type D^0 mesons yield different A' mass resolutions, Eq. (32) is evaluated using the appropriately weighted average of the two event categories.

relevant to the signal candidate and discarding the rest of the event information.

In terms of required bandwidth, in principle the resonant search can be carried out by storing just $m_{e^+e^-}$ for each signal candidate, though in practice one wants to store additional information to perform cross checks. At minimum, we expect that the four-momenta and various detector-response information can be kept for all particles in each signal candidate. More information can be kept for displaced candidates that are inconsistent with arising due to photon conversion, where the criteria used to select such candidates will be much looser than those applied to define our reach above. Finally, full events can be kept for a small fraction of signal candidates to permit detailed offline studies of the detector performance and background contributions. The experience gained by LHCb during Run 2 using reduced event storage will aid in determining how to optimize the data-storage strategy employed for the $D^{*0} \rightarrow D^0 A'$ search.

VII. POSSIBLE IMPROVEMENTS

The predicted reach in Fig. 2 covers most of the available A' parameter space for $m_{A'} \lesssim 100$ MeV, and one expects that LHCb could fully cover this region with modest improvements. That said, this reach is based on simulation, so despite our conservative approach, the actual reach may be less than predicted. Furthermore, the ultimate goal is to discover an A' boson and measure its properties. Discovery requires 5σ sensitivity, which our search provides over most of the relevant A' parameter space but not all of it. Precision measurement of the A' mass and lifetime may be vital to determining whether the A' decays invisibly into dark matter. For these reasons, it is worth investigating possible improvements.

There are various ways to improve the dark photon search at LHCb. At minimum, one could assume that the integrated luminosity accumulated by LHCb is larger than the 15 fb^{-1} baseline used in this study.¹⁴ One year of running at the nominal instantaneous luminosity corresponds to about 10 fb^{-1} ; therefore, our Run 3 estimate conservatively provides substantial LHC ramp up and LHCb commissioning time. LHCb also plans to take data in Run 4, collecting a minimum of 50 fb^{-1} in Runs 3 and 4 combined [51]. Furthermore, LHCb is currently investigating ways to increase the integrated luminosity in Run 4. By the end of Run 4, LHCb may collect 10–20 times more data than we used to estimate the reach in this study.

There are several ways to improve the $m_{e^+e^-}$ resolution. The mass-constrained fit we employed in Sec. III D

only makes use of kinematical information. For large flight distance, the vector formed from the pp collision to the A' decay vertex provides another measurement of the direction of flight. Thus, a more sophisticated fitting procedure would not only constrain energy and momentum, but also enforce a consistent decay topology [72], likely improving the resolution. In addition, during Run 1 LHCb managed to recover many of the bremsstrahlung photons emitted by electrons by reconstructing them in the calorimeter system [73]. These photons were then added back to the electron momenta. Since the calorimeter occupancy will be higher in Run 3, we chose to assume that this recovery procedure would not be possible in our analysis. If bremsstrahlung recovery is possible, though, it would improve the $m_{e^+e^-}$ resolution by $\mathcal{O}(10\%)$. We note that if a major upgrade of the calorimeter system [74] is installed for Run 4 that utilizes large-area picosecond photodetectors [75], this would significantly improve the $m_{e^+e^-}$ resolution. Such an upgrade would also permit using many more D^0 decays.

Installing additional tracking stations onto the face of the dipole magnet would provide LONG-track-level momentum resolution for UP tracks [74]. The hit resolution required is $\mathcal{O}(\text{mm})$, which means that such a tracking station would not be too costly. Since many A' decays are reconstructed with at least one UP track, installing such a system could greatly improve the $m_{e^+e^-}$ resolution. Furthermore, many D^0 decays produce UP tracks whose resolution would also be improved by installing a tracking system onto the face of the magnet.

For simplicity, our resonant search imposed a conversion veto, but since the conversion background is smooth, one could still perform a resonant search without any ℓ_T requirement at the expense of larger backgrounds. This is particularly relevant for small $m_{A'}$ where the D_T buffer requirement is most inefficient.

Finally, one could consistently combine both displaced searches and the resonant search following a strategy similar to Ref. [76]. This is particularly relevant if the A' can also decay into dark matter, since then the relationship between $\Gamma_{A'}$, $m_{A'}$, and ϵ^2 given in Eq. (20) will no longer hold and A' will have a shorter lifetime. To avoid introducing model dependence into the A' search, one can define exclusive prompt and displaced regions experimentally, and then optimally combine the information from both regions regardless of the A' lifetime [76]. A combined strategy might also help close the remaining gap between the resonant and displaced searches at large $m_{A'}$.

VIII. COMPARISON TO OTHER EXPERIMENTS

There is a rich planned program of dark boson searches [19], and various experimental proposals are sensitive to the same dark photon mass and coupling range targeted by this $D^{*0} \rightarrow D^0 A'$ search. Their anticipated reaches are

¹⁴ As stated in footnote 6, the first measurement of the $c\bar{c}$ cross section at 13 TeV suggests that the cross section in Eq. (6) is likely about 20% too small. If so, one can view our reach results as for 12 fb^{-1} .

