

Beam-dump ceiling and its experimental implications: The case of a portable experiment

Doojin Kim^{1,2,*}, Jaehoon Yu^{3,†}, Jong-Chul Park^{4,5,‡} and Hyunyong Kim^{6,2,§}

¹*Department of Physics, University of South Dakota, Vermillion, South Dakota 57069, USA*

²*Mitchell Institute for Fundamental Physics and Astronomy, Department of Physics and Astronomy, Texas A&M University, College Station, Texas 77843, USA*

³*Department of Physics, University of Texas, Arlington, Texas 76019, USA*

⁴*Department of Physics and Institute of Quantum Systems, Chungnam National University, Daejeon 34134, Republic of Korea*

⁵*Institute for Sciences of the Universe, Chungnam National University, Daejeon 34134, Republic of Korea*

⁶*Department of Physics and Astronomy, Seoul National University, Seoul 08826, Republic of Korea*



(Received 26 February 2024; accepted 21 April 2026; published 2 June 2026)

We generalize the nature of the so-called beam-dump “ceiling” beyond which the improvement on the sensitivity reach in the search for fast-decaying mediators dramatically slows down, and we point out its experimental implications that motivate tabletop-sized beam-dump experiments for the search. Light (bosonic) mediators are well-motivated new-physics particles, as they can appear in dark-sector portal scenarios and models to explain various laboratory-based anomalies. Due to their low mass and feebly interacting nature, beam-dump-type experiments, utilizing high-intensity particle beams, can play a crucial role in probing the parameter space of such visibly decaying mediators—in particular, the “prompt decay” region, where the mediators feature relatively large coupling and mass. We present a general and semianalytic proof that the ceiling effectively arises in the prompt-decay region of an experiment and show its insensitivity to data statistics, background estimates, and systematic uncertainties, considering a concrete example, the search for axion-like particles interacting with ordinary photons at three benchmark beam facilities: PIP-II at FNAL, and SPS and LHC-dump at CERN. We then identify optimal criteria to perform a cost-effective and short-term experiment to reach the ceiling, demonstrating that very short-baseline compact experiments enable access to the parameter space unreachable thus far.

DOI: [10.1103/5nst-v9lh](https://doi.org/10.1103/5nst-v9lh)

I. INTRODUCTION

Although the Standard Model (SM) provides a successful description of elementary particles and their mutual interactions, many issues remain unresolved. For example, the nonvanishing masses of neutrinos and the existence of dark matter are unexplained by the SM, demanding new physics beyond the SM. Furthermore, various experimental anomalies—e.g., ATOMKI [1–5], LSND [6], and Mini-BooNE [7–9], together with recent experimental investigations by MicroBooNE relevant to the MiniBooNE

excess (e.g., Refs. [10,11])—may hint at the existence of new physics.

Well-motivated explanations of these phenomena are deeply related to new particles that are very weakly or feebly interacting with the SM particles. In particular, new bosonic particles with MeV-range masses have been popular choices to address these observational anomalies and phenomena while satisfying the existing constraints. As a theoretical approach, the ideas of portal scenarios such as (pseudo)scalar and vector portals (see, e.g., Ref. [12]) promote these MeV-scale particles to the mediators through which the dark-sector particles, including dark matter, can communicate with the SM particles, and they often motivate MeV-scale light thermal dark matter. Therefore, searches for mediators of this sort can be regarded as an alternative avenue for dark-matter detection, especially in an accelerator-based experiment.

Considering the expected mass range of these particles (henceforth collectively called mediators), fixed-target experiments such as beam-dump-type experiments are better suited than the energy-frontier facilities like the

*Contact author: doojin.kim@usd.edu

†Contact author: jaehoon@uta.edu

‡Contact author: jcpark@cnu.ac.kr

§Contact author: hyunyong.kim@cern.ch

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

Large Hadron Collider (LHC) to probe MeV-scale physics, as their beam energy is as large as $\mathcal{O}(1\text{--}100\text{ GeV})$. Moreover, their high-instantaneous-intensity particle beams enable the production of feebly-interacting mediators. For these reasons, the search prospects of (MeV-scale) Higgs-portal scalars, axion-like particles (ALPs), light Z' , dark photons, etc., have been extensively investigated in the beam-dump-type experiments, especially in ongoing/upcoming/proposed experiments such as CCM, COHERENT, DUNE, FASER/FASER ν , ICARUS, MicroBooNE, SBND, and SHiP (see also, e.g., Refs. [12–16], and references therein).

A major goal of the experiments above is to expand the search as much as possible into the “prompt decay” region, which none of the existing laboratory-based experiments or astrophysical/cosmological considerations have been able to constrain; a mediator belonging to this region would decay rather earlier due to its (relatively) large coupling and/or large mass value. A recent ambitious proposal, DAMSA [17]—in which the detector system is placed extremely close to the signal source point (i.e., the beam target or dump)¹—showed the capability to start probing deeper into the prompt-decay region. Indeed, it has recently been demonstrated that for a given experimental configuration, there exists an inherent limitation dubbed the beam-dump “ceiling” beyond which the sensitivity improvement slows down dramatically, so that an ordinary beam-dump-type experiment would have difficulty probing without a nonsensical increase of data statistics [18]. The actual location of the ceiling crucially depends on the beam energy and the detector distance from the beam target [18].

