Physics Letters B 823 (2021) 136758

Contents lists available at ScienceDirect

Physics Letters B

www.elsevier.com/locate/physletb

Probing naturally light singlets with a displaced vertex trigger

Yuri Gershtein [a], Simon Knapen [b,*], Diego Redigolo [b,c]

[a] Department of Physics and Astronomy, Rutgers University, Piscataway, NJ 08854, USA
[b] CERN, Theoretical Physics Department, Geneva, Switzerland
[c] INFN, Sezione di Firenze and Department of Physics and Astronomy, University of Florence, Via G. Sansone 1, 50019 Sesto Fiorentino, Italy

ARTICLE INFO

Article history:
Received 13 January 2021
Received in revised form 14 September 2021
Accepted 25 October 2021
Available online 27 October 2021
Editor: B. Grinstein

ABSTRACT

We investigate the physics case for a dedicated trigger on a low mass, hadronic displaced vertex at the high luminosity LHC, relying on the CMS phase II track trigger. We estimate the trigger efficiency with a simplified simulation of the CMS track trigger and show that the L1 trigger rate from fake vertices, B meson decays and secondary interactions with the detector material can likely be brought down to the kHz level with a minimal set of cuts. While it would with any doubt be a severe experimental challenge to implement, we conclude that a displaced vertex trigger could open qualitatively new parameter space for exotic Higgs decays, exotic B decays and even direct production of light resonances. We parametrize the physics potential in terms of a singlet scalar mixing with the Standard Model Higgs and an axion-like particle with a coupling to gluons, and review a number or relevant models motivated by the hierarchy and strong CP problems, Dark Matter and baryogenesis.

1. Introduction

With its high luminosity upgrade, the LHC will be capable of delivering up to 7 times its current luminosity to ATLAS and CMS, which implies a corresponding increase of the number pile-up vertices per bunch crossing. Both ATLAS [1] and CMS [2] will undergo substantial upgrades to enable them to handle this increasingly challenging environment. One of the most important innovations is the introduction of tracking information into the level-1 (L1) trigger decision making. The CMS collaboration will accomplish this by making use of double layered sensors in the outer tracker, which detect track "stubs" rather than individual hits. By correlating both sensors, the tracker can then assign an approximate p_T to each stub, such that the computationally expensive track fitting can be restricted to stubs satisfying $p_T > 2$ GeV. This innovation makes track reconstruction feasible at the L1 trigger, though necessarily with a somewhat worse resolution as compared to off-line track reconstruction.

Interestingly, this method also enables the CMS track trigger to reconstruct *displaced* tracks with impact parameters up to about 10 cm [2,3], which can be used to search for displaced jets [3–5] and low mass, displaced dimuon resonances [6]. Both signatures are supported by very strong theory motivation, as low mass, long-lived particles (LLP) are rather generic ingredients for hidden sector models. Specifically, so far the displaced jet topology has been studied for exotic Higgs decays [3] and axion-like particles (ALPs) [7]. A displaced dimuon vertex trigger on the other hand can be a very powerful probe of low mass extended Higgs sectors, by leveraging exotic B-meson decays [6,8].

In this work, we combine both ideas and investigate the physics potential of a hypothetical trigger on a displaced, multi-track hadronic vertex. Of the L1 track trigger applications listed above, this would clearly be the most difficult to implement, given the combinatorical challenge when attempting to reconstruct such a vertex at the trigger level. It is therefore paramount to establish *(i)* a strong physics case for this type of trigger and *(ii)* that manageable background rates can be achieved, assuming the vertex reconstruction challenge can be met. Our goal with this paper is to address both questions and to motivate further experimental studies in this direction.

To parametrize the potential of a displaced vertex trigger, we study two simple benchmark models which encode the salient features of a vast class of beyond the Standard Model (SM) scenarios. In the first model we consider the most minimal extension of the SM Higgs sector by adding a single, real scalar field whose most general lagrangian reads

$$\mathcal{L}_S \supset -\frac{1}{2}\tilde{m}_S^2 S^2 - \mu S H^\dagger H - \frac{1}{2}\lambda_{SH} S^2 H^\dagger H - V_{\rm int}(S)\,, \qquad (1)$$

where $V_{\rm int}(S)$ contains the singlet self-interactions. As we detail in Sec. 2, the spontaneous and/or explicit breaking of the approximate $\mathbb{Z}_2$-symmetry $S \to -S$ (through the second term in Eq. 1 or

* Corresponding author.
E-mail address: smknapen@lbl.gov (S. Knapen).

https://doi.org/10.1016/j.physletb.2021.136758

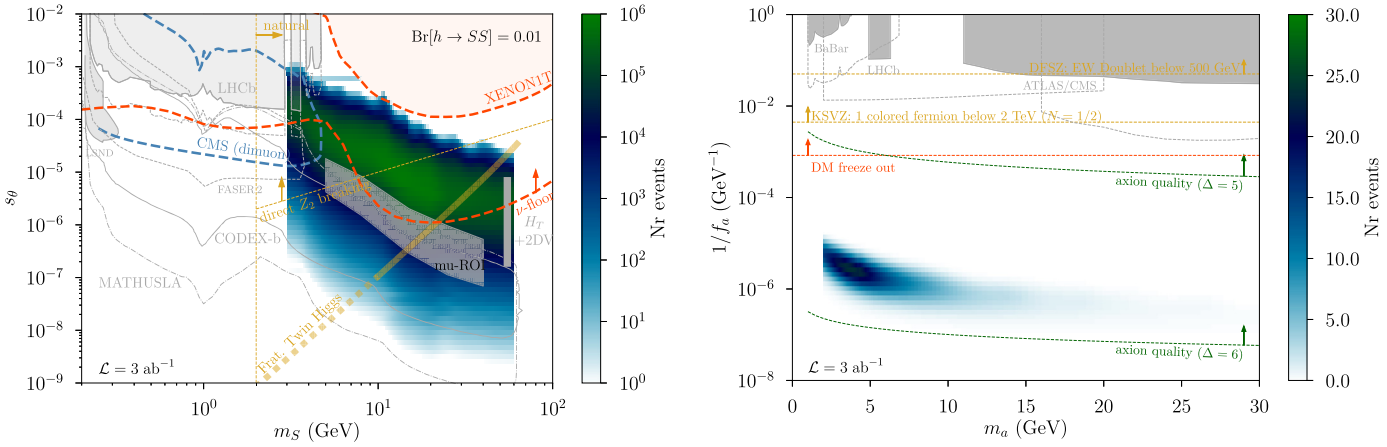

Fig. 1. Left: Projected number of events collected by a CMS L1 displaced vertex trigger at HL-LHC for the scalar model, fixing BR[$h \to SS$] = 1%. **shaded gray** regions represent existing limits from LSND [9], LHCb [10], the ATLAS [11] ("mu-ROI") and CMS ("H_T+2DV") [12]. The **dashed gray** curves are projections for Belle II [13], FASER 2 [14], MATHUSLA [15] and CODEX-b [16]. The **blue dashed** curve is the projected reach of the CMS track trigger in the dimuon channel [8]. The **yellow** lines represent theory priors from naturalness (Eq. 7 in Sec. 2.1), perturbativity (Eq. 9 in Sec. 2.1) and the Fraternal Twin Higgs scenario (Eq. 20 in Sec. 3.1). The two **red** lines represent the XENON1T limit [17] and the neutrino floor in an example Dark Matter model [18], where we fixed $m_\chi/m_S = 3$ (Sec. 3.3). **Right:** Projected number of events collected by a CMS L1 displaced vertex trigger at HL-LHC for the ALP model. **shaded gray** regions represent existing limits, assuming $E/N = 8/3$ (see Sec. 2.2), from diphoton searches at ATLAS and CMS [19], the diphoton cross section measurements at ATLAS and CMS [20] and at LHCb [21], the boosted dijets search at CMS [22] and the Υ decays at Babar [23]. The **dashed gray** curves are projections for Belle II [21], and HL-LHC [20]. The regions above the **dashed yellow** lines have colored fermions below 2 TeV in the KSVZ UV completion (Eq. 13 in Sec. 2.2) or heavy Higgses below 500 GeV in the DFSZ UV completion (Eq. 14 in Sec. 2.2). The region above the **dashed dark green** lines shows where the ALP can solve the strong CP problem even if the PQ symmetry is broken by operators with dimension $\Delta = 5, 6$ (Eq. 22 in Sec. 2.2). Above the **dashed dark red** line standard Dark Matter freeze-out can be realized in a perturbative model (Sec. 3.3).

through singlet potential $V_{\rm int}(S)$.) induces a mixing between the S and the SM Higgs, parametrized by the mixing angle θ. The latter sets the lifetime of the singlet, which has a dominant branching ratio into hadrons in most of the parameter space. The $\mathbb{Z}_2$-preserving quartic (λ_{SH}) controls the exotic Higgs decay $h \to SS$. Given that the SM Higgs is an exceptionally narrow particle, the branching ratio for this process can be large even for very small values of λ_{SH}.

As a second benchmark, we study an axion-like particle (ALP) which couples primarily to gauge bosons

$$\mathcal{L}_a \supset -\frac{1}{2}m_a^2 a^2 - \frac{\alpha_s}{8\pi}\frac{a}{f_a}\tilde{G}G + \frac{E}{N}\frac{\alpha_{\rm em}}{8\pi}\frac{a}{f_a}\tilde{F}F\,, \tag{2}$$

where coupling to massive gauge bosons are mostly irrelevant in the ALP mass range of interest. N and E are the anomaly coefficients for the gluon and photon couplings respectively. To simplify the notation in the crucial second term in Eq. 2, we have absorbed N in the definition of the f_a. The rational number E/N therefore sets the model dependent strength of the coupling to photon pairs relative to that to the gluons. The gluon coupling dominates the ALP width as long as $E/N \lesssim 8(\alpha_s/\alpha_{\rm em})^2 \simeq 4 \times 10^3$, where the factor of 8 accounts for the color factor enhancing the gluon partial width. This upper bound is certainly satisfied in model of grand unification where $E/N \simeq \mathcal{O}(1)$ (see for instance [24]). Moreover, for low $\sqrt{\hat{s}}$, the gluon-gluon luminosity at the LHC is so enormous that there is still an appreciable cross section for $pp \to a + X$ (where X stands for one or more jets) even for extremely high values of the decay constant f_a. In part of the parameter space, this means that a can decay through displaced hadronic vertex, while still being abundantly produced at the LHC in processes such as $pp \to a + j$.

The event yield for our proposed trigger is shown in Fig. 1 for both models, where we have fixed the quartic λ_{SH} by choosing the branching ratio of $h \to SS$ to be 1%. This benchmark value for the exotic Higgs branching ratio will be difficult to exclude just with precision measurements of the Higgs couplings at the HL-LHC [25–27]. We find that up to 10^6 events could be recorded, and the trigger could cover between 3 and 4 orders of magnitude in the mixing angle s_θ. Hunting directly for exotic Higgs decays in this manner provides then an experimental avenue that is very complementary with the Higgs precision program.

The event yield for the ALP model on the other hand is probably marginal, due to the still rather short lifetime of the ALP. This likely generalizes to other scenarios for which the LLP is heavier than the B-meson and for which both production and decay controlled by a single interaction. In this sense, the ALP model is likely already the most optimal benchmark for this class of models, because of the dominance of the gluon parton distribution functions for low partonic center of mass energies. On the other hand, the ALP model is still a useful proxy for less minimal models, in which the lifetime and production cross section of the LLP are not controlled by a single parameter. Indeed, both models are to be viewed as straw man models for more complete scenarios which address the naturalness of the electroweak scale, the strong CP problem, the origin of the Dark Matter or the present baryonic asymmetry. In Sec. 3 we explain how the phenomenology of these more complete frameworks can be mapped onto our simple benchmark models, hereby establishing the strong theory motivation for this experimental effort. In Secs. 4 and 5 we discuss respectively our simulation framework and background calculations, arguing that a rate well below 10 kHz can likely be achieved with a minimal set of cuts. We close with additional results and an outlook in Sec. 6.

2. Benchmark models

In this section we provide more details on both benchmark models, with a special focus on their most natural parameter space and their various production modes.

2.1. A light real singlet

The first model is the most minimal extension of the SM, with only one new, real degree of freedom S, one mass parameter (m_S) and two coupling constants (μ, λ_{HS}). The lagrangian is given in Eq. 1 and the resulting mass squared matrix of the Higgs-singlet system is

$$\mathcal{M} = \begin{pmatrix} m_h^2 & v(\mu + \lambda_{SH} s_0) \\ v(\mu + \lambda_{SH} s_0) & \tilde{m}_S^2 + \frac{1}{2}\lambda_{SH} v^2 + V''_{\rm int}(s_0) \end{pmatrix}, \tag{3}$$

where in the limit of small mixing, $\mathcal{M}_{11} \simeq m_h^2 = 2\lambda_H v^2$ and $\mathcal{M}_{22} \simeq m_S^2$ are the mass eigenstates corresponding to the SM Higgs and the singlet S. s_0 is the singlet vacuum expectation value (VEV) and $v = 246$ GeV is the VEV of the SM Higgs.

Here we are interested in the light singlet regime where $m_S < m_h/2$, which means that the mixed quartic coupling will induce the exotic Higgs decay $h \to SS$ with branching ratio

$$\text{BR}[h \to SS] \approx \frac{\Gamma_{h\to SS}}{\Gamma_{h\to b\bar{b}}} \approx \frac{\lambda_{SH}^2}{6 y_b^2 \lambda_H}, \tag{4}$$

where we expanded for $m_S \ll m_h$. $y_b \simeq 2 \times 10^{-2}$ and $\lambda_H \simeq 0.13$ are respectively the bottom Yukawa and SM Higgs quartic couplings. The extremely narrow width of the SM Higgs implies that $\text{BR}[h \to SS] = 0.01$, as assumed in Fig. 1, corresponds to a small quartic

$$\lambda_{SH} \approx 1.7 \times 10^{-3}. \tag{5}$$

The value of this quartic has two important theoretical consequences that we illustrate in turn.