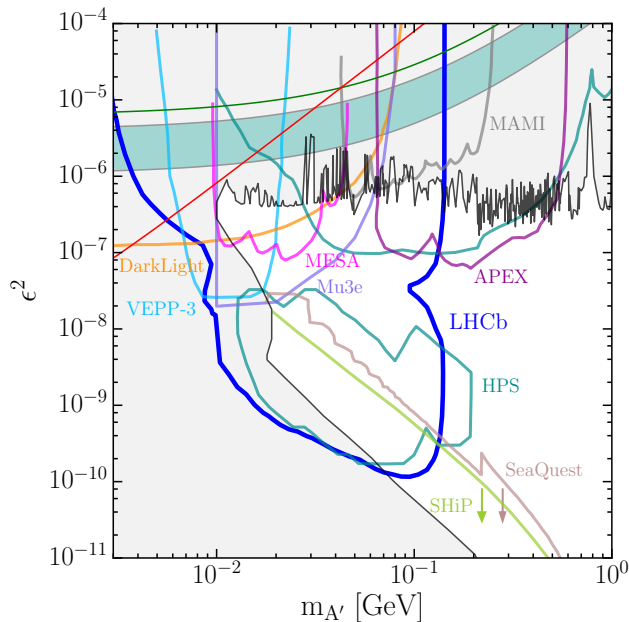


FIG. 9. Comparing the LHCb reach to other proposed dark photon experiments.

summarized in Fig. 9. Experiments like APEX [32, 77], MESA/MAMI [78], DarkLight [79, 80], VEPP-3 [23], and Mu3e [81] are high luminosity experiments that use a resonant search strategy.¹⁵ Experiments like SHiP [82] and SeaQuest [83] are beam-dump experiments that use a displaced strategy. See App. C for discussion of a possible $D^{*0} \rightarrow D^0 A'$ search at Belle-II.

The closest comparison to the LHCb $D^{*0} \rightarrow D^0 A'$ search is the HPS experiment [52]. Because HPS has a dedicated tracking and vertexing detector, it is able to search for both resonant and displaced A' signals, corresponding to the upper and lower HPS regions in Fig. 9. In terms of vertex performance and mass resolution, HPS is quite comparable to LHCb.¹⁶ It is therefore reasonable to ask why the resonant and displaced search regions overlap for LHCb but not for HPS.

There are three main advantages of LHCb over HPS, at least for $m_{A'} < 100$ MeV.

- *Parasitic running.* For a fixed mass resolution, the resonant search is limited only by the available statistics. The $D^{*0} \rightarrow D^0 A'$ search does not require any modifications to the standard LHC running environment, so it immediately benefits from the high data-taking rate (and long run times) already needed by other LHCb measurements. By

contrast, HPS is a dedicated experiment with an anticipated runtime of only 3 weeks to cover the 30-100 MeV mass range.

- *Access to smaller opening angles.* As a fixed-target experiment, HPS produces A' bosons in the very forward direction, effectively within the envelope of the beam pipe hole. This means that HPS has no sensitivity in the “dead zone” where the A' decay opening angle is less than 30 mrad. As a colliding beam experiment, LHC produces A' bosons with a sizable transverse momentum kick, such that the A' trajectory itself goes into the LHCb VELO. This allows LHCb to reconstruct much smaller opening angles down to around a few mrad, which helps the reach at low A' masses.
- *Larger Lorentz boosts.* The reach in the displaced search benefits from large γ factors (up to a point, see below). The median A' boost at the LHC is roughly three times larger than the maximum A' boost at HPS. Moreover, there is a tail of A' events at the LHC which extends to much higher boost factors, which can be exploited when combined with the high event rate.

Note that these last two bullet points are in direct conflict at a fixed-target experiment, since going to larger Lorentz boosts by using a higher beam energy means that the signal has smaller A' opening angles. At a colliding beam experiment, the opening angle resolution is limited only by the hit resolution, so one can in principle exploit larger Lorentz boosts up until the point that the displaced A' signal escapes the detector. From this we conclude that HPS is probably close to optimal for a fixed-target dark photon search in this mass range.

IX. SUMMARY

In this paper, we showed that in Run 3, LHCb can explore the entire dark photon parameter space between the prompt- A' and beam-dump limits for $m_{A'} \lesssim 100$ MeV using the decay $D^{*0} \rightarrow D^0 A'$. This reach is possible due to the large D^{*0} production rate and sizable A' Lorentz boost factor at the LHC, combined with the excellent vertex/mass resolution and planned triggerless-readout system of LHCb. The displaced and resonant strategies give complementary coverage of the A' parameter space. Even if the displaced vertex signature is absent due to a modified A' lifetime, there is still substantial coverage from the resonant search owing to the excellent $m_{e^+e^-}$ resolution.

Given the impressive reach below $\Delta m_D = 142$ MeV, one might wonder if a similar search could be performed at LHCb for dark photons with larger masses. The $D^{*0} \rightarrow D^0 A'$ search relies on the D^{*0} mass constraint to suppress backgrounds and to improve the e^+e^- invariant mass resolution. Without these handles, the dark photon

¹⁵ Strictly speaking, VEPP-3 uses a missing mass strategy.

¹⁶ One minor difference is that HPS uses strips for its tracking while LHCb will use pixels in Run 3. This means that HPS has worse hit resolution in the bending plane, so some of the topological requirements in App. B would not be helpful for HPS.

search becomes considerably more difficult, but there are a number of potential strategies to explore.

Above the dimuon threshold, a generic search could be performed for displaced $A' \rightarrow \mu^+\mu^-$ vertices that are inconsistent with originating from interactions with the detector material. The dimuon invariant mass resolution is sufficient without applying a mass constraint, and data samples of displaced dimuon vertices are nearly background free at LHCb (see, e.g., Ref. [84]). A similar search was performed by LHCb using Run 1 data that looked for the very rare decay $K_S^0 \rightarrow \mu^+\mu^-$ [85]. A key challenge for these generic searches, though, is converting a cross section bound into a bound on ϵ^2 , since pp collisions provide many potential sources of dark photons with uncertain production rates.