An important implication of the above observations with the beam-dump ceiling is that a collection of large amounts of data is no longer necessary to reach sensibly achievable sensitivity limits. In this context, we delve into the robustness of the sensitivity reaches near the beam-dump ceiling against the overall/instantaneous beam intensity, the detector geometry (size and angular coverage), and underlying background/systematic uncertainty assumptions, taking three benchmark beam facilities: PIP-II [19] at FNAL, and SPS [20] and LHC-dump [21] at CERN. From this investigation, we will demonstrate that a *tabletop-sized* detector system can probe the prompt-decay region of MeV-scale mediators rather *quickly*. Therefore, one can imagine experimenting at one beam facility for a short period until the ceiling has been reached and taking the same detector system to a different beam facility to perform another experiment to expand the ceiling. Inspired by DAMSA [17], now being proposed at PIP-II LINAC, we propose the concept of a portable DAMSA that could subsequently be performed at other facilities with various beam energies.

¹From now on, we will use *beam target* and *dump* interchangeably, as we will show that the target width does not affect our results.

The rest of this paper is structured as follows: In Sec. II, we discuss the behavior of the beam-dump ceiling, examining in turn its dependence on the total beam intensity, backgrounds and systematic uncertainties, detector geometry, and instantaneous beam intensity. In Sec. III, we validate the predicted ceiling behavior for the three benchmark beam facilities introduced above, focusing on ALPs coupled to the SM photon. We begin with a theory-level comparison in Sec. III A, adopting a common experimental configuration, including the beam-target and detector geometries. In Sec. III B, we present sensitivity projections for the SPS and LHC-dump configurations, considering feasible experimental setups—compatible with existing facility designs—and assuming a conservative level of reducible background suppression. Our conclusions appear in Sec. IV.

II. BEHAVIOR OF BEAM-DUMP “CEILING”

A. Dependence on total beam intensity

We begin with a very rough sketch of the existing constraints of a generic (bosonic) mediator—say, φ (e.g., ALPs, dark photon, and light Z')—in Fig. 1. Regions I and II are constrained by collider and beam-dump searches, respectively, whereas region III is constrained by various astrophysical and cosmological considerations including SN1987a, cosmic microwave backgrounds, and big-bang nucleosynthesis. Region IV is open and can be reached by increasing the statistics. Our primary interest, however, is to extend the beam-dump reach to region V—i.e., the prompt-decay region mentioned earlier.

In generic beam-dump or fixed-target experiments, if a mediator is produced inside the dump or target, it should reach and decay into visible particles within the detector fiducial volume in order for it to be detected. This requirement can be formulated as a detection probability P_{det} in terms of the exponential decay law:

$$P_{\text{det}} = \exp\left(-\frac{L_{\text{dist}}}{\tilde{l}_{\varphi}}\right) \left[1 - \exp\left(-\frac{\Delta L_{\text{det}}}{\tilde{l}_{\varphi}}\right)\right], \quad (1)$$




FIG. 1. A schematic diagram of the existing constraints on generic mediators.

where L_{dist} and ΔL_{det} denote the detector distance from the signal production point and the detector size along the direction of the mediator, respectively. $\tilde{l}_\varphi$ is the mean decay length of φ in the laboratory frame. Note that the smaller the L_{dist} , the larger the P_{det} . By definition,

$$L_{\text{dist}}, \Delta L_{\text{det}} \gg \tilde{l}_\varphi \quad (2)$$

in the prompt-decay region, resulting in

$$P_{\text{det}} \approx \exp\left(-\frac{L_{\text{dist}}}{\tilde{l}_\varphi}\right). \quad (3)$$

Suppose that for a given mass m_φ , an experiment can reach a coupling value g with N source particles inducing the φ production. If data statistics increases (via beam intensity or beam exposure duration) by a factor of X , the new sensitivity reach g' is given by [18]

$$g' \approx g \left(1 + \frac{\log X}{\log(N\langle P_{\text{prod}} \rangle / N_{\text{req}}) - 1}\right)^{1/2}, \quad (4)$$

where N_{req} is the number of signal events required to define the sensitivity reach g , and $\langle P_{\text{prod}} \rangle$ is the average mediator production probability at g for the given source particle. For a large background, $\log X$ can be replaced with $\log X^{1/2}$.

In ongoing/planned beam-dump or fixed-target experiments, particle beams are highly intensified, hence N is very large. In the prompt-decay region, the coupling is not too small, and so is $\langle P_{\text{prod}} \rangle$. Finally, unless an experiment suffers from poor background rejection or poor understanding of systematic uncertainties, N_{req} need not be too large. Therefore, the denominator in Eq. (4) is larger than the numerator, as long as X is at a sensible level. This implies that any feasible level of an increase in data collection would not allow for an appreciable improvement in the sensitivity reach. This conversely implies that decreasing data collection would not result in a considerable degradation in the sensitivity reach, as long as $|\log X| \ll \log(N\langle P_{\text{prod}} \rangle / N_{\text{req}}) - 1$. The overall decrease in data collection can be accomplished in two ways:

- (a) smaller beam exposure, and
- (b) smaller beam power.

