

Light Dark Matter eXperiment (LDMX): A Preliminary Design Study

Owen Colegrove,¹ Bertrand Echenard,² Norman Graf,³ Joshua Hiltbrand,⁴ David Hitlin,² Joseph Incandela,⁵ John Jaros,³ Robert Johnson,⁶ Gordan Krnjaic,⁷ Jeremiah Mans,⁴ Takashi Maruyama,³ Jeremy McCormick,³ Omar Moreno,³ Timothy Nelson,³ Philip Schuster,^{3,8} Natalia Toro,^{3,8} Nhan V Tran,⁷ and Andrew Whitbeck⁷

¹*University of California, Santa Barbara, Santa Barbara, CA 93106, USA*

²*California Institute of Technology, Pasadena, CA 91125, USA*

³*SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA*

⁴*University of Minnesota, Minneapolis, MN 55455, USA*

⁵*University of California at Santa Barbara, Santa Barbara, CA 93106, USA*

⁶*Santa Cruz Institute for Particle Physics,*

University of California at Santa Cruz, Santa Cruz, CA 95064, USA

⁷*Fermi National Accelerator Laboratory, Batavia, IL 60510, USA*

⁸*Perimeter Institute for Theoretical Physics, Waterloo ON N2L 2Y5, Canada*

(Dated: October 17, 2016 : Emailed to Department of Energy Office of High Energy Physics)

This note presents the physics motivation and a preliminary conceptual design study for an experiment to search for low-mass dark matter at the DASEL facility, now under study for the LCLS-II accelerator complex at SLAC. The experiment would use missing momentum to search for dark matter produced via “dark bremsstrahlung” by scattering electrons in a thin target. To identify rare signal events, the LDMX experiment would individually tag incoming beam-energy electrons, unambiguously associate them with low energy, moderate transverse-momentum recoils of the incoming electron, and establish the absence of an energetic forward photon or any additional forward-recoiling charged particles or neutral hadrons. Ultimately, LDMX aims to probe thermal dark matter over most of the viable sub-GeV mass range to a decisive level of sensitivity, with orders of magnitude more sensitivity than any previous or currently envisioned experiment. To integrate adequate statistics, LDMX requires a beamline that can deliver of order 10^8 multi-GeV electrons per second on target, which is compatible with the design parameters being proposed at SLAC with the DASEL facility. Furthermore, the experimental apparatus requires tracking and electromagnetic calorimetry that can perform well with such an event rate. The LDMX conceptual design makes use of a low-mass tracking system to tag incoming electrons with high purity, and to cleanly reconstruct electron recoils. A high-speed, high-granularity, Si-W calorimeter with MIP sensitivity is used to reject the high rate bremsstrahlung background at trigger level, and to work in tandem with a scintillator-based hadron calorimeter to veto rare photo-nuclear reactions. The LDMX conceptual design leverages new calorimeter technology under development for the HL-LHC, as well as existing tracking technology and experience from the HPS experiment. This note summarizes the current status of the LDMX conceptual design with layout and performance studies.

Contents

I. Introduction	4
II. Science Goals	5
A. Thermal Dark Matter: A Target of Opportunity	5
B. Primary Goal: Light Thermal Dark Matter (LDM)	6
C. A Benchmark Scenario for LDM	7
D. Searching for Light Dark Matter Production	10
III. Kinematics and Rate of Dark Matter Signal Production Process	11
IV. Detector Concept	15
A. Beamline	17
B. Overview of Tracking and Target Systems	18
C. Tagging Tracker	19
D. Recoil Tracker	21
E. Forward Electromagnetic Calorimeter	22
F. Hadronic Veto System	26
G. Trigger System	27
H. DAQ	28
1. Overall DAQ Architecture	29
2. Tracking DAQ	30
3. ECAL DAQ	31
4. HCAL and Target DAQ	32
V. Physics and Detector Simulation	34
A. Simulation of the Tagging and Recoil Trackers	34
1. Hit Reconstruction	34
2. Track Reconstruction	34
B. Simulation of the Calorimetry Systems	36
1. Simulation of the Electromagnetic Calorimeter	36
2. Digitization of the Hadronic Veto System	37
C. External Physics Generator for Signal Reaction (Dark Matter Production)	37
D. Photonuclear Model	38
VI. Performance Studies	39
A. Tagging Tracker Performance	39
B. Recoil Tracker Performance	42
C. Forward Electromagnetic Calorimeter Performance	47
D. Hadronic Veto System	50
E. Trigger Performance	53
VII. Conclusion	55
A. Light Dark Matter Theory	56
1. Thermal Origin	56
2. “WIMP” Dark Matter	57

3. Light Dark Matter (LDM)	58
4. Predictive LDM Targets	59
5. Current Bounds on LDM	60
6. Future Bounds on LDM	62
References	63

I. INTRODUCTION

Discovering the particle nature of Dark Matter is perhaps the most pressing challenge facing elementary particle physics today. Among the simplest possibilities is one in which dark matter arose as a thermal relic from the hot early Universe, which only requires small non-gravitational interactions between dark and familiar matter, and is robustly viable over the MeV to TeV mass range. Testing the hypothesis that the dark matter abundance arises from weak W- and Z- boson mediated interactions has been the primary focus of direct and indirect detection experiments to date, which are most sensitive to dark matter particles with masses ranging from a few GeV to a TeV. However, the lower mass range of MeV to GeV, where the most stable forms of ordinary matter are found, has remained stubbornly difficult to test, and is largely unexplored.

In the last several years, powerful ideas to probe “light dark matter” (LDM) in the sub-GeV mass range have emerged from efforts to test the well-motivated possibility that dark matter is part of a dark sector that is neutral under all Standard Model (SM) forces. Building on several years of experience with the first generation of dark sector experiments, the community has identified probing thermal dark matter in the sub-GeV mass range as a key science priority [1]. The most powerful experimental strategies for addressing this goal call for the use of electron beams to produce dark matter in fixed-target collisions, making use of missing mass, energy, or momentum to identify this process [2, 3]. Already, NA64 at CERN has carried out a test run for an experiment using missing energy as the identifying signature. That experiment promises to reach a dark matter sensitivity surpassing all existing constraints by 2020 [1, 4].

The “Light Dark Matter eXperiment” (LDMX) described in this note aims to explore dark matter interactions with sensitivity that extends four orders of magnitude beyond NA64’s eventual 2020 reach and so would be powerful enough to decisively test most predictive scenarios of thermal dark matter in the sub-GeV mass range. To achieve sufficient statistics, LDMX proposes to use a low-current (\sim pA) but high-repetition electron beam with multi-GeV energy. A beam with 10^8 electrons/second on target and energy in the 4 to 10 GeV range can explore most of the sub-GeV dark matter parameter space, while remaining below threshold for production of neutrinos, which are an irreducible background. The DASEL beamline at SLAC, discussed in a separate LOI [5], is ideal for this purpose. We envision running LDMX in SLAC’s End Station A.

An electron beam incident on a thin target can produce dark matter particles through a “dark bremsstrahlung” process, in which most of the incident electron’s energy is typically carried away by the invisible dark matter. To search for this process, LDMX reconstructs the kinematics of each beam electron both up- and down-stream of the target using low-mass tracking. The up-stream tracker tags the incoming beam electrons while the down-stream tracker selects the low-energy, moderate transverse-momentum recoils of the beam electrons. The calorimetry then vetoes events with an energetic forward photon or any additional forward-recoiling charged particles or neutral hadrons. For the first-phase experiment described here, the total required luminosity is a modest 0.2 pb^{-1} corresponding to 10^{14} tagged electrons on target. Nonetheless, because each electron passes through the detector, the experiment must contend with high event rates in the tracker and electromagnetic calorimeter. Therefore, LDMX requires low-mass tracking that provides high-purity tagging for incoming electrons and clean, efficient reconstruction of recoils in a high-rate environment. The calorimetry for LDMX must be simultaneously fast enough to support this high rate of background events, most of which are “easy” to reject based on their high electromagnetic energy deposition, and sensitive enough to reject rare but subtle processes where a hard bremsstrahlung photon undergoes a photo-nuclear reaction in the target or in the calorimeter itself. These simultaneous requirements call for a high-speed, high-granularity calorimeter with

MIP sensitivity to identify photo-nuclear products, used in conjunction with a hadron calorimeter that experiences much lower event rates. As described below, the LDMX concept plans to meet these challenges by leveraging technology under development for the HL-LHC and experience from the HPS experiment.

This note presents a self contained discussion of the dark matter science goals for LDMX (Section II); a summary of the signal kinematics (Section III); a description of the detector design concept and ongoing developments (Section IV); a summary of the simulation methodology (Section V) and performance results (Section VI) from early design studies. Finally, this document concludes with a detailed appendix on light dark matter theory. The primary design studies for LDMX are being carried about by groups at University of California Santa Barbara, University of California Santa Cruz, University of Minnesota, California Institute of Technology, Fermilab, and SLAC. The collaboration is still growing, and has garnered strong interest from international collaborators as well, particularly within INFN. Effort in FY17 will be focused on developing a full technical design for LDMX. It is very likely that some detailed aspects of the design will change from those presented in the preliminary conceptual design described in this note. This is particularly true of the target system and hadronic veto system.

II. SCIENCE GOALS

Dark matter is the dominant particle constituent of the Universe, yet its identity and origin will remain mysterious until we explore its non-gravitational interactions. The primary goal of LDMX is to search for dark matter particles in the sub-GeV mass range — a region that is simultaneously well motivated by the thermal freeze-out hypothesis for the origin of dark matter and experimentally open territory. In simple models, a thermal origin for dark matter implies a *minimum* interaction strength between dark and ordinary matter (as does any thermal contact between dark and ordinary matter in the early Universe) — the “thermal relic target”. While many recent experimental efforts aim to explore dark matter in this mass range, LDMX will have unprecedented sensitivity. In particular, for DM masses below 100 MeV LDMX aims in its first phase to *fully* explore the scalar thermal DM parameter space in its first phase, and the fermion thermal DM parameter space in its second phase. Exploring light dark matter to this level of sensitivity is a science priority recently highlighted in the Dark Sectors 2016 community report [1], and resonates with the dark matter priorities emphasized in the P5 report [6]. This section provides a self-contained discussion of thermal dark matter, realizations of sub-GeV thermal dark matter in particular, and the precise sensitivity targets that this general hypothesis motivates; a more comprehensive discussion of light dark matter may be found in Appendix A (see also [1, 3]). We also briefly survey the various approaches that have been proposed to explore this parameter space, and argue that the missing momentum approach used by LDMX is uniquely well suited to reach the sensitivity target motivated by thermal dark matter. Beyond dark matter science, LDMX will also be capable of supporting a range of important nuclear physics measurements, but we defer that discussion to future reports.