First, the quartic back-reacts on the SM Higgs potential and sets the cutoff Λ of a possible, natural UV completion for the lagrangian in Eq. 1. The most divergent contribution is the infamous mass correction to the SM Higgs mass

$$\delta m_H^2 \sim \frac{\lambda_{SH}}{16\pi^2}\Lambda^2, \tag{6}$$

where Λ is a proxy for any expected mass threshold of the unspecified UV theory. Setting $\delta m_H^2 \sim m_H^2$ implies $\Lambda \lesssim 38$ TeV which is easily outside the reach of the LHC. In other words, the hierarchy problem introduced by coupling S to the Higgs is much less severe then hierarchy problem corresponding to the SM top quark. A similar cut-off dependence arises in the analogous correction to $\tilde{m}_S$, but it can easily be UV completed without directly observable consequences.

Second, a non-zero λ_{SH} induces an irreducible, tree-level contribution to the singlet mass in Eq. 3. This means that fine tuning in the lower right block of Eq. 3 is needed if $\frac{1}{2}\lambda_{SH} v^2 \gg m_S^2$. Numerically, we can write m_S as

$$m_S \approx 2\text{ GeV} \times \left(\frac{\text{BR}[h \to SS]}{0.01}\right)^{1/4} \times \left(\frac{\Delta}{0.1}\right)^{1/2}, \tag{7}$$

where $\Delta \equiv m_S^2/(\frac{1}{2}\lambda_{SH} v^2)$ is a measure of the degree of fine tuning in $\mathcal{M}_{22}$. In other words, for $m_S \lesssim 2$ GeV, more than 10% fine tuning is needed given our benchmark value for λ_{SH}. This lower bound leaves most of the relevant parameter space in Fig. 1 open.

The remaining couplings of S to the SM are induced by the off-diagonal terms in Eq. 3, which mix the singlet with the SM Higgs with a small angle θ. The dependence of the mixing angle on the underlying theory parameters depends on whether the approximate $S \leftrightarrow -S$ parity is primarily broken by the μ parameter in Eq. 1, or by the dynamics of $V_{\rm int}(S)$. We will refer to these two cases as *induced* $\mathbb{Z}_2$ *breaking* and *direct* $\mathbb{Z}_2$ *breaking* respectively. Using the common shorthand notation $s_\theta = \sin\theta$, we can write

$$s_\theta \simeq \frac{|\mathcal{M}_{12}|}{m_h^2} = \begin{cases} \frac{\mu v}{m_h^2}\left|1 - \frac{\lambda_{SH} v^2}{2m_S^2}\right|, & s_0 = -\frac{\mu v^2}{2m_S^2} \\ \frac{\lambda_{SH}}{\sqrt{2}\lambda_S}\frac{v m_S}{m_h^2}, & s_0 = \frac{m_S}{\sqrt{2}\lambda_S} \end{cases}, \tag{8}$$

where the first and second line refer respectively to the induced and direct $\mathbb{Z}_2$ breaking. The second equation on the second line serves as the definition of the effective singlet quartic λ_S. By taking $\mu \ll v$, the induced $\mathbb{Z}_2$ breaking scenario allows for arbitrarily small values of the mixing angle s_θ, without fine tuning. This is simply a manifestation of the fact that μ is a technically natural parameter. In the direct $\mathbb{Z}_2$ breaking scenario on the other hand, the singlet quartic λ_S is bounded from above by perturbativity, which implies a lower bound[1] on s_θ

$$s_\theta \gtrsim 10^{-6} \times \left(\frac{m_S}{1\text{ GeV}}\right) \times \left(\frac{4\pi}{\sqrt{\lambda_S}}\right), \tag{9}$$

as indicated by the diagonal yellow line on Fig. 1.

Assuming no other dark sector decay channels are open, the mixing angle s_θ fixes the singlet width as

$$\Gamma_S(m_S) = s_\theta^2 \times \Gamma_h(m_S), \tag{10}$$

with $\Gamma_h(m_S)$ the width of the SM Higgs evaluated at mass m_S. Especially for $m_S \sim$ GeV, the form of $\Gamma_h(m_S)$ is very complicated and subject to large theory uncertainties. For $m_S \lesssim 1.0$ GeV, the width can be estimated using chiral perturbation theory and dispersion methods, as most recently updated by Winkler [29]. For $m_S \gtrsim 2$ GeV, perturbative calculations can be used, though still subject to sizable uncertainties. In the intermediate regime we must currently rely on an interpolation. In Fig. 2 we provide a minor update to the analysis in [29] by including the NLO QCD correction to the width and the running of quark masses, as calculated in [28]. This has an $\mathcal{O}(1)$ effect on the upper boundary condition of the interpolation region, and thus also has some impact on the regime where interpolation is needed.

Direct singlet production in association with a jet does not lead to interesting displaced jet signatures: The signal strength is only large enough for $s_\theta \gtrsim 10^{-4}$ and in most of this regime the singlet decays more or less promptly. We therefore find no substantial viable region associated with the $pp \to Sj$ process, in sharp contrast to the ALP case. Concretely, from the left-hand panel of Fig. 2 we see that the partial width to gluons is subdominant for $m_S \gtrsim 4$ GeV, which means that for a fixed gluon coupling (i.e. for a fixed direct production at the LHC) the scalar will decay more promptly than the ALP. This results in the small but non-zero event yield shown in the right-hand panel of Fig. 1. We will return to this argument in Sec. 2.2.

Below the B meson threshold, the light singlet S can be probed through exotic B decays. One possible channel is $B \to SX_s$, generated through an electroweak penguin [31–33] and with signal strength controlled by s_θ. A second channel is $B \to SSX_s$, also generated through an electroweak penguin [34] but with signal strength controlled by λ_{SH}. The branching ratios of the two processes are

$$\text{BR}(B \to SX_s) \simeq 3 \times 10^{-8}\left(\frac{s_\theta}{10^{-4}}\right)^2, \tag{11}$$

$$\text{BR}(B \to SSX_s) \simeq 3 \times 10^{-11}\left(\frac{\lambda_{\rm SH}}{10^{-3}}\right)^2, \tag{12}$$

where we neglected $\mathcal{O}(1)$ factors distinguishing between exclusive and inclusive channels [35]. For the values of the quartic considered here, the dominant channel is $B \to SX_s$ as long as $s_\theta \gtrsim 3 \times 10^{-6}$. As shown in Fig. 1, the exclusive searches for $B^+ \to K^+\mu^+\mu^-$ [10] and $B^0 \to K^*\mu^+\mu^-$ [36] at LHCb already exclude the region where $s_\theta > 10^{-3}$ and further improvements are

[1] Two upper bounds on s_θ can also be derived from *(i)* insisting that $\mathcal{M}$ in Eq. 3 is a positive definite matrix and *(ii)* from the naturalness of radiative correction to λ_S, induced by λ_{SH}. While neither is relevant for the parameter space considered here, they are important for very low mass scalars [30].

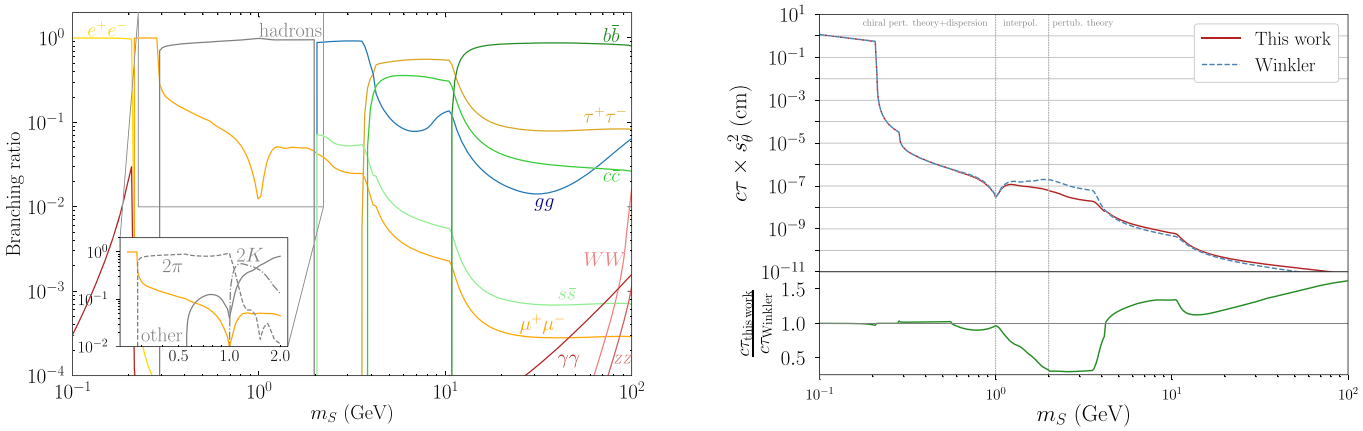

Fig. 2. Left: Estimated branching ratios for the scalar benchmark, derived from perturbative calculations in [28] and dispersive calculations in [29]. **Right:** Lifetime for the scalar benchmark, in comparison with the result in [29].

expected from the full LHCb dataset and from Belle2 [13,37]. The implementation of a displaced dimuon vertex trigger at CMS would moreover allow to probe the region further down to $s_\theta \simeq 10^{-4}$ [6,8]. Probing lower mixing angles would require dedicated experiments such as FASER 2 [14], MATHUSLA [15], CODEX-b [38] or SHiP [39]. (The latter, not shown on Fig. 1 is similar to the reach of MATHUSLA, but restricted to $m_S < m_B - m_K$.)

There also are a number of searches which can probe this model above the B meson threshold, for our specific choice λ_{SH}. (We refer to [40] for a more complete study of the full parameter space.) Concretely, there are two searches for displaced objects which do already have (limited) sensitivity: ATLAS performed a powerful search for anomalous activity in the muon chambers, using a dedicated trigger ("mu-ROI" in Fig. 1) [11]. In addition, the CMS search relying on an H_T cut plus two displaced vertices in the tracker ("H_T+2DV") [12] has some sensitivity for the higher end of the m_S range. For phase II, it has been argued that the timing detector and high granularity calorimeters will add substantial additional sensitivity [41,42]. A dedicated displaced vertex trigger would be complementary to all these strategies, as it can probe shorter lifetimes than the muon chambers and would have higher signal yield than triggers relying on MET, H_T or a VBF tag.

Aside from the aforementioned displaced searches, one may wonder whether various prompt searches already have sensitivity in the high s_θ region. In this regime, low mass dimuon resonances [43] can be sensitive to single S production through gluon fusion. Similarly, there exist a number of relevant searches for prompt exotic higgs decays, most notably in the $h \to b\bar{b}\mu\mu$ channel [44]. In all such cases one however must either pay the branching ratio of $S \to \mu\mu$ (see Fig. 2) at least once in the signal rate or content with very large hadronic backgrounds. As a result, we find that these searches are not yet constraining the parameter space in Fig. 1.

Before concluding this section it is important to note that the formulas derived in this section assume that the singlet couplings to the SM are exactly those of a lighter SM Higgs boson universally rescaled by the mixing angle s_θ. Large deviations from this assumption occur in models where the coupling of the singlet to the SM is driven purely by higher dimensional operators. A prominent example would be the one of a light dilaton [45,46] where the coupling to fermions can be suppressed compared to the one into gluons by carefully engineering how the SM yukawas are generated by the conformal sector dynamics [47,48].

2.2. A light ALP

The second benchmark model, in Eq. 2, describes the SM interactions of a pseudo-Nambu-Goldstone boson (pNGB) of a spontaneously broken Peccei-Quinn (PQ) symmetry which has a non-zero mixed anomaly with the QCD gauge group [49,50]. The spontaneous breaking of such a symmetry is a necessary requirement in axion solutions of the SM strong CP problem, and predicts an ALP with anomalous coupling to gluons. The anomalous ALP coupling to photons is controlled by the mixed anomaly with the $U(1)_{\text{em}}$ and will also be generically non-zero. As a consequence of the approximate shift symmetry acting on the ALP, the hierarchy $m_a \ll f_a$ is technically natural. We briefly discuss two different UV completions of the ALP lagrangian in Eq. 2, which differ by the type of states we expect to be present in the UV.

In the class of UV completions put forward by Kim, Shifman, Vainshtein and Zakharov (KSVZ) [51,52], the ALP is embedded in a complex scalar singlet Φ which couples to heavy, colored fermions. The latter are charged under the PQ symmetry such that the model is defined by

$$\mathcal{L}_{\text{KSVZ}} \supset g_* \Phi \tilde{\psi} \psi \quad , \quad \Phi = \frac{v_a + \varphi}{\sqrt{2}} e^{ia/v_a} \, , \tag{13}$$

where in this setup the singlet VEV (v_a) is related to the axion decay constant, defined in Eq. 2, by $v_a = 2Nf_a$. The anomaly coefficients N and E can be related to the multiplicity and the representation of the fermions: $N = q_{\text{PQ}} \sum_\psi C_3(R_\psi)$ and $E = q_{\text{PQ}} \sum_\psi Q^2(R_\psi)$, where C_3 is the color index of the fermion representation $\text{Tr}(R^a_\psi R^b_\psi) = C_3(R_\psi)\delta^{ab}$ and $Q(R_\psi)$ its electromagnetic charge. We can set the PQ charge $q_{\text{PQ}} = 1$ without loss of generality. As a well-motivated example, we take N_ψ flavors of heavy fermions in the $5 + \bar{5}$ representation of the $SU(5)$ GUT group, which implies $N = N_\psi/2$ and $E = 4/3N_\psi$. The coupling g_* sets the mass of the fermions $m_\psi = \sqrt{2}Ng_*f_a$.