B. Dependence on background and systematic uncertainty estimates

A similar exercise can be done for the case where a different number of signal events N''_{req} is required to determine the sensitivity reach. For example, one could incorrectly estimate backgrounds and/or systematics so that the true sensitivity reach would have been g'' . We find that g'' is expressed as

$$g'' \approx g \left(1 + \frac{\log(N''_{\text{req}}/N_{\text{req}})}{\log(N\langle P_{\text{prod}} \rangle / N_{\text{req}}) - 1}\right)^{-1/2}. \quad (5)$$

This relation suggests that unless an estimate of backgrounds or systematics is very poorly done, the numerator $|\log(N''_{\text{req}}/N_{\text{req}})|$ is much smaller than the denominator $\log(N\langle P_{\text{prod}} \rangle / N_{\text{req}}) - 1$, making g'' differ insignificantly from g .

All these observations demonstrate that the sensitivity near the beam-dump ceiling is robust and nearly insensitive to details of backgrounds, systematics, and beam intensity for a given experimental configuration unless they are grossly misestimated. We will verify this expectation shortly in a (semi)quantitative way, taking the ALP interacting with the SM photon as a concrete example.

C. Dependence on detector geometry

Many experiments that aim to detect mediator decay signals often accompany a finite-size decay volume or equivalent module immediately upstream of the particle detection modules (e.g., DAMSA [17], DarkQuest [22], NA62 [23], SHiP [24], etc). Obviously, the mediators residing near the ceiling are likely to decay promptly, and thus a long decay volume is unessential to probe this region. More specifically, decay events contributing to the sensitivity determination near the ceiling mostly happen at the beginning of the decay volume. Therefore, the sensitivity reaches in the prompt-decay region are nearly unaffected by the length of the decay volume.

Furthermore, the prompt mediators that reach and decay in the decay volume are likely to be highly boosted. Mediators with a large boost factor are induced by high-energetic, more forward-directed source particles. Therefore, wide angular coverage is less motivated, in order for a detector to explore the prompt-decay region.

Considering the above two geometry-wise observations, which can be summarized as

- (c) shorter length of the decay volume, and
- (d) smaller angular coverage of the detector system,

we expect that a tabletop-sized detector system can, in principle, be sufficient to probe the prompt-decay region and reach its own beam-dump ceiling. We will show the dependence of sensitivity reaches on these aspects again in the context of the ALP scenario.

D. Dependence on instantaneous beam intensity

The discussion thus far is relevant to any generic mediators, irrespective of whether their visible decay products involve some features that are distinguishable from potential (reducible) backgrounds. Depending on the final-state features in the mediator decay, an appropriate choice of instantaneous beam intensity (i.e., single-bunch intensity) would further reduce the required beam exposure to achieve the aimed sensitivity goal in the prompt-decay

region. For example, in the DAMSA proposal [17], the authors studied the ALP decaying to a photon pair and identified an accidental overlap of a pair of beam-related-neutron-induced photons that mimics the diphoton signal as the primary background.

One can expect that the number of background events will be reduced if the number of beam-related-neutron-induced photons in a single beam bunch length decreases. For a more quantitative calculation, suppose that a given protons-on-target (PoT) is delivered by n_{bunch} beam bunches, and each bunch results in n_γ neutron-induced photons (possibly after a set of cuts). Since the number of photon pairs within the given beam bunch is $\sim n_\gamma^2$, the total number of signal-faking accidental photon pairs in the given PoT is $\sim n_{\text{bunch}} n_\gamma^2$. Therefore, if the same PoT is delivered by f times more bunches, the number of beam-related-neutron-induced photons in a beam bunch is $\sim n_\gamma/f$, and the expected total number of backgrounds will be reduced by a factor of f to $\sim n_{\text{bunch}} n_\gamma^2/f$.

This simple algebra clearly shows that while backgrounds such as the beam-related-neutron-induced one could be formidable, an experiment to perform a search for featureful mediator signals could benefit greatly in reaching its ceiling at a lower beam exposure with (e) smaller instantaneous beam intensity.

As this is related to backgrounds, we study the sensitivity dependence on the instantaneous beam intensity by varying background estimates.

III. BENCHMARK STUDY

We are now in the position to validate the arguments thus far in the context of actual analyses for an example physics case. To this end, we consider ALPs interacting with the SM photon:

$$-\mathcal{L}_{\text{int}} \supset \frac{1}{4} g_{a\gamma\gamma} a F_{\mu\nu} \tilde{F}^{\mu\nu}, \quad (6)$$

where a , $F_{\mu\nu}$ ($\tilde{F}_{\mu\nu}$), and $g_{a\gamma\gamma}$ are the ALP field, the (dual-)field strength tensor of the SM photon, and the interaction strength between the ALP and the SM photon, respectively. Under this Lagrangian, an ALP is produced by the Primakoff scattering process of the photons [25] inside the beam target, and it decays into a photon pair inside the experimental system.