Other radiative charm decays may also offer viable A' search channels at LHCb. For example, the yield of $D^0 \rightarrow \bar{K}^{*0}\gamma$ decays, where both the kaon and pion from the $\bar{K}^{*0} \rightarrow K\pi$ decay are reconstructed as LONG tracks, will be $\mathcal{O}(10^{10})$ in Run 3. A search could then be performed for $D^0 \rightarrow \bar{K}^{*0}A'$, with $A' \rightarrow e^+e^-$, $\mu^+\mu^-$, and $\pi^+\pi^-$. This search is sensitive to $m_{A'} \lesssim 1$ GeV and the D^0 mass constraint would ensure good A' mass resolution. The A' lifetime is shorter at larger $m_{A'}$, however, so a displaced search would not be possible over most of the $m_{A'}$ range. Nevertheless, such a search is worth investigating. Similarly, the decay $D_{s1}(2460)^+ \rightarrow D_s^+A'$ may provide a viable search for $m_{A'} \lesssim 500$ MeV.

The common theme for all of the above search strategies is that they benefit greatly from triggerless readout with real-time data calibration. These and other searches provide strong physics motivation for these upgrades to LHCb in Run 3. We encourage more effort in exploring how best to exploit these advances in the hunt for dark photons and other new particles.

ACKNOWLEDGMENTS

We are indebted to Natalia Toro for extensive discussions on prompt backgrounds to the pre-module search. We thank Rouven Essig, Marat Freytsis, Tim Gershon, Yoni Kahn, Patrick Koppenburg, Gordan Krnjaic, Zoltan Ligeti, Jia Liu, Iain Stewart, and Vincenzo Vagnoni for providing helpful advice and feedback. We thank Yasmine Amhis, Thomas Bird, Chris Jones, Marie-Hélène Schune, Paul Seyfert, Edwige Tournefier, and Daniel Vieira for help determining the Run 3 LHCb performance. J.T. and W.X. are supported by the U.S. Department of Energy (DOE) under cooperative research agreement DE-SC-00012567. J.T. is also supported by the DOE Early Career research program DE-SC-0006389, and by a Sloan Research Fellowship from the Alfred P. Sloan Foundation. J.T. thanks the Galileo Galilei Institute for Theoretical Physics for the hospitality and the INFN for partial support during the completion of this work. M.W. and P.I. are supported by the U.S. National Science Foundation grant PHY-1306550.

Appendix A: Estimating LHCb Tracking Performance

LHCb does not provide public fast simulation software. Therefore in this appendix, we estimate the various performance numbers needed for our study from public documents.

We start with deriving the resolutions for LONG and UP tracks, as defined in Sec. III A. The resolution on σ_p/p for LONG tracks with $p > 10$ GeV as a function of p in the current LHCb detector is provided in Fig. 17 of Ref. [59]. The LONG tracks used in this analysis typically have momenta in the region where

$$\sigma_p^{\text{LONG}}/p \approx 0.5\% \quad (\text{A1})$$

(n.b. σ_p/p has no strong dependence on p for $p < 10$ GeV). The LONG track resolution is expected to improve by about 20% in Run 3, but to be conservative we assume the performance will be as measured in the current detector. The resolution for UP tracks with momenta typical of electrons coming from A' decays is

$$\sigma_p^{\text{UP}}/p \approx 12\%, \quad (\text{A2})$$

as can be seen in Fig. 2.5 in Ref. [65].

The resolution on the track direction is not explicitly published in any LHCb document to our knowledge, but we can extract it from published mass resolutions. LHCb measures σ_p/p for tracks using various two-body decays, e.g. $J/\psi \rightarrow \mu^+\mu^-$. For the case where $p_{\mu^+} \approx p_{\mu^-}$, the resolution on the track momentum is related to the resolution on the J/ψ mass m and decay opening angle α by

$$\left(\frac{\sigma_p}{p}\right)^2 \approx 2\left(\frac{\sigma_m}{m}\right)^2 - 2\left(\frac{p\sigma_\alpha}{m}\right)^2, \quad (\text{A3})$$

which is the same as Eq. (1) of Ref. [59] but without the additional factor of α in the denominator of the rightmost term (that is a typo in Ref. [59]). Taking $\sigma_m = 14$ MeV from Tab. 2, along with $\sigma_p/p = 0.5\%$ and $p \approx 15$ GeV from Fig. 17 of Ref. [59], we obtain $\sigma_\alpha \approx 0.3$ mrad, which gives a track polar-angle resolution of $\sigma_\theta \approx 0.2$ mrad in the $\theta_{J/\psi} \rightarrow 0$ limit.

The value $\sigma_\theta = 0.2$ mrad should be considered valid only for tracks with large momenta. For low-momentum tracks, multiple scattering dominates the resolution. We simulate a toy model of the Run 3 VELO taking the radiation length of the foil and detector modules from Ref. [64]. The resulting resolution is

$$\sigma_\theta \approx \left(0.2 + \frac{1.7 \text{ GeV}}{p}\right) \text{ mrad}, \quad \sigma_\phi \approx \sigma_\theta \cot \theta. \quad (\text{A4})$$

When the A' is constrained to originate from the pp collision, as assumed in deriving Fig. 5, $\sigma_\phi \approx \left(\sqrt{\sigma_{\text{hit}}^2 + \sigma_{pp,T}^2}\right)/6 \text{ mm} \approx 3 \text{ mrad}$ (the hit and transverse pp location resolution divided by the radial distance to the first hit).

