M^3: a new muon missing momentum experiment to probe $(g-2)_\mu$ and dark matter at Fermilab

Yonatan Kahn,^{a} Gordan Krnjaic,^{b} Nhan Tran^b and Andrew Whitbeck^b

^{a}*Princeton University,*
Princeton, 08544 NJ, U.S.A.

^{b}*Fermi National Accelerator Laboratory,*
Batavia, 60510 IL, U.S.A.

E-mail: ykahn@princeton.edu, krnjaicg@fnal.gov, ntran@fnal.gov,
awhitbe1@fnal.gov

ABSTRACT: New light, weakly-coupled particles are commonly invoked to address the persistent $\sim 4\sigma$ anomaly in $(g-2)_\mu$ and serve as mediators between dark and visible matter. If such particles couple predominantly to heavier generations and decay invisibly, much of their best-motivated parameter space is inaccessible with existing experimental techniques. In this paper, we present a new fixed-target, missing-momentum search strategy to probe invisibly decaying particles that couple preferentially to muons. In our setup, a relativistic muon beam impinges on a thick active target. The signal consists of events in which a muon loses a large fraction of its incident momentum inside the target without initiating any detectable electromagnetic or hadronic activity in downstream veto systems. We propose a two-phase experiment, M^3 (Muon Missing Momentum), based at Fermilab. Phase 1 with $\sim 10^{10}$ muons on target can test the remaining parameter space for which light invisibly-decaying particles can resolve the $(g-2)_\mu$ anomaly, while Phase 2 with $\sim 10^{13}$ muons on target can test much of the predictive parameter space over which sub-GeV dark matter achieves freeze-out via muon-philic forces, including gauged $U(1)_{L_\mu - L_\tau}$.

Contents

1 Introduction

Despite decades of dedicated searches, the particle nature of dark matter (DM) remains one of the greatest mysteries in all of physics (for a review see [1]). In response to null results from direct detection and collider experiments designed to probe weakly interacting massive particles, there has recently been a surge of activity in designing new techniques to explore "hidden sectors" in which DM is a Standard Model (SM) singlet with its own dynamics and cosmological history (see [2, 3] for a review). In a well-motivated class of such models, DM interacts predominantly with muons through a new force carrier, which enables thermal freeze-out in the early universe. This scenario generally features sizable couplings to visible particles and predictive experimental benchmarks, yet remains inaccessible to traditional DM search strategies, thereby motivating new techniques.

Such muon-specific forces are also well-motivated independently of any possible connection to dark matter. The persistent $\sim 3.6\sigma$ discrepancy in the anomalous magnetic moment of the muon remains one of the largest anomalies in particle physics [4]. It is well-known that light, weakly coupled particles can bring theoretical predictions into agreement with observations (see, for example, [5–8]). However in recent years, the most popular new-physics explanation, a light ($< \mathrm{GeV}$) dark photon with kinetic mixing $\epsilon \sim 10^{-3}$ and flavor-universal couplings, has been ruled out regardless of whether it decays visibly or invisibly [3]. The only remaining class of light new-physics explanations involves particles that couple predominantly (or exclusively) to muons, thereby evading the searches upon which the dark photon exclusions are based.[1] Thus, a robust test of the new-physics hypothesis requires improved sensitivity to muonic forces.

One promising approach to probing invisibly decaying muonic forces is the NA64 experiment, which aims to utilize the CERN SPS muon beam to perform a missing energy search [11]. The technique exploits the radiative production process $\mu N \to \mu N \not{E}$ shown in figure 1, where the final state missing energy $\not{E}$ arises from the production of a new muon-philic particle (S or V for scalar or vector, respectively). Here the signal is defined by the kinematics of the outgoing muon, which loses $\sim 50\%$ of its incident energy and momentum in a typical production event. Since the new radiated particle decays invisibly to DM or neutrinos, there is no additional visible energy, so events with additional activity in an ECAL/HCAL downstream of the target are vetoed as SM-induced backgrounds. The projected sensitivity for this approach can cover the entire region of parameter space over which a muonic force can reconcile the $(g-2)_\mu$ anomaly for invisibly decaying, muon-philic bosons [11]. This approach complements the efforts of the NA64 [12, 13] and LDMX [14, 15] experiments which are based on a similar setup, but utilize electron beams to probe invisibly-decaying particles which couple to electrons. The missing energy technique also complements proposed muon beam dump searches for *visibly decaying* muon-philic scalars proposed in [16], which inherit couplings to electron and photons through mixing with the Higgs, as well as searches at BaBar for mediators with mass greater than $2m_\mu$ which can decay back to muons [17].

---

[1]New physics explanations with heavier electroweak states also remain viable, but are under tension from null LHC results — see e.g. [9, 10].
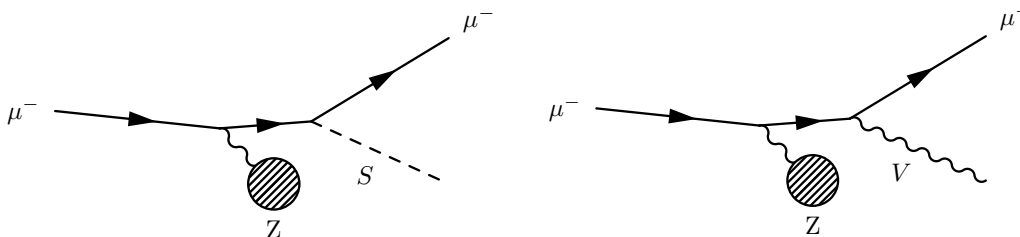
Figure 1. Dark bremsstrahlung signal process for simplified models with invisibly decaying scalar (*left*) and vector (*right*) forces that couple predominantly to muons. In both cases, a relativistic muon beam is incident on a fixed target and scatters coherently off a nucleus to produce the new particle as initial- or final-state radiation.

In this paper, we extend the discussion in [11] to propose a new fixed-target missing-momentum muon-beam experiment at Fermilab, which we dub Muon Missing Momentum or M^3, suitable for any muon-philic invisibly decaying particles. We also present thermal DM targets that such an effort could reach. Our study considers the new physics discovery potential for a 15 GeV muon beam, which is representative of the Fermilab muon beam capabilities, paired with a thick (~ 50 radiation length) target and a downstream detector similar to the planned LDMX setup [15] to veto SM backgrounds. Our detector design is illustrated in figure 2. The thicker active target compared to the ~ 0.1 radiation length in the nominal LDMX setup allows for a larger signal production rate while exploiting the fact that the muons will lose much less energy than electrons in a similarly-sized target. In analogy with similar processes involving electron beams, one can take advantage of the distinctive kinematics of the radiated massive scalar or vector particle S, V to distinguish signal from background (see figure 1). The Fermilab muon beam option provides several advantages over existing proposals for new physics searches with either electron beams or high-energy muon beams:

- *Bremsstrahlung backgrounds suppressed.* The principal reducible backgrounds for LDMX are dominated by hadronic processes initiated by a real bremsstrahlung photon. Relative to electron beams, the M^3 bremsstrahlung rate is suppressed by $(m_e/m_\mu)^2 \approx 2 \times 10^{-5}$, so background rejection becomes much simpler for muon beams for an equivalent target thickness.

- *Compact experimental design.* For $m_{S,V} \ll E_{\rm beam}$, the signal production cross section is largely independent of beam energy. However, compared to the CERN/SPS option [11], with $\sim 100 - 200\,{\rm GeV}$ beam muons, a lower-energy, e.g. 15 GeV, muon beam allows for greater muon track curvature and, therefore, a more compact experimental design. In particular, percent-level momentum resolution is possible in M^3 with the target placed in the magnetic field region, reducing acceptance losses from having the magnet downstream of the target.

We propose a two-phase experiment, each covering a well-motivated region of parameter space:

- *Phase 1: $(g-2)_\mu$ search.* With 10^{10} muons on target (MOT) and existing detector technology, we will show that our setup can probe the entire $(g-2)_\mu$ region not
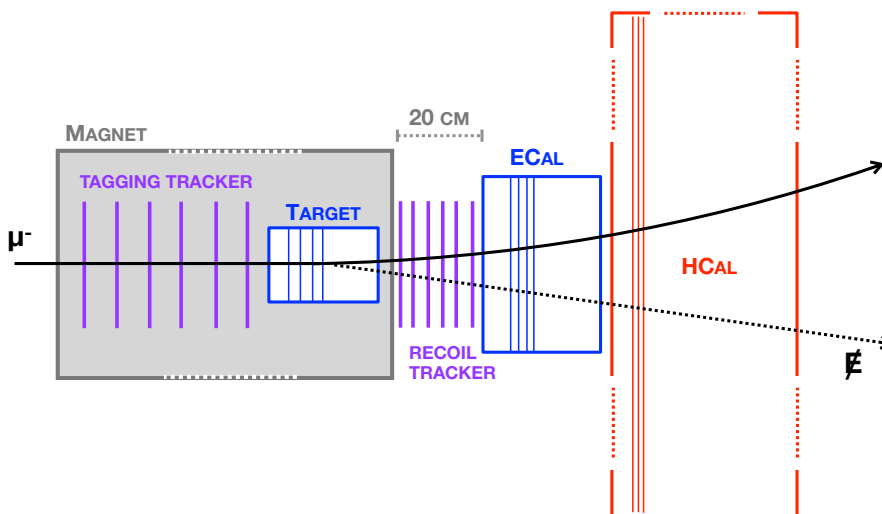
Figure 2. Experimental schematic. The incoming muon beam passes through a tagging tracker in the magnetic field region before entering the tungsten target. Outgoing muons are detected with a recoil tracker, with the magnet fringe field providing a momentum measurement. Electromagnetic and hadronic calorimeters veto on photons and hadrons produced in hard interactions in the target which could lead to significant muon energy loss.

currently excluded by experiments, for vectors with $m_V \lesssim 500$ MeV and scalars with $m_S \lesssim 100$ MeV which couple exclusively to muons and decay invisibly.[2] Here we are agnostic as to the UV completion of such a model, and we are simply aiming for an apples-to-apples comparison between a virtual S or V contributing to $(g-2)_\mu$ and a real S or V emitted from an initial- or final-state muon.

- *Phase 2: thermal muon-philic DM search.* With a larger flux of 10^{13} MOT and up-graded detector performance to reject backgrounds at the level of 10^{-13}, our setup can probe a significant portion of parameter space for which DM is thermally produced through $U(1)_{L_\mu - L_\tau}$ gauge interactions, and V is identified as the gauge boson of this new U(1). Such models are inaccessible with both traditional WIMP searches [21–27] and to most of the emerging sub-GeV dark matter search program, which consists of new direct detection [28–41] and fixed target experiments with electron [14, 15, 42–45] and proton beams [18, 46–53]; for a review and summary, see [3].

We emphasize that M^3 Phase 1 can be completed with minimal modifications to the Fermilab muon source and with as little as 3 months of data-taking. A null result would decisively exclude any new-physics explanation of the $(g-2)_\mu$ anomaly from invisibly-decaying muon-philic particles below 100 MeV. Phase 2 is comparable to the CERN SPS proposal, and in this paper we focus specifically on the advantages of pairing such an experiment with

---

[2]Models with a more complicated dark sector can fail our search criteria, for example an inelastic DM model $V \to \chi_1 \chi_2$ where the decay $\chi_2 \to \chi_1 e^+ e^-$ is prompt and proceeds through a different mediator which couples to electrons [18–20].

the lower-energy Fermilab muon beam, highlighting the relevance of this search to the thermal DM parameter space. Furthermore, both phases could be implemented as muon-beam reconfigurations of the proposed LDMX experiment with few additional modifications.

This paper is organized as follows. In section 2 we review the physics motivation for our benchmark models; in section 3 we discuss the characteristics of signal production; in section 4 we describe the basic experimental setup and relevant background processes; in section 5 we describe the necessary detector and beam properties; in section 6 we describe the projected sensitivities of our Phase 1 and Phase 2 proposals; finally, in section 7 we offer some concluding remarks.

2 Physics motivation

In this section we present the physics motivation for invisibly decaying muon-specific scalars S or vectors V. We begin by reviewing the contributions of vector and scalar particles to $(g-2)_\mu$, and then present a concrete benchmark model with a muon-philic gauge interaction which can be coupled to dark matter. Although this is not the only model that preferentially couples a new force carrier to muons, it serves as a representative example without loss of essential generality. Basic variations away from this example (i.e. the scalar force model in [16]) feature the same basic degrees of freedom and their signal characteristics are similar to what we consider below. When discussing a generic muon-specific vector mediator, we will use the notation V, reserving Z' for the gauge boson coupling the muon to dark matter in our representative model.