A. Theoretical benchmark comparison of the three facilities

The three benchmark beam facilities and their key specifications are summarized in Table I. The default bunch rate (henceforth denoted by R_b^*) can be calculated from PoT/bunch in this table. For PIP-II, a range of bunch repetition rates is possible [19], but we assume 100 MHz

TABLE I. Benchmark beam facilities: Key beam parameters (second through fourth columns) and our analysis cuts on photons at ECAL (fifth column).

Beam	E_p [GeV]	PoT/bunch	PoT/yr	$E_{\gamma,\text{cut}}$ [GeV]
PIP-II [19]	0.8	7.7×10^7	4×10^{23}	0.1
SPS [20,27]	400	1.7×10^5	4×10^{19}	5
LHC-dump [21]	7,000	2.2×10^{11}	1.1×10^{17}	20

for illustration. SPS features a slow beam spill of 4×10^{13} protons over 1.2 seconds. We assume that the beam is extracted in bunches evenly spaced at 200 MHz [26], corresponding to 1.7×10^5 PoT per bunch. By contrast, the proton beams are dumped to the LHC-dump 2,756 times per day, with 2.2×10^{11} PoT per bunch [21]. This corresponds to a yearly exposure of 1.1×10^{17} PoT assuming a 50% duty factor.

The photon flux expected in each of the beam facilities is estimated using GEANT4 [28]. We assume a tungsten dump for the three beam facilities for comparison purposes, but our conclusions do not depend on the choice of the dump as long as the material is the same. The selected target lengths are 1 meter (PIP-II), 1.2 meters (SPS), and 1.2 meters (LHC-dump), and our simulation shows that the mean photon production positions are at ~ 5 cm, ~ 15 cm, and ~ 20 cm after $E_{\gamma,\text{cut}}$ in Table I, respectively. We then envision a similar experimental setup as in the DAMSA proposal [17]: a certain length of vacuum decay chamber follows from the beam-target system, and an electromagnetic calorimeter (ECAL) is placed downstream of the decay chamber. We further assume a 100 MeV photon energy threshold and a 1° angular resolution of the detector, although the detailed ECAL capabilities hardly affect our conclusions. Finally, we follow the same data analysis scheme and beam-related-neutron-induced background suppression strategy as explicated in Ref. [17].

We first consider PIP-II, providing an 800-MeV proton beam, and we show the expected 90% CL sensitivity reaches under various experimental environments in Fig. 2(a). The reference or default parameters are defined as follows:

- (1) Detector distance from the mean photon production position in the target: $L_{\text{dist}} = 1$ meter.
- (2) Length of the decay volume: $\Delta L_{\text{det}} = 10$ meters.
- (3) Half the detector opening angle: $\theta_{\text{det}} = 0.5$ rad.
- (4) Beam bunch rate: R_b , to deliver the given PoT.
- (5) Beam exposure: $N_{\text{yr}} = 1$ year.

In addition to these parameter choices, we impose an energy cut on signal photons—i.e., $E_{\gamma,\text{cut}} = 100$ MeV—to mitigate the neutron-induced photon backgrounds and to reflect the reasonable detector capabilities. We find that a negligible number of accidental photon pairs can pass this hard energy cut. We do not vary the detector distance but




FIG. 2. Expected 90% CL sensitivity reaches for coupling $g_{\gamma\gamma}$ as a function of ALP mass m_a at (a) PIP-II, (b) SPS, and (c) LHC-dump. The inset in panel (a) shows the fractional change of the sensitivity with respect to the default case. We compare the three benchmark beam facilities for $\Delta L_{\text{det}} = 1$ meter, $\theta_{\text{det}} = 0.05$ rad, and $N_{\text{yr}} = 3$ months in panel (d), where the sensitivity reaches with the default setup are shown by dotted lines [i.e., the red solid curves in panels (a), (b), and (c)] for reference purposes. In panel (d), the black dot-dashed line represents other future prospects, including DarkQuest [22], DUNE [39], FASER/FASER2 [40], LDMX [41], and SHiP [27], for comparison.

fix it to be $L_{\text{dist}} = 1$ meter throughout our study, since we are primarily interested in the prompt-decay region.² The thin red solid curve shows our reference sensitivity reaches expected at the 800 MeV PIP-II. For comparison purposes, the current constraints compiled in, e.g., Refs. [12,16] are in the grey and yellow shaded area; they include Belle-II [29], CHARM [30], E137 [31], E141 [32], LEP [33], LHC (Pb) [34,35], NA64 [36], NuCal [37], and PrimEx [38].

The remaining four curves compare the sensitivity reaches with different experimental parameter choices against the reference one, keeping all others at default values. The magenta solid line is for a 10-times-increased beam exposure. It enables extending the sensitivity reaches toward the bottom and the higher-mass regions (region IV in Fig. 1) that are sensitive to statistics. The sensitivity

toward the prompt-decay region (region V in Fig. 1), however, does not change, since the limit has already been reached. A similar statement is relevant to the other cases: a 10-times-decreased polar-angular coverage with the blue dot-dashed line, a 100-times-increased bunch rate with the green dashed line, and a 10-times-decreased decay volume length with the orange dotted line. Again, for all of them, we lose or gain some sensitivity in the floor and higher-mass regions (region IV), but their impacts on the sensitivity in the prompt-decay region (region V) are negligible.