We now turn to the resolutions for F-type and P-type D^0 decays. For two-body F-type D^0 decays, both tracks are LONG and so one would naively expect $\sigma_p/p \approx 0.5\% \times \sqrt{2}$. However, applying the constraint of the known D^0 mass improves this by about a factor of two in our toy simulation. In the $\theta_D \rightarrow 0$ limit,

$$\sigma_\theta^{D^0_F} \approx \frac{\sigma_\alpha}{2} \approx 0.2 \text{ mrad}, \quad (\text{A5})$$

with

$$\sigma_\phi^{D^0_F} \approx \sigma_\theta \cot \theta \approx 2 \text{ mrad}. \quad (\text{A6})$$

For four-body F-type decays, the resolution is similar to the two-body case after applying the D^0 mass constraint, even if one or two of the tracks are of the UP type. For simplicity, the largest σ_p/p value obtained (when two tracks are UP)

$$\sigma_p^{D^0_F}/p \approx 0.5\% \quad (\text{A7})$$

is used for all F-type decays.

For P-type D^0 decays, the resolution depends strongly on the D^0 flight distance ℓ . Therefore, improved resolution comes at the expense of signal efficiency. LHCb achieves a proper lifetime resolution of $\sigma_t \approx 50$ fs in D^0 decays [86]. Using the relationship

$$\sigma_t^2 = \left(\frac{m}{p}\right)^2 \sigma_\ell^2 + (m\ell)^2 \left(\frac{\sigma_p}{p}\right)^2, \quad (\text{A8})$$

we can obtain a per-event estimate of the resolution on the flight distance. The vertex resolution in the plane transverse to the beam line ($x-y$) is much better than along the beam direction (z), e.g. in Ref. [59] $\sigma_z/\sigma_x \approx 5.5$ for the pp -collision vertex. Therefore, the resolution on the transverse flight distance is $\sigma_{\ell_T} \approx \sigma_\ell \sin \theta$, and $\sigma_\theta \approx (\sigma_{\ell_T}/\ell \cos \theta)(5.5/\sqrt{2}) \approx \tan \theta \sigma_\ell/4\ell$, where σ_ℓ is obtained from Eq. (A8) for each D^0 decay. Similarly, $\sigma_\phi \approx \sigma_\ell/4\ell$. Both expressions have the expected scaling with flight distance. Finally, applying D^0 selection criteria that are 50% efficient gives

$$\sigma_\theta^{D^0_P} \approx 2 \text{ mrad}, \quad \sigma_\phi^{D^0_P} \approx 30 \text{ mrad}. \quad (\text{A9})$$

In this toy study, we require $\sigma_\theta < 5$ mrad and assume additional requirements, e.g. on particle identification, are applied that are 90% efficient. In a more sophisticated analysis where the full covariance matrix from the vertex fit is available, one could choose to cut on σ_p/p as in Ref. [87].

The D^0 momentum in P-type decays is solved for as described in Sec. III B. For small ℓ it is common that p_{mis}^\perp is unphysical, i.e. it is larger than the expected total missing momentum in the D^0 rest frame due to the resolution on the D^0 decay vertex. In such cases, we vary the location of the D^0 vertex to find the position with the smallest vertex χ^2 that provides a physical p_{mis}^\perp value. After requiring $\sigma_\theta < 5$ mrad, we find

$$\sigma_p^{D^0_P}/p \approx 9\%. \quad (\text{A10})$$

It is likely that a more sophisticated selection based on σ_p/p would provide better resolution, though this has little impact on the search proposed in this paper.

Appendix B: Estimating LHCb Vertex Performance

The resolution on the IP and vertex location are key elements of the displaced searches. The one-dimensional track IP resolution expected in Run 3 is well approximated by [64]

$$\sigma_{\text{IP}} = \left(11.0 + \frac{13.1 \text{ GeV}}{p_T}\right) \mu\text{m}. \quad (\text{B1})$$

In the $\theta_{A'} \rightarrow 0$ limit, this is the IP resolution both in and orthogonal to the A' decay plane. We can use σ_{IP} to cross check our estimate of σ_θ for tracks. Assuming that the location of the pp collision is perfectly measured, then $\sigma_\theta \approx \sigma_{\text{IP}} \sin \theta/6$ mm, where 6 mm is the mean radial distance of the inner-most track hit [64]. This gives $\sigma_\theta \approx 2.2 \text{ GeV}/p$, which, for typical electron momenta in $D^{*0} \rightarrow D^0 A'$ decays, agrees to within about 10% with our estimate in Eq. (A4) based on mass resolution and a toy VELO simulation.

The resolution on ℓ_T in the pre-module search is well approximated by

$$\sigma_{\ell_T} \approx \frac{\sin \theta}{\alpha_{e^+e^-}} \sqrt{\sigma_{e^+\text{IP}}^2 + \sigma_{e^-\text{IP}}^2}, \quad (\text{B2})$$

where $\alpha_{e^+e^-}$ is the A' decay opening angle and the A' is constrained to originate from the pp collision. The resolution on ℓ_T as a function of $m_{A'}$ is shown in Fig. 10, as is the required displacement $n\sigma_{\ell_T}$ to achieve ≈ 1 background event in a given A' signal mass window. The $m_{A'}$ dependence is driven by the dependence of $\alpha_{e^+e^-}$ on $m_{A'}$.

In our pre-module displaced search, we require a consistent decay topology. These requirements are as follows:

- the A' decay vertex is downstream of the pp collision;
- the distance of closest approach between the e^+ and e^- tracks is consistent with zero;
- the angle between $\vec{p}_{A'}$ and the vector formed from the pp collision to the A' decay vertex is consistent with zero;
- the IP out of the A' decay plane—defined by the pp collision point and the first hits on the e^+ and e^- tracks—for each electron is consistent with zero.

In each case, we define consistent with zero as having a p -value greater than 1%. Therefore, the efficiency on a true displaced A' decay is close to 100%.