A. Thermal Dark Matter: A Target of Opportunity

Perhaps the simplest, most physically grounded possibility for the origin of dark matter is that its abundance arises from non-gravitational interactions with Standard Model particles in the hot early Universe. This compelling scenario is largely model independent and only requires the DM-




FIG. 1: The allowed mass range over which DM can thermalize with the SM in the early universe and yield the observed relic abundance via annihilation. For masses below $\lesssim 10$ keV, DM is too hot to form the observed structure of the universe on large scales [7] and for masses above $\gtrsim 10$ TeV, a perturbative annihilation rate cannot achieve the correct relic abundance in simple models [8].

SM interaction rate to exceeds the Hubble expansion rate at some point in the early Universe. If such a condition is satisfied, dark matter automatically reaches thermal equilibrium with visible matter and a residual DM abundance is guaranteed. This mechanism for generating the DM abundance is attractive and important for many reasons. (See Appendix A for more detail.)

First, this mechanism is simple and generically realized because equilibrium is hard to avoid even for tiny couplings between dark and visible matter. Therefore, most *discoverable* models of DM with non-gravitational interactions fall into this category. Second, this mechanism implies a minimum annihilation rate in order to avoid producing an overabundance of dark matter at “freeze-out”. This minimum annihilation rate defines a minimum cross section which must be experimentally probed to rule out dark matter of thermal origins. Finally, the thermal DM mass range is well defined – see Fig. 1. The most natural mass range is comparable to that of Standard Model matter, so that DM annihilation into SM final states is generically possible. These features provide strong motivations and clear targets for exploration of thermally produced dark matter.

B. Primary Goal: Light Thermal Dark Matter (LDM)

If DM is realized in the upper half of the thermal mass window $\sim \text{GeV} - 10 \text{ TeV}$, it can be a Weakly Interacting Massive Particle (WIMP) charged under the familiar electroweak force. This has been the traditional focus of dark matter direct and indirect detection experiments, driven in part by the well known connection between WIMPs and supersymmetry (SUSY), whose DM candidate naturally realizes this paradigm. However, despite decades of dedicated searches for DM near the electroweak scale, there has been no discovery in this mass range and the simplest WIMP scenarios are now ruled out by several orders of magnitude. Furthermore, years of null LHC searches for SUSY [9] motivate a broad discovery effort that explores the full range of masses compatible with a thermal origin, which is a powerful, physically grounded organizing principle with minimal model dependence.

The lower half of the thermal mass window, $\sim 10 \text{ keV} - \text{GeV}$, has remained stubbornly difficult to test with traditional experiments designed to probe WIMPs and is not well explored. Indeed, most of the constraints in this mass range derive from theorists’ reinterpretations of existing data gathered for other purposes (see Appendix A). This is particularly unfortunate because the sub-GeV mass range is well-motivated by “hidden sector” (or “dark sector”) scenarios in which dark matter is simply a particle with its own forces and interactions, neutral under the Standard Model, but with sufficient coupling to visible matter that thermal equilibrium is achieved in the early Universe. This mass range is also independently motivated by asymmetric dark matter scenarios,

in which dark matter carries a net particle number in analogy with the baryon asymmetry observed in visible matter. The particle physics community has highlighted these scenarios as among the most important to test, in both the P5 report [6] and in the recent Dark Sectors 2016 community report [1].

The primary science goal of LDMX is to explore dark matter interactions with electrons to a level of sensitivity needed to decisively test most predictive thermal dark matter scenarios over nearly the entire sub-GeV mass range. To this end, LDMX aims to become the first ever dedicated experiment to search for sub-GeV dark matter with sensitivity extending all the way down to the lowest DM masses for which a thermal origin is empirically viable.

C. A Benchmark Scenario for LDM

In the keV-GeV mass range, viable models of LDM have the following properties:

- **Light Forces:** There have to be comparably light force carriers to mediate an efficient annihilation rate for thermal freeze-out (this follows from a simple generalization of the Lee-Weinberg bound [10]).
- **Neutrality:** Both the DM and the mediator must be singlets under the full SM gauge group; otherwise they would have been produced and detected at LEP or at hadron colliders.

These properties single out the hidden sector scenario highlighted in [1, 6], which is the focus of considerable experimental activity. Thus, for the remainder of this note, we will use one of the simplest and most representative hidden sector model in the literature – a dark matter particle charged under a $U(1)$ gauge field (i.e. “dark QED”). We note in passing that LDMX is also sensitive to DM produced through several other vector-mediated interactions, as long as they couple to the electron – see Appendix A for a comprehensive discussion.

We define the LDM particle to be χ and the $U(1)$ gauge boson (popularly called a “dark photon”) A' with lagrangian

$$\mathcal{L} = i\bar{\chi}\not{\partial}\chi + m_\chi\bar{\chi}\chi + g_D A'_\mu \bar{\chi}\gamma^\mu\chi - \frac{1}{4}F'_{\mu\nu}F'^{\mu\nu} + \frac{\epsilon}{2}F'_{\mu\nu}F^{\mu\nu} + \frac{m_{A'}^2}{2}A'_\mu A'^\mu, \quad (1)$$

where $F'_{\mu\nu} = \partial_\mu A'_\nu - \partial_\nu A'_\mu$, $\epsilon \ll 1$ is the kinetic mixing parameter, which controls A' mixing with the SM photon, and $g_D \equiv \sqrt{4\pi\alpha_D}$ is the A' coupling to the DM. Although we consider fermionic DM above for concreteness, LDMX sensitivity to scalar DM is essentially identical.

After diagonalizing the kinetic mixing interaction, the dark photon A' acquires a coupling to the SM electromagnetic current

$$\mathcal{L} \rightarrow i\bar{\chi}\not{\partial}\chi + m_\chi\bar{\chi}\chi + A'_\mu \left(g_D \bar{\chi}\gamma^\mu\chi + \epsilon e \sum_f Q_f \bar{f}\gamma^\mu f \right) - \frac{1}{4}F'_{\mu\nu}F'^{\mu\nu} + \frac{m_{A'}^2}{2}A'_\mu A'^\mu, \quad (2)$$

where f is a SM fermion and Q_f is its electromagnetic charge.

This model permits two annihilation scenarios depending on the A' and χ masses.

- **Secluded Annihilation:** For $m_{A'} < m_\chi$, DM annihilates predominantly into A' pairs (Fig. 41, left). This annihilation rate is independent of the SM- A' coupling ϵ . While this makes direct A' or DM production difficult in laboratory experiments, the simplest version of this

scenario is robustly constrained by CMB data [11], which rules out DM masses below $O(10)$ GeV for simple secluded annihilation models. More complex secluded models remain viable for low DM masses; these are potentially discoverable by LDMX but are not our primary focus.

- **Direct Annihilation:** For $m_{A'} > m_\chi$, annihilation proceeds via $\chi\chi \rightarrow A'^* \rightarrow ff$ to SM fermions f through a virtual mediator. This scenario is quite predictive, because the SM- A' coupling ϵ must be large enough, and the A' mass small enough, in order to achieve the thermal relic cross-section. No robust constraint on this case can be extracted from CMB data. Therefore, the observed DM abundance implies a *minimum* DM production rate at accelerators. The detailed phenomenology depends on the ratio $m_{A'}/m_\chi$. If $m_\chi < m_{A'} < 2m_\chi$, the mediator decay to DM is kinematically forbidden so that visible dark photon searches and DM searches like LDMX are both relevant. If instead $m_\chi < m_{A'}/2$, the mediator decays primarily to DM, making searches for DM production even more sensitive and absolutely essential. This case represents most of the parameter space for direct annihilation, so we focus on it from now on.

Following conventions in the literature (see [1]), we introduce the dimensionless interaction strength y as

$$\sigma v(\chi\chi \rightarrow A'^* \rightarrow ff) \propto \epsilon^2 \alpha_D \frac{m_\chi^2}{m_{A'}^4} = \frac{y}{m_\chi^2} \quad , \quad y \equiv \epsilon^2 \alpha_D \left(\frac{m_\chi}{m_{A'}} \right)^4 \quad (3)$$

This is a convenient variable for quantifying sensitivity because for each choice of m_χ there is a unique value of y compatible with thermal freeze-out independently of the individual values of α_D , ϵ and $m_\chi/m_{A'}$. Reaching experimental sensitivity to this benchmark for masses between 10 keV – GeV would provide decisive coverage of these scenarios.

All existing constraints, collected in the y vs m_χ parameter space, are depicted in Fig. 2 alongside projections for LDMX and for other proposed experiments (see A for a discussion). Also shown are the thermal targets for fermion and scalar LDM candidates, which are invariant in this parameter space regardless of the values of ϵ , α_D , and $m_\chi/m_{A'}$ separately. Direct detection constraints are also naturally expressed as functions of y and m_χ , but accelerator-based constraints are not — for example, collider production bounds depend on $m_{A'}$ and ϵ . However, within the predictive framework of direct DM annihilation, both α_D and $m_\chi/m_{A'}$ are bounded from above, accelerator-based constraints typically become *stronger* as either α_D or $m_\chi/m_{A'}$ is decreased (specifically, lowering α_D expands the constrained regions downward in Figure 2 while lowering $m_\chi/m_{A'}$ moves them downward and to the left). Therefore, in the lower panel we plot conservative versions of these constraints with α_D and $m_\chi/m_{A'}$ near their upper limits to reveal all gaps in the DM parameter space consistent with the assumption of direct annihilation.

The upper panel of Figure 2 includes only missing mass/momentum/energy experiments, whose sensitivity can be directly compared to LDMX irrespective of α_D and $m_\chi/m_{A'}$. In particular, these are the only experiments sensitive to LDM in the limit of small α_D . This plot illustrates the large, orders of magnitude gaps in coverage between existing and projected constraints, and the thermal relic contours for LDM. In order to decisively cover thermal LDM in the direct annihilation regime, a dedicated effort is required. To this end, we also show the projections for LDMX at DASEL, which is the only proposed effort to probe the thermal target for both scalars and fermions down to the MeV range.