At fixed f_a, a larger multiplicity N_ψ or a strong coupling g_* will make the colored fermions heavier, leaving the light ALP as the main phenomenological target of this setup. In the right-hand panel of Fig. 1 we indicate where the colored fermions would lie below 2 TeV for g_* fixed at its unitary bound $g_* = 4\pi/\sqrt{N_\psi}$, assuming $N_\psi = 1$. As we see, ALP searches into diphotons will be able to probe a portion of the parameter space where the colored fermions are out of reach of the LHC.

In models following the Dine-Fischler-Srednicki-Zhitnitsky (DFSZ) setup [53,54] the SM Higgs sector is extended to a two Higgs doublet model (H_u, H_d), plus a complex scalar singlet (Φ). The singlet couples to the two electroweak Higgs doublets through

$$\mathcal{L}_{\text{DFSZ}} \supset \lambda_* H_u H_d \Phi^2 \, , \tag{14}$$

fixing again the PQ charge of the singlet to be $q_{\text{PQ}} = 1$ the PQ charges of the two Higgses are fixed in this model: $q_{H_u} =$

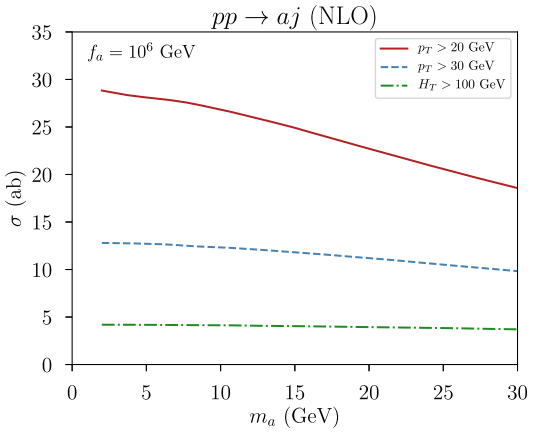

Fig. 3. Production cross section for an ALP in association with a jet, computed with Madgraph@NLO [56,57] for different fiducial cuts. The axion decay constant is fixed here to $f_a = 10^6$ GeV which is the typical size of the decay constant required for the ALP to decay displaced (see Eq. 17).

$-2\cos^2\beta$, $q_{H_d} = -2\sin^2\beta$, where $\tan\beta \equiv v_u/v_d$ with $v_{u,d}$ the VEVs of $H_{u,d}$. The DFSZ ALP couples to the SM fermions through their couplings to the Higgs doublets. Here we take a type II two Higgs doublet as an example. The SM fermions that are heavier than the ALP can be integrated out and contribute to the gluon and the photon anomaly, while the couplings to the light quarks and leptons are a new phenomenological feature of this setup. For $m_b < m_a < m_t$ we have $N = \cos\beta^2$ and $E = 8/3\cos\beta^2$ from the top contribution. The axion decay constant f_a is directly related to the mass of a heavy, CP-even Higgs bosons in this setup through $m_H \sim 2\sqrt{\lambda_*} f_a$. The LHC bounds on these extra Higgs doublets are mild relative to the bounds on new, colored fermions, which means that the DFSZ ALPs can generally have a lower decay constant than KSVZ ALPs. In Fig. 1 we indicate the regime where the CP Higgs bosons would roughly lay below 500 GeV, again saturating the unitarity bound $\lambda_* = 16\pi^2$.

The ALP width into two gluons and two photons is given by

$$\Gamma_{gg} = K_{gg}\frac{\alpha_s^2}{32\pi^3}\frac{m_a^3}{f_a^2} \quad , \quad \Gamma_{\gamma\gamma} = \frac{E^2}{N^2}\frac{\alpha_{\rm em}^2}{256\pi^3}\frac{m_a^3}{f_a^2} \, , \tag{15}$$

where K_{gg} is the k-factor accounting for NNLO corrections to the gluon width [55]

$$K_{gg} = 1 + \frac{\alpha_s}{\pi}E_A + \left(\frac{\alpha_s}{\pi}\right)^2 E_A\left(\frac{3}{4}E_A + \frac{\beta_1}{\beta_0}\right) \, , \tag{16}$$

where $E_A = 97/4 - 7N_f/6$, $\beta_0 = 11/4 - N_f/6$ and $\beta_1 = 51/8 - 19N_f/24$ and α_s is evaluated at m_a. N_f is the number of quarks with mass below m_a. Numerically, the k-factor ranges from 3.5 at $m_a = 5$ GeV to 2.5 for $m_a = 20$ GeV. From Eq. 15 we see explicitly that the gluon width dominates the total width as long as $E/N \lesssim K_{gg} 8(\alpha_s/\alpha_{\rm em})^2 \simeq 10^4$. The axion lifetime therefore scales as

$$c\tau_a \simeq 0.2 \text{ cm}\left(\frac{f_a}{10^6 \text{ GeV}}\right)^2\left(\frac{10 \text{ GeV}}{m_a}\right)^3 \, , \tag{17}$$

where in this simplified formula we neglected the k-factor dependence on the ALP mass. We see that for a decay constant as large as 10^6 GeV the ALP can produce a displaced vertex.

The signal events for the ALP model where generated at leading order with Madgraph [56] and decayed with Pythia 8 [58]. We normalized this sample to fixed order[2] NLO cross section for $pp \to a + j$, as computed with Madgraph@NLO [56,57]. This is shown in Fig. 3, where for reference we also show the cross sections for the fiducial cuts used in [7]. The corresponding k-factor for this process is $\simeq 2$ in the full mass range under consideration.[3] Fig. 3 and Eq. 17 make it clear that at most $\mathcal{O}(100)$ long-lived ALPs could be produced at the HL-LHC, which explains the limited signal yield in Fig. 1.

We now comment on exotic Higgs decays into a pair of ALPs. These are induced by the dimension six operator

$$\mathcal{L}_{\text{h-ALP}} = \frac{c_{\rm aH}}{N^2} H^\dagger H \frac{(\partial_\mu a)^2}{f_a^2} \, , \tag{18}$$

which respects the ALP shift-symmetry. The branching ratio is

$$\text{BR}[h \to aa] \approx \frac{\Gamma_{h\to aa}}{\Gamma_{h\to b\bar{b}}} = \frac{2c_{\rm aH}^2}{3N^4}\frac{v^4}{f_a^4}\frac{\lambda_H}{y_b^2} \, , \tag{19}$$

which implies that $\text{BR}[h \to aa] \simeq 0.01$ can be obtained for $f_a \simeq 1$ TeV and $c_{aH} \simeq \mathcal{O}(1)$. For values of the ALP decay constant which leads to displaced jet signatures (see Eq. 17), the signal rate from Higgs decays is very suppressed unless the decay constant controlling the gluon width is taken to be very different from the one controlling the operator in Eq. 18. Achieving $c_{\rm aH} \gg N^2$ is in principle possible since this hierarchy is protected by a symmetry, however it is fair to say that vanilla ALP scenarios would not lead to Higgs decay to long-lived ALPs. It is worth noting that this is very different from the singlet case, because the ALP couples to the SM mainly derivatively and its mixing with the Higgs is therefore parametrically suppressed by m_a^2/f_a^2. In practical terms, this means that a substantial model building effort is needed to enhance $c_{\rm aH}/N^2$ in Eq. 18. In the singlet case on the other hand, we have shown in the previous section that a large hierarchy between λ_{SH} and s_θ is completely natural.

Finally, the ALP could also be produced in B decays, in particular in DFSZ style models, see e.g. [59,60]. (In KSVZ style models this process takes place either at two loops [61] or through the ALP-WW coupling [62], and the reach is somewhat less promising.) Due to the B meson's large cross section and exceptionally small width, it is possible to produce a sizable sample of long-lived ALPs at the LHC through an exotic B meson decay. A displaced, hadronic vertex trigger would have the advantage that one does not need to pay the branching ratio of the ALP to muons, which tends to be very small in most models. We leave this case for a future study.

3. Theory motivation

In this section we review the theory motivation for the simplified benchmark models in the previous section, and low mass, displaced vertices in general. Some of the models we review make a concrete prediction in the parameter space shown in Fig. 1.

3.1. Hierarchy problem

Neutral naturalness models The non-observation of colored top partners at the LHC has revived the interest in a class of models where the top partners are charged under a dark $SU(3)$, rather than under the SM color group. The great ancestors of these models are the Twin Higgs [63,64], folded supersymmetry [65] and the quirky little higgs [66], which can all be seen as examples of a broader class of "neutral naturalness" models [67–70]. All these

[2] We verified that matching up to two jets does not lead to relevant differences in efficiency and worked for simplicity with the fixed order cross section.

[3] Notice that our k-factor differs from the one of Ref. [7], which was taken from the NLO corrections to $\sigma(pp \to a)$. The latter process has a very small efficiency and it is less well suited to assign a reliable k-factor.

models relax the fine-tuning of the electroweak scale by introducing an extended Higgs sector together with non-trivial dynamics in a hidden sector. As a consequence, an experimental smoking gun of these scenarios are exotic higgs decays into a dark parton shower, followed by a dark hadronization process. The resulting spray of dark sector hadrons can decay to the standard model through displaced vertices. In this sense, neutral naturalness models are themselves examples of a broader class of a hidden valley models [71,72]. (See [73] for a recent review of dark shower phenomenology.)

The phenomenology of the fraternal Twin Higgs model [69] in particular was studied in great detail [69,74] and can be summarized as follows: The SM Higgs decays to dark sector bottom quarks (b') with branching ratio $\mathrm{Br}[h \to b'b'] \approx v^2/f^2$, where f is the VEV of the twin Higgs.[4] The b' subsequently shower and hadronize, and eventually produce a number of dark glueballs. The lightest glueball is known to be a CP-even scalar [76] and can be identified with S in the simplified model in Sec. 2.1. Because most $h \to b'b'$ decays are expected to produce one or more 0^{++} glueballs, we roughly identify $\mathrm{Br}[h \to b'b'] \sim \mathrm{Br}[h \to SS]$. With this identification, the slice of the parameter space plotted in Fig. 1 corresponds to $v/f = 1/10$, which is considered to be rather fine tuned region of the Twin Higgs construction already. In other words, the event rate in the natural parameter space of the fraternal Twin Higgs is likely higher than the rate we assumed in Fig. 1.

The glueball lifetime is controlled by its mixing with the SM Higgs, which is given by [69]

$$s_\theta \simeq \frac{3.06 m_0^3 v}{24\pi^2 f^2 m_h^2}\,, \tag{20}$$

so that at fixed f, the mixing is predicted in terms of the mass of the glueball (m_0). This relation is shown by the yellow line in Fig. 1. Along the dotted line, m_0 is too low to be accommodated in the fraternal twin Higgs without substantial fine-tuning. It can be realized however by adding more matter to the dark sector, as in the vector-like twin Higgs [70].

Crucially, the presence of a dark shower decouples the glueball's production and decay rates, such that its lifetime can be parametrically enhanced without suppressing its production rate. Our identification of this rather complicated scenario with the simple benchmark model in Sec. Eq. 2.1 necessarily neglects $\mathcal{O}(1)$ theory uncertainties in the event yield. However these uncertainties do not affect the qualitative result that a displaced vertex trigger would be a major asset in the search for neutral naturalness models.

Composite Higgs models In composite Higgs models, the Higgs boson is a pseudo Nambu-Goldstone boson (pNGB) of a global symmetry, which is broken spontaneously by a strong sector in the UV. Depending on the size of the symmetry group in the model, additional light pNGB's may be present, which one can identify with the ALP benchmark. The ALP decay constant f_a is of the same order as the compositeness scale of the model, which must be $\sim$ few TeV to preserve naturalness. As we have seen in Sec. 2, for $f_a \sim 1$ TeV, the ALP is copiously produced at the LHC, though it typically decays promptly. Searches for prompt ALPs decaying to $\gamma\gamma$ or $\tau\tau$ therefore tend to be the most promising, be it in direct production [20,21,77] or in the exotic $h \to aa$ decay [78].

On the other hand, arguments have been presented in favor of pushing the compositeness scale to $\gtrsim 10$ TeV, which reduces tension with electroweak precision bounds, at the cost of moderate amount of fine tuning [79]. In this case, any additional pNGBs aside the Higgs could have macroscopic lifetimes [80].

Supersymmetry A generic pNGB arising in low energy SUSY-breaking is the R-axion related to the spontaneous breaking of the $U(1)$ R-symmetry [81,82]. The R-symmetry is universally associated to the $N = 1$ SUSY algebra as the only abelian symmetry that does not commute with the supercharges and has to be broken to realize a phenomenologically viable SUSY spectrum with (large) Majorana gaugino masses. The R-axion can behave like a DFSZ ALP or a KSVZ ALP, depending on the R-charge assignment of the MSSM fields [82]. Moreover, the values of the R-axion decay constant can be related to the mass of the messengers, which can easily be of the order of 10^3 TeV in vanilla gauge mediation scenarios [83,84].