2.1 Simplified models for $(g-2)_\mu$

The current discrepancy in the anomalous magnetic moment is characterized by $a_\mu \equiv \frac{1}{2}(g-2)_\mu$, the observed value of which differs from the SM theoretical prediction by an amount [4]:

$$\Delta a_\mu \equiv a_\mu(\text{obs}) - a_\mu(\text{SM}) = (28.8 \pm 8.0) \times 10^{-10}. \quad (2.1)$$

It is well known that for a light scalar S or vector V coupling to the muon,

$$g_S S \overline{\mu} \mu \quad \text{(scalar)} , \quad g_V V_\alpha \overline{\mu} \gamma^\alpha \mu \quad \text{(vector)} \quad (2.2)$$

the leading-order contribution to the muon anomalous magnetic moment from scalars is [7]

$$\Delta a_\mu^S = \frac{g_S^2}{16\pi^2} \int_0^1 dz \frac{m_\mu^2 (1-z)(1-z^2)}{m_\mu^2 (1-z)^2 + m_S^2 z} \simeq 9.5 \times 10^{-11} \left(\frac{g_S}{10^{-4}}\right)^2 \quad (m_S \ll m_\mu), \quad (2.3)$$

and the corresponding expression for vector particles is

$$\Delta a_\mu^V = \frac{g_V^2}{4\pi^2} \int_0^1 dz \frac{m_\mu^2 z(1-z)^2}{m_\mu^2 (1-z)^2 + m_V^2 z} \simeq 1.3 \times 10^{-10} \left(\frac{g_V}{10^{-4}}\right)^2 \quad (m_V \ll m_\mu). \quad (2.4)$$

Note that pseudoscalar and axial-vector couplings contribute to a_μ with opposite sign, pushing the theoretical value farther away from the measured value. As a result we only consider parity-even mediators, scalars S or vectors V.

For the purposes of our Phase 1 search, we make no attempt to describe a complete theoretical model of S or V, and take the observed discrepancy in a_μ as positive evidence for a new particle which can be probed in beam dump experiments. For Phase 2, we can define a well-motivated region of parameter space and make a connection to thermal dark matter in a particular representative model, which we describe below.

2.2 A complete vector model: $\mathrm{U}(1)_{L_\mu - L_\tau}$

For our Phase 2 study, we extend the SM to include the anomaly-free $\mathrm{U}(1)_{L_\mu - L_\tau}$ gauge group under which μ, τ and their corresponding neutrino flavors couple to a new gauge boson Z'. The Lagrangian for this scenario is

$$\mathcal{L} = \mathcal{L}_{\mathrm{SM}} - \frac{1}{4} F'^{\alpha\beta} F'_{\alpha\beta} + \frac{m_{Z'}^2}{2} Z'^\alpha Z'_\alpha - Z'_\alpha J^\alpha_{\mu - \tau}, \tag{2.5}$$

where $F'_{\alpha\beta} \equiv \partial_\alpha Z'_\beta - \partial_\beta Z'_\alpha$ is the field strength tensor and $m_{Z'}$ is the gauge boson mass; we assume that the gauge symmetry is spontaneously broken in the IR, but that states responsible for that breaking are sufficiently decoupled that their effects are negligible at the GeV scale. The $\mu - \tau$ current in eq. (2.5) is

$$J^\alpha_{\mu - \tau} = g_{\mu - \tau} \left(\bar{\mu} \gamma^\alpha \mu + \bar{\nu}_\mu \gamma^\alpha P_L \nu_\mu - \bar{\tau} \gamma^\alpha \tau - \bar{\nu}_\tau \gamma^\alpha P_L \nu_\tau \right), \tag{2.6}$$

where $g_{\mu - \tau}$ is the gauge coupling and $P_L \equiv \frac{1}{2}(1 - \gamma^5)$ is the left projection operator. The rest frame partial widths for Z' decays are

$$\Gamma(Z' \to ff) = \frac{\alpha_{\mu - \tau} m_{Z'}}{3} \left(1 + \frac{2m_f^2}{m_{Z'}^2} \right) \sqrt{1 - \frac{4m_f^2}{m_{Z'}^2}} \quad (f = \mu, \tau) \ ,$$

$$\Gamma(Z' \to \nu_i \nu_i) = \frac{\alpha_{\mu - \tau} m_{Z'}}{6}, \tag{2.7}$$

where $\alpha_{\mu - \tau} \equiv g_{\mu - \tau}^2 / 4\pi$. In the absence of additional decay channels, for $m_{Z'} < 2m_\mu$ the Z' will always decay invisibly to μ and τ neutrinos and for $m_{Z'} > 2m_\mu$ the invisible branching ratio is approximately $1/3$. If the Z' can also decay to lighter hidden sector states (see section 2.3), it is possible to have a large branching fraction to invisible states throughout the parameter space if the hidden sector particles couple more strongly to the mediator.

Since this model modifies neutrino interactions, so there are strong limits from neutrino trident scattering via the process $\nu N \to \nu N \mu^+ \mu^-$ which exclude this interpretation of the $(g - 2)_\mu$ anomaly for $m_{Z'} \gtrsim 400 \, \mathrm{MeV}$. Furthermore, for $m_{Z'} \lesssim$ few MeV, there are strong bounds from cosmology since a Z' in this mass range decays to neutrinos after the latter have decoupled from the photons of the SM thermal bath. Such a late decay heats the neutrinos relative to the photons, thereby predicting $\Delta N_{\mathrm{eff}} \sim 3$. Thus, the current Planck limit of $N_{\mathrm{eff}} = 3.15 \pm 0.23$ [54] rules out this mass range unless the Z' is too feebly coupled to ever thermalize with SM particles in the early universe. However, avoiding thermalization requires $g_{\mu - \tau} \ll 10^{-5}$, which falls far short of resolving the $(g-2)_\mu$ anomaly, so our primary focus for the remainder of this work will be on the range $\mathrm{MeV} \lesssim m_{Z'} \lesssim 400 \, \mathrm{MeV}$.

Note that, because gauged $\mathrm{U}(1)_{L_\mu - L_\tau}$ breaks lepton flavor universality, naively it seems that this gauge extension could lead to dangerous flavor-violating processes after spontaneous symmetry breaking. However, as noted in [55], due to $\mathrm{U}(1)_{L_\mu - L_\tau}$ gauge invariance at high energies, the charged lepton Yukawa couplings are already diagonal in the interaction basis prior to symmetry breaking; in the UV the electron, muon and tau are distinct particles and do not mix with each other. Thus, after the Z' acquires its mass, there is no need to perform any additional lepton flavor rotations and there are no tree-level flavor-changing neutral currents in this model. Furthermore, although this model imposes no barrier to generating successful neutrino masses, since the muon and tau neutrinos now carry additional gauge charges, the Weinberg operator $\sim (LH)^2$ is not a $\mathrm{U}(1)_{L_\mu - L_\tau}$ gauge singlet, so the sector responsible for generating neutrino masses must couple to the source of symmetry breaking. A full discussion of neutrino masses and mixings in this extension is beyond the scope of this work, but see [55, 56] for a discussion of this issue.

2.3 Including dark matter

In addition to the model in eq. (2.5), we can also include dark matter charged under $\mathrm{U}(1)_{L_\mu - L_\tau}$, which adds the additional interaction $\mathcal{L} \supset -Z'_\alpha J^\alpha_\chi$ to the Lagrangian in eq. (2.5). For a variety of benchmark DM candidates χ with mass m_χ, the current coupled to the Z' is

$$
J^\mu_\chi = g_\chi \times \begin{cases} i\chi^* \partial_\mu \chi + \text{h.c.} & \text{Complex Scalar} \\ \overline{\chi}_1 \gamma^\mu \chi_2 + \text{h.c.} & \text{Pseudo--Dirac Fermion} \\ \frac{1}{2} \overline{\chi} \gamma^\mu \gamma^5 \chi & \text{Majorana Fermion} \\ \overline{\chi} \gamma^\mu \chi & \text{Dirac Fermion} \end{cases} \tag{2.8}
$$

where $g_\chi = g_{\mu-\tau} q_\chi$ is the Z' coupling to DM and q_χ accounts for the possibility that DM does not carry unit charge under $\mathrm{U}(1)_{L_\mu - L_\tau}$. As long as the DM is vector-like under the gauge extension, q_χ is a free parameter, so we adopt the convention from eq. (2.6) which assigns unit charge to μ, τ, and their neutrinos. Generically, loops of μ and τ will generate kinetic mixing between the Z' and the photon (this scenario and its implications for electron beam experiments are explored further in refs. [55, 57]), but this can always be cancelled by another UV contribution to the kinetic mixing parameter. In this work we assume that *no* kinetic mixing is present in the IR, and the Z' does not couple to charged particles other than muons and taus.

In the early universe, the $L_\mu - L_\tau$ gauge boson Z' mediates interactions that maintain thermal equilibrium between dark and visible matter. In the absence of a particle asymmetry in the dark sector, the relic abundance is determined by the dark-visible interaction rate at the freeze-out temperature $T_{\text{fo}} \sim m_\chi / 20$. For $m_\chi > m_{Z'}$ the abundance is set primarily by $\chi\chi \to Z'Z'$ *secluded* annihilation which depends only on g_χ and is independent of the Z' coupling to SM particles. In the opposite regime $m_\chi < m_{Z'}$, the abundance is determined by direct annihilation, $\chi\chi \to Z'^* \to \mu^+\mu^-, \tau^+\tau^-, \nu_i\nu_i$, the cross section for which is proportional to the combination $(g_\chi g_{\mu-\tau})^2$. In this case, even for $g_\chi \sim 1$ near the perturbative unitarity bound, there is still a minimum $g_{\mu-\tau}$ required to yield the observed abundance at late times.
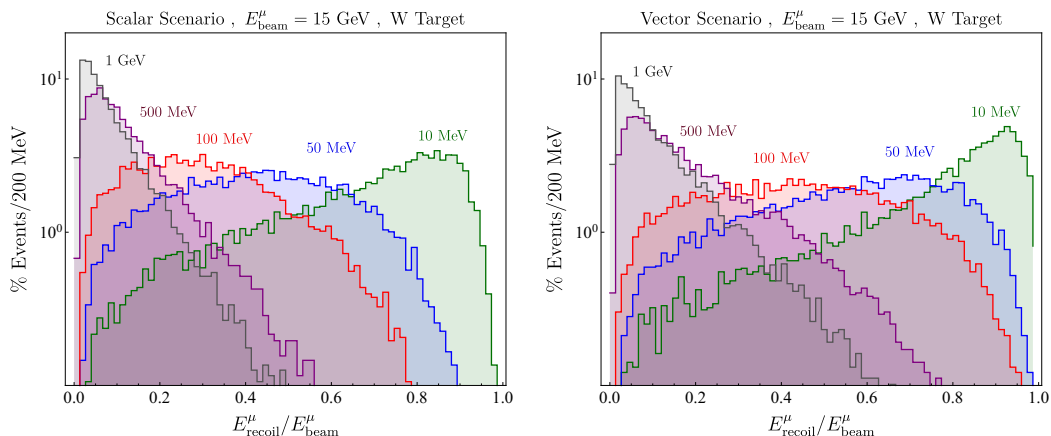
Figure 3. Muon energy distributions for signal events for simplified phenomenological models with invisibly decaying scalar S (*left*) and vector V (*right*) particles that couple only to muons. Each histogram corresponds to a different choice of scalar or vector mass.
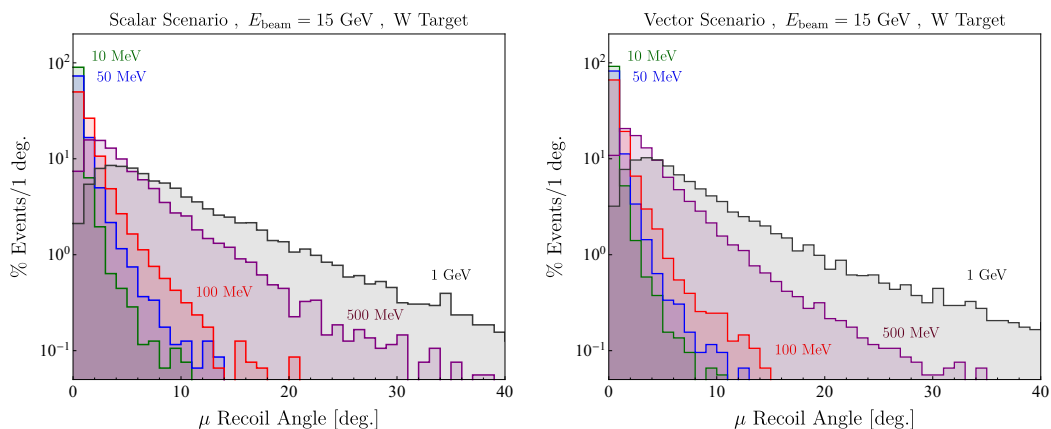

Figure 4. Muon angular distributions for signal events for phenomenological simplified models with invisibly decaying scalar S (*left*) and vector V (*right*) particles coupled to muons.