All these comparisons suggest that sensitivity estimates near the beam-dump ceiling (below $m_a \lesssim 300$ MeV for PIP-II) are robust as predicted earlier, changing at most within the 5%–10% level below $m_a = 300$ MeV, as shown in the inset. More quantitatively speaking, for example, the comparison between the red and green curves essentially represents the case where the required number of signal events to determine the sensitivity differs by an order of

²Strictly speaking, signal production points differ from event to event, hence L_{dist} varies. We consider this as an average distance.

magnitude (i.e., a 1,000% change) due to a mismodeling of systematics and/or an incorrect background estimation.

Figures 2(b) and 2(c) display our sensitivity estimates at SPS and LHC-dump, respectively, assuming a fixed baseline of $L_{\text{dist}} = 1$ m. While facility-specific studies are necessary to determine the optimal baselines in these cases, we adopt this value for illustrative purposes to enable a uniform comparison across different beam facilities. Unlike PIP-II, more energetic photons would come out of the target in addition to the neutron-induced photons, such as those from processes like $K_L^0 \rightarrow 3\pi^0 \rightarrow 6\gamma$, due to higher proton beam energies. Our GEANT4 simulation suggests that the backgrounds from photons exiting the target are significantly suppressed with $E_{\gamma,\text{cut}}$ and PoT/bunch in Table I; more quantitatively, there are $\sim 1.2 \times 10^9$ and $\sim 4.2 \times 10^{12}$ accidental photon pairs per year at SPS and LHC-dump, respectively. Assuming a similar-level suppression factor of $\sim 10^8$ – 10^9 as in Ref. [17] along with posterior cuts (e.g., invariant mass window) therein, we expect ~ 12 and $\sim 42,000$ diphoton backgrounds annually. The effectiveness of the posterior cuts, such as the distance-of-closest-approach cut, arrival-time-difference cut, fiducial-volume cut, trace-back cut, and invariant-mass-window cut, as developed in Ref. [17], relies on the detector’s capabilities. For illustration, we assume a similar level of detector performance to achieve comparable background suppression.

Similarly to the PIP-II case, the maximum sensitivity reaches around the beam-dump ceiling are rather robust and nearly insensitive to beam parameters, detector geometry, statistics, etc. As pointed out earlier, the impact of the detector angular coverage is much less pronounced in SPS and the LHC than that in PIP-II, since the signal is more highly forward-directed.

Along these lines, our study further enables the identification of key experimental requirements for exploring the “prompt-decay” parameter space, suggesting that a compact “tabletop-sized” detector is sufficient to reach the beam-dump ceiling of a given experiment with a *moderate beam exposure*. Therefore, multiple, short-term experiments with the same compact detector system can be performed at different beam facilities, providing complementary information for probing the parameter space, especially toward the prompt-decay regime. For example, the solid curves in Fig. 2(d) show the sensitivity prospects of the three beam facilities with $\Delta L_{\text{det}} = 1$ meter, $\theta_{\text{det}} = 0.05$ radians, and $N_{\text{yr}} = 3$ months. We employ an ECAL with just a 10 cm-scale radius for these sensitivity estimates. For comparison, we also show the sensitivity curves corresponding to the default configurations by dotted lines. These examples clearly support the case of compact, portable experiments that can be performed in different facilities, complementing each other in expanding reachable parameter space, possibly beyond other future prospects, including DarkQuest [22], DUNE [39],

FASER/FASER2 [40], LDMX [41], and SHiP [27], represented by the black dot-dashed line.

We stress that Fig. 2(d) is not a facility-optimized projection, but a controlled benchmark in which key experimental inputs are held fixed to isolate the impact of beam energy on production kinematics and acceptance. Because the achievable (instantaneous) beam exposure, beam-time structure, and background environment are highly facility-dependent, the relative ordering of reaches can shift once these effects are included. Therefore, Fig. 2(d) should be read as demonstrating complementarity in kinematics/acceptance rather than as evidence for a universal optimal beam energy.

B. Practical comparison of the SPS and LHC-dump facilities

As mentioned earlier, the study in the previous section is largely theoretical for the SPS and LHC-dump facilities and will require more detailed investigations to assess the ECAL’s exposure to unwanted particles, including escaping hadrons (e.g., beam-related neutrons) and muons. While we leave a full feasibility study of this type for future publications, including the DAMSA Pathfinder work [42], we present here sensitivity estimates based on more practical configurations that could plausibly be achieved at the two facilities:

- (1) *SPS*: We leverage the planned SHiP experimental setup [27] to test the DAMSA proposal at the SPS. SHiP employs a ~ 1.4 m TZM (titanium-zirconium-doped molybdenum) $\oplus$ tungsten target, followed by a ~ 40 m-scale magnetic field area to bend away the muons escaping from the target. The collaboration’s dedicated muon-shielding simulation indicates the existence of a nearly muon-free “sweet spot” located about 15 m downstream from the front of the shielding region [43,44]. Given the mean photon production point, the effective baseline is ~ 16 meters.