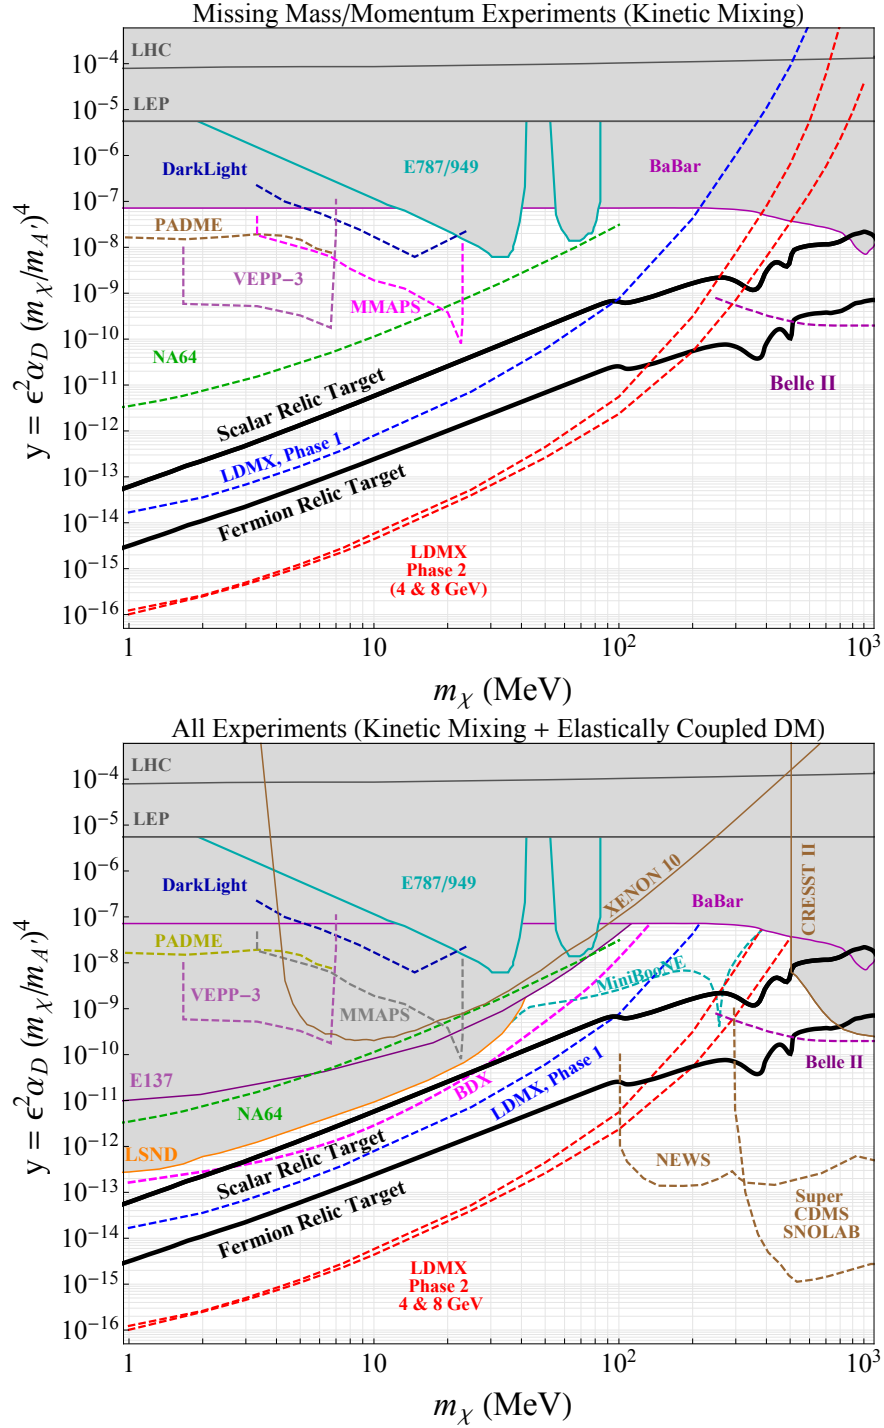


FIG. 2: The parameter space for LDM and future experimental projections in the y vs. m_χ plane plotted against the thermal relic targets for scalar and fermion DM – see text for a discussion. The top compares only constraints and projections based on missing mass/momentum techniques, which do not depend on any assumption about the nature of dark matter provided that the mediator decays invisibly on characteristic experimental length scales. The bottom panel shows the same parameter space on the assumption that the mediator couples to either scalar or fermion dark matter which scatters elastically off SM particles.




FIG. 3: Feynman diagram for radiation of an LDM particle-antiparticle pair off a beam electron as it scatters off a target nucleus. The DM production is mediated by a kinetically mixed dark photon (see (5)). Searching for this process is the primary science goal of LDMX.

D. Searching for Light Dark Matter Production

As discussed above, the primary science goal for LDMX is to search for the process depicted in Figure 3, wherein a DM particle-antiparticle pair is radiated off a beam electron as it scatters off a target nucleus. The DM production is mediated by a kinetically mixed dark photon – see Eq. (5).

Depending on the dark photon and DM masses, the leading contribution to this process may be *decay* of a dark photon into a DM particle-antiparticle pair ($m_{A'} > 2m_\chi$) or pair-production through a virtual dark photon. In either case, a constraint on the DM particle production rate can be used to infer a robust bound on the interaction strength y , which can in turn be directly compared to the targets from thermal freeze-out shown in Figure 2. This note will focus, for concreteness, on the former case, though the final-state kinematics is extremely similar for either on- or off-shell dark photons and the same search strategies apply to both.

LDMX will search for this process using the *energy-angle kinematics* of the recoiling electron, or “missing momentum” approach. As discussed in detail in Section III, this kinematics is distinctive, with the recoiling electron typically carrying a small fraction of the beam energy (the rest is carried by the DM and hence invisible) and receiving an appreciable transverse kick from the DM production process. Thus, the experimental signature for the signal comprises three basic features: (i) a reconstructed recoiling electron with energy substantially less than beam energy but also (ii) detectable, with measurable transverse momentum, and (iii) an absence of any other activity in the final state.

This search strategy has distinct advantages over other approaches that have been used to detect DM production:

- **Missing mass** (as in BaBar, MMAPS*, and VEPP-3* – * indicates proposed experiments) relies on full reconstruction of *all* recoiling particles, is only practical in e^+e^- collisions, and requires a much lower luminosity, greatly reducing production yield and hence sensitivity.
- **DM re-scattering** in a detector downstream of the production point (as in E137, MiniBoone, and BDX*) one can use very intense beams of either protons or electrons, but the low probability of DM scattering weakens sensitivity relative to what is possible in a kinematic search

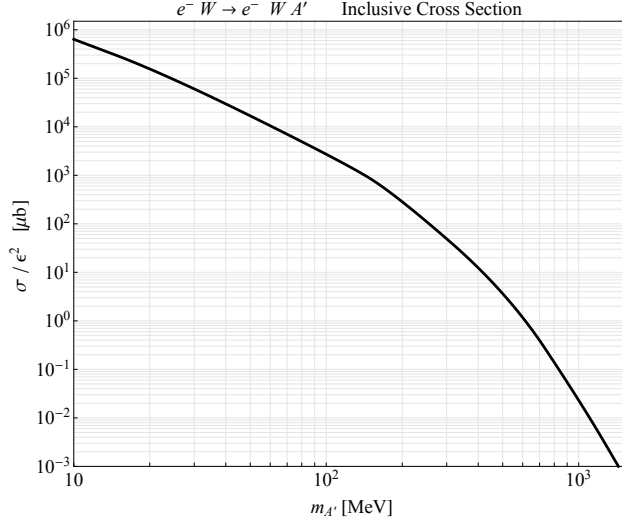


FIG. 4: Inclusive cross-section normalized to $\epsilon = 1$ for dark matter pair production off a tungsten target.

– whereas the kinematic signals of DM production scale as the square of the weak DM coupling, re-scattering signals scale as the fourth power. As a result, even the most aggressive proposals with intense beams fall short of the anticipated LDMX reach.

- **Missing energy** (as in NA64), reconstructing only the energy (not the angle), is closely related to the missing-momentum approach but with fewer kinematic handles to reject SM backgrounds and measure veto inefficiencies *in situ*. In addition, missing energy experiments lack the ability to distinguish final-state electrons from one or more photons, introducing irreducible neutrino backgrounds to high-rate missing energy experiments [12]).

However, reaching the full potential of this technique places demanding constraints on the experiment and beamline supporting it. A high repetition rate of electrons is required (~ 50 M e/sec on target for phase 1 and as much as ~ 1 G e/sec on target for phase 2), and so also a fast detector that can individually resolve the energies and angles of electrons incident on the detector, while simultaneously rejecting a variety of background processes varying in rate over many orders of magnitude. The remainder of this note describes a design that we expect to be suitable for Phase I with minor modifications, with technologies that are also expected to be useful for a Phase II experiment.

III. KINEMATICS AND RATE OF DARK MATTER SIGNAL PRODUCTION PROCESS

This section briefly summarizes cross-sections and kinematic features of the DM production process in the case of an on-shell dark photon. More detail can be found in [2, 12]. The DM production cross-section σ_{DM} scales as the square of the kinetic mixing parameter ϵ , therefore, the quantity σ_{DM}/ϵ^2 and the kinematic distributions of the recoiling electron depend only on particle masses. Throughout this section, we assume representative kinematics $m_{A'} = 3m_\chi$.

The normalized DM production cross-section on a tungsten target, as a function of DM mass, is shown in Figure 4. For low masses, the production cross-section scales as $1/m_{A'}^2$; for higher masses, it is further suppressed by the steeply falling form-factor for scattering off the nucleus with high enough momentum exchange to produce an on-shell mediator.

The kinematics of the electron in DM production also depends on the A' mass. In general, as long as the A' or $\chi\bar{\chi}$ pair that's produced is heavy relative to the electron mass, the differential cross-section for DM production is peaked in the phase space where the DM carries away the majority of the beam energy and the electron carries relatively little. This is, of course, the exact opposite of the structure for the kinematics of the vast majority of bremsstrahlung events (with the tiny remainder constituting the main LDMX backgrounds, as discussed in the next section), and the primary kinematic handle for a missing momentum search. This qualitative behavior is accentuated for larger χ and A' masses, and less dramatic for lower masses comparable to m_e , as shown in the top left panel of Figure 5. The selection efficiency of a requirement $E_e < E_{cut}$ is shown on the bottom left panel as a function of E_{cut} (the fraction of events accepted saturates below 100% for some mass points because events with backward-going recoil electrons are always rejected). A nominal value $E_{cut} = 0.3E_{beam} = 1.2$ GeV is assumed as a signal selection throughout this note. This cut keeps 75-90% of signal events over the full range of masses considered. Reducing it by a factor of 2 would only slightly degrade signal efficiency. In the remainder of this section we also impose a cut $E_e > 50$ MeV motivated by the degradation of tracking acceptance at lower energies.

The top right panel of Figure 5 shows the electron p_T distribution for events that pass this $E_e < 1.2$ GeV cut for a range of DM and A' masses. These are to be contrasted with the sharply falling p_T distribution from bremsstrahlung, which (even after accounting for multiple scattering in a 10% r.l. target) falls off as $1/p_T^3$ for $p_T \gtrsim 4$ MeV. As we will see below, bremsstrahlung-originated events dominate the background, so this kinematic difference is a powerful one.