To define relic targets for the experimental program, we consider in the $m_{Z'} \gg m_\chi$ for a Dirac χ particle. Defining the dimensionless variable $y \equiv g_\chi^2 g_{\mu-\tau}^2 (m_\chi/m_{Z'})^4$, the cross section and abundance have the approximate scaling

$$\langle \sigma v \rangle \simeq \frac{3 g_\chi^2 g^2 m_\chi^2}{\pi^2 m_{Z'}^2} = \frac{3y}{\pi m_\chi^2} \implies \Omega_\chi h^2 \sim 0.1 \left(\frac{3 \times 10^{-9}}{y} \right) \left(\frac{m_\chi}{\text{GeV}} \right)^2, \tag{2.9}$$

where the factor of 3 accounts for annihilation channels and the requisite y from this estimate agrees well with the full numerical results presented in figure 10. As we will show in section 6, there is a region of parameter space where $\text{U}(1)_{L_\mu - L_\tau}$ can explain both thermal DM and the $(g-2)_\mu$ anomaly, but the parameter space for thermal DM is a priori unconstrained by $(g-2)_\mu$.
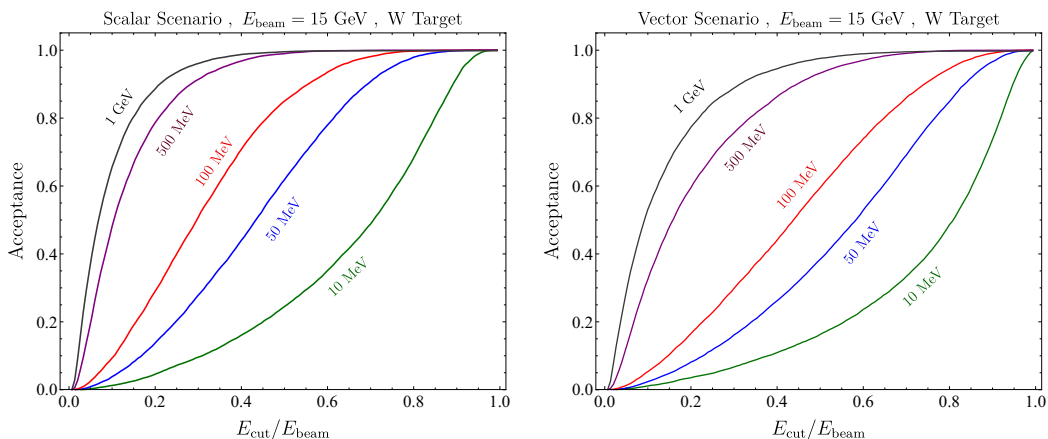
Figure 5. Signal acceptance for radiated scalar S (*left*) and vector V (*right*) particles as a function of the fractional cut on the maximum muon recoil energy for various choices of S/V masses.

3 Signal kinematics

In this section we discuss the physics of the new mediator radiated from a muon beam, noting some important features of the signal production cross section and kinematics. To develop intuition for the signal kinematics, we will use the generalized Weizsacker-Williams (WW) approximation for the cross section for mediator production, appropriate for the limit where the beam energy is much greater than both $m_{S,V}$ and m_μ, where we also keep track of the mass of the beam particle [16, 58–60]. It is known that the WW approximation may give unreliable results in certain regions of parameter space, and thus for all of our reach calculations, we make no approximations and calculate the full diagrams in figure 1 with `Madgraph`.

For scalar S production we have

$$\frac{d\sigma}{dx}\bigg|_S \simeq \frac{g_S^2 \alpha^2}{4\pi} \chi_S \beta_S \beta_\mu \frac{x^3 \left[m_\mu^2(3x^2 - 4x + 4) + 2m_S^2(1-x)\right]}{\left[m_S^2(1-x) + m_\mu^2 x^2\right]^2}, \tag{3.1}$$

and a corresponding expression for V production

$$\frac{d\sigma}{dx}\bigg|_V \simeq \frac{g_V^2 \alpha^2}{4\pi} \chi_V \beta_V \beta_\mu \frac{2x \left[x^2 m_\mu^2(3x^2 - 4x + 4) - 2m_V^2(x^3 - 4x^2 + 6x - 3)\right]}{\left[m_V^2(1-x) + m_\mu^2 x^2\right]^2}, \tag{3.2}$$

where $x = E_{S,V}/E_{\text{beam}}$, $\beta_{S,V} = \sqrt{1 - m_{S,V}^2/(xE_{\text{beam}})^2}$, $\beta_\mu = \sqrt{1 - m_\mu^2/E_{\text{beam}}^2}$, and

$$\chi_{S,V} = \int_{t_{\min}}^{t_{\max}} dt \, \frac{t - t_{\min}}{t^2} G_2(t) \simeq \int_{m_{S,V}^4/(4E_{\text{beam}}^2)}^{m_{S,V}^2 + m_\mu^2} \frac{t - m_{S,V}^4/(4E_{\text{beam}}^2)}{t^2} G_2(t), \tag{3.3}$$

with G_2 the target form factor.

As has been noted in previous studies, and as shown here in figure 3, the energy spectrum differs markedly depending on whether $m_{S,V} < m_\mu$ (in which case S/V is soft and the spectrum closely resembles ordinary QED bremsstrahlung) or $m_{S,V} > m_\mu$ (in
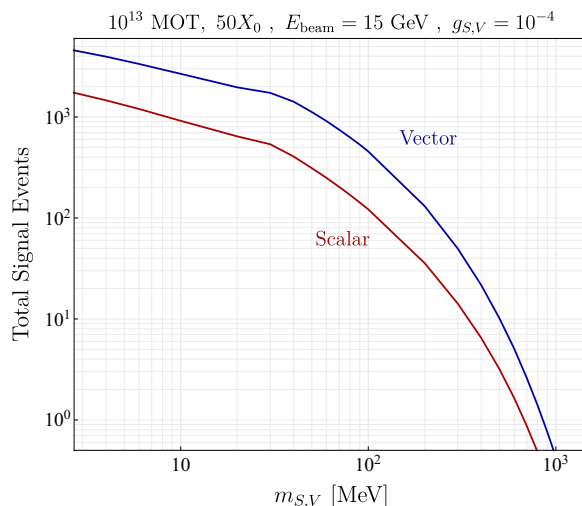
Figure 6. Signal event yield for radiative S or V production in muon-tungsten coherent scattering with a 15 GeV incident beam energy, 10^{13} muons on target, and a target of thickness 50 X_0. Here we have normalized to a reference coupling $g_{V,S} = 10^{-4}$ which is roughly the value required to explain the $(g-2)_\mu$ anomaly; see section 2.1 and figure 9.

which case S/V takes most of the beam energy, as is familiar from electron beam studies). The same is true to a lesser extent for the angular spectrum (figure 4), where the outgoing muon is peaked in the forward direction for light S/V but can be emitted at wider angles for heavier S/V. However, we also note here that both the shape of the energy spectrum and the inclusive cross section are largely independent of beam energy for small $m_{S,V}$. This can be seen directly from eqs. (3.1)–(3.2), where $d\sigma/dx$ is only a function of the recoil energy fraction x and not E_{beam} in the limit $m_{S,V} \to 0$ and $m_\mu \ll E_{\text{beam}}$; we have also confirmed this with `Madgraph` simulations.[3] Therefore, the reach of a muon beam experiment at low masses will be determined almost entirely by the cut on recoil energy carried by the outgoing muon, required to keep backgrounds to a tolerable level. In section 5 we will find that $E_{\text{cut}} = 9$ GeV with $E_{\text{beam}} = 15$ GeV allows us to achieve zero background, while still corresponding to $\mathcal{O}(1)$ acceptance over much of the desired mass range. Figure 5 shows the signal acceptance as a function of recoil cut, and figure 6 shows the inclusive cross section for a beam energy of 15 GeV, for both S and V. The vector cross section is roughly a factor of 3 larger for the same value of the coupling, due almost entirely to the spin sum over the 3 polarization states for a massive vector.

4 Experimental setup

In this section, we introduce the basic experimental concept which takes advantage of the missing momentum technique [61]. We describe the general layout of the detector setup, leaving a dedicated discussion of the detector technology for section 5. After introducing the experimental technique, we discuss the main background processes.

---

[3]We thank Chien-Yi Chen and Yiming Zhong for providing us with the modified `Madgraph` files which implement fixed-target scattering on tungsten targets.

4.1 Missing momentum technique

The missing momentum technique probes the dark bremsstrahlung production of new, invisibly decaying particles produced in $\mu N \to \mu N \not{E}$ processes, as illustrated in figure 1, from an incident muon beam on a fixed nuclear target. The particle of interest (S or V) is radiated off of the muon as initial- or final-state radiation via interaction with the nucleus and, by assumption, promptly decays to dark matter or neutrinos, thereby yielding missing energy and momentum. The technique requires identifying and measuring the momentum of individual muons; the experimental signature is an outgoing muon emerging from the target with significant momentum loss compared to the incident muon, with no additional visible energy deposited in the target material or in the downstream ECAL/HCAL, as illustrated in figure 2.

The signal production cross section depends on the incoming beam energy and the number of radiation lengths, X_0, of the target. Good muon momentum resolution for both the incoming and outgoing tracking system is crucial. A high-field dipole magnet, along with its fringe field, surrounds the tracking region in order to provide a good momentum measurement. The target should be active to detect background processes from hard muon scattering in the target. Downstream of the target, we have additional electromagnetic and hadronic calorimeters to capture muon energy loss from standard QED bremsstrahlung and possible photon-nucleon and muon-nucleon hard processes which produce hadrons. The angular spectrum of the signal process, shown in figure 4, informs the geometry of the tracker and calorimeters.

Two key features separate this experimental proposal from previous studies [11, 16]: the energy of the incident muons (15 GeV as opposed to 150 GeV in previous studies), and the technology, based on the proposed LDMX experiment, which is well-suited to this energy scale and does not rely on measurement of visible decay products of S or V. This design allows for a more compact experimental setup while still achieving the necessary detector performance. In this paper we consider an incident 15 GeV muon beam as a typical beam energy for the experiment, but a range of comparable muon energies is reasonable and our setup does not rely crucially on the precise value of the beam energy.

4.2 Backgrounds

The relevant background processes for electron beams were enumerated in [14], and for muon beams in [11]. Here, we summarize the various categories of backgrounds and how their effects are modified in the case of a muon beam. The main difference between muon and electron beams is that muons are minimum ionizing particles (MIPs), and will continuously lose energy passing through a thick target. Energy loss of the muon without detectable signals in the target from electronic interactions, typically referred to as dE/dx, for a 15 GeV muon through a radiation length 50 X_0 of tungsten absorber material is roughly 550 MeV. Large fluctuations of the energy loss can cause a signal-like event topology. However, these events are often accompanied by other particles produced in a hard process, and can be vetoed. A study of this effect and discussion of vetoes in our setup is given in section 5.1.

4.2.1 Beam-related backgrounds

- *Muon energy spread:* the Fermilab muon source derives from decays in flight of a monoenergetic 32 GeV pion beam. Prior to delivery to the target, the beam passes through an iron shield where most of the surviving pions are absorbed. Muons lose energy and scatter in the shield, resulting in a spread of muon momenta exiting the shield, from 10 to 30 GeV with a significant low-energy tail below 3 GeV from hard bremsstrahlung in the shield [62]. In our proposal, the tagging tracker is designed to measure the momentum of each incoming muon. While it will be difficult to trigger on this momentum and neglect any events with low-energy muons, such events can be discarded in the analysis phase. The tagging tracker must be designed such that the momentum measurement is sufficiently accurate to identify muons with energy below the desired recoil threshold to an accuracy of 10^{-10} for Phase 1 and 10^{-13} for Phase 2; we discuss the performance of the tracker in section 5.1.

- *Pion contamination:* pions which decay in the target region can fake a signal because the pion mass is so close to the muon mass, and the muon beam is not mono-energetic. Specifically, the initial beam particle momentum is measured solely by track curvature in the tagging tracker region, and pions of slightly lower energy are degenerate with muons of slightly higher energy. Pions which decay in the bulk of the tagging tracker region can be distinguished from muons by the kink in the track curvature resulting from the lower-energy muon decay product. The background events will be pions which decay between the second-to-last tagging tracking layer and the first recoil tracking layer. The target thickness, including all material related to the active elements, is $\ell = 38.8$ cm. Assuming the pions all have energy 32 GeV, $\gamma c\tau = 1790$ m, and so the probability of decay in the target region is $P_{\text{decay}} \approx \ell/\gamma c\tau = 2 \cdot 10^{-4}$. Thus the pion contamination must be less than $5 \cdot 10^{-7}$ for Phase 1, and less than $5 \cdot 10^{-10}$ for Phase 2. The current design of the test beam facility includes 12 feet of iron, which should be sufficient to provide a beam purity of $5 \cdot 10^{-7}$ for Phase 1 (see section 5.2.1). For Phase 2, an additional hadron absorber can be placed between the beam source and the tagging tracker, as suggested in [11].

4.2.2 Reducible backgrounds

- *Single bremsstrahlung:* $\mu N \to \mu N \gamma$. This background results from a hard photon escaping detection in the calorimeter. For a sufficiently thick calorimeter, the single-photon bremsstrahlung $eN \to eN\gamma$ was shown to be negligible for LDMX, and will be further suppressed by $(m_e/m_\mu)^2$ for muon beams.