Some SUSY scenarios also predict exotic Z^0 decays. One of the most straightforward ones arises in R-parity violating (RPV) supersymmetry with a light, bino-like lightest supersymmetric particle (LSP). If the colored superpartners are inaccessible at the LHC, the bounds on such a $\tilde{\chi}_1^0$ can be very mild, and even $m_{\tilde{\chi}_1^0} \sim 1$ GeV is still allowed. However, to prevent such a light $\tilde{\chi}_1^0$ from overclosing the universe, R-parity must be violated, rendering $\tilde{\chi}_1^0$ unstable [85]. The decay of $\tilde{\chi}_1^0$ will go through a heavy, off-shell sfermion, and therefore often occurs displaced, even for relatively large RPV couplings. In addition, the current LHC bounds permit a higgsino as light as ~ 150 GeV [86], as long as the wino is kinematically inaccessible. In this scenario, the mixing of the $\tilde{\chi}_1^0$ with the higgsino can be large enough to induce an observable branching ratio for $Z^0 \to \tilde{\chi}_1^0 \tilde{\chi}_1^0$ [87,88]. Given that the final states could be fully hadronic and that the partonic center-of-mass energy is merely $\sim m_{Z^0}$, this scenario will benefit greatly from a dedicated displaced vertex trigger.

Cosmological solutions We first discuss the relaxion mechanism, which is the prototypical cosmological solution to the hierarchy problem [89]. In this setup, the VEV of the SM Higgs is driven towards a small value by cosmological evolution of a light scalar field, the relaxion. Relaxion models behaves parametrically as the singlet models with induced $\mathbb{Z}_2$-symmetry breaking discussed in Sec. 2. Moreover, a single scale μ_b controls both the quartic coupling with the Higgs and the relaxion mass, such that $m_\phi^2 \sim \lambda_{SH} v^2$ with $\lambda_{SH} = \mu_b^2/f^2$ where f is the relaxion decay constant. For the value of λ_{SH} chosen in Fig. 1, the relaxion mass will be roughly around ~ 10 GeV. Moreover, requiring the relaxion to address the (little) hierarchy problem one obtains

$$\sin\theta = \frac{m_\phi^2}{m_h^2}\frac{f}{v} \gtrsim 2.6 \times 10^{-1} \left(\frac{m_\phi}{10\ \mathrm{GeV}}\right)^2 \left(\frac{f}{10\ \mathrm{TeV}}\right), \tag{21}$$

where f needs to be larger than the UV cutoff of the Higgs sector, which we took to be 10 TeV in this expression. The above equation shows that a successful relaxion mechanism is a fortiori associated to a promptly decaying relaxion for $m_\phi \gtrsim 1$ GeV. This feature is even more pronounced if one accounts for the suppression of the relaxion mass compare to its mixing with the Higgs [30]. On the other hand, for $m_\phi \lesssim 1$ GeV CMS may have some sensitivity in the displaced dimuon channel, by leveraging the production through exotic B-meson decay.

A different cosmological solution of the hierarchy problem has recently been proposed in Ref. [90]. Here, the idea is that the Higgs mixes with a dilaton of a conformal field theory (CFT) which has a large, negative vacuum energy. The flat dilaton potential, gets modified by the Higgs dynamics if the Higgs VEV is small enough and develops a new vacuum where the cosmological evolution is standard. Hubble patches with a large Higgs VEV on the other hand

[4] The Twin Higgs itself could also be within the reach of the LHC, depending the value of the twin quartic [75].

will crunch immediately under the influence of the CFT's large negative vacuum expectation value, and thus provide a form of anthropic selection. The existence of our Universe today can then be related to the smallness of the Higgs VEV. The light dilaton in [90] can be accessible at colliders, and its preferred mass range lies exactly where a displaced vertex trigger could probe interesting new parameter space.

3.2. Strong CP problem

Here we comment on the possibility that the ALP phenomenology described in Sec. 2 could arise in QCD axion models addressing the strong CP problem. While they are elegant and predictive solutions to the strong CP problem, axion models generally suffer from what is known as the *axion quality problem*. The quality problem states that any sources of explicit breaking of the PQ symmetry must be extremely small compared to the contribution from the QCD sector. In particular, quantum gravity is expected to break the PQ symmetry at the Planck scale ($M_{\rm Pl}$) with an order one coefficient. A successful solution of the strong CP problem thus implies an upper bound on the axion decay constant

$$f_a \lesssim \left(10^{-10} m_a^2 M_{\rm Pl}^{\Delta-4}\right)^{\frac{1}{\Delta-2}} , \tag{22}$$

with Δ the dimension of the leading operator which breaks the PQ symmetry. Using the relation $m_a \simeq \Lambda_{\rm QCD}^2/f_a$ this implies

$$f_a \lesssim 10^{-\frac{10}{\Delta}} \left(\frac{\Lambda_{QCD}}{M_{\rm Pl}}\right)^{\frac{4}{\Delta}} M_{\rm Pl} \lesssim 10^5\,{\rm GeV}, \tag{23}$$

where in the last inequality we assumed $\Delta = 6$, in analogy to baryon number violation in the SM. This regime however suffers from the rather unfortunate problem that it is already experimentally excluded by rare meson decays [91] and the non-observation of anomalous stellar cooling [92–94]. This leaves us with two possible avenues to salvage the axion solution to the strong CP problem: The most common approach is build more elaborate models in the UV which forbid or suppress all dangerous operators (see e.g. [95–98]), effectively raising Δ. This is the idea underlying all searches for the low mass QCD axions, e.g. using high quality resonance cavities.

The second option is to attempt to raise m_a by breaking the $m_a \simeq \Lambda_{\rm QCD}^2/f_a$ relation. This can be accomplished by adding a heavy dark sector which also contributes to the axion mass. This is still a non-trivial model building exercise, as the phase of the new contribution to the axion potential must be perfectly aligned with that of the QCD contribution. Perhaps the simplest possibility is to introduce a mirror QCD sector [99–103]. The possibility of probing this particular heavy axion model at the LHC was recently studied in Ref. [7]. Alternative models modifies the running of the QCD coupling constant at high energies [104–109] or embed the color group in a larger product group [110]. Once accomplished however, the strong CP problem can be solved with $m_a \gtrsim$ GeV and an f_a which could be within reach of the LHC. The green curves in Fig. 1 indicate the upper bounds on f_a from Eq. 22, assuming dimension $\Delta = 5$ and $\Delta = 6$.

3.3. Dark matter

The diversity of the SM forces can be taken as a motivation to go beyond the standard WIMP paradigm by introducing new forces that could control the Dark Matter freeze-out process. Both the ALP and singlet benchmark model are frequently used as new force mediators in this context, see e.g. [18,21,111–114]. If the mediator can decay to the Dark Matter, the signal at the LHC will typically be MET, associated with jets, leptons, photons or electroweak gauge bosons, depending on the model. In this case, a robust prediction of the signal rate is possible. If the mediator is lighter than the Dark Matter, it does typically decay back to the SM but there is no longer a direct connection between the mediator's coupling to SM and the Dark Matter relic density. This is the scenario towards which a displaced vertex trigger can contribute.

As long as the mediator is in thermal contact with the SM during the Dark Matter freeze-out epoch, it has the potential to play a role in setting the Dark Matter relic density. For the singlet benchmark, the most important process is $SS \leftrightarrow f\bar{f}$ scattering through an off-shell Higgs boson, with f the SM fermions. Concretely, we find that S always maintains thermal equilibrium with the SM as long as there is a fermion species f for which

$$m_\phi \gtrsim \frac{4\pi m_h^4}{\lambda_{SH}^2 m_f^2 M_{\rm Pl}} \qquad \text{and} \qquad m_\phi > m_f , \tag{24}$$

with $M_{\rm Pl}$ the Planck mass. This is easily satisfied everywhere in the parameter space in Fig. 1, regardless of the value of s_θ. The ALP model is able to maintain thermal equilibrium with the SM down to $T \sim m_a$ as long as

$$f_a \lesssim \frac{\alpha_s}{2\pi}\sqrt{m_a M_{\rm Pl}} , \tag{25}$$

which is also satisfied everywhere in Fig. 1.

An exhaustive review of all relevant models that can explain the Dark Matter abundance is well beyond our scope, and we instead conclude this general discussion by mentioning a few selected examples. An elegant and minimal example was provided by Evans et al. [18], where they extend the scalar benchmark in Eq. 1 with a Majorana Dark Matter particle (χ) which couples to the singlet S as

$$\mathcal{L}_{S\chi} \supset \frac{1}{2} y_\chi S \chi\chi , \tag{26}$$

and obtains its mass from the vacuum expectation value of S. Their model has 4 independent parameters, m_χ, m_S, y_χ and s_θ, as they set $\mu = 0$ in Eq. 1. This corresponds to the direct $\mathbb{Z}_2$ breaking scenario in Sec. 2.1. Here we will slightly generalize their setup by allowing for $\mu \neq 0$, and thus gain λ_{SH} as a 5$^{\rm th}$ independent parameter. One can fix y_χ by requiring that the $\chi\chi \to SS$ process reproduces the correct Dark Matter relic abundance, using the calculations in [18]. In Fig. 1 we have also fixed $m_\chi/m_S = 3$. λ_{SH} was fixed by our choice of BR$[h \to SS]$, as explained in Sec. 2.1. The lower and upper orange curves saturate respectively the XENON1T limit [17] and the neutrino floor. The orange shaded region in Fig. 1 is therefore ruled out for this model with these particular choice for m_S/m_χ. Interestingly, this also implies the region that will be probed by the next generation of large Dark Matter detectors coincides with the region that could be accessible to CMS. While the model does not predict that the Dark Matter must live in this specific region of parameter space, the prospect of a possible double discovery is nevertheless exciting.

In a similar spirit, the ALP model can be easily extended to a minimal model of Dark Matter freeze-out. For instance, we may add a Dirac Dark Matter candidate, charged under the $U(1)_{\rm PQ}$ in the KSVZ model of Eq. 13 [21]. The interaction is then

$$\mathcal{L}_{a\chi} \supset y_\chi \Phi \tilde{\chi}\chi + \text{h.c.} \tag{27}$$

where the mass of the Dark Matter is $m_\psi = \sqrt{2} y_\chi N f_a$ with N the anomaly coefficient[5] as defined in Sec. 2.2. The dominant

[5] The appearance of N in this equation is an artefact of the normalization we chose for f_a in Eq. 2.

Dark Matter annihilation processes are $\chi\chi \to aa$ and $\chi\chi \to gg$. Whichever dominates, depends on the value of N. For $N = 1/2$, corresponding to one heavy flavor, and fixing $y_\chi < 4\pi$, we indicate in Fig. 1 the range of f_a where the correct Dark Matter relic abundance can be obtained.

Both benchmarks can also play a role in more exotic freeze-out mechanisms: For example, D'Angolo et al. have shown that sub-GeV thermal Dark Matter candidates can exist through the coannihilation [113] or coscattering [114] of a compressed multiplet of dark sector particles, for which they used the singlet benchmark model as a mediator. These mechanisms are also well known in the context of for higher mass Dark Matter candidates, and their displaced signatures at the LHC were studied in detail in [115–117]. Given the compression of the spectrum, a displaced vertex trigger may well contribute substantially for Dark Matter candidates well above 100 GeV, something we leave for a future study.

3.4. Baryogenesis

The presence of out-of-equilibrium dynamics in the early universe is necessary condition for generating the observed baryon asymmetry. It is therefore natural to consider a relatively small coupling in a hidden sector, which can be responsible for a particle species decaying out-of-equilibrium before the onset of Big Bang Nucleosynthesis (BBN). While a baryogenesis mechanisms at very high energy scales can be fairly easily constructed, there are a number of schemes which are explicitly tied to energy scales that are accessible to the LHC. In those models, displaced vertices are a generic expectation.

For example, in WIMP baryogenesis [118,119] an electroweak state freezes out in the early universe, much like a classic WIMP. Unlike the standard WIMP however, it is allowed to decay to a SM state carrying baryon number and another stable particle, which will be the Dark Matter. If this decay occurs sufficiently slowly, the baryon asymmetry can be generated. As it turns out, the necessary lifetimes correspond to displaced decays at the LHC. So far the main focus has been on high mass scenarios, partially inspired by supersymmetry, for which the MET and/or H_T triggers are adequate. A displaced vertex trigger could however open a complementary, low mass parameter space.

In a more recent example, the baryon asymmetry is generated through the CP-violating oscillations of heavy flavor SM baryons [120,121] or mesons [122,123]. (See [124] for a supersymmetric implementation.) Concretely, the first ingredient is a massive, non-relativistic particle, which decays out-of-equilibrium to the SM b quarks at a low temperature. The b quarks subsequently hadronize to heavy flavor mesons and hadrons, which start to oscillate due to the CP violating phase in the SM CKM matrix. Finally, a new, exotic decay mode of the B-mesons into SM baryons and a dark sector then generates the asymmetry. The dark sector states can be stable, or decay back to the SM, depending on the specific model. In the latter case, its decay must however occur through an operator with rather high dimension, which predicts a macroscopic lifetime. This scenario is likely difficult to detect, since the unstable hidden sector state must be lighter than m_B. On the other hand, the experimental constraints on such a decay are currently very limited and given the huge $b\bar{b}$ cross section, a very large number of events could be collected by a displaced vertex trigger, even if the trigger efficiency itself is relatively small.

4. Simulation framework

For this study we rely on a toy simulation of the CMS L1 track trigger, as developed and described in [3,6], with minor modifications, which we describe below.