The optimal strategy for using angular/ p_T information in a DM search will depend on the ultimate rejection power of the experiment's visible-particle veto, and will likely vary depending on the DM mass of interest. As the detector design is currently still evolving, we have not yet attempted to do this optimization, but note some general considerations below. If the veto inefficiency is low enough that no bremsstrahlung particles are expected, then the optimal strategy is to loosen signal selection to maximize acceptance. In this case one would not use the electron p_T in signal selection, but it is still a very useful handle for demonstrating the new-physics origin of events that pass the energy requirement and veto. If a substantial number of events pass the energy selection and veto, then the recoil electron p_T measurement can be used either to *reject* these events or to measure the bremsstrahlung component. Requiring a *minimum* recoil electron p_T of 20–50 MeV would offer substantial kinematic rejection of bremsstrahlung background, but substantially degrades the experiment's acceptance for A' masses above ~ 100 MeV. At lower masses, a 2D selection using both energy and angle information, or a multi-component signal+background fit to the 2D p_T -energy distribution may be the more effective use of this measurement.

Another very important kinematic feature is the angular/transverse separation between the recoiling electron and the direction of the beam. The larger this separation, the better the spatial separation between the recoil electron's shower and other final-state particles that one would like to veto. Figure 6 (left) shows the angle of the recoiling electron at the target for events passing the energy cut, while Figure 6 (right) shows, for the same events, the transverse separation of the recoiling electron from the beam spot at the face of the ECAL for the detector geometry described in Section IV.

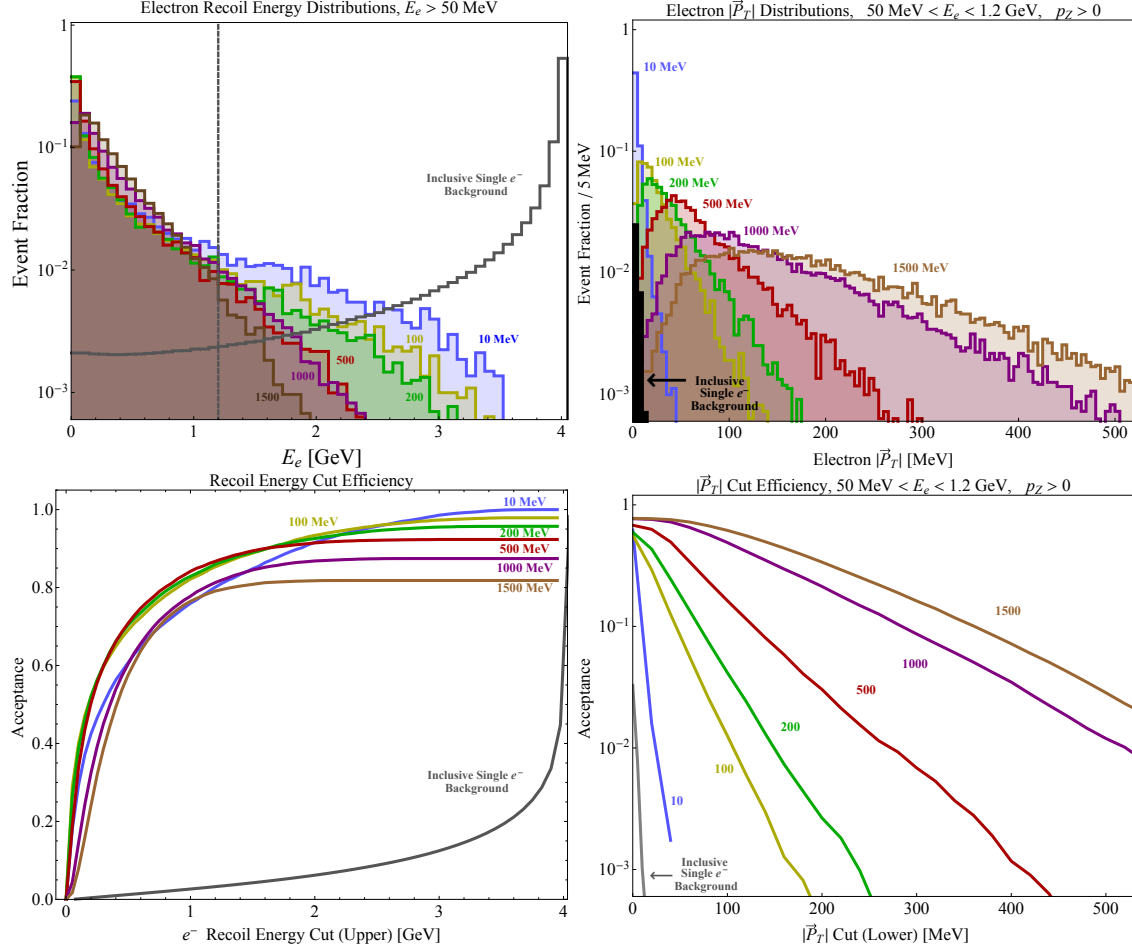


FIG. 5: Top: Electron energy (left) and p_T (right) spectra for DM pair radiation process, at various dark matter masses. Bottom Left: Selection efficiency for energy cut $E_e < E_{cut}$, as a function of E_{cut} , on inclusive signal events, The nominal cut is $E_{cut} = 0.3E_{beam}$. Bottom Right: Selection efficiency for p_T cut $p_{T,e} > p_{T,cut}$, as a function of $p_{T,cut}$, on events with $50 \text{ MeV} < E_e < E_{cut}$. In all panels, the numbers next to each curve indicate A' mass. Also included in each plot is the corresponding inclusive single electron background distribution.




FIG. 6: Left: Angular distribution of signal events passing the energy cut $E_e < 1.2 \text{ GeV}$. The numbers next to each curve indicate A' mass. Right: Transverse separation of the recoiling electron from the beam spot at the face of the ECAL for the detector geometry described in Section IV.

IV. DETECTOR CONCEPT

As noted above, the signature of DM production that is used by LDMX involves (i) substantial energy loss by the electron (e.g. recoil with $\lesssim 30\%$ of incident energy), (ii) a potentially large transverse momentum kick, and (iii) the absence of any additional visible final-state particles which could carry away energy lost by the electron. In the first phase of LDMX we aim to perform a search for events with this signature in a sample of $4 \cdot 10^{14}$ incident electrons, delivered onto a 10% radiation-length (X_0) target with a 46 MHz bunch spacing and average charge of 1 e^- per bunch. The target thickness has not been precisely optimized, but represents a reasonable compromise between event rate and transverse momentum resolution. Performing the measurement of each kinematic component (i and ii above) on every incident electron requires a tracker and electromagnetic calorimetry (ECAL) downstream of the target, with a sensitive area extending onto the beam axis itself. Placing the tracker in a weak magnetic field allows us to use the most striking of these events, where a soft and wide-angle recoil electron doesn't penetrate into the ECAL. In addition, stray low-energy particles from the beam halo must be rejected with very low inefficiency, as these could mimic the DM signal. This motivates another tracker (with larger B-field) upstream of the target that precisely measures the momentum and trajectory of the incoming beam electrons.

While the energy and angle resolution requirements for LDMX are modest, the experiment poses two main challenges. First, every incident electron passes through the trackers and showers in the ECAL; the detectors placed directly in the beam line must therefore contend with high radiation doses and a large event rate. Second, the requirement of a veto to take advantage of (iii), that's robust enough to handle a variety of rare backgrounds is an important physics performance driver for the tracking and ECAL, and in fact calls for the addition of a dedicated hadron veto system surrounding and behind the ECAL.

These four detector systems – the “tagging tracker” upstream of the target, “recoil tracker” downstream, a radiation hard forward ECAL, and the hadronic veto, form the majority of the LDMX experimental concept. To keep the detector compact and the field in the ECAL minimal, we place the tagging tracker inside the bore of a dipole magnet and the recoil tracker in its fringe field. The layout for LDMX is illustrated in Figure 7. The tracker and ECAL must be able to contend with a high rate of events with one of several dominant topologies:

1. **Electrons that do not interact in the target.** These electrons experience some straggling in the trackers and target (~ 46 MHz), but do not lose appreciable energy. These events feature a hard track through both trackers and a high-energy (≈ 4 GeV) shower in the ECAL.
2. **Hard bremsstrahlung.** These occur in the target (or, less frequently, in a tracking layer), such that the electron track through the recoil tracker falls below $E_{cut} = 1.2$ GeV, with a frequency of $\mathcal{O}(1)$ MHz. These events have two showers in the ECAL, with combined shower energy ≈ 4 GeV, separated by 1 – 2 cm or more depending on the electron energy.
3. **Trident events.** These occur in the target (or, less frequently, in a tracking layer) and produce two or three tracks (depending on kinematics) through the recoil tracker and two or three well-separated ECAL showers with combined shower energy ≈ 4 GeV. These are a higher-order process, and so occur at a frequency of $\mathcal{O}(100)$ kHz.

These necessitate radiation-hard, fast detectors and are the principal event types that must be rejected by the trigger.

The tracking proposed for LDMX borrows heavily from the Silicon Vertex Tracker (SVT) of the HPS experiment [13], a fixed target search for visibly decaying dark photons using the CEBAF beam at the Thomas Jefferson National Accelerator Facility (JLab). The HPS SVT was designed to provide high-purity, high-precision tracking for low-momentum ($\lesssim 3$ GeV) electrons at high occupancies (up to 4 Mhz/mm²) with the nearly-continuous (2 ns bunch spacing) CEBAF beam. The SVT meets these requirements with low-mass construction (0.7% X_0 per 3d hit), excellent hit-time resolution (~ 2 ns) and operation inside of the beam vacuum. Given very similar requirements, the solutions developed for HPS are also highly optimized for LDMX.