- *Bremsstrahlung-initiated hadronic events:* $\mu N \to \mu N \gamma$, $\gamma N \to$ hadrons. Many hadronic events will deposit a significant amount of energy into the active target or will produce easily identifiable secondaries. Rare processes such as $\gamma p \to \pi^+ n$, $\gamma N \to (\rho, \omega, \phi)N \to \pi^+\pi^- N$, and $\gamma n \to n\bar{n}n$ can fake signal events depending on the probability of the hadronic calorimeter missing the final-state hadrons. Again, since these processes are photon-initiated, they are suppressed by $(m_e/m_\mu)^2$ compared to the case of electron beams, and thus can be mitigated with relatively mild

background rejection; with veto inefficiencies $\lesssim 0.01$ for charged pions and $\lesssim 0.1$ for neutrons, these processes will give less than 1 fake event per 10^{15} incident muons. This informs the detector performance requirements for Phase 2.

- *Muon pair production:* $\mu N \to \mu N \mu^+ \mu^-$ or $\mu N \to \mu N \gamma$, $\gamma N \to N \mu^+ \mu^-$. This rate is small for Phase 1 luminosity, but is the dominant reducible background for the Phase 2 muon beam. Indeed, the kinematics of the Bethe-Heitler "trident" diagram are peaked in the region of phase space where one muon is produced collinear with the incoming muon [63], but this background is reducible since the majority of the time the remaining muon can be tagged, and/or the presence of two MIPs can be seen in the tracker. We discuss the performance of the tracker further in section 5.1.2, and include a detailed discussion of the pair production background in appendix B.

4.2.3 Irreducible backgrounds

- *Irreducible neutrino pair production:* $\mu N \to \mu N \nu \overline{\nu}$. The only single process resulting in real missing energy relevant for muon beams is Z-mediated neutrino pair production, but the small cross section at $15\,\text{GeV}$ beam energies renders this rate negligible even for Phase 2.

- *Moller + CCQE:* $\mu e \to \mu e$, $ep \to n \nu_e$. For this process, the incident muon scatters a target electron, which acquires a large fraction of the incident beam energy and then undergoes a charged-current quasi-elastic (CCQE) scatter off a proton. The resulting final state contains one low energy muon and a majority of the beam energy carried away by the n and ν. If the neutron is not vetoed, this process is an irreducible background, whose probability per muon is conservatively estimated to be

$$P_{\mu e \to \mu e + \text{CCQE}} = (\alpha \, n_e G_F \ell)^2 \frac{m_p}{2 m_e} \log \frac{E_\mu}{E_{\text{cut}}} \simeq 5 \times 10^{-15}, \qquad (4.1)$$

where, for a tungsten target, the electron density is $n_e = 4.6 \times 10^{24}\,\text{cm}^{-3}$, the total length of tungsten is $X_0 = 17.5\,\text{cm}$, and $E_{\text{cut}} = E_\mu/2$ is signal trigger threshold — see appendix C for a derivation and discussion. Thus, this background is irrelevant for the luminosities we consider in this paper.

Finally, we mention a class of background processes which are irreducible for electron beams but negligible for muon beams. CCQE with final-state neutrinos, $eN \to N' \nu_e$, accompanied by a γ or π^0 in the final state, can pose a significant background when the γ or π^0 fakes an electron. However, the analogous process for muon beams, $\mu N \to N' \nu_\mu + (\gamma, \pi^0)$, would be easily vetoed by the target and ECAL signals.

4.2.4 Target background rates for Phases 1 and 2

The dominant background processes and relative rates are shown in figure 7, along with the maximum veto inefficiencies for Phase 1 and Phase 2 required to achieve the required sensitivity. Because of the reduced bremsstrahlung rate for muon beams and overall smaller total integrated muons on target, there are milder requirements on the hadronic vetoes than
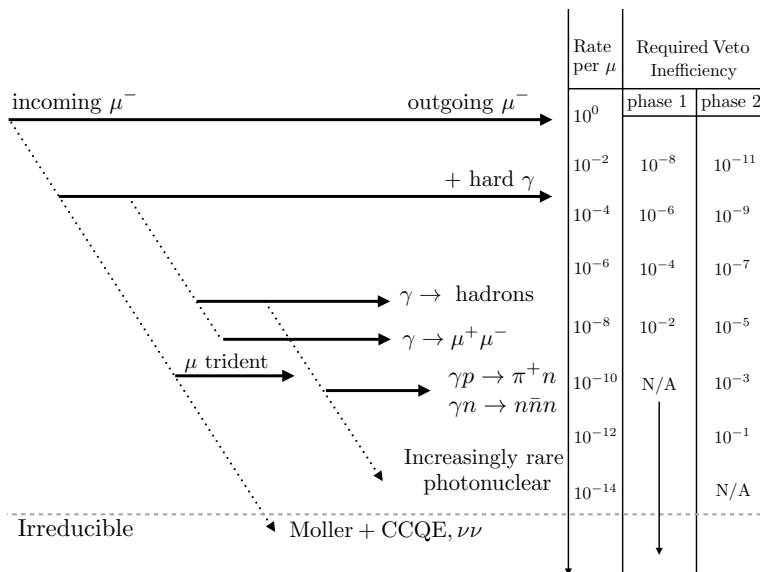
Figure 7. Relative rates of various background processes and the required veto inefficiencies for a zero background experiment for phases 1 and 2.

for electron beams. In summary, we expect that a muon beam experiment based on the LDMX detector concept and technology can be made background-free with 10^{13} muons. While detailed studies of the experimental backgrounds are beyond the scope of this work, we have performed a GEANT [64] simulation of muons on an example target geometry to verify the rates of the rarer reactions, such as photonuclear and muon-nuclear interactions. Discussions of these studies are given in section 5.1.2.

5 Detector and beam parameters

The experimental setup, described in the previous section, requires the ability to identify and track individual muons. Thus a low-current, high-repetition rate muon beam would be ideal, with per-bucket occupancies of $\mathcal{O}(\text{few})$ muons. The sensitivity of the experiment to $g_{S,V}$ is proportional to $\sqrt{N_\mu}$, so acquiring as many muons as possible on target is paramount. To maintain an experimental setup similar to LDMX, which is relatively compact and therefore low-cost, the beam energy should be $\mathcal{O}(\text{several})$ GeV. This drives the recoil muon p_T resolution.

Given these general beam requirements we discuss a specific realistic example in order to define benchmark scenarios for the physics reach of such an experiment. We take as our example the possibility of a muon beam at the Fermilab Accelerator complex [65].

5.1 An LDMX-like detector

We now discuss the most important differences between an LDMX-like detector for a muon beam experiment and its original design with an electron beam. More details of the LDMX design can be found in [15].

5.1.1 LDMX design

The LDMX detector is very similar to the schematic laid out in figure 2 except that the target area is much thinner, approximately 0.1 X_0. The tagging tracker upstream of the target measures the incoming electron momentum. The recoil tracker, along with the electromagnetic calorimeter (ECAL), measures the outgoing electron momentum. The tagging and recoil trackers employ a silicon microstrip detector with strip readout pitch of 60 μm. The tagging tracker and target are inside the bore of a 1 T dipole magnet while the recoil tracker is in the fringe field of the magnet. The electromagnetic calorimeter is a high granularity tungsten-silicon sample calorimeter with a hexagonal geometry and hexagon cell size of 1 cm. The silicon-based calorimeter is required primarily for high electron energy resolution and radiation tolerance as the electromagnetic calorimeter is effectively an instrumented beam dump. The hadronic calorimeter is a steel-scintillator sampling calorimeter with a high sampling fraction for high neutral hadron detection efficiency. The detector is designed to handle incoming electron rates of 40-200 MHz and requires readout electronics accordingly.

The tracking system is based on technology in use at the HPS experiment [66, 67] and the electromagnetic calorimeter is based on technology being developed for the CMS ECAL endcap HL-LHC upgrade [68].

5.1.2 M³ modifications

The primary change to an LDMX-like detector paired with a muon beam is the target area, which now must be much thicker. Let us nominally consider a 50 X_0 thick target region. The most important design feature for such a thick target is that it is active in order to detect muon energy loss from SM interactions within the target. We propose, as a simple extension of the detector technology, to use the same modules for the target in M³ that are employed in the ECAL in LDMX. The high granularity and high sampling fraction make it a good candidate technology for high-efficiency detection of hard bremsstrahlung and photonuclear interactions. Other technologies such as LYSO scintillator are viable options, but here we consider the high granularity silicon-based calorimeter as the primary option as it is a minimal change to the detector technologies needed for the experiment. Compactness is important for fitting inside the magnet, both in the transverse and longitudinal directions. One particular challenge is accomodating the cooling, power, and electronics infrastructure for the active target.

In M³, the ECAL no longer provides information on the muon momentum, so the muon momentum is measured exclusively by the recoil tracker. As a result, the recoil track momentum resolution is paramount to the experiment. A simple estimate for the momentum resolution in the recoil tracker gives

$$\frac{\sigma_p}{p} \approx \frac{\sqrt{8/n}\sigma_x}{h}\frac{p}{qLB_y}, \tag{5.1}$$

where p is the muon momentum, n is the number of measurements by the tracker (nominally 6), σ_x is the hit resolution ($60/\sqrt{12}$ μm), LB_y is the length in the (fringe) magnetic field
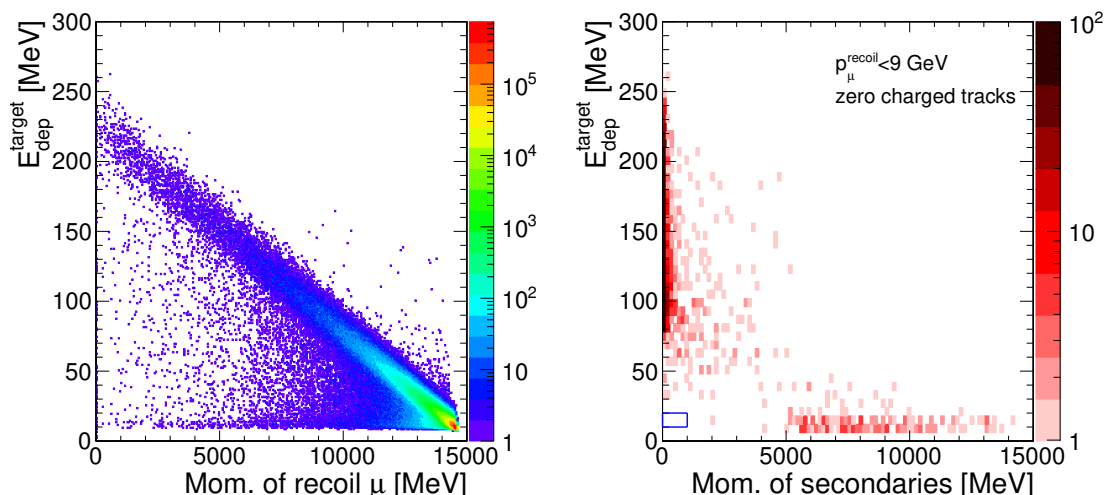
Figure 8. *Left*: distribution of recoil momentum of muons leaving the target versus energy deposited in active elements of the target. *Right*: distribution of energy deposited in the active elements of the target versus non-muon (secondary) momentum leaving the target for events with recoil muon momentum less than 9 GeV and zero charge particle leaving the target at an angle less than 40°. The blue box demarcates the simple signal selection described in the text.

($\sim$0.1 T-m) and h is the lever arm. For a lever arm of 0.25 m, this gives approximately 0.8% momentum resolution. This resolution drives the cut on the muon recoil momentum. The primary signal selection is defined by this cut combined with a minimal amount of energy loss in the active target from the muon (see below). The recoil track momentum resolution can be improved fairly simply by a factor of 2–4 by increasing the recoil track lever arm (at a loss of angular coverage) and/or with a more powerful magnet (2 T).

In order to find a reasonable selection for the signal, we study the energy loss of a muon traversing a 50 X_0 active target modeled after the LDMX electromagnetic calorimeter. The modeling is performed using the LDMX software suite [69] based on the GEANT simulation framework [64]. We model a sample of 10^7 muons with 15 GeV of incident energy on a tungsten-silicon sampling calorimeter where tungsten absorber layers are 1 X_0 per layer. We then examine the energy of the outgoing muon compared to the amount of energy deposited in the silicon. We can cut on the amount of energy deposited in the silicon to eliminate events where the muon underwent a hard interaction in the calorimeter, causing significant muon energy loss. The results are shown in figure 8.

In figure 8 (left), one can see that the incoming 15 GeV muon typically loses about 500 MeV of energy traversing the target. The important metric is the fluctuations of energy loss in the target without any significant energy deposited in the silicon in the target. The silicon energy deposit is shown on the y-axis, and we see that below 12.5 GeV of recoil momentum, there are relatively few events with less than 20 MeV energy deposited. Because all events have a minimum ionizing track traversing the target, very few events populate the region below 10 MeV. The signature of signal events, which consist of a single minimum ionizing muon, will be $10 - 20$ MeV of energy deposited in the active layers of

the target. Given the muon momentum resolution of $\sim 1\%$, we define a signal region which requires a muon outgoing momentum of < 9 GeV, 40 standard deviations away from the incoming beam energy. We combine this with a simplistic selection on the energy of $10 - 20$ MeV deposited in the silicon.