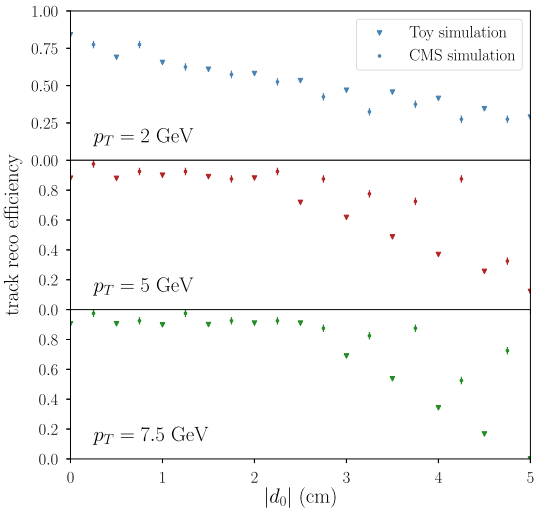

Fig. 4. The track reconstruction efficiency of our toy detector simulation as a function of the transverse impact parameter of the track ($|d_0|$), compared with the efficiency found by the CMS collaboration [4]. Tracks were generated with a particle gun, uniform in $|\eta| < 1$, and were required to have 5 reconstructed stubs. The error bars on the CMS result represent the bin size in which the efficiency was reported.

4.1. Track reconstruction

Our toy tracking detector consists out of 6 cylindrical layers located at the radii of the CMS phase II outer tracker [125], for $|\eta| < 2.4$. A track is propagated in the 3.8 T magnetic field, starting from the location of the decay vertex from which the track originated. In a realistic experimental setup, a charged particle would undergo multiple scattering in the detector material, slightly deflecting its trajectory and hereby degrading the vertex resolution. To model this, we deflect each track when passing through a layer with an angle drawn from a gaussian distribution centered around zero with width $4 \cdot 10^{-4}/p_T$ [126,127]. For this purpose, also the 4 layers of the inner tracker have been included, though they are otherwise "dark" from the point of view of the L1 track trigger. See [6] for a quantitive estimate of the impact of multiple scattering on the vertex reconstruction in our toy simulation.

Once the track has been propagated, the intersections with all 6 layers of the outer tracker are found, and for each intersection the azimuthal offset is computed. Only "stubs" for which the offset is consistent with a prompt, $p_T > 2$ GeV track are kept. We demand that at least 5 stubs per track pass this selection. All such surviving stubs are subsequently smeared with the expected resolution for the layer in question [126]. Finally, the hits are fit to a 5 parameter helix, allowing the track to originate away from the beam line. The reconstruction efficiency is most impacted by the p_T of the track, its transverse impact parameter (d_0) and the number of stubs required in the track.

In Fig. 4 we used a particle gun to compare the track reconstruction efficiency obtained with our toy simulation with the efficiency reported by the CMS collaboration [4]. The agreement is excellent for low impact parameters, though the toy simulation is somewhat underperforming for high impact parameters. The signal trigger efficiencies we report in Sec. 6 are in this sense conservative estimates. Moreover, the background which is most sensitive to the tracking efficiency is the B meson background (Sec. 5.2). This background is however restricted to very low d_0, for which the toy simulation is performing well, as we will see in Fig. 6.

4.2. Vertex reconstruction

We employ a very simplified vertexing algorithm, deliberately avoiding any χ^2-minimizations and/or numerical solutions for intersecting trajectories, as these operations are likely computationally prohibitive at the L1 track trigger. After the tracks have been reconstructed, our simplified vertex algorithm proceeds as follows:

i) The two hardest tracks of the candidate vertex are selected to seed the algorithm.
ii) For these two tracks, their intersections in the transverse plane are found, if they exist. If the tracks do not intersect in the transverse plane, the transverse location of the candidate vertex is chosen to be the point on the line connecting the centers of the circles for which the distance to each circle is equal. The distance of a track to the candidate vertex in the transverse plane is recorded in the variable Δ_{xy}, which is set to 0 if intersections are found.
iii) For each candidate vertex in the transverse plane, the distance along the z-direction between both tracks is computed. If two intersections were found in step (ii), the one with the smallest distance in the z-direction is chosen, with the z-coordinate of the candidate vertex being the midpoint between both tracks. The distance in the z-direction between a tracks and the vertex candidate is recorded in the variable Δ_z.[6]
iv) Having defined the location of the candidate vertex by intersecting the two hardest tracks, we compute subsequently Δ_{xy} and Δ_z for the remaining tracks, without updating the vertex location. Only tracks satisfying $\Delta_{xy} < 0.1$ cm and $\Delta_z < 0.5$ cm are assigned to the vertex. We further demand that all tracks satisfy $|d_0| > 0.1$ cm, to remove tracks originating from the beam line.
v) Only vertices with 4 or more tracks are being kept.

Better performance can likely be obtained by numerically finding the most optimal vertex location using the information of all tracks associated with the vertex candidate, rather than just the two hardest tracks. While speed is of course not an issue for our simple toy simulation, it would however likely be the most important bottleneck in a real life experimental implementation. We therefore deliberately refrain from using numerical optimization routines, and for all steps laid out above, closed form analytic expressions exist as a function of the track parameters, substantially speeding up the algorithm.

With this toy vertexing algorithm we have however left one important question unanswered: How can one efficiently identify the set of tracks corresponding to a candidate vertex? This is important given that a sizable number of fake, displaced tracks is expected in each event, and the rapid combinatorical growth of the possible combinations could be a major obstruction to implementing even a primitive vertex algorithm on the trigger. A possible approach could be to select the hardest displaced track(s) in the event and define a corresponding region of interest in ϕ and η around the hardest track, excluding all other tracks in the event. In the necessary detail, this question is highly non-trivial, and can only be addressed adequately through a more detailed study within the CMS collaboration. It is not a priori obvious that a workable solution exists, subject to the latency constraints of the L1 track trigger. The purpose of our toy simulation is therefore to provide the motivation needed to justify allocating CMS resources to perform such a study.

[6] Note that in [6], the definitions of Δ_{xy} and Δ_z differ with a factor of 2 from the ones we use here, as both variables defined as the distances between the two tracks defining the vertex, rather than the distance of the track to the vertex.

5. Background rate

We consider three backgrounds which we expect to drive most of the rate for a displaced vertex trigger:

- Secondary vertices from long-lived SM hadrons interacting with the detector material,
- true displaced vertices from SM hadron decays and
- fake vertices from randomly crossing fake tracks.

Each background requires dedicated modeling, as we describe below. Given the total L1 bandwidth of ∼750 kHz after the phase II upgrade, we consider rates $\lesssim 1$ kHz as "acceptable" in this work.

5.1. Material interactions

In the current configuration of CMS, about 5% of all pions with $p_T \gtrsim 5$ GeV were found to create a secondary vertex in the tracker [128]. Extrapolating to HL-LHC conditions, this corresponds to a daunting rate of roughly 30 MHz. The selection criteria for the tracks in this measurement were however less stringent than those which will be imposed by the p_T-modules in the L1 track trigger: In [128] a vertex was required to have two (off-line) reconstructed tracks with $p_T > 0.5$ GeV each, while our selections demand at least *four* tracks with $p_T > 2$ GeV each. A reliable extrapolation from this data is therefore not feasible, and we must perform a simplified GEANT4 [129] simulation instead.

Concretely, we set up a particle gun where we fired π^+ particles at a slab of silicon with thickness of 1 cm, using the standard `FTFP BERT` physics list, based on the Fritiof [130–133] and Bertini intra-nuclear cascade [134–136] models. The particle gun was then used to compute the probability for the resulting secondary vertex to have at least four hard tracks, as a function of the incident pion energy. For a given pseudo-rapidity (η), this can then be converted in a p_T-dependent efficiency by rotating the system and rescaling the GEANT4 output to account for the true, η-dependent material budget of the detector [126]. The resulting efficiency curve is shown in the left-hand panel of Fig. 5 for two values of η and three p_T cuts on the outgoing tracks. For the loosest p_T requirement, the efficiency sharply turns on around 15 GeV, and plateaus around 30 GeV, where the differences between the η-ranges can be attributed to the differences in material budget. The plateau at high p_T arises because the relevant process is deep inelastic scattering on the nuclei. The cross section is therefore essentially geometric in this regime and thus independent of the momentum of the incoming pion.

To estimate the total trigger rate, we simulate the rate for stable SM hadrons with Pythia 8 [58,137], in the form of a weighted dijet sample with a minimal $p_T > 5$ GeV cut. The resulting spectrum was conservatively normalized to a cross section of 68 mbn, the total inelastic cross section at $\sqrt{s}$ =13 TeV [138]. The combined rate for all stable or long-lived SM hadrons is shown by the dashed histogram in the right-hand panel of Fig. 5.

To estimate rate at which SM hadrons produce secondary vertices which pass our selection cuts, we fold the efficiency obtained by the GEANT4 particle gun against the hadron spectrum. The resulting rates are shown in the solid histograms in the right-hand panel of Fig. 5. For a minimal p_T cut of 2 GeV per track, the integrated rate for this background is roughly 25 kHz. While this is likely somewhat higher than what one may be able to accommodate comfortably, we stress that this estimate is very conservative and should be understood as an estimated upper bound on the rate. In particular, *(i)* for this background, no track reconstruction efficiencies were accounted for, as this require the full CMS detector simulation. Instead, we assumed that every track above the threshold was reconstructed with 100% efficiency, *regardless of its*

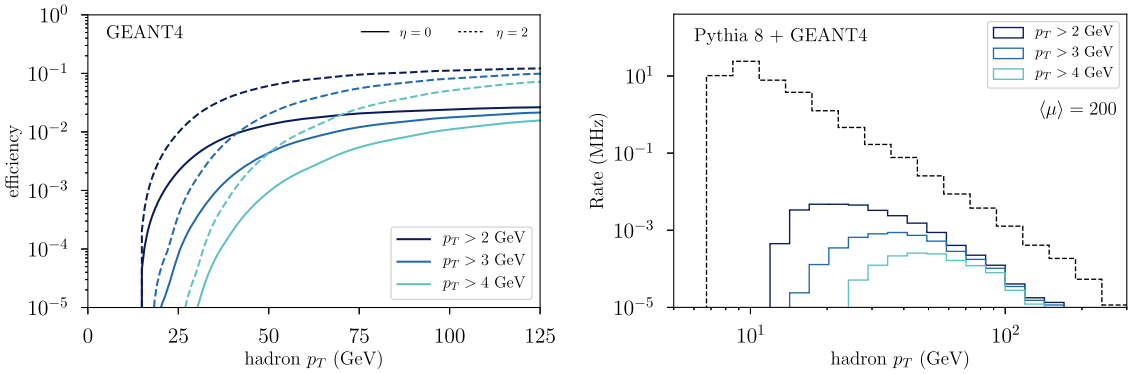

Fig. 5. *Left:* Efficiency for a hadron to create a secondary vertex with at least 4 tracks subject to the listed p_T cuts, as estimated with a GEANT4 particle gun. The dashed and solid lines represent different example values for the pseudo-rapidity of the incoming hadron, to illustrate the effect of the η-dependence of the material budget of the detector. (See text for details.) *Right:* Inclusive rate for long-lived SM hadrons with $|\eta| < 2.4$ (black dashed) and estimated rate for secondary vertices from a hadron interacting with detector material (solid) for different p_T cuts on the outgoing tracks. The integrated rates for $p_T > 2$ GeV, 3 GeV and 4 GeV are respectively 25 kHz, 5 kHz and 1 kHz.

impact parameters. We moreover assume that the vertex quality requirements in Sec. 4.2 are always satisfied for these background events. *(ii)* No material veto[7] was attempted and *(iii)* from Fig. 5 one can see that the rate is primarily driven by hadrons with $p_T \gtrsim 15$ GeV. Such hard hadrons will typically not be isolated, providing an additional handle to reduce this background if needed. Finally, the rate drops substantially when the threshold per track is raised slightly, to 5 kHz (1 kHz) for $p_T > 3$ GeV ($p_T > 4$ GeV). We therefore present our subsequent results for those 3 different threshold choices.

5.2. SM meson decays

First, we simulate an inclusive sample of K_S and K_L samples with Pythia 8 by generating a p_T-weighted dijet sample, where we subject the kaons to the fiducial cuts of $p_T > 8$ GeV and $|\eta| < 2.4$. Conservatively normalizing this sample to the total measured inelastic cross section of 68 mbn [138], this yields a rate of roughly 21 MHz for each species, which corresponds roughly one Kaon per event. The dominant decay modes which feature 4 charged final states are $K_S \to \pi^+\pi^- e^+e^-$ with branching ratio 4.79×10^{-5} and $K_L \to \pi^\pm e^\mp \nu e^+ e^-$ with branching ratio 1.26×10^{-5} [139]. Both modes therefore yield rates $<$ kHz, even before other fiducial and reconstruction efficiencies are factored in. We therefore focus on B-meson decays as the most dangerous hadron background.

We use Pythia 8 to simulate an inclusive sample of B-mesons, which are subsequently processed through the toy detector simulation described in Sec. 4. The sample is normalized to the inclusive $b\bar{b}$ cross section, as calculated with FONLL [140–143]. Demanding at least 4 tracks forming a vertex, the resulting rate is approximately 50 kHz. Most tracks originating from B decays however have small impact parameters, and a cut of $|d_0| > 1$ mm for at least 4 tracks in the vertex further drops the rate to roughly 130 Hz for the most loose p_T cut (see Fig. 6). The rate moreover drops substantially if the p_T cuts are tightened, as shown in Table 1. If needed, an additional cut on the transverse radius of the vertex (L_{xy}) is extremely effective at reducing this background (see e.g. [8]), at a minimal cost to the signal.

[7] Given that the pixel layers will not be available to the L1 trigger, the spacial resolution on the vertex location is expected to be relatively poor [6]. A material veto would therefore be less effective than in a full, offline analysis and we therefore chose to not rely on it in our study.