Similarly, a natural solution to the granularity, timing, and radiation hardness required of the ECAL is to exploit the silicon calorimetry designed for the CMS Phase-II forward calorimeter upgrade. The primary physics trigger will also require a positive signal in a scintillator pad adjacent to the tracker (to avoid triggering on empty buckets) coincident with *low* (or no) energy deposition in the ECAL in order to reject the first two classes of event types described above. The classes of rare, but potential backgrounds that will often pass the trigger, but must be rejected offline, are dominated by:

- **Hard bremsstrahlung + photo-nuclear reaction in ECAL.** These are events where the hard bremsstrahlung photon does not convert in the ECAL but instead undergoes a photo-nuclear reaction in one of the first layers of the ECAL (typically in a tungsten layer). The photo-nuclear cross-section is roughly 1000x smaller than the conversion cross-section, so these events will occur with sub-*kHz* rate. Photo-nuclear processes initiated by 2.8 – 4 GeV photons can result in a wide range of final states. When pions are produced in the interaction and escape the nucleus, they typically give rise to either “tracks” (π^\pm) or substantial energy deposition (π^0) in the ECAL. In many cases, a large number of low-energy protons and neutrons are liberated from a heavy nucleus; some of the protons may deposit energy in the ECAL, and some of the neutrons can be detected by the HCAL. A rare but important class of these events is characterized by only two or three $\mathcal{O}(1$ GeV), moderate-angle neutrons escaping from the nucleus. These events must be rejected by the HCAL. Though not strictly photo-nuclear, processes such as muon pair production are vetoed on the basis of the reconstruction of the “tracks” they produce in passing through the calorimeter.
- **Hard bremsstrahlung + photo-nuclear reaction in target/tracker.** Similar to the class above, these events have 20x lower rate because of the thin target material, and are more difficult to detect because the effective angular coverage of the HCAL is reduced.
- **Electro-nuclear interactions in target.** These are kinematically similar to the in-target photo-nuclear background, but at a still lower rate and broader electron p_T distribution.
- **Incident low-energy particles/beam impurities.** This potential background arises if the tagging tracker mis-measures a low-energy beam impurity as carrying close to the beam energy. Rejecting such potential backgrounds drives the performance requirements of the tagging tracker.
- **Neutrino backgrounds.** Neutrino production (through either charged or neutral currents) is entirely negligible for a Phase I experiment with $4 \cdot 10^{14}$ incident electrons on a 10% target.

In the following sub-sections, we describe the LDMX beamline and detector sub-systems in more detail.




FIG. 7: A cutaway overview of the LDMX detector showing, from left to right, the trackers and target inside a vacuum chamber in the spectrometer dipole, the forward ECAL, and the HCAL. Although not yet studied, the HCAL may be extended forward (transparent region surrounding the ECAL) to provide better coverage for wide-angle bremsstrahlung and neutral hadrons originating from photo-nuclear reactions in the target and the front of the ECAL.

A. Beamline

The LDMX beamline is relatively simple, consisting only of an analyzing magnet and a vacuum chamber into which the tagging and recoil trackers are installed, as shown in Figure 7. The analyzing magnet is a common 18D36 dipole magnet with a 14-inch vertical gap and operated at a central field of 1.5 T. A stainless-steel vacuum chamber with 1/2" thick walls fits just inside the magnet bore supporting the tracking detectors and their readout electronics. As with HPS, the upstream end of the vacuum chamber is closed by an additional vacuum box to accommodate feedthroughs for power, readout, and cooling, and the upstream end of this vacuum box is closed by a plate with a 6" conflat flange for connection to the DASEL beampipe. The downstream end of the vacuum chamber is closed by a thin vacuum window in front of the ECAL face, which sits 20 cm downstream from the target at roughly the same z position as the outer face of the magnet coils. The magnet is rotated by approximately 100 mrad about the vertical axis with respect to the upstream beamline so that the incoming 4 GeV beam follows the desired trajectory to the target, with the incoming beam arriving at normal incidence to, and centered on, the target, which is laterally centered in the vacuum chamber at $z = 40$ cm relative to the center of the magnet. Although the specific dimensions differ, this arrangement (aside from the vacuum window) is very similar to the magnet and vacuum chamber employed by the HPS experiment at JLab.

A number of 18D36 magnets, not currently in use, are in hand at SLAC, along with the steel required to adjust the magnet gap, if required to suit our purposes. These include a magnet that is already assembled with the 14-inch gap planned for LDMX. This magnet was tested to 1.0 T in 1978, at which point 199 kW of power was dissipated. Based on the current capacity of the

other similar magnets with smaller gaps, it is expected that this magnet can be operated at 1.5 T, resulting in a power dissipation of approximately 450 kW and requiring an approximately 55 gpm flow of cooling water. If this magnet proves to be suitable for LDMX, it will simply have to be split, cleaned up, and reassembled before testing and carefully mapping the field in the tracking volumes.

Since End Station A (ESA) is relatively distant from the critical areas in the Beam Switchyard, the vacuum requirements in ESA are quite modest: roughly 10^{-4} Torr or better. Therefore, standard vacuum fabrication techniques for the vacuum chamber and vacuum window will suffice. Furthermore, since HPS has achieved a vacuum of better than 10^{-5} Torr at JLab, it is clear that the construction techniques developed there will work as well for LDMX. However, as with HPS, a system of monitoring the beam vacuum, and interlocks capable of closing remotely operable valves will be required to ensure that the liquid cooling system of the trackers cannot spoil the beam vacuum.

Although the final location of LDMX has not been determined, the upstream end of the ESA beamline has the most room and best access for locating an apparatus of this size. In particular, there is a suitable concrete pedestal for mounting the dipole with sufficient room on all sides to support the calorimeters and accommodate the required beamline elements.

B. Overview of Tracking and Target Systems

Although the tagging and recoil trackers function as two distinct systems, they use common technologies and share the same support structures and data acquisition hardware. In particular, the first four layers of the recoil tracker are identical to the layers of the tagging tracker and share the same support and cooling structure, as shown in Figure 8. The key element of this




FIG. 8: An overview of the tracking systems and target inside the LDMX vacuum chamber.

upstream support structure is a vertically-oriented aluminum plate onto which the stereo modules

are mounted. To provide cooling, a copper tube through which coolant flows is pressed into a machined groove in the plate. This support plate slides from the upstream end of the vacuum chamber into precision kinematic mounts in a support box that is aligned and locked in place inside the vacuum chamber. Another similar plate slides into the support frame on the positron side of the chamber and hosts the Front End Boards (FEBs) that distribute power and control signals from the DAQ and digitize raw data from the modules for transfer to the external DAQ. The last two layers of the recoil tracker, being much larger, are supported on another structure: a cooled support ring onto which the single-sided, axial-only modules are mounted. This support ring is installed from the downstream end of the chamber, engaging precision kinematic mounts in the support box for precise alignment to the upstream stereo modules. The cooling lines for all three cooled structures—the upstream and downstream tracker supports and the FEB support—are routed to a cooling manifold at the upstream end of the vacuum chamber which, in turn, connects to a cooling feedthrough with dielectric breaks on the bottom of the vacuum box.

Overall, this design is similar to that of the HPS tracker, although with some important simplifications. First, because the radiation dose in LDMX is modest, cooling is needed only to remove heat from the readout electronics and not to keep the silicon itself cold. Therefore, cooling water that is close to room temperature can be used and there are no significant issues of differential thermal expansion to be concerned with. Second, the LDMX detector is in no danger from the nominal beam, so it does not need to be remotely movable, in contrast to HPS. Finally, the LDMX vacuum chamber is of a very generous size to achieve the desired acceptance in the ECal and recoil tracker, so it is not crowded, in contrast to the HPS chamber. In summary, the tracking systems are very similar to, but are a significant simplification of, the tracking system built for the HPS experiment. For this reason, the risk associated with this project is small.

It is the interposition of the target between the last layer of the tagging tracker and the first layer of the recoil tracker that clearly distinguishes between the target and recoil tracking systems. The target is a 350 micron tungsten sheet, comprising 10% of a radiation length. This choice of thickness provides a good balance between signal rate and transverse momentum transfer due to multiple scattering, which limits the utility of using transverse momentum as a signal discriminator. The tungsten sheet is glued to a stack of two 1 mm sheets of PVT scintillator, which provides a fast level-0 signal to the ECAL trigger, indicating the passage of a beam electron in a given beam pulse. The scintillator-tungsten sandwich is mounted in an aluminum frame that is inserted into the upstream tracker support plate from the positron side to simplify the process of replacing or swapping the target, as shown in Figure 9. With the scintillator mounted on the downstream face of the target, it can also be used to help identify events in which a bremsstrahlung photon has undergone a photonuclear reaction in the target. The structure and readout of the scintillator are discussed below as part of the trigger system.

C. Tagging Tracker

The tagging tracker is designed to unambiguously identify incoming beam electrons that have the correct energy, as well as precisely measure their momentum, direction, and impact position at the target. The key elements of the design are determined by this goal. First, the long, narrow layout of the tagging tracker accepts only beam electrons that have roughly the momentum and trajectory expected for the incoming beam, to eliminate off-energy beam electrons. Second the layers have low mass and are spaced far apart in a significant magnetic field to ensure a precise measurement of momentum. Finally, a large number of layers ensures the redundancy required to ensure high-purity pattern recognition.




FIG. 9: The LDMX target, indicated by the red arrow, installed between the last tagging tracker layer and first recoil tracker layer.

The layout and resolution of the tagging tracker are summarized in Table I. It consists of six

TABLE I: The layout and resolution of the tagging tracker.

Layer	1	2	3	4	5	6	7
z -position, relative to target (mm)	-607.5	-507.5	-407.5	-307.5	-207.5	-107.5	-7.5
Stereo Angle (mrad)	-100	100	-100	100	-100	100	-100
Bend plane (horizontal) resolution (μm)	≈ 6	≈ 6	≈ 6	≈ 6	≈ 6	≈ 6	≈ 6
Non-bend (vertical) resolution (μm)	≈ 60	≈ 60	≈ 60	≈ 60	≈ 60	≈ 60	≈ 60

double-sided modules of silicon microstrips arranged at 10 cm intervals upstream of the target, with the first module centered at $z = -7.5$ mm. The modules are positioned laterally within the vacuum chamber so that they are centered along the path of incoming 4 GeV beam electrons. Each module places a pair of $4 \text{ cm} \times 10 \text{ cm}$ sensors back to back: one sensor with vertically oriented strips for the best momentum resolution and the other at ± 100 mrad stereo angle to improve pattern recognition and provide three-dimensional tracking. The sensors are standard p^+ -in- n type silicon microstrip sensors identical to those used for HPS. These sensors have $30(60) \mu\text{m}$ sense(readout) pitch to provide excellent spatial resolution at high S/N ratios and are operable to at least 350 V bias for radiation tolerance.

The sensors are read out with CMS APV25 ASICs operated in multi-peak mode, which provides reconstruction of hit times with a resolution of approximately 2 ns. At very high occupancies, six-sample readout can be used to distinguish hits that overlap in time down to 50 ns, which limits the readout rate, and therefore the trigger rate, to approximately 50 kHz. However, at the low hit occupancies anticipated in LDMX, three-sample readout may suffice, enabling a maximum trigger rate approaching 100 kHz. These sensors are mounted on standard FR4 hybrid circuit boards, which provide the power conditioning and I²C control for the APV25, as well as a thermal path to

the cooled support structure.