After defining this signal selection, we illustrate in figure 8 (right) the amount of non-muon momentum exiting the back of the target at an angle of less than $40°$. This is plotted against the energy deposited in the silicon on the y-axis. For the region in which muons are depositing a small amount of energy in the target, the secondaries that are exiting the target are mostly a mixture of electrons/photons, e.g. from leakage of electromagnetic showers, and high energy handrons from photonuclear and muon-nuclear reactions in the target. Photons and electrons can be identified with the downstream ECAL by vetoing events with ECAL energies significantly above those expected from muons. Hadrons can also be identified by the ECAL, if their showers start developing early enough, but the downstream HCAL is expected to provide the best sensitivity to hadrons. For hadrons with energies above a GeV, veto rates for pions, protons, and neutrons of up to 10^{-5} are achievable [70] depending on the size and of the HCAL and the absorber thickness per layer. Thus, the events with high energy secondaries in figure 8 (right) are expected to be easily vetoed by requiring an upper bound on the energy measured in the downstream calorimeters. This signal region requirement is represented by the blue box which denotes only energy from a minimum-ionizing particle deposited in the target and no non-muon momentum exiting the back of the target. The plot shows no background events in the signal region for the statistics generated. The effect of these downstream veto on signal efficiency is expected to be of the order of the rate of hard bremsstrahlung and is concidered to be negligible.

Going beyond this simple signal selection, we could use additional information on MIP track shape and anomalous cell hits to further reject muons undergoing significant energy loss and not leaving much energy in the target. While we do not consider them in this work, these are important additional handles that can be used to reject background processes. Beyond the target and the recoil tracking systems, not much modification is needed to the nominal LDMX design. The ECAL and HCAL system will continue to be important in rejecting debris from photon interactions occurring late in the target or producing hadrons, though (as described in section 4.2) these backgrounds are suppressed with respect to the electron beam scenario. The trigger system may need to be modified slightly for muon rates above ~ 100 kHz, which is the current LDMX event rate capacity; this is most relevant for Phase 2, and we discuss this situation further below. In addition, the current detector configuration should be able to handle multiple muons per event, however, more than $\mathcal{O}(\text{several})$ would result in hit confusion from combinatorics and reduced signal efficiency.

5.2 Fermilab accelerator complex

We now discuss a potential accompanying beamline for the M^3 detector at the Fermilab Accelerator Complex. The experiment imagines a two-phase setup where, with lower total integrated luminosity (muons on target), the first phase will cover the $(g - 2)_\mu$ parameter space and the second will probe the thermal DM parameter space.

5.2.1 Phase 1: MTest beamline

The primary Main Injector beamline provides a proton beam of 120 GeV with an RF frequency of 53 MHz [65]. The Fermilab Test Beam Facility (FTBF) [71] beam is extracted from the Main Injector beam via resonant extraction over 4.2 s spills. The beam is then sent to the MTest or MCenter facilities which receive beams of various energies and particle types.

At the energies we require, the MTest beamline can produce a 10–30 GeV muon beam which is produced from pion decays at 32 GeV. For a pion beam of 32 GeV, the beamline has a > 80% composition of charged pions. In muon mode, pions are converted to muons where the final expected rate of 10^5 muons per spill [72]. Pion contamination of the MTest beamline is controlled by an upstream hadronic absorber. For the M^3 experiment, charged pion contamination should be at the level of 10^{-6} which is achievable with the available space in the beamline at MTest. The transverse beam size is roughly a few centimeters in x and y. The time between spills is approximately one minute. Over one week of continuous running, this sums to approximately 1×10^9 muons on target (MOT). Such a performance is achievable with little modification to the Fermilab Accelerator Complex. Therefore we consider the MTest beamline an excellent candidate for Phase 1 of the experiment which probes a large portion of the allowed region for $(g-2)_\mu$. In Phase 1, we plan for the following experimental parameters:

- Muon Energy = 15 GeV

- Muon Intensity = 10^{10} MOT (muons on target)

- Target Thickness = 50 X_0 (about 25 cm for tungsten/silicon)

This program could be achieved on a timescale of roughly a few months. There are no additional detector requirements beyond what is described in section 5.1. Given the estimated background rates in figure 7 and the muon repetition rate at MTest, a zero-background experiment is achievable.

5.2.2 Phase 2: neutrino (NM4) beamline

In order to achieve higher incident pion rates, improved radiation shielding for the MTest target area is required. A candidate would be a similar beamline, called "Neutrino Line", which similarly performs resonant extraction of beam from the Main Injector but sends the protons down another beamline which currently hosts the SeaQuest experiment. This beamline is specified to handle $10^6 - 10^7$ muons per spill which could greatly increase the sensitivity of the M^3 technique. However, the current Neutrino Line would require modifications to beamline itself to enable muon production. With this higher integrated luminosity of muons on target, we can now achieve enough MOT to probe the thermal DM phase space.

In Phase 2, we plan for the following experimental parameters:
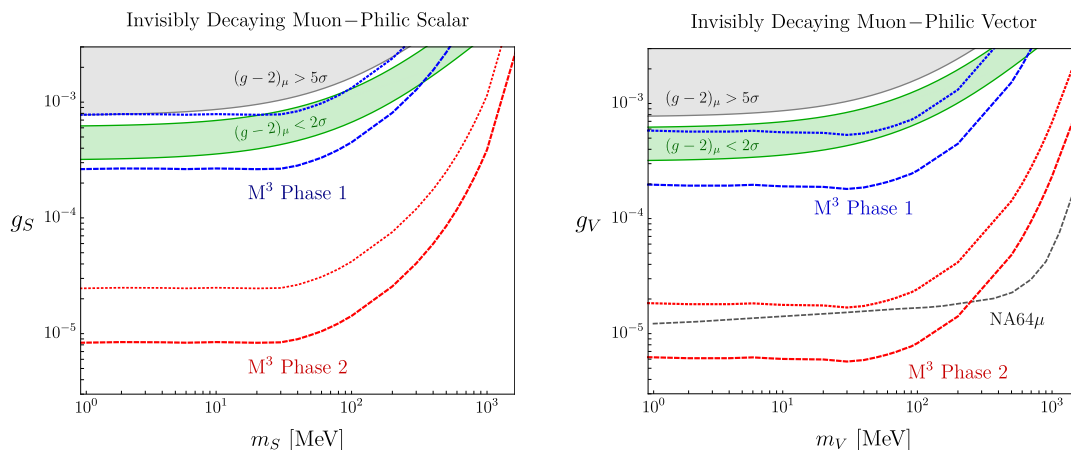
- Muon Energy = 15 GeV

Figure 9. Parameter space for a muon-philic scalar S (*left*) or vector V (*right*) particle as described in section 2.1. The green bands represent the parameter space for which such particles can reconcile the $(g-2)_\mu$ anomaly to within 2σ of the measured value. Also shown are the M^3 projections for Phases 1 and 2 involving the Fermilab test beam facility and the Neutrino (NM4) beamline respectively (see section 5.2 for more details). Here the for both Phases 1 and 2 the dashed (dotted) lines represent the 2σ statistical sensitivity for $0(100)$ background events, respectively. For the vector plot, the gray dashed curve is the NA64μ projection for an invisibly decaying vector particle, which uses the projections in [11], but rescales the region $m_V > 2m_\mu$ to ensure $\mathrm{BR}(V \to \text{invisible}) = 1$; a comparable analysis for scalars at NA64μ would also cover new parameter space. Note also that for $m_{V,S} < 1\,\mathrm{MeV}$, the new invisibly decaying particles will be in thermal equilibrium with the Standard Model during BBN and increase N_{eff}, so we do not consider this regime.

- Muon Intensity = 10^{13} MOT

- Target Thickness = $50\ X_0$

This program would be achieved in a longer timescale than Phase 1 depending on the achievable muon instantaneous luminosity. With a higher muon instantaneous luminosity, a trigger is needed to reduce the event rate to a more manageable 100 kHz. This will require additional R&D into the hardware for fast track reconstruction in order to keep the readout rate manageable. However, even in this scenario, background rates are expected to be manageable.

6 Projected sensitivities

In this section, we discuss the sensitivity of the experimental proposal. We draw limit curves assuming the 95% c.l. Poisson upper limit of 3.0 events. For our signal projection, we assume a 50 X_0 tungsten target and generate a sample of $\mu N \to \mu N(S,V)$ simulated events for various choices of S, V masses using a modified version of `MadGraph` which includes nuclear target form factors (see section 3). We also impose energy and angular cuts to keep only those events in which the outgoing muon passes through the forward tracking layers ($p_z > 0$ to avoid backwards scattering) and satisfies $E < E_{\mathrm{cut}} = 9\,\mathrm{GeV}$. For most values of $m_{S,V}$ that we consider, these cuts yield an order-one signal acceptance as shown in figure 5.
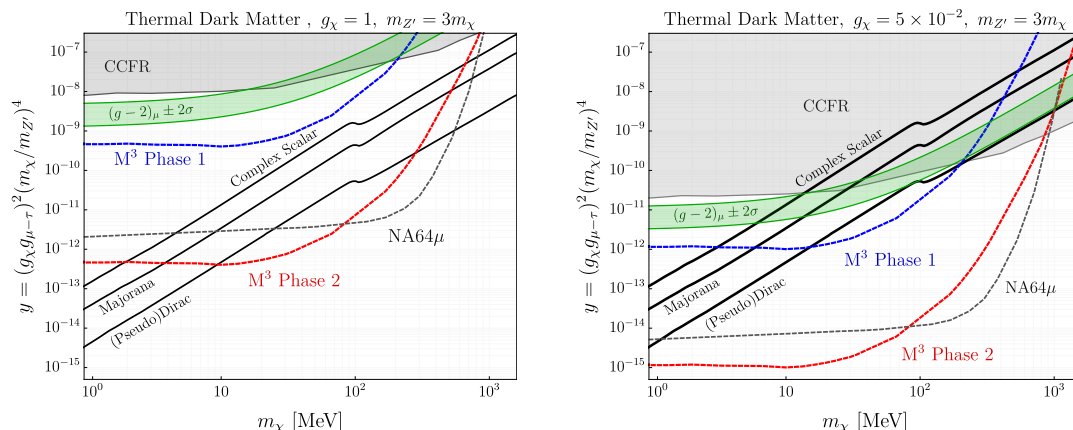
Figure 10. Parameter space for predictive thermal DM charged under $U(1)_{L_\mu - L_\tau}$, for DM charges near the perturbativity limit (*left*) or smaller such that the $(g-2)_\mu$ region overlaps with the thermal relic curves (*right*). Here the relic abundance arises through direct annihilation to SM particles via s-channel Z' exchange. The vertical axis is the product of couplings that sets the relic abundance for a given choice of DM mass and spin (see appendix A). Also plotted are constraints from the neutrino trident process from the CCFR experiment [8, 73] and projected limits from NA64 [13]. Note that there are also bounds on $m_\chi = \mathcal{O}(\text{MeV})$ from $\Delta N_{\text{eff.}}$ that arise from $\chi\bar{\chi} \to \nu\nu$ annihilation during BBN; these bounds differ depending on the choice of DM candidate spin [74, 75] and are not shown here. For the pure Dirac scenario, the annihilation process $\chi\bar{\chi} \to \mu^+\mu^-$ is s-wave, so this process is ruled out by CMB energy injection bounds for $m_\chi > m_\mu$ [54]. Here the M^3 projections correspond to 2σ statistical sensitivity assuming a zero background events; see section 4.2 for a discussion of background rates.

6.1 Phase 1: $(g-2)_\mu$ sensitivity

Figure 9 shows the target parameter space for light new physics contributions to $(g-2)_\mu$. The grey region is excluded because the contribution to $(g-2)_\mu$ is greater than 5σ from the measured value, while green region would resolve the $(g-2)_\mu$ anomaly to within $\pm 2\sigma$. In the zero background case, we see that Phase 1 can completely exclude all new-physics explanations for $(g-2)_\mu$ at mass scales below $100\,\text{MeV}$, and higher for vector mediators and/or looser cuts which could still allow for a background-free search. For comparison, we also show the reach for Phase 2, which can extend the exclusion region to $1\,\text{GeV}$ for both scalars and vectors.

In figure 9, we also illustrate the experimental sensitivity assuming 100 background events. While the initial estimates in this paper shows minimal background contamination, we present the case with non-negligible background levels to illustrate the sensitivity when accounting for non-negligible background levels. In rare process searches, unknown sources of backgrounds could affect the sensitivity of the experiment. For example, there could be rare backgrounds perhaps from difficult to model nuclear interactions in a particular phase space, detector mismodeling, and yet unconsidered upstream muon beam condition variations.

6.2 Phase 2: $U(1)_{L_\mu - L_\tau}$ thermal DM sensitivity

Figure 10 shows the target parameter space for thermal relic DM with a $L_\mu - L_\tau$ mediator. The vertical axis plots the dimensionless variable $y = g_\chi^2 g_{\mu-\tau}^2 (m_\chi/m_{Z'})^4$ which controls

the DM annihilation rate, and the black curves represent the unique value of y for each m_χ which results in the correct DM relic abundance (see appendix A), for DM a complex scalar, Majorana fermion, or (pseudo)-Dirac fermion (see section 2.3). The left panel shows the scenario $g_\chi = 1$ near the perturbativity limit, which corresponds to the weakest possible bounds on this model, while the right panel shows the case $g_\chi = 5 \times 10^{-2}$. In the latter case, there is a region of parameter space compatible with both thermal dark matter and $(g - 2)_\mu$, which can be probed by Phase 1, with the entire viable parameter space for thermal DM probed by Phase 2.[4] Even for the pessimistic case $g_\chi = 1$, a large portion of the parameter space is accessible to Phase 2. We emphasize that muon beam experiments like M^3 are the only terrestrial experiments which can probe such a muon-philic model of DM; direct detection signals are absent, and high-energy collider production cross sections are too small.