An important point to keep in mind is that virtually all fake highly displaced e^+e^- vertices that satisfy these consistency requirements are reconstructed with $\alpha_{e^+e^-}$ much

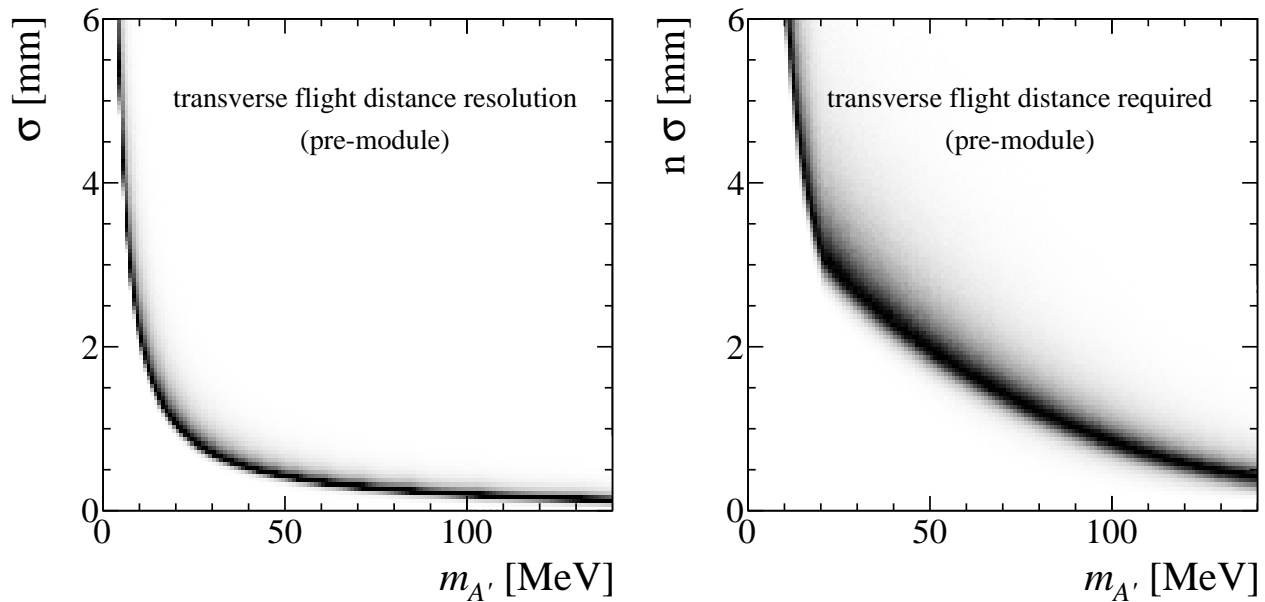


FIG. 10. Distribution of (left) resolution σ_{ℓ_T} and (right) required displacement $n\sigma_{\ell_T}$ versus $m_{A'}$ for pre-module A' decays.

larger than its true value. Therefore, the reconstructed value of $m_{e^+e^-}$ is larger than its true value in the absence of substantial bremsstrahlung, and the dominant background to the A' signal comes from (more copious) e^+e^- pairs at smaller true invariant mass.

The opening angle in our search for $m_{A'} = 100$ MeV is on average about the same as for the HPS experiment. One can see from Fig. 10 that the pre-module search requirement at 100 MeV is $\ell_T \gtrsim 1$ mm or $\ell \gtrsim 25$ mm. This value is similar to the flight distance requirement used by HPS to determine the reach of their displaced search, leading to a comparable high-side reach at that mass value (see Fig. 9).

For the post-module search, we define IP_i with respect to the point where the A' trajectory intersects the i^{th} VELO module, which is the one directly upstream of the A' decay vertex location. The resolution on IP_i is given by

$$\sigma_{\text{IP}_i} \approx \sqrt{(\Delta Z_i \sigma_\theta)^2 + (12 \mu\text{m})^2}, \quad (\text{B3})$$

where ΔZ_i is the separation of VELO modules i and $i+1$, and $12 \mu\text{m}$ is the hit resolution. For simplicity, we take $\Delta Z_i = 25$ mm for all modules. In reality, about 40% of the modules traversed by an A' candidate have a larger ΔZ_i . Such cases have slightly better resolution since the relative effect of the hit resolution is reduced. The resolution on the z position of the A' decay vertex in the post-module search is

$$\sigma_{z_i} \approx \frac{1}{\alpha_{e^+e^-}} \sqrt{\sigma_{e^+\text{IP}_i}^2 + \sigma_{e^-\text{IP}_i}^2}, \quad (\text{B4})$$

where again the A' is constrained to originate from the pp collision. The resolution on z_i is shown in Fig. 11, as is

the required displacement $n\sigma_{z_i}$. The consistency requirements imposed for the post-module displaced search are the same as those used in the pre-module search, with the following exceptions: the A' decay vertex is required to be downstream of the i^{th} VELO module, and IP_i is used in place of IP for the out-of-decay-plane requirement.

Appendix C: Belle-II Reach

The search strategy in this paper can be applied to any experiment with a large D^{*0} production rate. As a key example, the Belle-II experiment [88] will collect a data sample corresponding to about 50 ab^{-1} starting in 2017. Indeed, Belle-II plans to collect about 10 ab^{-1} prior to the start of Run 3 at the LHC, so they may be able to set initial bounds on $D^{*0} \rightarrow D^0 A'$ (or make a discovery) prior to LHCb.