The sensors and hybrids are assembled into half-modules, which are the smallest non-rebuildable units of the tagger tracker. Each half module consists of a single sensor, which is glued to the hybrid at one end with conductive epoxy to provide bias voltage and support. A thin ceramic plate is glued to the other end of the sensor for support. A pair of half modules is attached to either side of an aluminum module support with screws to form a module. A vacuum-compatible thermal compound applied between the half-module and the support ensures a good cooling path. Finally these double-sided modules are attached to the upstream support plate described in Section IV B to position them along the beamline and provide cooling for the hybrids as shown in Figure 10. This support arrangement places only silicon in the tracking volume and




FIG. 10: A beams' eye view of the tagging and recoil trackers.

no material on the side of the tagger towards which electrons bend, in order to avoid secondaries that would be generated by degraded beam electrons hitting material close to the tracking planes.

D. Recoil Tracker

The recoil tracker is designed to identify low-momentum (50 MeV to 1.2 GeV) recoil electrons and precisely determine their momentum, direction, and impact position at the target. In addition, it must work together with the calorimeters to correctly distinguish low-momentum signal recoils

from scattered beam electrons and multi-particle backgrounds. The key elements of the design are determined by this goal. First, the recoil tracker is placed at the end of the magnet in the beginning of the fringe field to optimize tracking for particles up to two orders of magnitude softer than the beam-energy electrons measured by the tagging tracker. Second, the recoil tracker is short and wide for good acceptance in angle and momentum and to minimize the distance from the target to the calorimeters to improve their angular coverage. Finally, the recoil tracker provides 3-d tracking near the target for both direction and impact parameter resolution but emphasizes low mass density over the longest possible lever arm further downstream to deliver the best possible momentum resolution. This design delivers good momentum resolution for both multiple-scattering limited, low-momentum tracks and beam energy electrons, which are nearly straight in the fringe field.

The layout and resolution of the recoil tracker are summarized in Table II. It consists of four stereo layers located immediately downstream of the target and two axial layers at larger intervals in front of the ECAL. The stereo layers are double-sided modules of silicon microstrips arranged

TABLE II: The layout and resolution of the recoil tracker.

Layer	1	2	3	4	5	6
z -position, relative to target (mm)	+7.5	+22.5	+37.5	+52.5	+90	+180
Stereo Angle (mrad)	100	-100	100	-100	-	-
Bend plane (horizontal) resolution (μm)	≈ 6	≈ 6	≈ 6	≈ 6	≈ 6	≈ 6
Non-bend (vertical) resolution (μm)	≈ 60	≈ 60	≈ 60	≈ 60	-	-

at 15 mm intervals downstream of the target, with the first module centered at $z = +7.5$ mm. These modules are laterally centered on the target and the center of the vacuum chamber and are identical to the modules of the tagging tracker that are mounted upstream on the same support plate.

The thinner axial-only layers, at $z = +90$ mm and $z = +180$ mm, are mounted on a separate support structure at the downstream end of the vacuum chamber as described in Section IV B. They have a somewhat different module design as shown in Figure 11. Each consists of a pair of sensors, glued end-to-end, with an APV25-based FR4 hybrid circuit board at each end of this structure to read out the two sensors. The sensors are standard p^+ -in- n silicon microstrip sensors, but are somewhat shorter and wider than those used in the stereo modules and therefore require six APV25 chips to read out each sensor instead of five. The modules are supported at both ends by screw attachment of the hybrids to castellated support blocks attached to the cooled support structure. A vacuum compatible thermal compound is applied to ensure good thermal conductivity between the hybrid and the support structure.

E. Forward Electromagnetic Calorimeter

A high granularity, Si-W sampling calorimeter will be used for the ECAL of LDMX. The ECAL will draw heavily on the forward calorimeter upgrade of the CMS experiment for the HL-LHC [14]. In particular, the hexagonal sensors, front-end readout electronics and front-end trigger architecture will be the same. This is facilitated by the fact that the CMS experiment has an instantaneous event rate of 40 MHz while that of LDMX will be 46 MHz.

An ECAL based on silicon pad sensors is well-suited to the LDMX experiment whose principal task for Phase I, with $4 \cdot 10^{14}$ electrons on target (EOT), will be the identification of ~ 1 to 4 GeV photons with extraordinarily high efficiency and very good energy resolution. The ECAL will also need to play a role in the identification of hadrons, especially those that are contained




FIG. 11: The target and recoil tracker.

within it. However, a hadronic veto calorimeter, described below, will be used to identify hadrons from photo-nuclear reactions that escape the ECAL without interacting. The high granularity of the ECAL provides the ability to track charged hadrons that cross multiple silicon layers and to identify isolated charged hadrons that range out in a single layer, depositing charge equivalent to that expected for multiple minimum ionizing particles (MIP). Detection of stable neutral hadrons is more difficult but may be enhanced in the ECAL by carefully introducing some layers of appropriate materials at various places in the stack that have higher cross sections for neutrons to interact and produce charged hadrons. Studies that vary absorber materials, thicknesses, and geometrical layout are underway. We expect to settle on final specifications for the ECAL in FY17.

For LDMX Phase II, we envision a total of 10^{16} EOT at an energy of 8 GeV, which will require the identification of an average of 20-30 electrons every 25 ns. While Phase I is the main topic of this note, we mention Phase II in the context of the ECAL because high granularity is needed even more in that high rate environment and because a silicon based calorimeter can sustain the associated higher fluences.

With a reasonably good transverse spread of the beam electrons, the granularity of the apparatus can be used to help identify overlapping showers. The central stack of hexagons could use the highest granularity CMS modules with 512 channels on a sensor made from an 8" wafer. The hexagonal pads on these sensors have sides of length 4.3 mm. It may be possible to double the pad count to as high as 1024, with pads having 3 mm sides, if that brings a significant improvement in performance. This would require a dedicated order of some 50 sensors and the production of a corresponding front-end readout board. Note that the majority of modules in LDMX could use the most common sensor design in CMS with 384 hexagonal cells having 5 mm sides. We plan to study the impact of higher granularity in the LDMX ECAL next.

The choice of silicon is important for handling the impact of fairly substantial radiation doses expected in LDMX Phase II, and even in the central stack of the ECAL in Phase I. Simulations of the radiation environment performed with FLUKA [15] indicate that, while electromagnetic inter-

actions are the dominant source of irradiation in the LDMX environment, a substantial hadronic component is also present. LDMX plans a diffuse beam to spread the electrons over a cross-sectional region of order 10 cm^2 in the central stack. The effective 1 MeV neutron-equivalent fluence at shower max is shown in Fig. 12. For 10^{14} EOT, the fluence is seen to peak at around $3 \times 10^{13} \text{ n/cm}^2$. This can be linearly projected to $3 \times 10^{15} \text{ n/cm}^2$ for 10^{16} EOT. This is certainly high enough to necessitate changes in the biasing of the silicon and require the silicon to be operated cold. It is, however, substantially lower than the worst case expectation for the CMS HGC. The CMS experiment has performed radiation studies for 300, 200 and 120 μm sensors exposed to maximum fluences of roughly 9×10^{14} , 4×10^{15} and 1.5×10^{16} 1 MeV equivalent neutrons per cm^2 , respectively. The full depletion voltage of the sensor, V_d , is proportional to the square of the thickness, which can lead to a reduced operational lifetime for thicker sensors in a high radiation environment. In the 120 μm case, the noise is much higher for irradiated sensors, reducing the signal-to-noise ratio to about 2. For the fluences expected in LDMX, the CMS studies suggest that a sensor thickness of 300 μm would have a worst-case signal-to-noise ratio of about 5. The radiation environment of the LDMX experiment is, however, low enough to allow the use of 500 μm sensors, which would boost the signal by 67%. For 500 μm thick sensors, one expects roughly 10 % improvement in the energy resolution of 4 GeV electrons over that obtained with 300 μm sensors. It may be possible to have even thicker sensors from a radiation-tolerance standpoint but this has the disadvantage of increasing the space between absorbers, thereby increasing the Molière radius.





FIG. 12: Effective 1 MeV neutron equivalent fluence at shower max after 10^{14} electrons on target. The 4 GeV electron beam is assumed to be spread out in order to distribute impacting electrons over a region of order 10 cm^2 . The maximum effective fluence is found to be of order $3 \times 10^{13} \text{ n/cm}^2$.

This minimum signal-to-noise ratio expected at the end of operation of LDMX, even for 300 μm sensors, will be adequate to identify MIPs. Moreover, as noted earlier, smaller pads could be used for the LDMX central stack. This would increase the signal-to-noise ratio by up to about 50%,



HGC Module Components
and completed module

FIG. 13: Prototype CMS HGC 6 inch, 128 channel module with (1) a W/Cu (75%/25%) baseplate, to which (2) a thin Gold-on-Kapton sheet is glued for insulation and to enable biasing of (3) a Si sensor glued directly on top of it. Finally, (4) a PCB is glued on the sensor. Connections to each set of 3 adjacent pads on the sensor are made by wirebonds through 4 mm diameter holes.

assuring excellent MIP detection capability for the life of the experiment.

Module components and the first completed CMS prototype HGC, a nominal 6-inch, 128 channel module, are shown in Fig. 13. We are targeting a design with roughly 30 Silicon layers, each comprised of a central module surrounded by a ring of 6 modules. As for CMS, the hexagonal sensors will likely be produced on ‘8-inch’ wafers, yielding a span of 170 mm between parallel sides. The basic readout unit is a hexagonal module containing a single hexagonal silicon sensor. The module is built on a metal baseplate that serves as electromagnetic absorber, and so has been chosen to be predominantly tungsten. It must also transport heat from the module electronics and sensor to a cooling/support plane. In order to enhance thermal transport without introducing an unacceptably large mis-match in thermal expansion coefficients between the silicon and the baseplate, CMS has chosen an alloy of 25% Copper and 75% Tungsten. In LDMX we have so far considered a pure W baseplate for the module, which however could be affixed to a Cu cooling/support plane similar to the CMS design.

Note that in the module shown in Fig. 13, the corners of the sensor and PCB have been removed to give access to mounting holes on the baseplate that are used to position and attach the modules precisely on the cooling plate. In CMS the cooling surface is a large 5 mm–thick planar Cu sheet with embedded CO₂ cooling lines, where the Cu contributes to the absorber material budget. For LDMX we are planning to use a simpler, liquid cooling system, since the dark currents and thermal load are lower given the less severe radiation environment. We have begun to study ECAL performance with a very thin and lightweight cooling layer, made up of small diameter, thin-wall stainless steel cooling lines embedded in carbon fiber. This could be useful in the upstream layers of the ECAL, where we find that it is beneficial for the rejection of some relatively rare types of

photo-nuclear interactions to start with a thin absorber, increasing the thickness to a nominal value near $1 X_0$ over some 5 layers or so.