Intriguingly, we also find that both Phase 1 and Phase 2 have sensitivity to a class of DM explanations for the $\sim 3.8\sigma$ anomaly reported by the EDGES collaboration [79]. It has been shown that a $\sim 1\%$ subcomponent of DM with a QED millicharge of order $\sim 10^{-3}e$ can cool the SM gas temperature at redshift $z \sim 20$ and thereby account for the magnitude of the observed absorption feature [80]. However, ref. [81] pointed out that such a scenario generically requires dark forces to deplete the millicharge abundance in the early universe to account for the $\sim 1\%$ fraction needed to resolve the anomaly. One viable possibility for generating the appropriate abundance is that the millicharges also carry additional U(1)$_{L_\mu - L_\tau}$ charge, which allows $\chi\chi \to \nu\nu, \mu\mu$ annihilation processes to set the appropriate DM fraction at late times. In figure 11 we present the viable parameter space for this scenario along with various constraints. The solid and dotted black curves are contours of constant χ abundance, and the purple region labeled "EDGES" represents the favored parameter space for appreciable hydrogen gas cooling at $z = 20$ if DM also carries a $\sim 10^{-3}e$ millicharge (see [81] for more details).

7 Conclusion

Light, weakly-coupled particles with muon-philic couplings arise in various extensions of the Standard Model, yet much of their best-motivated parameter space remains experimentally elusive. Of particular interest are scenarios in which such particles resolve the persistent $(g - 2)_\mu$ anomaly and/or serve as mediators that couple to both dark and visible matter. In this paper we have proposed a new muon-beam missing-momentum experiment, M^3, to greatly improve experimental coverage of invisibly-decaying, muon-philic forces with existing beams and detector technology. This technique complements other proposals for probing such forces involving muon beams at CERN [11] and proposed searches of muon-philic forces using rare meson decays with muonic final states [82].

Our setup consists of a relativistic $\sim 15\,\text{GeV}$ muon beam impinging on a fixed, active target with downstream electromagnetic and hadronic calorimeters to veto SM particles that emerge undetected from the target. Upstream and downstream of the target, tracking

---

[4]See also [76, 77] for other models relating thermal DM to $(g - 2)_\mu$, and [78] for a model where axial couplings can explain $B \to K$ anomalies and which could be probed at Phase 2.
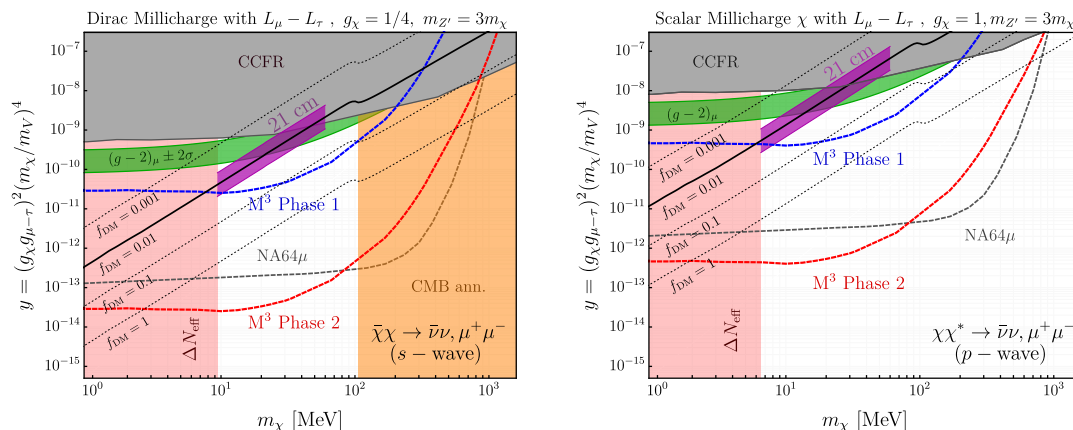
Figure 11. Parameter space and projections for fermion (*left*) and scalar (*right*) millicharge particles which constitute a subcomponent of the total dark matter abundance. If such particles also carry appreciable charge under $U(1)_{L_\mu - L_\tau}$, they can achieve their subdominant relic abundance via $\chi\chi \to \nu\nu, \mu\mu$ annihilation mediated by a virtual s-channel Z'. The black contours represent different fractional DM abundances f_{DM} and the purple region represents the parameter space for which such a DM subcomponent can resolve the $\sim 3.8\sigma$ anomaly in the EDGES 21 cm absorption signal [81]. Also plotted are CCFR constraints from the neutrino trident process [8, 73], $\Delta N_{\text{eff.}}$ bounds on light DM species [74, 75], CMB annihilation bounds for the Dirac fermion scenario with s-wave annihilation [54], and projected sensitivity for the NA64 experiment [13]. For more details about the DM interpretation of the 21 cm signal, see [81]. Here the M^3 projections correspond to 2σ statistical sensitivity assuming a zero background events; see section 4.2 for a discussion of background rates.

layers measure the incident and outgoing muon momenta, with the signal region consisting of evens for which the incident muon loses an order-one fraction of its initial momentum with no detectable energy deposited in the target or veto system. A test phase of our proposal with 10^{10} muons on target can exhaustively test all scenarios in which light, invisibly-decaying muon-philic particles resolve the $(g-2)_\mu$ anomaly. A mature realization of our concept with 10^{13} muons on target can test nearly all predictive sub-GeV dark matter scenarios in which muon-philic forces are responsible for thermal freeze-out. Both phases correspond to presently-feasible experimental options currently available at Fermilab.

A moderate-energy muon beam missing-momentum experiment has several appealing features. First, the cross sections for most background processes (due to SM particles produced, but not observed) are typically smaller than those of an electron-beam experiment with identical target thickness because most such events are initiated by photonuclear interactions with beam-initiated bremsstrahlung radiation. Since the bremsstrahlung rate for a muon beam is reduced by a factor of $(m_e/m_\mu)^2 \sim 2 \times 10^{-5}$ compared to electron beams, our proposal can be implemented with a much thicker target than LDMX, which increases the effective luminosity of the experiment and compensates for the overall reduction of available muons (relative to an ideally-executed electron beam effort). Second, by operating at $\sim 15\,\text{GeV}$, a tenth of the nominal muon beam energy of NA64 at the CERN SPS, our proposal allows for much greater beam deflection for recoiling muons as

they pass through the tracking layers surrounding the active target. This feature allows the experiment to operate on a much smaller ($\sim$ few m) length scale. Furthermore, lower beam energies require less powerful magnetic fields to achieve comparable muon angular deflections and hence comparable momentum resolution.

In this paper, we have primarily motivated our new technique by considering the experimental reach for invisibly decaying muon-philic particles produced in muon-nucleus collisions inside the target. Indeed, a compelling interpretation of this signature is that the new particles decay to dark matter, thereby enabling the technique to probe the couplings that are responsible for setting the relic abundance via thermal freeze-out. However, our case studies do not exhaust the full physics search potential of the proposed setup (and possible variations thereof). In particular, it would be interesting to study the possibility of cascade signatures or displaced vertices that involve combinations of missing and visible energy. Such processes are qualitatively distinct from background processes initiated in the target and are not covered by the analysis presented here. Furthermore, our setup may also be useful for better understanding rare muon-nucleus scattering processes, which could help constrain neutrino-nucleus interactions of relevance for the emerging long baseline neutrino oscillation program. We leave these (and other) questions for future work.

A   Relic density calculation

To compute the relic density curves presented in figures 10 and 11 in the context of a gauged $L_\mu - L_\tau$ model, we follow the thermal freeze-out prescription in [83]. For a DM particle χ with λ_χ internal degrees of freedom, the thermally averaged annihilation rate can be parametrized as $\langle \sigma |v| \rangle \equiv \sigma_0 x^{-n}$ where σ_0 is a constant, $x \equiv m_\chi / T$, and $n = 0$ and 1 for s and p-wave annihilation, respectively. The asymptotic χ abundance at late times is given by

$$\Omega_\chi h^2 = 1.07 \times 10^9 \frac{c(n+1)x_f^{n+1} \text{ GeV}^{-1}}{(g_{*,S}/g_*^{1/2})m_{Pl}\sigma_0} \ , \tag{A.1}$$

where $c = 1$ if the DM is its own antiparticle and $c = 2$ otherwise, g_* and $g_{*,S}$ are respectively the effective relativistic and entropic degrees of freedom, $m_{\text{Pl}} = 1.22 \times 10^{19}$ GeV is the Planck mass, and the value of x at freeze-out, x_f, satisfies

$$x_f \simeq \log a - \left(n + \frac{1}{2}\right) \log(\log a) \quad , \quad a = 0.038(n+1)(\lambda_\chi/g_*^{1/2})m_{Pl}m_\chi\sigma_0 . \quad (A.2)$$

Since σ_0 cannot be extracted analytically from eqs. (A.1) and (A.2), the system must be solved numerically. The leading order $\chi\bar{\chi} \to f\bar{f}$ annihilation cross sections can be straightforwardly extracted using the procedure in [84]. In the representative models we consider in this paper, $\chi\chi \to \bar{f}f$ annihilation where f is a charged SM fermion, we have

$$\sigma_0 = \frac{g_\chi^2 g_f^2 (m_f^2 + 2m_\chi^2)}{k\pi\left[(m_{Z'}^2 - 4m_\chi^2)^2 + m_{Z'}^2\Gamma_{Z'}^2\right]}\sqrt{1 - \frac{m_f^2}{m_\chi^2}}, \quad (A.3)$$

where $g_{\chi,f}$ is the mediator's coupling to χ or f particles and the parameters (n, c, k) are $(0, 2, 2)$ for a (pseudo-) Dirac fermion, $(1, 2, 12)$ for a complex scalar and $(1, 1, 6)$ for a Majorana fermion. For $\chi\chi \to \nu\nu$ annihilation, the Z' neutrino interaction in eq. (2.6) contains an extra insertion of P_L, so σ_0 is smaller by a factor of 2. In the $m_{Z'} \gg m_\chi \gg m_f$ limit away from the $m_{Z'} \sim 2m_\chi$ resonance, appropriate for most of the parameter space considered in this paper,

$$\sigma_0 \propto \frac{g_\chi^2 g_f^2 m_\chi^2}{m_{Z'}^4} \equiv \frac{y}{m_\chi^2} \quad , \quad y \equiv (g_\chi g_f)^2 \left(\frac{m_\chi}{m_{Z'}}\right)^4, \quad (A.4)$$

so for each value of m_χ, there is a unique value of the dimensionless parameter y which yields thermal freeze-out. The definition of y used in this context is adapted from the notation used in [85, 86]. Note that in our numerical results, the appropriate σ_0 includes all available annihilation channels, whereas eq. (A.3) represents only one possible final state. Note also that for the gauged $L_\mu - L_\tau$ results in figure 10, we take the SM fermions to carry unit charge and identify $g_f \to g_{\mu-\tau}$, but allow $g_\chi \neq g_{\mu-\tau}$. Because the DM is vectorlike under the gauge group, all charge assignments are consistent with anomaly cancellation with no additional matter required.

B   Reducible trident background

A potential reducible background is muon pair creation. Two diagrams contribute to this process, referred to as the radiative trident process and the Bethe-Heitler trident process, shown on the left and right of figure 12, respectively. As discussed in detail in ref. [63], the Bethe-Heitler trident dominates over the radiative trident, and moreover has a forward singularity where one muon is energetic and collinear with the recoiling muon, and the second muon is soft. This can potentially mimic the signal of a single outgoing muon with significant energy loss if the collinear muon is the same sign as the recoiling muon and has nearly identical momentum, thus producing only a single track to within the angular and spatial resolution of the tracker, and the second muon is soft and/or wide-angle. Indeed, ref. [11] finds that the trident is a dominant background in their proposal.
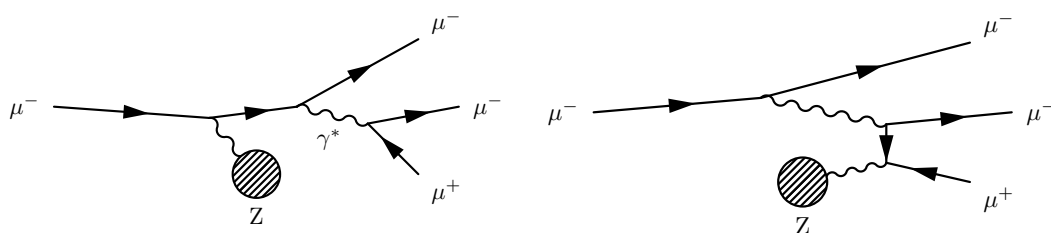
Figure 12. Feynman diagrams for muon pair creation: radiative trident (*left*) and Bethe-Heitler trident (*right*). If this process takes place inside the target, it can mimic the signal if the opposite-sign muon is not detected and the same-sign pair are collinear with similar momenta (and thus not resolved as separate particles based on deflection).