Again neglecting secondary D^{*0} production from b hadron decays, the total number of prompt $D^{*0} \rightarrow D^0 \gamma$ decays produced at Belle-II will be approximately

$$N(D^{*0} \rightarrow D^0 \gamma) \approx 510 \text{ pb} \times 0.381 \times 50 \text{ ab}^{-1} \approx 10^{10}, \quad (\text{C1})$$

where the first number is the inclusive prompt $e^+e^- \rightarrow D^{*0}$ cross section and the second number is the $D^{*0} \rightarrow D^0 \gamma$ branching ratio. This is about 500 times less than at LHCb in Run 3, though Belle-II will likely be able to make use of a larger fraction of D^0 decays than LHCb. Rescaling the limits derived in this paper, we anticipate Belle-II will not be able to probe unexplored parameter space using the resonant A' strategy.

The main challenge for a Belle-II displaced search is that the electrons produced in $D^{*0} \rightarrow D^0 A'$ decays will have very low momenta. Using PYTHIA, we estimate that

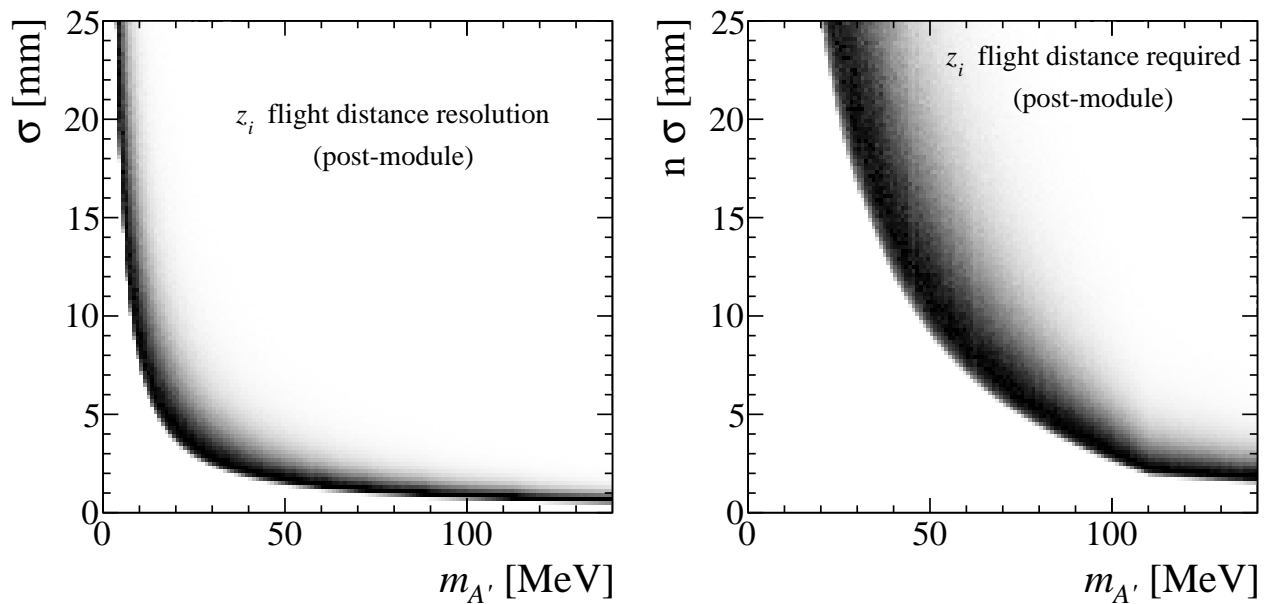


FIG. 11. Distribution of (left) resolution σ_{z_i} and (right) required displacement $n\sigma_{z_i}$ versus $m_{A'}$ for post-module A' decays.

the median value of the lower-momentum electron is only 60 MeV. This is likely to reduce the tracking efficiency, the electron-momentum resolution, and the e^+e^- vertex resolution. The Lorentz γ factors at Belle-II will also be much lower resulting in shorter flight distances. We

estimate that Belle-II could have sensitivity to displaced A' decays for small $m_{A'}$. Given that such sensitivity may occur prior to Run 3 at LHCb or elsewhere, we encourage a detailed study by the Belle-II collaboration to assess the discovery potential.