Initial simulated ECAL performance studies make use of a layout with 42 layers of silicon for which the first 10 layers have absorbers of $0.8X_0$ and the remaining 32 layers having absorbers of $1.0X_0$. We chose this geometry to allow us to understand shower containment and energy resolution very well, and we have found that at least for a 4 GeV electron beam, we can manage with 10 fewer sensing layers. We have therefore begun studies and estimated costs for a 30 layer, $40 X_0$ device. This would have 210 installed silicon modules. The active region of the ECAL would then be roughly 30 cm in depth and about 51 cm wide.

For the electronics, we plan to use the CMS HGCROC front-end chip and PCB design. The CMS development schedule is such that the final components will be produced in substantial quantities in time for the construction of the LDMX detector. The HGCROC, for instance, is expected to be ready for mass-production by the end of CY19. The PCB will be finalized even earlier, based on pin-compatible prototype front-end chips.

Given the small physical extent of the ECAL, it is our belief that there are no optical transmitters needed in the active region. Our plan is to bring low voltage in on relatively heavy Cu lines and to use micro-coaxial cables for the signal lines from each module. It may be that we can use the current CMS architecture with a data concentrator on each module, but full readout would also be relatively straightforward for this relatively small system. It would require one or two fast links, each with its own signal line, for each HGCROC chip. The data would be pushed to FPGAs for triggering and for full readout. The data acquisition is described in more detail later in this note.

The cooling planes would be stacked in an enclosed and thermally insulated box through which dry nitrogen gas or dry air would flow. The support planes with modules attached could be suspended like hanging file folders from a box frame. This simple approach was recently used for the CMS test-beam studies of the first HGC modules at FNAL and CERN. The box and support structure could be relatively simple and constructed from hard aluminum or stainless steel. Feed-throughs would be needed for the signal, clock, control, and power cables, as well as the cooling and gas lines. Fortunately, the density of services is relatively low in each plane. With appropriate care and attention in assembly and installation, we believe it will be possible to achieve an excellent thermal and gas enclosure.

F. Hadronic Veto System

It is important to extend the calorimeter to veto neutral hadrons being produced in photo-nuclear reactions. As exclusive reactions of this type are rare, the rates and radiation doses will be much lower, as the HG calorimeter is designed to fully contain the electromagnetic showers. This makes silicon detectors unnecessary in the hadronic calorimeter section. A scintillator-based sampling calorimeter is a natural solution in this situation.

The goal of the hadronic veto system is to identify neutral hadrons in the energy range from above approximately 100 MeV to several GeV with high efficiency, which requires a hadronic calorimeter with at least 5 nuclear interactions of depth in order to fully contain the most energetic neutrons with greater than 99% probability. Simultaneously, in order to detect lower energy neutrons, absorbing layers cannot be so thick that neutrons of hundreds of MeV are captured by them. Therefore, a steel-scintillator (polystyrene) calorimeter of approximately 15 layers and totaling 5 nuclear interaction lengths is proposed. Each layer is structured as 50 mm of steel, 2 mm air gap, 6 mm of scintillator, and 2 mm of air gap, where the air gaps are left for detector services. The transverse size of each layer is $1\text{m} \times 1\text{m}$ to cover the solid angle of the signal

acceptance. Transverse granularity of the system is not desired, due to the lower rates expected in the hadronic system and in order to maintain high efficiency for neutral hadron detection. An illustration of the hadronic veto system is given in Fig. 14. A particularly challenging background from photo-nuclear reactions occurs when a few neutral hadrons are produced and scatter at wide angles (~ 45 degrees or more). Pending further study of this background, it may be necessary to add a wide angle component to the hadronic veto system that would wrap around all but the front of the ECAL system. Studies are in progress to understand the size of this type of background as well as the optimal orientation and specifications for a wide angle veto component.




FIG. 14: Sketch of the hadronic veto system. Light areas represent steel absorbers and yellow areas represent plastic scintillators. An example design of grooves for wavelength shifting optical fibers can be seen in the last active layer.

Fast readout is also required of the hadronic calorimeter system in order to coincide with electromagnetic calorimeter information. Readout will be based on the CMS Phase 1 upgrade HCAL system, which has fast readout capabilities intrinsically at frequencies of 40 MHz and would be sufficient for the DASEL beam structure. Scintillating light is read out by wavelength shifting fibers with silicon photomultipliers (SiPMs) as the photodetectors. SiPM technology is chosen due to high gain and low noise capabilities. Each layer is read out by approximately 6 wavelength shifting fibers, and the light is optically summed into one photodetector. In addition to the SiPMs, readout modules (48/64 channels per) and charge integration electronics designed for the CMS detector can be re-purposed for the LDMX experiment with minimal changes, taking good advantage of the experience of CMS collaborators on LDMX. The time scale for commissioning of the readout electronics in CMS are 2017/2018 and thus are appropriate for the LDMX timeline.

G. Trigger System

The LDMX trigger system is designed to reduce the typical beam particle arrival rate of 46 MHz to a rate of 5 kHz for storage and analysis. The selection is performed by a combination of dedicated hardware and software running on general-purpose computers.

The first stage trigger is implemented in hardware and allows the selection of both candidate events for dark matter production and important samples for calibration and detector performance monitoring. The overall trigger management is provided by a trigger manager board, which receives inputs from the various triggering subsystems including the silicon calorimeter and the scintillator calorimeter. The latency requirements of the trigger calculation are set by the tracker readout ASIC to $2 \mu\text{s}$.




FIG. 15: Drawing of the concept for the target and trigger scintillator system, showing the relationship of the SiPMs to the scintillator planes and strips.

The primary physics trigger is based on the silicon calorimeter and is designed to select events with energy deposition significantly lower than the full beam energy. The silicon calorimeter HGCROC calculates the total energy in 2×2 fundamental cells for every 46 MHz bucket. The energy information is transferred by digital data link to the periphery of the calorimeter, where sums are made over larger regions and transferred by optical link to the trigger electronics. The total energy is then used to select the events. The details of the electronics for the ECAL trigger are discussed along with the DAQ for the ECAL below.

The use of a calorimeter trigger requiring energy below a threshold also requires a beam-presence measurement to avoid very high trigger rates during crossings in which there is no beam electron present. The trigger scintillator satisfies this requirement. As discussed above, two thin (1 mm) layers of plastic scintillator are connected to the downstream side of the tungsten target. The first layer is a continuous plane, while the second layer is segmented at 2.5 mm intervals. The two layers are separated by a reflective foil, and a second layer of foil covers the second layer. The structure is shown in Fig. 15. The first layer, directly in contact with the target, is important for identifying the short-range fragmentation products from photonuclear interactions in the target. The second layer provides information on the position of the interaction within the target.

The trigger scintillator is read out using an array of 2×2 mm SiPMs connected to QIE11 electronics developed for the CMS HCAL upgrade. The SiPMs are coupled directly to the edge of the scintillator stack, with each SiPM observing both layers. The spatial information can be extracted by comparison between SiPMs to determine the relative contributions of the two layers.

The SiPMs and QIE11 readout cards will be located in the vacuum vessel along with the tracker front end electronics. The SiPMs will be cooled in common with the tracker elements to achieve very low thermal noise levels in the trigger scintillator readout. The readout electronics will continuously digitize the SiPM signal, providing an integrated charge as well as time-of-arrival measurement for the pulse with an LSB of 500 ps. Both amplitude and timing information can be provided to the trigger, allowing correction of the calorimeter amplitude for timewalk effects already at trigger level.

H. DAQ

The off-detector trigger and readout electronics are assumed to share a single ATCA crate, using the Reconfigurable Cluster Element (RCE) generic computational building block developed by

SLAC. The RCE is based on the Xilinx Zynq 7 System-On-Chip technology, running a dual core 1 GHz ARM processor with 1 GB of DDR3 memory tightly integrated with on-chip programmable logic (FPGA). The RCE blocks are integrated into Data Processing Modules (DPM), each of which hosts two RCE blocks. The DPMs receive data from rear transition modules (RTM), which are customized to adapt to the specific portion of the experiment that is being read out. Up to four DPMs can be mounted on a single Cluster-On-Board (COB) ATCA blade. The modular ATCA design permits LDMX to re-use architecture and functionality from other DAQ systems whose components are similar to those used by LDMX.

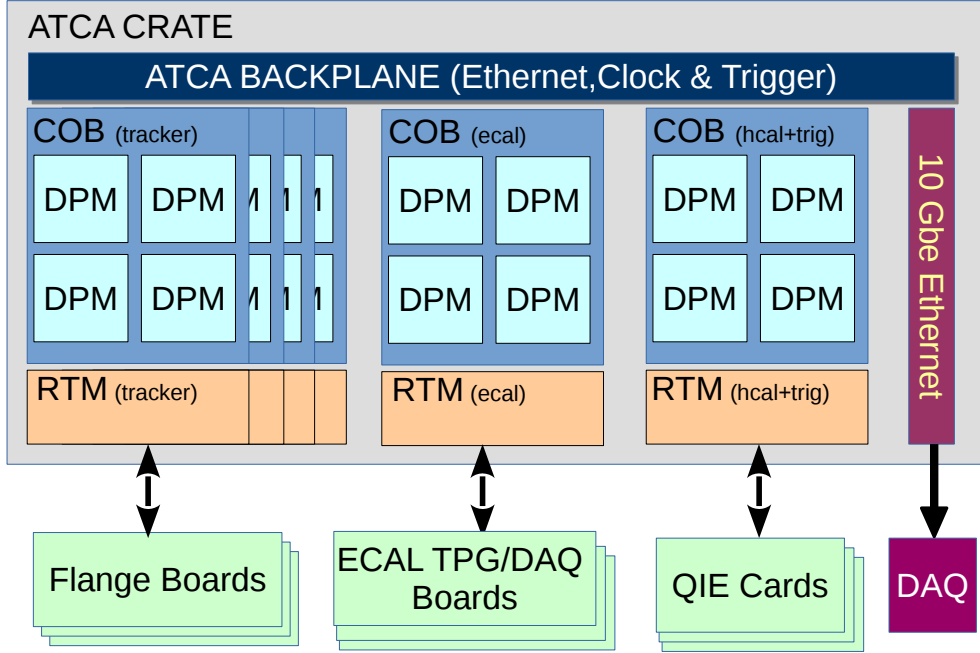


FIG. 16: Schematic block diagram of the RCE-based DAQ.

1. Overall DAQ Architecture

One of the DPMs functions as the trigger interface, which receives trigger signals from the optical fiber module on the RTM, distributes clock and trigger signals, and handles communication with the software trigger supervisor. Specific interfaces to beamline and infrastructure services are provided through a custom RTM. Final DAQ is carried out through the 10 Gbps Ethernet readout path available separately from each COB. The COBs in a single ATCA crate are also interconnected by 10 Gbps readout and contain ARM processors that can be used to execute high-level trigger algorithms.