Following the method of ref. [63], we simulated the Bethe-Heitler diagram only (ignoring interference effects) in `MadGraph`, selecting events for which two same-sign muons were collinear to within angular resolution σ_θ and the momenta were matched to resolution σ_p. Figure 13 (left) shows the event rate for this potential background as a function of the signal energy threshold for various choices of the energy and angular resolutions. The background rate saturates above a cut at 50% of the beam energy because if the two collinear muons have nearly identical energy, each must carry less than half the beam energy, thus trivially satisfying the cut. For fairly conservative resolutions $\sigma_p = 5\%$, $\sigma_\theta = 0.5°$, the maximum trident rate is about 3 events for 10^{10} MOT. Moreover, if the third muon is not too soft or wide-angle, it can be detected in the tracker and vetoed. Figure 13 (right) shows the energy spectrum of the non-collinear muon.

Crucially, our proposal has another strong handle on vetoing the trident background. Even if the two same-sign muon tracks are exactly parallel and have exactly the same momentum, both muons are minimum-ionizing particles, and thus would deposit twice the expected energy in the tracker as would a single particle. In each layer of the HCAL, $\mathcal{O}(10)$ photoelectrons will be produced per MIP per layer. With 10 layers, a single muon would yield $\mathcal{O}(100)$ photoelectrons while two parallel muons would yield $\mathcal{O}(200)$; assuming Poisson statistics, this results in a fake rate of less than 10^{-4}. The HCAL will likely have many more than 10 layers, but this is already sufficient to bring the Bethe-Heitler rate below 10^{-13} per MOT, rendering this background negligible for both Phase 1 and Phase 2.

C Irreducible CCQE background

The signal considered in this paper is $\mu(E_{\text{beam}})N \to \mu(E \ll E_{\text{beam}})N(S, V)$, followed by an invisible S or V decay. The main irreducible background for this process is $\mu e \to \mu e$ in which the recoiling electron receives a large fraction of the incident muon energy and undergoes an $ep \to n\nu$ CCQE conversion before it can be vetoed as an electromagnetic shower inside the active target.

The differential cross section for fixed target $\mu e \to \mu e$ scattering is

$$\frac{d\sigma}{dE_e} = \frac{2\pi\alpha^2}{m_e E_e^2} \left(1 - \frac{E_e}{E_\mu} + \frac{E_e^2}{2E_\mu^2}\right), \tag{C.1}$$
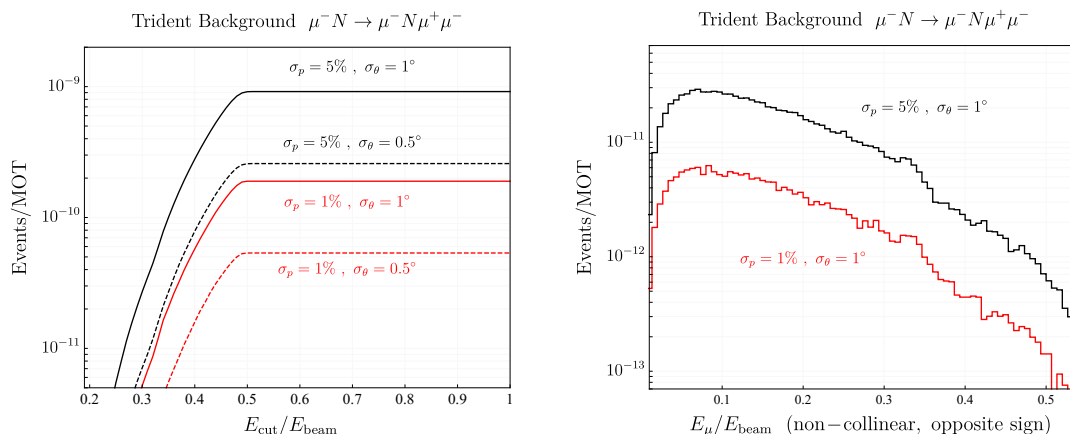
Figure 13. *Left*: total muon trident background events $\mu N \rightarrow \mu N \mu \mu$ for which the two same-sign muons are collinear to within various energy σ_E and angular σ_θ resolutions plotted as a function of the signal energy threshold $E_{\rm cut}$. *Right*: differential energy spectrum for the opposite sign, non-collinear muon for $E_{\rm cut} = 8\,{\rm GeV}$, $\sigma_\theta = 1^\circ$ and different choices of σ_p.

where E_μ is the incident muon beam energy, E_e is electron recoil energy, and both are evaluated in the lab frame. The probability of a CCQE background event with e recoil energies above threshold $E_{\rm cut}$ in a material of electron density n_e and infinitesimal thickness dz is

$$dP_{\mu e \rightarrow \mu e + {\rm CCQE}} = dz \, n_e \int_{E_{\rm cut}}^{E_\mu} dE_e \frac{d\sigma}{dE_e} P_{\rm CCQE} \, \Theta\big(E_\mu - E_\mu^{\rm min}\big) \,, \tag{C.2}$$

where the Θ function enforces energy conservation,

$$E_\mu^{\rm min} = \frac{E_e}{2} \left[1 + \sqrt{\left(1 + \frac{2m_e}{E_e}\right)\left(1 + \frac{2m_\mu^2}{m_e E_e}\right)} \right] \tag{C.3}$$

is the minimum beam energy required to induce an electron recoil with energy E_e, and

$$P_{\rm CCQE} = n_p \sigma_F (\ell - z) \,\,,\,\,\, \sigma_F = \frac{G_F^2 m_p E_e}{2\pi} \,\,, \tag{C.4}$$

is the probability of a CCQE conversion event for an electron with energy E_e passing through a target with proton density $n_p = n_e$ and longitudinal length $(\ell - z)$, corresponding to the amount of material an electron encounters after forward scattering at position $z \in (0, \ell)$ in the one-dimensional approximation. Keeping only the leading term of eq. (C.1) and integrating over z in eq. (C.2), the probability of a CCQE event per incident muon in a tungsten target is

$$P_{\mu e \rightarrow \mu e + {\rm CCQE}} \approx (\alpha \, n_e G_F \ell)^2 \frac{m_p}{2m_e} \log \frac{E_\mu}{E_{\rm cut}} \simeq 5 \times 10^{-15}, \tag{C.5}$$

where, in the last step we have taken $\ell = 50 X_0 = 17.5$ cm, $n_e = 4.6 \times 10^{24}\,{\rm cm}^{-3}$, and $E_{\rm cut} = E_\mu/2$. This estimate is conservative because we have neglected the kinematic

restriction imposed by the Θ function in eq. (C.2), which reduces the overall number of scattering events; evaluating the full expression gives $P_{\mu e \to \mu e + \mathrm{CCQE}} = 3 \times 10^{-16}$ for the same setup.

Furthermore, unlike LDMX with an electron beam, which requires a thin, passive target, using a muon beam with a thicker, active target makes it possible to veto events with electromagnetic showers inside the much larger target. As such, the above estimate is overly conservative by integrating the contribution over the total target length; in practice this process could only happen for lengths of order $\sim X_0$ without being vetoed. Thus, we conclude that the CCQE background is not relevant even for the high luminosity Phase 2 run with a benchmark luminosity of 10^{13} MOT.

Open Access. This article is distributed under the terms of the Creative Commons Attribution License (CC-BY 4.0), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

References

[1] G. Bertone and D. Hooper, *A history of dark matter*, submitted to *Rev. Mod. Phys.*, arXiv:1605.04909 [INSPIRE].

[2] J. Alexander et al., *Dark sectors 2016 workshop: community report*, arXiv:1608.08632 [INSPIRE].

[3] M. Battaglieri et al., *US cosmic visions: new ideas in dark matter 2017: community report*, arXiv:1707.04591 [INSPIRE].

[4] PARTICLE DATA GROUP collaboration, C. Patrignani et al., *Review of particle physics*, *Chin. Phys.* **C 40** (2016) 100001 [INSPIRE].

[5] S.N. Gninenko and N.V. Krasnikov, *The muon anomalous magnetic moment and a new light gauge boson*, *Phys. Lett.* **B 513** (2001) 119 [hep-ph/0102222] [INSPIRE].

[6] E. Ma, D.P. Roy and S. Roy, *Gauged L_μ-L_τ with large muon anomalous magnetic moment and the bimaximal mixing of neutrinos*, *Phys. Lett.* **B 525** (2002) 101 [hep-ph/0110146] [INSPIRE].

[7] M. Pospelov, *Secluded* U(1) *below the weak scale*, *Phys. Rev.* **D 80** (2009) 095002 [arXiv:0811.1030] [INSPIRE].

[8] W. Altmannshofer, S. Gori, M. Pospelov and I. Yavin, *Neutrino trident production: a powerful probe of new physics with neutrino beams*, *Phys. Rev. Lett.* **113** (2014) 091801 [arXiv:1406.2332] [INSPIRE].

[9] M. Endo, K. Hamaguchi, S. Iwamoto and T. Yoshinaga, *Muon g − 2 vs. LHC in supersymmetric models*, *JHEP* **01** (2014) 123 [arXiv:1303.4256] [INSPIRE].

[10] T. Moroi, *The muon anomalous magnetic dipole moment in the minimal supersymmetric standard model*, *Phys. Rev.* **D 53** (1996) 6565 [*Erratum ibid.* **D 56** (1997) 4424] [hep-ph/9512396] [INSPIRE].

[11] S.N. Gninenko, N.V. Krasnikov and V.A. Matveev, *Muon g − 2 and searches for a new leptophobic sub-GeV dark boson in a missing-energy experiment at CERN*, *Phys. Rev.* **D 91** (2015) 095015 [arXiv:1412.1400] [INSPIRE].

[12] NA64 collaboration, D. Banerjee et al., *Search for invisible decays of sub-GeV dark photons in missing-energy events at the CERN SPS*, *Phys. Rev. Lett.* **118** (2017) 011802 [arXiv:1610.02988] [INSPIRE].

[13] S.N. Gninenko, N.V. Krasnikov, M.M. Kirsanov and D.V. Kirpichnikov, *Missing energy signature from invisible decays of dark photons at the CERN SPS*, *Phys. Rev.* **D 94** (2016) 095025 [arXiv:1604.08432] [INSPIRE].

[14] E. Izaguirre, G. Krnjaic, P. Schuster and N. Toro, *Testing GeV-scale dark matter with fixed-target missing momentum experiments*, *Phys. Rev.* **D 91** (2015) 094026 [arXiv:1411.1404] [INSPIRE].

[15] LDMX collaboration, J. Mans, *The LDMX experiment*, *EPJ Web Conf.* **142** (2017) 01020.

[16] C.-Y. Chen, M. Pospelov and Y.-M. Zhong, *Muon beam experiments to probe the dark sector*, *Phys. Rev.* **D 95** (2017) 115005 [arXiv:1701.07437] [INSPIRE].

[17] BaBar collaboration, J.P. Lees et al., *Search for a muonic dark force at BABAR*, *Phys. Rev.* **D 94** (2016) 011102 [arXiv:1606.03501] [INSPIRE].

[18] E. Izaguirre, Y. Kahn, G. Krnjaic and M. Moschella, *Testing light dark matter coannihilation with fixed-target experiments*, *Phys. Rev.* **D 96** (2017) 055007 [arXiv:1703.06881] [INSPIRE].

[19] E. Izaguirre, G. Krnjaic, P. Schuster and N. Toro, *Physics motivation for a pilot dark matter search at Jefferson Laboratory*, *Phys. Rev.* **D 90** (2014) 014052 [arXiv:1403.6826] [INSPIRE].

[20] E. Izaguirre, G. Krnjaic and B. Shuve, *Discovering inelastic thermal-relic dark matter at colliders*, *Phys. Rev.* **D 93** (2016) 063523 [arXiv:1508.03050] [INSPIRE].

[21] DAMIC collaboration, A. Aguilar-Arevalo et al., *Search for low-mass WIMPs in a 0.6 kg day exposure of the DAMIC experiment at SNOLAB*, *Phys. Rev.* **D 94** (2016) 082006 [arXiv:1607.07410] [INSPIRE].

[22] LUX collaboration, D.S. Akerib et al., *Results from a search for dark matter in the complete LUX exposure*, *Phys. Rev. Lett.* **118** (2017) 021303 [arXiv:1608.07648] [INSPIRE].

[23] XENON collaboration, E. Aprile et al., *First dark matter search results from the XENON1T experiment*, *Phys. Rev. Lett.* **119** (2017) 181301 [arXiv:1705.06655] [INSPIRE].

[24] SuperCDMS collaboration, R. Agnese et al., *Results from the Super Cryogenic Dark Matter Search experiment at Soudan*, *Phys. Rev. Lett.* **120** (2018) 061802 [arXiv:1708.08869] [INSPIRE].

[25] PandaX collaboration, M. Xiao, *Status and recent dark matter search results from the PandaX experiment*, in the proceedings of the 12th *Patras Workshop on Axions, WIMPs and WISPs (PATRAS 2016)*, June 20–24, Jeju Island, South Korea (2017).