Table 1
Estimated rates for various backgrounds, for different choices of the minimum track p_T. The rate for the secondaries is to be understood as a conservative upper bound. See text for details.

min track p_T	2 GeV	3 GeV	4 GeV
secondaries (kHz)	25	5	1
B-mesons (kHz)	0.13	0.04	0.01
fake vertices (kHz)	0.04	0.01	0.004

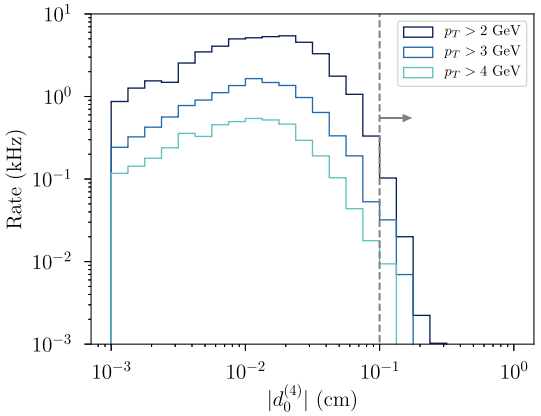

Fig. 6. Estimated L1 trigger rate from B-meson backgrounds, as a function of the transverse impact parameter of the 4$^{\text{th}}$ reconstructed track ($|d_0^{(4)}|$), where the tracks were ordered according to decreasing $|d_0|$. We require $|d_0^{(4)}| > 0.1$ cm, as indicated by the dashed vertical line.

5.3. Fake vertices

Finally, to estimate the probability of 4 fake tracks forming a vertex, we generated 10^7 pairs of fake tracks, uniformly distributed in the track parameters. (See [6] for details.) The pairs of fake tracks where processed through the toy vertex reconstruction algorithm and discarded if the Δ_{xy} and Δ_z parameters of the candidate vertex failed our vertex quality requirements in Sec. 4.2. The remaining two-track vertices are then pairwise compared with one another, where we select only those pairs for which the distance between both vertices in the transverse and longitudinal directions is compatible with the respective vertex resolution cuts. Finally, all remaining candidate four-track vertices are again processed through the algorithm in Sec. 4.2. With this method, the probability of 4 random fake tracks forming a vertex is found to be approximately 3×10^{-11}.

Based on the expected occupancy of the trigger system [144], we assume on average 30 fake tracks per event. We thus compute the average combinatorical enhancement factor per event to be

$$\left\langle \frac{n_{\text{fakes}}!}{(n_{\text{fakes}}-4)!4!} \right\rangle \approx 3.4 \times 10^4 \,, \tag{28}$$

where the brackets indicates the Poisson weighted average, assuming $\langle n_{\text{fakes}} \rangle = 30$. Combining this with the fake vertex probability, we find the rate to about 40 Hz or lower (See Table 1), giving a substantial buffer in case our estimate of 30 tracks per event proves to be insufficiently conservative.

Finally, one may be concerned that a significant contribution from the rate may come from real, 3-track vertices to which a fourth, fake track is erroneously assigned. With the assumptions laid out in Sec. 5.1 we find that the rate for a 3-track vertex from hadronic interactions with the detector material is roughly 50 kHz for a p_T cut on the tracks of 2 GeV. Following Sec. 5.2, the rate for a 3-track vertex from the $K^\pm \to \pi^\pm\pi^+\pi^-$ process is bounded from above by roughly 30 kHz. We hereby imposed a truth-level $p_T > 6$ GeV cut for the $K^\pm$ and required it to decay within 30 cm of the beamline, but no requirements were made on the kinematics of the individual pions. With $\mathcal{O}(30)$ fake tracks per event, the rate for both cases therefore remains below 1 kHz if the probability for a fake track to be associated with a real displaced vertex is below $\mathcal{O}(10^{-3})$. This appears plausible, but in particular for the vertices from material interactions, we cannot accurately estimate this probability without the full CMS detector simulation. Our estimates of the rate for real, 3-track vertices are moreover conservative, and in practice we expect that a $\mathcal{O}(10^{-2})$ suppression could be sufficient. Should this process nevertheless prove to be problematic, it can be suppressed further by tightening the p_T cuts on the tracks: In particular, by increasing the cut to 3 GeV (4 GeV), the rate for 3-track material interactions drops to 13 kHz (5 kHz). Similarly, increasing the p_T cut on the $K^\pm$ to 9 GeV (12 GeV) reduces the rate to 5 kHz (0.8 kHz).

6. Discussion and outlook

As we showed in the previous section, if a displaced vertex could be reconstructed at the trigger level, the $H_T > 100$ GeV requirement used in [3–5,7] may no longer be needed to bring the background rate down to a manageable level. This in turn substantially increases the signal efficiency for scenarios where the final states are rather soft.

Fig. 7 illustrates this point for the ALP benchmark, where we normalized all efficiencies to a sample with a 20 GeV cut on the truth-level p_T of the ALP. The trigger efficiency is $\mathcal{O}(10\%)$, which we consider fairly decent for a signal this soft. We find that a displaced vertex trigger would roughly be a factor of 3 improvement over the displaced jet trigger with $H_T > 100$ GeV, as studied Hook et al. [7].[8] Hook et al. also provided a second, more speculative selection consisting out of a single, displaced jet with $p_T > 30$ GeV. This selection unsurprisingly outperforms a displaced vertex trigger, as the former only requires 3 displaced tracks vs 4 displaced tracks for the latter. Given the high rate of $p_T > 30$ GeV jets, secondary vertices from material interactions and fake tracks, we however suspect that the background rate for this proposal could be prohibitive. In sum, Fig. 7 makes it clear that the reason for the rather low signal yield in the right-hand panel of Fig. 1 is not related to the low reconstruction efficiency per se. Rather, both the lifetime and the production cross section of the ALP are controlled by f_a and the lifetime is simply very short for values of f_a which yield interesting cross sections (see Eq. 17 and Fig. 3).

[8] We thank Soubhik Kumar for providing us with this unpublished data.

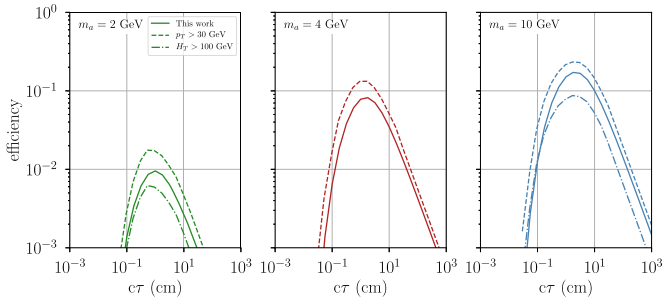

Fig. 7. Estimated L1 trigger efficiency for a CMS displaced vertex trigger for the ALP model (solid). Also shown are the estimates based on the $H_T > 100$ GeV and $p_T > 30$ GeV selections from Hook et al. [7] All efficiencies were normalized to a truth-level cross section with $p_T > 20$ GeV cut on the ALP. (See text for details).

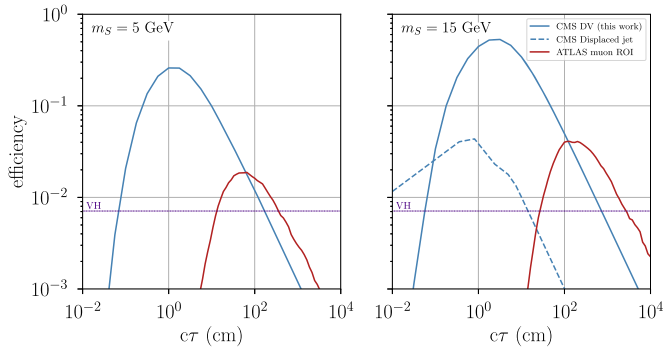

Fig. 8. Estimated L1 trigger efficiency for a CMS displaced vertex trigger (solid blue). Also shown is the projected efficiency for CMS L1 track jet trigger for displaced jets [3,4], assuming 5 kHz rate (dashed blue) and the efficiency of the existing ATLAS muon region of interest trigger (solid red) [11]. The dashed purple line indicates the efficiency to trigger on an associated lepton in VH production, normalized against the gluon fusion cross section.

This is in sharp contrast with the scalar model, where the lifetime and production rate are controlled by *independent* parameters, s_θ and λ_{SH} respectively. As a result, a displaced vertex trigger could open a large portion of the parameter space where the mass of the singlet is natural with respect to its quartic coupling with the SM Higgs. We showed that the values of the mixing angles that could be explored are further motivated by simple models of Dark Matter freeze-out and neutral naturalness scenarios.

The model independent trigger efficiency for the $h \to SS$ topology is shown Fig. 8, where a displaced vertex trigger would improve on the displaced jet + H_T trigger by more than an order of magnitude. For reference, we also include the efficiency for the existing ATLAS trigger which relies on anomalous activity in a region of interest in the muon chamber ("muon ROI") [11]. As expected, both approaches are highly complimentary, as a trigger on a displaced vertex in the tracker could cover shorter lifetimes.

We conclude by commenting on another opportunity for the track trigger, which we intend explore in upcoming work [145]. Traditionally, Dark Matter searches rely heavily on the MET trigger, which has low efficiency for low mass Dark Matter or soft production modes. This is particularly so for inelastic Dark Matter models, where the Dark Matter resides in a narrowly split multiplet. Inelastic Dark Matter models predict a soft displaced vertex with a moderate amount of MET, typically well below the current trigger thresholds. Models of heavy neutral leptons produce a similar signature. By identifying the displaced tracks with the track trigger, CMS may be able to substantially lower the MET requirement, while maintaining a manageable background rate. The implementation of such a trigger may allow CMS to probe a large portion

of the parameter space of these models, complementing existing proposals [146,147].

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

We are grateful to Brando Bellazzini, Matthew Citron, Soubhik Kumar, Jared Evans, Sam Junius, Zhen Liu, Gaia Lianfranchi, Alberto Mariotti, Matthew McCullough, Toby Opferkuch, Simone Pagan Griso, Michele Papucci, Dean Robinson, Filippo Sala, Ennio Salvioni, Jessie Shelton and Scott Thomas for useful discussions. We thank especially Gauthier Durieux, Davide Pagani, Pier Monni and Marco Zaro for important clarifications on the QCD corrections to the cross section and for help with *MadGraph5@NLO*. YG was supported by the National Science Foundation under grant nr. 1913356.