A full simulation incorporating both the beam target and the entire shielding region is computationally prohibitive. Instead, we simulate the target followed by an iron yoke with four representative lengths—0.25 m, 0.5 m, 1 m, and 2 m—using 10^6 incident protons in each configuration. We find that the number of photons exiting the yoke with $E_{\gamma} > 5$ GeV decreases rapidly with increasing yoke length, and the trend is well described by a straight line in log space. Extrapolating to a 15 m iron yoke yields an expected photon rate below 10^{-22} per incident proton. Given the PoT per bunch and the annual beam intensity, this corresponds to $\sim 10^{-8}$ photons per bunch, implying ~ 0.01 accidental photon pairs per year. Therefore, under this setup, a nearly background-free analysis is feasible without imposing additional *a posteriori* cuts.

(2) *LHC-dump* [21]: Once the luminosity decays below an operational limit, the fill is intentionally terminated, and the circulating protons are extracted to the ~ 8 m-scale graphite beam-dump block, followed by a ~ 3 m-thick concrete to prepare for the next injection cycle. A few-meter-long empty space is located downstream of the concrete module. We envision placing an additional 1.5 m tungsten dump in this region, followed by the DAMSA detector. Our GEANT4 simulation indicates that most dumped-proton interactions occur at ~ 2.5 meters, so the effective baseline is 10 meters under this setup.

Our simulation further suggests that, per proton bunch containing 2.2×10^{11} protons, 17 photons with energies above 20 GeV enter the DAMSA detector region. This results in 136 accidental photon pairs, some of which could potentially mimic the ALP signal events, yielding approximately 6.8×10^7 diphoton-like events annually.

We are now ready to present the sensitivity estimates for the SPS and LHC-dump configurations, assuming the experimental setup described above. For comparison, we adopt the same detector size as in the $\theta_{\text{det}} = 0.05$ rad case: a calorimeter region with a length of order 1 m and a radius of order 10 cm. With this geometry, the DAMSA detector subtends angles of 5.9 mrad and 9.1 mrad with respect to the beam axis for the SPS and LHC-dump cases, respectively.

The blue solid curve in Fig. 3 shows our SPS sensitivity estimates assuming a one-year exposure, while the green solid curve corresponds to the LHC-dump case with no cuts applied other than $E_\gamma > 20$ GeV. The black dot-dashed

line represents other future prospects as before. Since the LHC-dump case (green solid curve) is overwhelmed by a formidable amount of reducible backgrounds, its full potential is not adequately reflected. Indeed, we observe that, in the experimental configuration under consideration, the LHC-dump case would not reach its own “ceiling,” since the beam intensity does not appear to be sufficient for the experiment to become limited by the beam-dump ceiling rather than by backgrounds. In other words, reducing the backgrounds would allow one to probe deeper into the prompt-decay region, potentially up to the ceiling. We therefore also show the maximal reach, indicated by the green dashed curve, which could be achieved with an appropriate set of cuts. For example, the set of kinematic cuts proposed in Ref. [17] can suppress accidental diphoton-like backgrounds by roughly 8–9 orders of magnitude, yielding a nearly background-free environment. However, validating this expectation requires a careful and dedicated analysis, which we defer to future work as mentioned earlier.

IV. CONCLUSIONS

In this paper, we presented a study of the (visibly) decaying signals of new mediators at beam-dump-type experiments, taking ALPs with couplings to photons for illustration. We clearly showed that once a given experiment reaches its inherent ceiling, the sensitivity near the beam-dump ceiling, the maximum sensitivity reach that an experiment can accomplish in the prompt-decay region, is robust and insensitive to detector geometry, beam specifications, and background/systematics estimates, regardless of the detector baseline. Moreover, we provided means to assess whether the projected sensitivity of an experiment hits the ceiling in Eqs. (2)–(5).

We also considered more realistic background levels and experimental configurations compatible with existing SPS and LHC-dump facility designs. Our sensitivity estimates still indicate that a tabletop-sized DAMSA detector can probe new regions of ALP parameter space, potentially extending beyond other future prospects. These results clearly motivate further dedicated studies, including the exploration of additional physics cases at these facilities.

We hope that DAMSA can be realized in the near future [42], probing unexplored and challenging regions of dark-sector parameter space and potentially uncovering hints of new physics.

ACKNOWLEDGMENTS

For facilitating portions of this research, the authors wish to acknowledge the Center for Theoretical Underground Physics and Related Areas (CETUP*), the Institute for Underground Science at the Sanford Underground Research Facility (SURF), and the South Dakota Science and Technology Authority for hospitality and financial support, as well as for providing a stimulating environment.




FIG. 3. Expected 90% CL sensitivity reaches for coupling $g_{\gamma\gamma}$ as a function of ALP mass m_a at SPS (blue curve) and LHC-dump (green curves), assuming the experimental configuration described in the text. To illustrate the maximal potential of the LHC-dump configuration, we also show a sensitivity estimate (green dashed curve) assuming a nearly background-free environment. The black dot-dashed line represents other future prospects, as in Fig. 2.