-
- [1] L. Okun, Sov. Phys. JETP **56**, 502 (1982).
[2] P. Galison and A. Manohar, Phys. Lett. **B136**, 279 (1984).
[3] B. Holdom, Phys. Lett. **B166**, 196 (1986).
[4] M. Pospelov, A. Ritz, and M. B. Voloshin, Phys. Lett. **B662**, 53 (2008), arXiv:0711.4866 [hep-ph].
[5] N. Arkani-Hamed, D. P. Finkbeiner, T. R. Slatyer, and N. Weiner, Phys. Rev. **D79**, 015014 (2009), arXiv:0810.0713 [hep-ph].
[6] J. D. Bjorken, R. Essig, P. Schuster, and N. Toro, Phys. Rev. **D80**, 075018 (2009), arXiv:0906.0580 [hep-ph].
[7] G. Bernardi, G. Carugno, J. Chauveau, F. Dicarolo, M. Dris, *et al.*, Phys. Lett. **B166**, 479 (1986).
[8] R. Meijer Drees *et al.* (SINDRUM I collaboration), Phys. Rev. Lett. **68**, 3845 (1992).
[9] P. Astier *et al.* (NOMAD collaboration), Phys. Lett. **B506**, 27 (2001), arXiv:hep-ex/0101041 [hep-ex].
[10] S. Gninenko, Phys. Rev. **D85**, 055027 (2012), arXiv:1112.5438 [hep-ph].
[11] P. Adlarson *et al.* (WASA-at-COSY collaboration), Phys. Lett. **B726**, 187 (2013), arXiv:1304.0671 [hep-ex].
[12] G. Agakishiev *et al.* (HADES collaboration), Phys. Lett. **B731**, 265 (2014), arXiv:1311.0216 [hep-ex].
[13] A. Adare *et al.* (PHENIX collaboration), Phys. Rev. **C91**, 031901 (2015), arXiv:1409.0851 [nucl-ex].
[14] J. Batley *et al.* (NA48/2 collaboration), (2015), arXiv:1504.00607 [hep-ex].
[15] F. Bergsma *et al.* (CHARM collaboration), Phys. Lett. **B166**, 473 (1986).
[16] S. Gninenko, Phys. Lett. **B713**, 244 (2012), arXiv:1204.3583 [hep-ph].
[17] F. Archilli *et al.* (KLOE-2 collaboration), Phys. Lett. **B706**, 251 (2012), arXiv:1110.0411 [hep-ex].
[18] D. Babusci *et al.* (KLOE-2 collaboration), Phys. Lett. **B720**, 111 (2013), arXiv:1210.3927 [hep-ex].
[19] R. Essig *et al.*, in *Community Summer Study 2013: Snowmass on the Mississippi (CSS2013) Minneapolis, MN, USA, July 29-August 6, 2013* (2013) arXiv:1311.0029 [hep-ph].
[20] B. Aubert *et al.* (BaBar collaboration), in *Proceedings, 34th International Conference on High Energy Physics (ICHEP 2008)* (2008) arXiv:0808.0017 [hep-ex].
[21] B. Batell, M. Pospelov, and A. Ritz, Phys. Rev. **D80**, 095024 (2009), arXiv:0906.5614 [hep-ph].
[22] P. deNiverville, M. Pospelov, and A. Ritz, Phys. Rev. **D84**, 075020 (2011), arXiv:1107.4580 [hep-ph].
[23] B. Wojtsekhowski, D. Nikolenko, and I. Rachek, (2012), arXiv:1207.5089 [hep-ex].
[24] Y. Kahn and J. Thaler, Phys. Rev. **D86**, 115012 (2012), arXiv:1209.0777 [hep-ph].
[25] E. Izaguirre, G. Krnjaic, P. Schuster, and N. Toro, Phys. Rev. **D88**, 114015 (2013), arXiv:1307.6554 [hep-ph].
[26] B. Batell, R. Essig, and Z. Surujon, Phys. Rev. Lett. **113**, 171802 (2014), arXiv:1406.2698 [hep-ph].
[27] Y. Kahn, G. Krnjaic, J. Thaler, and M. Toups, Phys. Rev. **D91**, 055006 (2015), arXiv:1411.1055 [hep-ph].