References

[1] K. Einsweiler, et al., ATLAS collaboration, Technical Design Report for the ATLAS Inner Tracker Pixel Detector, Tech. Rep. CERN-LHCC-2017-021. ATLAS-TDR-030, CERN, Geneva, 2017.
[2] CMS Collaboration, The Phase-2 Upgrade of the CMS Level-1 Trigger, Tech. Rep. CERN-LHCC-2020-004. CMS-TDR-021, CERN, Geneva, 2020, final version.
[3] Y. Gershtein, C.M.S. Hardware, Track trigger: new opportunities for long-lived particle searches at the HL-LHC, Phys. Rev. D 96 (2017) 035027, arXiv:1705.04321 [hep-ph].
[4] CMS Collaboration, First Level Track Jet Trigger for Displaced Jets at High Luminosity LHC, Tech. Rep. CMS-PAS-FTR-18-018, Geneva, 2018.
[5] B. Bhattacherjee, S. Mukherjee, R. Sengupta, P. Solanki, Triggering long-lived particles in HL-LHC and the challenges in the first stage of the trigger system, J. High Energy Phys. 08 (2020) 141, arXiv:2003.03943 [hep-ph].
[6] Y. Gershtein, S. Knapen, Trigger strategy for displaced muon pairs following the CMS phase II upgrades, Phys. Rev. D 101 (2020) 032003, arXiv:1907.00007 [hep-ex].
[7] A. Hook, S. Kumar, Z. Liu, R. Sundrum, High quality QCD axion and the LHC, Phys. Rev. Lett. 124 (2020) 221801, arXiv:1911.12364 [hep-ph].
[8] J.A. Evans, A. Gandrakota, S. Knapen, H. Routray, Searching for exotic B meson decays with the CMS L1 track trigger, arXiv:2008.06918 [hep-ph], 2020.
[9] S. Foroughi-Abari, A. Ritz, LSND constraints on the Higgs portal, Phys. Rev. D 102 (2020) 035015, arXiv:2004.14515 [hep-ph].
[10] R. Aaij, et al., LHCb collaboration, Search for long-lived scalar particles in $B^+ \to K^+ \chi(\mu^+\mu^-)$ decays, Phys. Rev. D 95 (2017) 071101, arXiv:1612.07818 [hep-ex].
[11] M. Aaboud, et al., ATLAS, Search for long-lived particles produced in pp collisions at $\sqrt{s} = 13$ TeV that decay into displaced hadronic jets in the ATLAS muon spectrometer, Phys. Rev. D 99 (2019) 052005, arXiv:1811.07370 [hep-ex].
[12] A.M. Sirunyan, et al., CMS, Search for long-lived particles using displaced jets in proton-proton collisions at $\sqrt{s} = 13$ TeV, arXiv:2012.01581 [hep-ex], 2020.
[13] A. Kachanovich, U. Nierste, I. Nišandžić, Higgs portal to dark matter and $B \to K^{(*)}$ decays, Eur. Phys. J. C 80 (2020) 669, arXiv:2003.01788 [hep-ph].
[14] J.L. Feng, I. Galon, F. Kling, S. Trojanowski, Dark Higgs bosons at the ForwArd search ExpeRiment, Phys. Rev. D 97 (2018) 055034, arXiv:1710.09387 [hep-ph].
[15] C. Alpigiani, et al., MATHUSLA, An update to the letter of intent for MATHUSLA: search for long-lived particles at the HL-LHC, arXiv:2009.01693 [physics.ins-det], 2020.
[16] G. Aielli, et al., Expression of interest for the CODEX-b detector, arXiv:1911.00481 [hep-ex], 2019.
[17] E. Aprile, et al., XENON, Dark matter search results from a one ton-year exposure of XENON1T, Phys. Rev. Lett. 121 (2018) 111302, arXiv:1805.12562 [astro-ph.CO].
[18] J.A. Evans, S. Gori, J. Shelton, Looking for the WIMP next door, J. High Energy Phys. 02 (2018) 100, arXiv:1712.03974 [hep-ph].
[19] G. Aad, et al., ATLAS, Search for scalar diphoton resonances in the mass range $65 - 600$ GeV with the ATLAS detector in pp collision data at $\sqrt{s} = 8\ TeV$, Phys. Rev. Lett. 113 (2014) 171801, arXiv:1407.6583 [hep-ex].
[20] A. Mariotti, D. Redigolo, F. Sala, K. Tobioka, New LHC bound on low-mass diphoton resonances, Phys. Lett. B 783 (2018) 13, arXiv:1710.01743 [hep-ph].
[21] X. Cid Vidal, A. Mariotti, D. Redigolo, F. Sala, K. Tobioka, New axion searches at flavor factories, J. High Energy Phys. 01 (2019) 113, Erratum: J. High Energy Phys. 06 (2020) 141, arXiv:1810.09452 [hep-ph].
[22] A.M. Sirunyan, et al., CMS, Search for low mass vector resonances decaying into quark-antiquark pairs in proton-proton collisions at $\sqrt{s} = 13$ TeV, J. High Energy Phys. 01 (2018) 097, arXiv:1710.00159 [hep-ex].
[23] J. Lees, et al., BaBar, Search for hadronic decays of a light Higgs boson in the radiative decay $\Upsilon \to \gamma A^0$, Phys. Rev. Lett. 107 (2011) 221803, arXiv:1108.3549 [hep-ex].
[24] G.F. Giudice, R. Rattazzi, A. Strumia, Unificaxion, Phys. Lett. B 715 (2012) 142, arXiv:1204.5465 [hep-ph].
[25] P. Bechtle, S. Heinemeyer, O. Stål, T. Stefaniak, G. Weiglein, Probing the standard model with Higgs signal rates from the tevatron, the LHC and a future ILC, J. High Energy Phys. 11 (2014) 039, arXiv:1403.1582 [hep-ph].
[26] G. Belanger, B. Dumont, U. Ellwanger, J. Gunion, S. Kraml, Global fit to Higgs signal strengths and couplings and implications for extended Higgs sectors, Phys. Rev. D 88 (2013) 075008, arXiv:1306.2941 [hep-ph].
[27] C. Frugiuele, E. Fuchs, G. Perez, M. Schlaffer, Relaxion and light (pseudo)scalars at the HL-LHC and lepton colliders, J. High Energy Phys. 10 (2018) 151, arXiv:1807.10842 [hep-ph].
[28] M. Spira, QCD effects in Higgs physics, Fortschr. Phys. 46 (1998) 203, arXiv:hep-ph/9705337.
[29] M.W. Winkler, Decay and detection of a light scalar boson mixing with the Higgs boson, Phys. Rev. D 99 (2019) 015018, arXiv:1809.01876 [hep-ph].
[30] A. Banerjee, H. Kim, O. Matsedonskyi, G. Perez, M.S. Safronova, Probing the relaxed relaxion at the luminosity and precision frontiers, J. High Energy Phys. 07 (2020) 153, arXiv:2004.02899 [hep-ph].
[31] R.S. Willey, H.L. Yu, Decays $K^\pm \to \pi^\pm l^+ l^-$ and limits on the mass of the neutral Higgs boson, Phys. Rev. D 26 (1982) 3287.
[32] R.S. Chivukula, A.V. Manohar, Limits on a light Higgs boson, Phys. Lett. B 207 (1988) 86.
[33] B. Grinstein, L.J. Hall, L. Randall, Do B meson decays exclude a light Higgs?, Phys. Lett. B 211 (1988) 363.
[34] C. Bird, P. Jackson, R.V. Kowalewski, M. Pospelov, Search for dark matter in b −> s transitions with missing energy, Phys. Rev. Lett. 93 (2004) 201803, arXiv:hep-ph/0401195.
[35] I. Boiarska, K. Bondarenko, A. Boyarsky, V. Gorkavenko, M. Ovchynnikov, A. Sokolenko, Phenomenology of GeV-scale scalar portal, J. High Energy Phys. 11 (2019) 162, arXiv:1904.10447 [hep-ph].
[36] R. Aaij, et al., LHCb, Search for hidden-sector bosons in $B^0 \to K^{*0}\mu^+\mu^-$ decays, Phys. Rev. Lett. 115 (2015) 161802, arXiv:1508.04094 [hep-ex].
[37] A. Filimonova, R. Schäfer, S. Westhoff, Probing dark sectors with long-lived particles at BELLE II, Phys. Rev. D 101 (2020) 095006, arXiv:1911.03490 [hep-ph].
[38] V.V. Gligorov, S. Knapen, M. Papucci, D.J. Robinson, Searching for long-lived particles: a compact detector for exotics at LHCb, Phys. Rev. D 97 (2018) 015023, arXiv:1708.09395 [hep-ph].
[39] S. Alekhin, et al., A facility to search for hidden particles at the CERN SPS: the SHiP physics case, Rep. Prog. Phys. 79 (2016) 124201, arXiv:1504.04855 [hep-ph].
[40] E. Fuchs, O. Matsedonskyi, I. Savoray, M. Schlaffer, Collider searches of scalar singlets across lifetimes, arXiv:2008.12773 [hep-ph], 2020.
[41] J. Liu, Z. Liu, L.-T. Wang, Enhancing long-lived particles searches at the LHC with precision timing information, Phys. Rev. Lett. 122 (2019) 131801, arXiv:1805.05957 [hep-ph].
[42] J. Liu, Z. Liu, L.-T. Wang, X.-P. Wang, Enhancing sensitivities to long-lived particles with high granularity calorimeters at the LHC, J. High Energy Phys. 11 (2020) 066, arXiv:2005.10836 [hep-ph].
[43] Search for a narrow resonance decaying to a pair of muons in proton-proton collisions at 13 TeV, Tech. Rep. CMS-PAS-EXO-19-018, CERN, Geneva, 2019.
[44] M. Aaboud, et al., ATLAS, Search for Higgs boson decays into a pair of light bosons in the $bb\mu\mu$ final state in pp collision at $\sqrt{s} = 13$ TeV with the ATLAS detector, Phys. Lett. B 790 (2019) 1, arXiv:1807.00539 [hep-ex].
[45] B. Bellazzini, C. Csaki, J. Hubisz, J. Serra, J. Terning, A naturally light dilaton and a small cosmological constant, Eur. Phys. J. C 74 (2014) 2790, arXiv:1305.3919 [hep-th].
[46] F. Coradeschi, P. Lodone, D. Pappadopulo, R. Rattazzi, L. Vitale, A naturally light dilaton, J. High Energy Phys. 11 (2013) 057, arXiv:1306.4601 [hep-th].
[47] B. Bellazzini, C. Csaki, J. Hubisz, J. Serra, J. Terning, A higgslike dilaton, Eur. Phys. J. C 73 (2013) 2333, arXiv:1209.3299 [hep-ph].
[48] Z. Chacko, R.K. Mishra, Effective theory of a light dilaton, Phys. Rev. D 87 (2013) 115006, arXiv:1209.3022 [hep-ph].
[49] R.D. Peccei, H.R. Quinn, Constraints imposed by CP conservation in the presence of instantons, Phys. Rev. D 16 (1977) 1791.
[50] R.D. Peccei, H.R. Quinn, CP conservation in the presence of instantons, Phys. Rev. Lett. 38 (1977) 1440.
[51] J.E. Kim, Weak interaction singlet and strong CP invariance, Phys. Rev. Lett. 43 (1979) 103.
[52] M.A. Shifman, A. Vainshtein, V.I. Zakharov, Can confinement ensure natural CP invariance of strong interactions?, Nucl. Phys. B 166 (1980) 493.