This article was supported by the computing resources of the Global Science Experimental Data Hub Center (GSDC) at the Korea Institute of Science and Technology Information (KISTI). We thank Wooyoung Jang of the University of Texas at Arlington for valuable discussions. The work of D. K. is supported in part by U.S. Department of Energy Grant No. DE-SC0010813. The work of J.-C. P. is supported by the National Research Foundation of Korea (NRF) grant funded by the Ministry of Science and ICT [No. RS-2024-00356960] and by the Ministry of Education [No. RS-2025-25442707]. The work of J. Y. is supported

by the University of Texas at Arlington, and by the U.S. Department of Energy under Grant No. DE-SC0011686 and NRF Grant No. RS-2024-00350406. The work of H. K. is supported by the U.S. Department of Energy under Grant No. DE-SC0010813 and NRF Grants No. RS-2026-25493297 and No. RS-2024-00350406.

DATA AVAILABILITY

The data are not publicly available. The data are available from the authors upon reasonable request.

-
- [1] A. J. Krasznahorkay *et al.*, Observation of anomalous internal pair creation in ^8Be : A possible indication of a light, neutral boson, *Phys. Rev. Lett.* **116**, 042501 (2016).
 - [2] A. J. Krasznahorkay *et al.*, New evidence supporting the existence of the hypothetical X17 particle, [arXiv:1910.10459](#).
 - [3] A. J. Krasznahorkay, M. Csatlós, L. Csige, J. Gulyás, A. Krasznahorkay, B. M. Nyakó, I. Rajta, J. Timár, I. Vajda, and N. J. Sas, New anomaly observed in ^4He supports the existence of the hypothetical X17 particle, *Phys. Rev. C* **104**, 044003 (2021).
 - [4] A. J. Krasznahorkay *et al.*, New anomaly observed in ^{12}C supports the existence and the vector character of the hypothetical X17 boson, *Phys. Rev. C* **106**, L061601 (2022).
 - [5] A. J. Krasznahorkay, A. Krasznahorkay, M. Csatlós, L. Csige, J. Timár, M. Begala, A. Krakó, I. Rajta, and I. Vajda, Observation of the X17 anomaly in the decay of the Giant Dipole Resonance of ^8Be , [arXiv:2308.06473](#).
 - [6] A. Aguilar-Arevalo *et al.* (LSND Collaboration), Evidence for neutrino oscillations from the observation of $\bar{\nu}_e$ appearance in a $\bar{\nu}_\mu$ beam, *Phys. Rev. D* **64**, 112007 (2001).
 - [7] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), Unexplained excess of electron-like events from a 1-GeV neutrino beam, *Phys. Rev. Lett.* **102**, 101802 (2009).
 - [8] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), Significant excess of electronlike events in the MiniBooNE short-baseline neutrino experiment, *Phys. Rev. Lett.* **121**, 221801 (2018).
 - [9] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), Updated MiniBooNE neutrino oscillation results with increased data and new background studies, *Phys. Rev. D* **103**, 052002 (2021).
 - [10] P. Abratenko *et al.* (MicroBooNE Collaboration), Search for an anomalous production of charged-current νe interactions without visible pions across multiple kinematic observables in MicroBooNE, *Phys. Rev. Lett.* **135**, 081802 (2025).
 - [11] P. Abratenko *et al.* (MicroBooNE Collaboration), Inclusive search for anomalous single-photon production in MicroBooNE, [arXiv:2502.06064](#).
 - [12] B. Batell *et al.*, Dark sector studies with neutrino beams, in *Snowmass 2021* (2022); [arXiv:2207.06898](#).
 - [13] P. Ilten, Y. Soreq, M. Williams, and W. Xue, Serendipity in dark photon searches, *J. High Energy Phys.* **06** (2018) 004.
 - [14] M. Bauer, P. Foldenauer, and J. Jaeckel, Hunting all the hidden photons, *J. High Energy Phys.* **07** (2018) 094.
 - [15] M. Fabbrichesi, E. Gabrielli, and G. Lanfranchi, The dark photon, [10.1007/978-3-030-62519-1](#) (2020).
 - [16] J.-F. Fortin, H.-K. Guo, S. P. Harris, D. Kim, K. Sinha, and C. Sun, Axions: From magnetars and neutron star mergers to beam dumps and BECs, *Int. J. Mod. Phys. D* **30**, 2130002 (2021).
 - [17] W. Jang, D. Kim, K. Kong, Y. Kwon, J.-C. Park, M. S. Ryu, S. Shin, R. G. Van de Water, U.-K. Yang, and J. Yu, Search prospects for axionlike particles at rare nuclear isotope accelerator facilities, *Phys. Rev. D* **107**, L031901 (2023).
 - [18] B. Dutta, D. Kim, and H. Kim, Proposal to use LHC general-purpose detectors in “beam-dump” measurements for long-lived particles, *Phys. Lett. B* **861**, 139246 (2025).
 - [19] R. Ainsworth *et al.*, An upgrade path for the Fermilab accelerator complex, [arXiv:2106.02133](#).
 - [20] C. C. Ahlida *et al.*, *SPS Beam Dump Facility—Comprehensive Design Study*, edited by M. Calviani, B. Goddard, R. Jacobsson, and M. Lamont, CERN Yellow Reports: Monographs, Vol. 2/2020 (2020); [arXiv:1912.06356](#).
 - [21] J. Maestre *et al.*, Design and behaviour of the large hadron collider external beam dumps capable of receiving 539 MJ/dump, *J. Instrum.* **16**, P11019 (2021).
 - [22] A. Apyan *et al.*, DarkQuest: A dark sector upgrade to SpinQuest at the 120 GeV Fermilab main injector, in *Snowmass 2021* (2022); [arXiv:2203.08322](#).
 - [23] E. Cortina Gil *et al.* (NA62 Collaboration), The beam and detector of the NA62 experiment at CERN, *J. Instrum.* **12**, P05025 (2017).
 - [24] S. Alekhin *et al.*, A facility to search for hidden particles at the CERN SPS: The SHiP physics case, *Rep. Prog. Phys.* **79**, 124201 (2016).
 - [25] H. Primakoff, Photoproduction of neutral mesons in nuclear electric fields and the mean life of the neutral meson, *Phys. Rev.* **81**, 899 (1951).
 - [26] C. Ahlida *et al.* (SHiP Collaboration), The SHiP experiment at the proposed CERN SPS beam dump facility, *Eur. Phys. J. C* **82**, 486 (2022).
 - [27] R. Albanese *et al.* (SHiP Collaboration), *BDF/SHiP at the ECN3 High-Intensity Beam Facility*, Technical Report (CERN, Geneva, 2023).