- [28] E. Izaguirre, G. Krnjaic, P. Schuster, and N. Toro, *Phys. Rev.* **D91**, 094026 (2015), arXiv:1411.1404 [hep-ph].
- [29] E. Izaguirre, G. Krnjaic, P. Schuster, and N. Toro, (2015), arXiv:1505.00011 [hep-ph].
- [30] H. Merkel *et al.* (A1 collaboration), *Phys. Rev. Lett.* **106**, 251802 (2011), arXiv:1101.4091 [nucl-ex].
- [31] H. Merkel *et al.*, *Phys. Rev. Lett.* **112**, 221802 (2014), arXiv:1404.5502 [hep-ex].
- [32] S. Abrahamyan *et al.* (APEX collaboration), *Phys. Rev. Lett.* **107**, 191804 (2011), arXiv:1108.2750 [hep-ex].
- [33] A. Konaka, K. Imai, H. Kobayashi, A. Masaike, K. Miyake, *et al.*, *Phys. Rev. Lett.* **57**, 659 (1986).
- [34] E. Riordan, M. Krasny, K. Lang, P. De Barbaro, A. Bodek, *et al.*, *Phys. Rev. Lett.* **59**, 755 (1987).
- [35] J. Bjorken, S. Ecklund, W. Nelson, A. Abashian, C. Church, *et al.*, *Phys. Rev.* **D38**, 3375 (1988).
- [36] A. Bross, M. Crisler, S. H. Pordes, J. Volk, S. Errede, *et al.*, *Phys. Rev. Lett.* **67**, 2942 (1991).
- [37] M. Davier and H. Nguyen Ngoc, *Phys. Lett.* **B229**, 150 (1989).
- [38] C. Athanassopoulos *et al.* (LSND collaboration), *Phys. Rev.* **C58**, 2489 (1998), arXiv:nucl-ex/9706006 [nucl-ex].
- [39] S. Adler *et al.* (E787 collaboration), *Phys. Rev.* **D70**, 037102 (2004), arXiv:hep-ex/0403034 [hep-ex].
- [40] A. V. Artamonov *et al.* (BNL-E949 collaboration), *Phys. Rev.* **D79**, 092004 (2009), arXiv:0903.0030 [hep-ex].
- [41] R. Essig, R. Harnik, J. Kaplan, and N. Toro, *Phys. Rev.* **D82**, 113008 (2010), arXiv:1008.0636 [hep-ph].
- [42] J. Blumlein and J. Brunner, *Phys. Lett.* **B701**, 155 (2011), arXiv:1104.2747 [hep-ex].
- [43] J. Blumlein and J. Brunner, *Phys. Lett.* **B731**, 320 (2014), arXiv:1311.3870 [hep-ph].
- [44] D. Hanneke, S. Fogwell, and G. Gabrielse, *Phys. Rev. Lett.* **100**, 120801 (2008), arXiv:0801.1134 [physics.atom-ph].
- [45] G. F. Giudice, P. Paradisi, and M. Passera, *JHEP* **11**, 113 (2012), arXiv:1208.6583 [hep-ph].
- [46] M. Endo, K. Hamaguchi, and G. Mishima, *Phys. Rev.* **D86**, 095029 (2012), arXiv:1209.2558 [hep-ph].
- [47] M. Pospelov, *Phys. Rev.* **D80**, 095002 (2009), arXiv:0811.1030 [hep-ph].
- [48] J. B. Dent, F. Ferrer, and L. M. Krauss, (2012), arXiv:1201.2683 [astro-ph.CO].
- [49] D. Kazanas, R. N. Mohapatra, S. Nussinov, V. L. Teplitz, and Y. Zhang, *Nucl. Phys.* **B890**, 17 (2014), arXiv:1410.0221 [hep-ph].
- [50] K. A. Olive *et al.* (Particle Data Group), *Chin. Phys.* **C38**, 090001 (2014).
- [51] “LHCb Trigger and Online Technical Design Report,” (2014), LHCb-TDR-016.
- [52] O. Moreno, (2013), arXiv:1310.2060 [physics.ins-det].
- [53] T. Sjöstrand *et al.*, *Comput. Phys. Commun.* **191**, 159 (2015), arXiv:1410.3012 [hep-ph].
- [54] R. Aaij *et al.* (LHCb collaboration), *Nucl. Phys.* **B871**, 1 (2013), arXiv:1302.2864 [hep-ex].
- [55] R. Aaij *et al.* (LHCb collaboration), “LHCb-PAPER-2015-041, in preparation,”.
- [56] G. A. Miller and P. Singer, *Phys. Rev.* **D37**, 2564 (1988).
- [57] I. W. Stewart, *Nucl. Phys.* **B529**, 62 (1998), arXiv:hep-ph/9803227 [hep-ph].
- [58] A. A. Alves Jr. *et al.* (LHCb collaboration), *JINST* **3**, S08005 (2008).
- [59] R. Aaij *et al.* (LHCb collaboration), *Int. J. Mod. Phys.* **A30**, 1530022 (2015), arXiv:1412.6352 [hep-ex].
- [60] R. Aaij *et al.*, *JINST* **9**, P09007 (2014), arXiv:1405.7808 [physics.ins-det].
- [61] M. Adinolfi *et al.*, *Eur. Phys. J.* **C73**, 2431 (2013), arXiv:1211.6759 [physics.ins-det].
- [62] A. A. Alves Jr. *et al.*, *JINST* **8**, P02022 (2013), arXiv:1211.1346 [physics.ins-det].
- [63] R. Arink *et al.*, *JINST* **9**, P01002 (2014), arXiv:1311.3893 [physics.ins-det].
- [64] “LHCb VELO Upgrade Technical Design Report,” (2013), LHCb-TDR-013.
- [65] “LHCb Tracker Upgrade Technical Design Report,” (2014), LHCb-TDR-015.
- [66] R. Aaij *et al.* (LHCb collaboration), *JHEP* **07**, 014 (2014), arXiv:1405.2797 [hep-ex].
- [67] R. a. Aaij (LHCb collaboration), *JHEP* **10**, 005 (2014), arXiv:1408.1299 [hep-ex].
- [68] H. W. Koch and J. W. Motz, *Rev. Mod. Phys.* **31**, 920 (1959).
- [69] “LHCb PID Upgrade Technical Design Report,” (2013), LHCb-TDR-014.
- [70] HPS collaboration, “Proposal for 2014-2015 run,”.
- [71] A. Borsellino, *Phys. Rev.* **89**, 1023 (1953).
- [72] W. D. Hulsbergen, *Nucl. Instrum. Meth.* **A552**, 566 (2005), arXiv:physics/0503191 [physics].
- [73] R. Aaij *et al.* (LHCb collaboration), *JHEP* **04**, 064 (2015), arXiv:1501.03038 [hep-ex].
- [74] S. Stone, “private communication,”.
- [75] LAPPD collaboration, “http://psec.uchicago.edu,”.
- [76] M. Williams, *JINST* **10**, P06002 (2015), arXiv:1503.04767 [hep-ex].
- [77] R. Essig, P. Schuster, N. Toro, and B. Wojtsekhowski, *JHEP* **02**, 009 (2011), arXiv:1001.2557 [hep-ph].
- [78] T. Beranek, H. Merkel, and M. Vanderhaeghen, *Phys. Rev.* **D88**, 015032 (2013), arXiv:1303.2540 [hep-ph].
- [79] M. Freytsis, G. Ovanessian, and J. Thaler, *JHEP* **01**, 111 (2010), arXiv:0909.2862 [hep-ph].
- [80] J. Balewski, J. Bernauer, W. Bertozzi, J. Bessuille, B. Buck, *et al.*, (2013), arXiv:1307.4432.
- [81] B. Echenard, R. Essig, and Y.-M. Zhong, *JHEP* **01**, 113 (2015), arXiv:1411.1770 [hep-ph].
- [82] S. Alekhin *et al.*, (2015), arXiv:1504.04855 [hep-ph].
- [83] S. Gardner, R. J. Holt, and A. S. Tadepalli, (2015), arXiv:1509.00050 [hep-ph].
- [84] R. Aaij *et al.* (LHCb collaboration), (2015), arXiv:1508.04094 [hep-ex].
- [85] R. Aaij *et al.* (LHCb collaboration), *JHEP* **01**, 090 (2012), arXiv:1209.4029 [hep-ex].
- [86] R. Aaij *et al.* (LHCb collaboration), *Phys. Rev. Lett.* **110**, 101802 (2013), arXiv:1211.1230 [hep-ex].
- [87] R. Aaij *et al.* (LHCb collaboration), *Nature Phys.* **11**, 17 (2015), arXiv:1504.01568 [hep-ex].
- [88] Belle collaboration, “Belle II technical design report,” KEK Report 2010-1.