[53] A. Zhitnitsky, On possible suppression of the axion hadron interactions, Sov. J. Nucl. Phys. 31 (1980) 260 (in Russian).
[54] M. Dine, W. Fischler, M. Srednicki, A simple solution to the strong CP problem with a harmless axion, Phys. Lett. B 104 (1981) 199.
[55] K. Chetyrkin, B.A. Kniehl, M. Steinhauser, W.A. Bardeen, Effective QCD interactions of CP odd Higgs bosons at three loops, Nucl. Phys. B 535 (1998) 3, arXiv:hep-ph/9807241.
[56] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H.S. Shao, T. Stelzer, P. Torrielli, M. Zaro, The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations, J. High Energy Phys. 07 (2014) 079, arXiv:1405.0301 [hep-ph].
[57] P. Artoisenet, et al., A framework for Higgs characterisation, J. High Energy Phys. 11 (2013) 043, arXiv:1306.6464 [hep-ph].
[58] T. Sjostrand, et al., An introduction to PYTHIA 8.2, Comput. Phys. Commun. 191 (2015) 159, arXiv:1410.3012 [hep-ph].
[59] B. Batell, M. Pospelov, A. Ritz, Multi-lepton signatures of a hidden sector in rare B decays, Phys. Rev. D 83 (2011) 054005, arXiv:0911.4938 [hep-ph].
[60] M. Freytsis, Z. Ligeti, J. Thaler, Constraining the axion portal with $B \to Kl^+l^-$, Phys. Rev. D 81 (2010) 034001, arXiv:0911.5355 [hep-ph].
[61] D. Aloni, Y. Soreq, M. Williams, Coupling QCD-scale axionlike particles to gluons, Phys. Rev. Lett. 123 (2019) 031803, arXiv:1811.03474 [hep-ph].
[62] E. Izaguirre, T. Lin, B. Shuve, Searching for axionlike particles in flavor-changing neutral current processes, Phys. Rev. Lett. 118 (2017) 111802, arXiv:1611.09355 [hep-ph].
[63] Z. Chacko, H.-S. Goh, R. Harnik, The twin Higgs: natural electroweak breaking from mirror symmetry, Phys. Rev. Lett. 96 (2006) 231802, arXiv:hep-ph/0506256 [hep-ph].
[64] Z. Chacko, H.-S. Goh, R. Harnik, A twin Higgs model from left-right symmetry, J. High Energy Phys. 01 (2006) 108, arXiv:hep-ph/0512088 [hep-ph].
[65] G. Burdman, Z. Chacko, H.-S. Goh, R. Harnik, Folded supersymmetry and the LEP paradox, J. High Energy Phys. 02 (2007) 009, arXiv:hep-ph/0609152 [hep-ph].
[66] H. Cai, H.-C. Cheng, J. Terning, A quirky little Higgs model, J. High Energy Phys. 05 (2009) 045, arXiv:0812.0843 [hep-ph].
[67] N. Craig, S. Knapen, P. Longhi, The orbifold Higgs, J. High Energy Phys. 03 (2015) 106, arXiv:1411.7393 [hep-ph].
[68] N. Craig, S. Knapen, P. Longhi, Neutral naturalness from orbifold Higgs models, Phys. Rev. Lett. 114 (2015) 061803, arXiv:1410.6808 [hep-ph].
[69] N. Craig, A. Katz, M. Strassler, R. Sundrum, Naturalness in the dark at the LHC, J. High Energy Phys. 07 (2015) 105, arXiv:1501.05310 [hep-ph].
[70] N. Craig, S. Knapen, P. Longhi, M. Strassler, The vector-like twin Higgs, J. High Energy Phys. 07 (2016) 002, arXiv:1601.07181 [hep-ph].
[71] M.J. Strassler, K.M. Zurek, Echoes of a hidden valley at hadron colliders, Phys. Lett. B 651 (2007) 374, arXiv:hep-ph/0604261 [hep-ph].
[72] T. Han, Z. Si, K.M. Zurek, M.J. Strassler, Phenomenology of hidden valleys at hadron colliders, J. High Energy Phys. 07 (2008) 008, arXiv:0712.2041 [hep-ph].
[73] J. Alimena, et al., Searching for long-lived particles beyond the standard model at the large hadron collider, J. Phys. G 47 (2020) 090501, arXiv:1903.04497 [hep-ex].
[74] D. Curtin, C.B. Verhaaren, Discovering uncolored naturalness in exotic Higgs decays, J. High Energy Phys. 12 (2015) 072, arXiv:1506.06141 [hep-ph].
[75] S. Alipour-Fard, N. Craig, S. Gori, S. Koren, D. Redigolo, The second Higgs at the lifetime frontier, J. High Energy Phys. 07 (2020) 029, arXiv:1812.09315 [hep-ph].
[76] C.J. Morningstar, M.J. Peardon, The glueball spectrum from an anisotropic lattice study, Phys. Rev. D 60 (1999) 034509, arXiv:hep-lat/9901004.
[77] G. Cacciapaglia, G. Ferretti, T. Flacke, H. Serodio, Revealing timid pseudoscalars with taus at the LHC, Eur. Phys. J. C 78 (2018) 724, arXiv:1710.11142 [hep-ph].
[78] M. Bauer, M. Neubert, A. Thamm, Collider probes of axion-like particles, J. High Energy Phys. 12 (2017) 044, arXiv:1708.00443 [hep-ph].
[79] J. Barnard, P. Cox, T. Gherghetta, A. Spray Long-Lived, Colour-triplet scalars from unnaturalness, J. High Energy Phys. 03 (2016) 03, arXiv:1510.06405 [hep-ph].
[80] D. Curtin, et al., Long-lived particles at the energy frontier: the MATHUSLA physics case, Rep. Prog. Phys. 82 (2019) 116201, arXiv:1806.07396 [hep-ph].
[81] H-S. Goh, M. Ibe, R-axion detection at LHC, J. High Energy Phys. 03 (2009) 049, arXiv:0810.5773 [hep-ph].
[82] B. Bellazzini, A. Mariotti, D. Redigolo, F. Sala, J. Serra, *R*-axion at colliders, Phys. Rev. Lett. 119 (2017) 141804, arXiv:1702.02152 [hep-ph].
[83] Z. Komargodski, D. Shih, Notes on SUSY and R-symmetry breaking in Wess-Zumino models, J. High Energy Phys. 04 (2009) 093, arXiv:0902.0030 [hep-th].
[84] G.F. Giudice, R. Rattazzi, Theories with gauge mediated supersymmetry breaking, Phys. Rep. 322 (1999) 419, arXiv:hep-ph/9801271.
[85] P. Bechtle, et al., Killing the cMSSM softly, Eur. Phys. J. C 76 (2016) 96, arXiv:1508.05951 [hep-ph].
[86] M. Aaboud, et al., ATLAS, Search for electroweak production of supersymmetric states in scenarios with compressed mass spectra at $\sqrt{s} = 13$ TeV with the ATLAS detector, Phys. Rev. D 97 (2018) 052010, arXiv:1712.08119 [hep-ex].
[87] A. Bartl, W. Majerotto, N. Oshimo, On the production of neutralinos at the z and w and their decay into Higgs bosons, Phys. Lett. B 216 (1989) 233.
[88] J.C. Helo, M. Hirsch, Z.S. Wang, Heavy neutral fermions at the high-luminosity LHC, J. High Energy Phys. 07 (2018) 056, arXiv:1803.02212 [hep-ph].
[89] P.W. Graham, D.E. Kaplan, S. Rajendran, Cosmological relaxation of the electroweak scale, Phys. Rev. Lett. 115 (2015) 221801, arXiv:1504.07551 [hep-ph].
[90] C. Csáki, R.T. D'Agnolo, M. Geller, A. Ismail, Crunching dilaton, hidden naturalness, arXiv:2007.14396 [hep-ph], 2020.
[91] H. Georgi, D.B. Kaplan, L. Randall, Manifesting the invisible axion at low energies, Phys. Lett. B 169 (1986) 73.
[92] G. Raffelt, A. Weiss, Red giant bound on the axion - electron coupling revisited, Phys. Rev. D 51 (1995) 1495, arXiv:hep-ph/9410205.
[93] N. Viaux, M. Catelan, P.B. Stetson, G. Raffelt, J. Redondo, A.A.R. Valcarce, A. Weiss, Neutrino and axion bounds from the globular cluster M5 (NGC 5904), Phys. Rev. Lett. 111 (2013) 231301, arXiv:1311.1669 [astro-ph.SR].
[94] M.M. Miller Bertolami, B.E. Melendez, L.G. Althaus, J. Isern, Revisiting the axion bounds from the Galactic white dwarf luminosity function, J. Cosmol. Astropart. Phys. 10 (2014) 069, arXiv:1406.7712 [hep-ph].
[95] L. Randall, Composite axion models and Planck scale physics, Phys. Lett. B 284 (1992) 77.
[96] M. Redi, R. Sato, Composite accidental axions, J. High Energy Phys. 05 (2016) 104, arXiv:1602.05427 [hep-ph].
[97] L. Di Luzio, E. Nardi, L. Ubaldi, Accidental Peccei-Quinn symmetry protected to arbitrary order, Phys. Rev. Lett. 119 (2017) 011801, arXiv:1704.01122 [hep-ph].
[98] M. Duerr, K. Schmidt-Hoberg, J. Unwin, Protecting the axion with local Baryon number, Phys. Lett. B 780 (2018) 553, arXiv:1712.01841 [hep-ph].
[99] V. Rubakov, Grand unification and heavy axion, JETP Lett. 65 (1997) 621, arXiv:hep-ph/9703409.
[100] Z. Berezhiani, L. Gianfagna, M. Giannotti, Strong CP problem and mirror world: the Weinberg-Wilczek axion revisited, Phys. Lett. B 500 (2001) 286, arXiv:hep-ph/0009290.
[101] A. Hook, Anomalous solutions to the strong CP problem, Phys. Rev. Lett. 114 (2015) 141801, arXiv:1411.3325 [hep-ph].
[102] H. Fukuda, K. Harigaya, M. Ibe, T.T. Yanagida, Model of visible QCD axion, Phys. Rev. D 92 (2015) 015021, arXiv:1504.06084 [hep-ph].
[103] S. Dimopoulos, A. Hook, J. Huang, G. Marques-Tavares, A collider observable QCD axion, J. High Energy Phys. 11 (2016) 052, arXiv:1606.03097 [hep-ph].
[104] B. Holdom, M.E. Peskin, Raising the axion mass, Nucl. Phys. B 208 (1982) 397.
[105] K. Choi, C. Kim, W. Sze, Mass renormalization by instantons and the strong CP problem, Phys. Rev. Lett. 61 (1988) 794.
[106] B. Holdom, Strong QCD at high-energies and a heavy axion, Phys. Lett. B 154 (1985) 316, Erratum: Phys. Lett. B 156 (1985) 452.
[107] M. Dine, N. Seiberg, String theory and the strong CP problem, Nucl. Phys. B 273 (1986) 109.
[108] J.M. Flynn, L. Randall, A computation of the small instanton contribution to the axion potential, Nucl. Phys. B 293 (1987) 731.
[109] K. Choi, H.D. Kim, Small instanton contribution to the axion potential in supersymmetric models, Phys. Rev. D 59 (1999) 072001, arXiv:hep-ph/9809286.
[110] P. Agrawal, K. Howe, Factoring the strong CP problem, J. High Energy Phys. 12 (2018) 029, arXiv:1710.04213 [hep-ph].
[111] G. Krnjaic, Probing light thermal dark-matter with a Higgs portal mediator, Phys. Rev. D 94 (2016) 073009, arXiv:1512.04119 [hep-ph].
[112] S. Matsumoto, Y.-L.S. Tsai, P.-Y. Tseng, Light fermionic WIMP dark matter with light scalar mediator, J. High Energy Phys. 07 (2019) 050, arXiv:1811.03292 [hep-ph].
[113] R.T. D'Agnolo, C. Mondino, J.T. Ruderman, P.-J. Wang, Exponentially light dark matter from coannihilation, J. High Energy Phys. 08 (2018) 079, arXiv:1803.02901 [hep-ph].
[114] R.T. D'Agnolo, D. Pappadopulo, J.T. Ruderman, P.-J. Wang, Thermal relic targets with exponentially small couplings, Phys. Rev. Lett. 124 (2020) 151801, arXiv:1906.09269 [hep-ph].
[115] A. Filimonova, S. Westhoff, Long live the Higgs portal!, J. High Energy Phys. 02 (2019) 140, arXiv:1812.04628 [hep-ph].
[116] A. Bharucha, F. Brümmer, N. Desai, Next-to-minimal dark matter at the LHC, J. High Energy Phys. 11 (2018) 195, arXiv:1804.02357 [hep-ph].
[117] N. Nagata, H. Otono, S. Shirai, Probing Bino-Wino coannihilation at the LHC, J. High Energy Phys. 10 (2015) 086, arXiv:1506.08206 [hep-ph].
[118] Y. Cui, R. Sundrum, Baryogenesis for weakly interacting massive particles, Phys. Rev. D 87 (2013) 116013, arXiv:1212.2973 [hep-ph].
[119] Y. Cui, B. Shuve, Probing baryogenesis with displaced vertices at the LHC, J. High Energy Phys. 02 (2015) 049, arXiv:1409.6729 [hep-ph].
[120] D. McKeen, A.E. Nelson, CP violating baryon oscillations, Phys. Rev. D 94 (2016) 076002, arXiv:1512.05359 [hep-ph].
[121] K. Aitken, D. McKeen, T. Neder, A.E. Nelson, Baryogenesis from oscillations of charmed or beautiful baryons, Phys. Rev. D 96 (2017) 075009, arXiv:1708.01259 [hep-ph].
[122] G. Elor, M. Escudero, A. Nelson, Baryogenesis and dark matter from *B* mesons, Phys. Rev. D 99 (2019) 035031, arXiv:1810.00880 [hep-ph].
[123] A.E. Nelson, H. Xiao, Baryogenesis from B meson oscillations, Phys. Rev. D 100 (2019) 075002, arXiv:1901.08141 [hep-ph].

[124] G. Alonso-Álvarez, G. Elor, A.E. Nelson, H. Xiao, A supersymmetric theory of baryogenesis and sterile sneutrino dark matter from B mesons, J. High Energy Phys. 03 (2020) 046, arXiv:1907.10612 [hep-ph].
[125] The Phase-2 Upgrade of the CMS Level-1 Trigger, Tech. Rep. CERN-LHCC-2020-004. CMS-TDR-021, CERN, Geneva, 2020, final version.
[126] K. Klein, et al., CMS collaboration, The Phase-2 Upgrade of the CMS Tracker, Tech. Rep. CERN-LHCC-2017-009. CMS-TDR-014, CERN, Geneva, 2017.
[127] M. Tanabashi, et al., Particle Data Group, Review of particle physics, section 33.3 Phys. Rev. D 98 (2018) 030001.
[128] CMS Collaboration, Studies of Tracker Material, Tech. Rep. CMS-PAS-TRK-10-003, 2010.
[129] S. Agostinelli, et al., Geant4—a simulation toolkit, Nucl. Instrum. Methods Phys. Res., Sect. A, Accel. Spectrom. Detect. Assoc. Equip. 506 (2003) 250.
[130] B. Andersson, G. Gustafson, B. Nilsson-Almqvist, A model for low-pt hadronic reactions with generalizations to hadron-nucleus and nucleus-nucleus collisions, Nucl. Phys. B 281 (1987) 289.
[131] B. Andersson, A. Tai, B.-H. Sa, Final state interactions in the (nuclear) FRITIOF string interaction scenario, Z. Phys. C, Part. Fields 70 (1996) 499.
[132] B. Nilsson-Almqvist, E. Stenlund, Interactions between hadrons and nuclei: the Lund Monte Carlo, FRITIOF version 1.6, Comput. Phys. Commun. 43 (1987) 387.
[133] B. Ganhuyag, V. Uzhinsky, Modified FRITIOF code: negative charged particle production in high energy nucleus nucleus interactions, in: Relativistic Heavy-Ion Physics. Proceedings, International School/Workshop for Young Physicists, Prague, Czech Republic, September 2-6, 1996, Czechoslov. J. Phys. 47 (1997) 913.
[134] M.P. Guthrie, R.G. Alsmiller, H.W. Bertini, Calculation of the capture of negative pions in light elements and comparison with experiments pertaining to cancer radiotherapy, Nucl. Instrum. Methods 66 (1968) 29.
[135] H.W. Bertini, M.P. Guthrie, News item results from medium-energy intranuclear-cascade calculation, Nucl. Phys. A 169 (1971) 670.
[136] V.A. Karmanov, Light front wave function of relativistic composite system in explicitly solvable model, Nucl. Phys. B 166 (1980) 378.
[137] T. Sjostrand, S. Mrenna, P.Z. Skands, PYTHIA 6.4 physics and manual, J. High Energy Phys. 05 (2006) 026, arXiv:hep-ph/0603175 [hep-ph].
[138] M. Aaboud, et al., ATLAS, Measurement of the inelastic proton-proton cross section at $\sqrt{s} = 13$ TeV with the ATLAS detector at the LHC, Phys. Rev. Lett. 117 (2016) 182002, arXiv:1606.02625 [hep-ex].
[139] P.A. Zyla, et al., Review of particle physics, Prog. Theor. Exp. Phys. 2020 (2020) 083C01.
[140] M. Cacciari, M. Greco, P. Nason, The P(T) spectrum in heavy flavor hadroproduction, J. High Energy Phys. 05 (1998) 007, arXiv:hep-ph/9803400.
[141] M. Cacciari, S. Frixione, P. Nason, The p(T) spectrum in heavy flavor photoproduction, J. High Energy Phys. 03 (2001) 006, arXiv:hep-ph/0102134.
[142] M. Cacciari, S. Frixione, N. Houdeau, M.L. Mangano, P. Nason, G. Ridolfi, Theoretical predictions for charm and bottom production at the LHC, J. High Energy Phys. 10 (2012) 137, arXiv:1205.6344 [hep-ph].
[143] M. Cacciari, M.L. Mangano, P. Nason, Gluon PDF constraints from the ratio of forward heavy-quark production at the LHC at $\sqrt{S} = 7$ and 13 TeV, Eur. Phys. J. C 75 (2015) 610, arXiv:1507.06197 [hep-ph].
[144] D. Contardo, M. Klute, J. Mans, L. Silvestris, J. Butler, Technical Proposal for the Phase-II Upgrade of the CMS Detector, Tech. Rep. CMS-TDR-15-02, 2015.
[145] Y. Gershtein, S. Junius, S. Knapen, A. Mariotti, D. Redigolo, 2021, in preparation.
[146] E. Izaguirre, G. Krnjaic, B. Shuve, Discovering inelastic thermal-relic dark matter at colliders, Phys. Rev. D 93 (2016) 063523, arXiv:1508.03050 [hep-ph].
[147] A. Berlin, F. Kling, Inelastic dark matter at the LHC lifetime frontier: ATLAS, CMS, LHCb, CODEX-b, FASER, and MATHUSLA, Phys. Rev. D 99 (2019) 015021, arXiv:1810.01879 [hep-ph].