- [28] S. Agostinelli *et al.* (GEANT4 Collaboration), GEANT4—a simulation toolkit, *Nucl. Instrum. Methods Phys. Res., Sect. A* **506**, 250 (2003).
- [29] F. Abudinén *et al.* (Belle-II Collaboration), Search for axion-like particles produced in e^+e^- collisions at Belle II, *Phys. Rev. Lett.* **125**, 161806 (2020).
- [30] F. Bergsma *et al.* (CHARM Collaboration), Search for axion like particle production in 400-GeV proton-copper interactions, *Phys. Lett.* **157B**, 458 (1985).
- [31] J. D. Bjorken, S. Ecklund, W. R. Nelson, A. Abashian, C. Church, B. Lu, L. W. Mo, T. A. Nunamaker, and P. Rassmann, Search for neutral metastable penetrating particles produced in the SLAC beam dump, *Phys. Rev. D* **38**, 3375 (1988).
- [32] E. M. Riordan *et al.*, A search for short lived axions in an electron beam dump experiment, *Phys. Rev. Lett.* **59**, 755 (1987).
- [33] G. Abbiendi *et al.* (OPAL Collaboration), Multiphoton production in e^+e^- collisions at $\sqrt{s} = 181\text{-GeV}$ to 209-GeV , *Eur. Phys. J. C* **26**, 331 (2003).
- [34] A. M. Sirunyan *et al.* (CMS Collaboration), Evidence for light-by-light scattering and searches for axion-like particles in ultraperipheral PbPb collisions at $\sqrt{s_{NN}} = 5.02\text{ TeV}$, *Phys. Lett. B* **797**, 134826 (2019).
- [35] G. Aad *et al.* (ATLAS Collaboration), Measurement of light-by-light scattering and search for axion-like particles with 2.2 nb^{-1} of Pb + Pb data with the ATLAS detector, *J. High Energy Phys.* **03** (2021) 243; **11** (2021) 050(E).
- [36] D. Banerjee *et al.* (NA64 Collaboration), Search for axion-like and scalar particles with the NA64 experiment, *Phys. Rev. Lett.* **125**, 081801 (2020).
- [37] J. Blumlein *et al.*, Limits on neutral light scalar and pseudoscalar particles in a proton beam dump experiment, *Z. Phys. C* **51**, 341 (1991).
- [38] D. Aloni, C. Fanelli, Y. Soreq, and M. Williams, Photo-production of axionlike particles, *Phys. Rev. Lett.* **123**, 071801 (2019).
- [39] V. Brdar, B. Dutta, W. Jang, D. Kim, I. M. Shoemaker, Z. Tabrizi, A. Thompson, and J. Yu, Axionlike particles at future neutrino experiments: Closing the cosmological triangle, *Phys. Rev. Lett.* **126**, 201801 (2021).
- [40] J. L. Feng, I. Galon, F. Kling, and S. Trojanowski, Axionlike particles at FASER: The LHC as a photon beam dump, *Phys. Rev. D* **98**, 055021 (2018).
- [41] A. Berlin, N. Blinov, G. Krnjaic, P. Schuster, and N. Toro, Dark matter, millicharges, axion and scalar particles, gauge bosons, and other new physics with LDMX, *Phys. Rev. D* **99**, 075001 (2019).
- [42] P. Bhattarai *et al.* (DAMSA Collaboration), DAMSA experiment conceptual design white paper, [arXiv:2601.15255](https://arxiv.org/abs/2601.15255).
- [43] A. De Roeck, The ship experiment: Search for hidden particles (2025), presented on June 25, (2025) https://indico.cern.ch/event/1441321/contributions/6542403/attachments/3078461/5451176/deroeck_LL2025_SHIP_v2-2.pdf.
- [44] R. Jacobsson (private communication).