[26] CRESST collaboration, F. Petricca et al., *First results on low-mass dark matter from the CRESST-III experiment*, talk given at the 15th *International Conference on Topics in Astroparticle and Underground Physics (TAUP 2017)*, July 24–28, Sudbury, Ontario, Canada (2017), arXiv:1711.07692 [INSPIRE].

[27] DarkSide collaboration, P. Agnes et al., *Low-mass dark matter search with the DarkSide-50 experiment*, *Phys. Rev. Lett.* **121** (2018) 081307 [arXiv:1802.06994] [INSPIRE].

[28] R. Essig, J. Mardon and T. Volansky, *Direct detection of sub-GeV dark matter*, *Phys. Rev.* **D 85** (2012) 076007 [arXiv:1108.5383] [INSPIRE].

[29] P.W. Graham, D.E. Kaplan, S. Rajendran and M.T. Walters, *Semiconductor probes of light dark matter*, *Phys. Dark Univ.* **1** (2012) 32 [arXiv:1203.2531] [INSPIRE].

[30] R. Essig et al., *First direct detection limits on sub-GeV dark matter from XENON10*, *Phys. Rev. Lett.* **109** (2012) 021301 [arXiv:1206.2644] [INSPIRE].

[31] Y. Hochberg, Y. Zhao and K.M. Zurek, *Superconducting detectors for superlight dark matter*, *Phys. Rev. Lett.* **116** (2016) 011301 [arXiv:1504.07237] [INSPIRE].

[32] S.K. Lee, M. Lisanti, S. Mishra-Sharma and B.R. Safdi, *Modulation effects in dark matter-electron scattering experiments*, *Phys. Rev.* **D 92** (2015) 083517 [arXiv:1508.07361] [INSPIRE].

[33] R. Essig et al., *Direct detection of sub-GeV dark matter with semiconductor targets*, *JHEP* **05** (2016) 046 [arXiv:1509.01598] [INSPIRE].

[34] Y. Hochberg, M. Pyle, Y. Zhao and K.M. Zurek, *Detecting superlight dark matter with Fermi-degenerate materials*, *JHEP* **08** (2016) 057 [arXiv:1512.04533] [INSPIRE].

[35] Y. Hochberg et al., *Directional detection of dark matter with two-dimensional targets*, *Phys. Lett.* **B 772** (2017) 239 [arXiv:1606.08849] [INSPIRE].

[36] S. Derenzo et al., *Direct detection of sub-GeV dark matter with scintillating targets*, *Phys. Rev.* **D 96** (2017) 016026 [arXiv:1607.01009] [INSPIRE].

[37] R. Essig, T. Volansky and T.-T. Yu, *New constraints and prospects for sub-GeV dark matter scattering off electrons in Xenon*, *Phys. Rev.* **D 96** (2017) 043017 [arXiv:1703.00910] [INSPIRE].

[38] SENSEI collaboration, J. Tiffenberg et al., *Single-electron and single-photon sensitivity with a silicon Skipper CCD*, *Phys. Rev. Lett.* **119** (2017) 131802 [arXiv:1706.00028] [INSPIRE].

[39] Y. Hochberg et al., *Detection of sub-MeV dark matter with three-dimensional Dirac materials*, *Phys. Rev.* **D 97** (2018) 015004 [arXiv:1708.08929] [INSPIRE].

[40] DARKSIDE collaboration, P. Agnes et al., *Constraints on sub-GeV dark matter-electron scattering from the DarkSide-50 experiment*, arXiv:1802.06998 [INSPIRE].

[41] SENSEI collaboration, M. Crisler et al., *SENSEI: first direct-detection constraints on sub-GeV dark matter from a surface run*, *Phys. Rev. Lett.* **121** (2018) 061803 [arXiv:1804.00088] [INSPIRE].

[42] E. Izaguirre, G. Krnjaic, P. Schuster and N. Toro, *New electron beam-dump experiments to search for MeV to few-GeV dark matter*, *Phys. Rev.* **D 88** (2013) 114015 [arXiv:1307.6554] [INSPIRE].

[43] B. Batell, R. Essig and Z. Surujon, *Strong constraints on sub-GeV dark sectors from SLAC beam dump E137*, *Phys. Rev. Lett.* **113** (2014) 171802 [arXiv:1406.2698] [INSPIRE].

[44] BDX collaboration, M. Battaglieri et al., *Dark matter search in a Beam-Dump eXperiment (BDX) at Jefferson Lab*, arXiv:1406.3028 [INSPIRE].

[45] BDX collaboration, M. Battaglieri et al., *Dark matter search in a Beam-Dump eXperiment (BDX) at Jefferson Lab*, arXiv:1607.01390 [INSPIRE].

[46] B. Batell, M. Pospelov and A. Ritz, *Exploring portals to a hidden sector through fixed targets*, *Phys. Rev.* **D 80** (2009) 095024 [arXiv:0906.5614] [INSPIRE].

[47] B. Batell et al., *Leptophobic dark matter at neutrino factories*, *Phys. Rev.* **D 90** (2014) 115014 [arXiv:1405.7049] [INSPIRE].

[48] P. deNiverville, M. Pospelov and A. Ritz, *Observing a light dark matter beam with neutrino experiments*, *Phys. Rev.* **D 84** (2011) 075020 [arXiv:1107.4580] [INSPIRE].

[49] MiniBooNE collaboration, R. Dharmapalan et al., *Low mass WIMP searches with a neutrino experiment: a proposal for further MiniBooNE running*, arXiv:1211.2258 [InSPIRE].

[50] MiniBooNE collaboration, A.A. Aguilar-Arevalo et al., *Dark matter search in a proton beam dump with MiniBooNE*, *Phys. Rev. Lett.* **118** (2017) 221803 [arXiv:1702.02688] [InSPIRE].

[51] P. Coloma, B.A. Dobrescu, C. Frugiuele and R. Harnik, *Dark matter beams at LBNF*, *JHEP* **04** (2016) 047 [arXiv:1512.03852] [InSPIRE].

[52] C. Frugiuele, *Probing sub-GeV dark sectors via high energy proton beams at LBNF/DUNE and MiniBooNE*, *Phys. Rev.* **D 96** (2017) 015029 [arXiv:1701.05464] [InSPIRE].

[53] B.A. Dobrescu and C. Frugiuele, *GeV-scale dark matter: production at the main injector*, *JHEP* **02** (2015) 019 [arXiv:1410.1566] [InSPIRE].

[54] Planck collaboration, P.A.R. Ade et al., *Planck 2015 results. XIII. Cosmological parameters*, *Astron. Astrophys.* **594** (2016) A13 [arXiv:1502.01589] [InSPIRE].

[55] M. Bauer, P. Foldenauer and J. Jaeckel, *Hunting all the hidden photons*, *JHEP* **07** (2018) 094 [arXiv:1803.05466] [InSPIRE].

[56] J. Heeck and W. Rodejohann, *Kinetic and mass mixing with three abelian groups*, *Phys. Lett.* **B 705** (2011) 369 [arXiv:1109.1508] [InSPIRE].

[57] S.N. Gninenko and N.V. Krasnikov, *Probing the muon $g_\mu - 2$ anomaly, L_μ-L_τ gauge boson and Dark Matter in dark photon experiments*, *Phys. Lett.* **B 783** (2018) 24 [arXiv:1801.10448] [InSPIRE].

[58] Y.-S. Liu, D. McKeen and G.A. Miller, *Validity of the Weizsäcker-Williams approximation and the analysis of beam dump experiments: Production of a new scalar boson*, *Phys. Rev.* **D 95** (2017) 036010 [arXiv:1609.06781] [InSPIRE].

[59] Y.-S. Liu and G.A. Miller, *Validity of the Weizsäcker-Williams approximation and the analysis of beam dump experiments: production of an axion, a dark photon, or a new axial-vector boson*, *Phys. Rev.* **D 96** (2017) 016004 [arXiv:1705.01633] [InSPIRE].

[60] S.N. Gninenko, D.V. Kirpichnikov, M.M. Kirsanov and N.V. Krasnikov, *The exact tree-level calculation of the dark photon production in high-energy electron scattering at the CERN SPS*, *Phys. Lett.* **B 782** (2018) 406 [arXiv:1712.05706] [InSPIRE].

[61] J.D. Bjorken, R. Essig, P. Schuster and N. Toro, *New fixed-target experiments to search for dark gauge forces*, *Phys. Rev.* **D 80** (2009) 075018 [arXiv:0906.0580] [InSPIRE].

[62] D. Jensen and M. Rominsky, personal communication.

[63] R. Essig, P. Schuster, N. Toro and B. Wojtsekhowski, *An electron fixed target experiment to search for a new vector boson A' decaying to e^+e^-*, *JHEP* **02** (2011) 009 [arXiv:1001.2557] [InSPIRE].

[64] GEANT4 collaboration, S. Agostinelli et al., *GEANT4 — a simulation toolkit*, *Nucl. Instrum. Meth.* **A 506** (2003) 250 [InSPIRE].

[65] B.C. Brown et al., *The Fermilab main injector: high intensity operation and beam loss control*, *Phys. Rev. ST Accel. Beams* **16** (2013) 071001 [arXiv:1307.2934] [InSPIRE].

[66] HPS collaboration, A. Celentano, *The heavy photon search experiment at Jefferson Laboratory*, *J. Phys. Conf. Ser.* **556** (2014) 012064 [arXiv:1505.02025] [InSPIRE].

[67] HPS collaboration, P. Hansson Adrian, *The Silicon Vertex Tracker for the heavy photon search experiment*, *IEEE Nucl. Sci. Symp. Med. Imag. Conf.* (2016) 7581862 [arXiv:1511.07844] [INSPIRE].

[68] D. Contardo et al., *Technical proposal for the Phase-II upgrade of the CMS detector*, CMS-TDR-15-02 (2015).

[69] *Light dark matter experiment*, https://confluence.slac.stanford.edu/display/ MME/Light+Dark+Matter+Experiment (2016).

[70] LDMX collaboration, T. Åkesson et al., *Light Dark Matter eXperiment (LDMX)*, arXiv:1808.05219 [INSPIRE].

[71] *Fermilab test beam facility*, http://ftbf.fnal.gov.

[72] D. Denisov, V. Evdokimov, S. Lukić and P. Ujić, *Muon beam at the Fermilab test beam area*, arXiv:1608.03453 [INSPIRE].

[73] CCFR collaboration, S.R. Mishra et al., *Neutrino tridents and WZ interference*, *Phys. Rev. Lett.* **66** (1991) 3117 [INSPIRE].

[74] C. Boehm, M.J. Dolan and C. McCabe, *A lower bound on the mass of cold thermal dark matter from Planck*, *JCAP* **08** (2013) 041 [arXiv:1303.6270] [INSPIRE].

[75] K.M. Nollett and G. Steigman, *BBN and the CMB constrain neutrino coupled light WIMPs*, *Phys. Rev.* **D 91** (2015) 083505 [arXiv:1411.6005] [INSPIRE].

[76] K. Kowalska and E.M. Sessolo, *Expectations for the muon g − 2 in simplified models with dark matter*, *JHEP* **09** (2017) 112 [arXiv:1707.00753] [INSPIRE].

[77] L. Calibbi, R. Ziegler and J. Zupan, *Minimal models for dark matter and the muon g − 2 anomaly*, *JHEP* **07** (2018) 046 [arXiv:1804.00009] [INSPIRE].

[78] F. Sala and D.M. Straub, *A new light particle in B decays?*, *Phys. Lett.* **B 774** (2017) 205 [arXiv:1704.06188] [INSPIRE].

[79] J.D. Bowman et al., *An absorption profile centred at 78 MegaHertz in the sky-averaged spectrum*, *Nature* **555** (2018) 67 [INSPIRE].

[80] R. Barkana, *Possible interaction between baryons and dark-matter particles revealed by the first stars*, *Nature* **555** (2018) 71 [arXiv:1803.06698] [INSPIRE].

[81] A. Berlin, D. Hooper, G. Krnjaic and S.D. McDermott, *Severely constraining dark matter interpretations of the 21 cm anomaly*, *Phys. Rev. Lett.* **121** (2018) 011102 [arXiv:1803.02804] [INSPIRE].

[82] M. Ibe, W. Nakano and M. Suzuki, *Constraints on L_μ-L_τ gauge interactions from rare kaon decay*, *Phys. Rev.* **D 95** (2017) 055022 [arXiv:1611.08460] [INSPIRE].

[83] E.W. Kolb and M.S. Turner, *The Early universe*, *Front. Phys.* **69** (1990) 1 [INSPIRE].

[84] J.D. Wells, *Annihilation cross-sections for relic densities in the low velocity limit*, hep-ph/9404219 [INSPIRE].

[85] E. Izaguirre, G. Krnjaic, P. Schuster and N. Toro, *Analyzing the discovery potential for light dark matter*, *Phys. Rev. Lett.* **115** (2015) 251301 [arXiv:1505.00011] [INSPIRE].

[86] G. Krnjaic, *Probing light thermal dark-matter with a Higgs portal mediator*, *Phys. Rev.* **D 94** (2016) 073009 [arXiv:1512.04119] [INSPIRE].