To establish a conservative design for the DAQ, we consider a target DAQ bandwidth of 25 kHz (trigger rejection factor of 5×10^{-4}). This provides a factor of five safety compared with the combined expectation for the performance of the hardware trigger and for the data volume in triggered events. The breakdown of data volumes by subsystem is shown in Table III and sums to an estimated total of 2.5 kilobytes per event, for a total DAQ bandwidth of 62.5 MBps. This bandwidth is within the capability of the 10 Gbps Ethernet bandwidth with a comfortable safety factor of ≈ 10 . The 10 Gbps ethernet connections from the COBs will be connected to one or

more data storage and data-quality-monitoring PCs. To estimate the potential total data volume, we assume there is no high-level trigger, in which case the full sample of 2×10^{11} events would require 500 TB of storage.

TABLE III: Estimated DAQ data volume per channel.

Subsystem	kbytes/event	Estimate Notes
Trigger	0.5	Includes event accounting information (32B), trigger sums from each module of the ECAL (210 B), trigger-counter readout (144B), and readout of trigger information from the HCAL (128B).
Tracker	1.0	Each raw hit requires 20B for id and data, and we expect 50 hits/event, including noise hits.
ECAL	0.8	Each hit with TDC is 10B including 4B for data and 6B for addressing and overhead, and each (low-amplitude) hit without TDC is 8B. Assuming a trigger threshold of < 400 MIP total in the calorimeter, we expect 45 hits, which we assume are all TDC hits for the estimate. We assume a noise cut at 2.5σ , which results in 44 noise hits, all assumed to be non-TDC hits.
HCAL	0.2	Each channel's readout is 20B for id and data, and we budget for 10 channels in a typical event.
Total	2.5	

2. Tracking DAQ

The data acquisition system for the tracking detectors supports the continuous 46 MHz readout and processing of signals from the silicon-strip sensors. It also selects and transfers those events that were identified by the trigger system to the back-end DAQ for further event processing at rates approaching 50 kHz. The system adopted is an evolution of the DAQ used for the HPS SVT and consists of APV25 readout ASICs hosted on hybrid circuit boards integrated into the tracking modules, a set of Front End Boards (FEBs) within the vacuum volume for power distribution and signal digitization, flange boards that transmit power and data across the vacuum barrier, and remotely located ATCA-based data acquisition hardware.

The APV25 ASIC, initially developed for the Compact Muon Solenoid silicon tracker at the Large Hadron Collider at CERN, was chosen because it provides excellent signal to noise, analog output for optimal spatial resolution, and signal pulse-shape sampling capability for good hit time resolution. Each hybrid board has five or six analog output lines (one for each of the APV25 ASICs) that are sent to an FEB using LVDS signals over about 50 cm of twisted-pair magnet wire. At the front-end readout board, a preamplifier scales the APV25 differential current output to match the range of a 14-bit Analog to Digital Converter (ADC). Each front-end board services four hybrids. The ADC operates at the system clock frequency of 46 MHz. The digitized output from the front-end board is sent through compact 8-pair mini-SAS cables to the vacuum flanges to connect to the external DAQ, which resides outside the vacuum chamber. The front-end readout board houses an FPGA and buffers signals to allow for control of the distribution of clock, trigger and I2C communication with the APV25 ASICs. To further simplify the services and minimize cabling that enter through the vacuum flanges, it contains linear regulators to distribute and regulate three low voltage power lines to each of the APV25 ASICs, as well as passing through the high voltage bias. Figure 17 shows a schematic layout of this part of the readout chain.

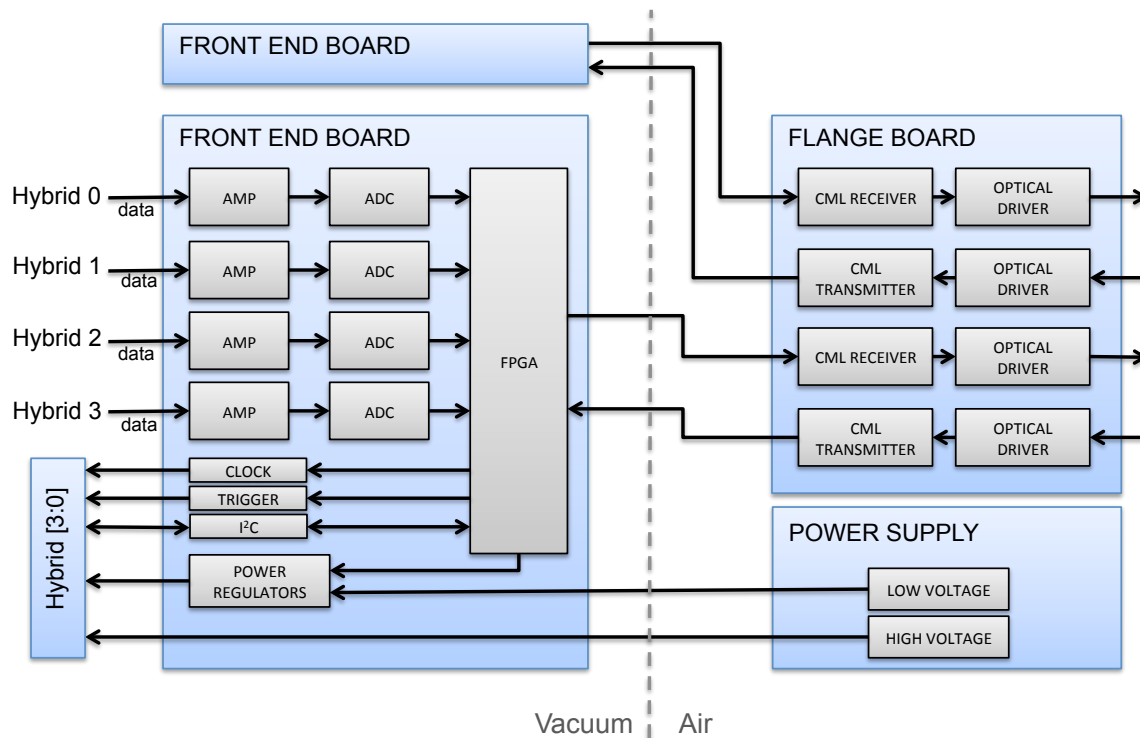


FIG. 17: Schematic view of the downstream part of the tracker DAQ.

The digitized signals are converted to optical signals just outside the vacuum flange on custom-built flange boards. Each flange board houses optical drivers to handle the electrical-optical conversion and to transmit the optical signals over fibers to the ATCA crate. The flange board also interfaces the low and high voltage power transmission from the Wiener MPOD power supplies to the front-end boards located inside the vacuum chamber.

The optical signals from four hybrids on one third of a flange board are received at one of four sections of the Rear Transition Module (RTM) board in the ATCA crate. Each section of the RTM connects to one of four DPMs. In order to minimize the complexity of the system inside the vacuum chamber, all signal processing is done at the DPM. Each data DPM receives the digitized signals from the RTM, applies thresholds for data reduction and organizes the sample data into Ethernet frames. Four COBs are sufficient to handle all of the hybrids of the trackers. The maximum readout rate of the tracking system is approximately 50 kHz, limited by the APV25 readout rate for six-sample, multi-peak readout.

3. ECAL DAQ

As described above, signals in the Si/W ECAL are digitized in the HGCROC and transferred on 368 MHz DC-coupled copper data links to the periphery of the calorimeter. At the edge of the detector, the copper links are received by FPGAs and processed before transmission on optical links off the detector. The architecture of the trigger primitive and readout electronics of the ECAL is shown in Fig. 18.

For trigger purposes, groups of four cells are summed, and the highest amplitude cells are

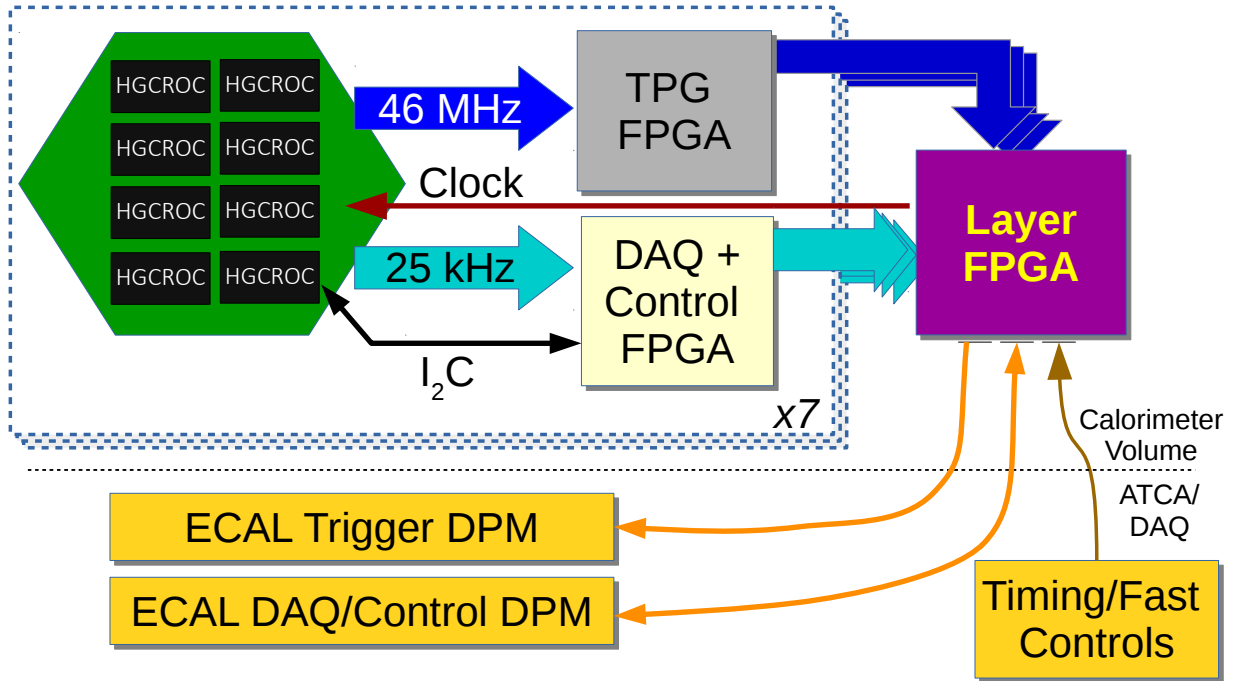


FIG. 18: Schematic view of the trigger and readout for the ECAL.

transferred on the electrical links to an FPGA at the periphery but inside the cooled volume of the ECAL. Thirteen electrical trigger links are required for each HGCROC chip, for a total of 104 per module. These trigger links are processed by a module-sum FPGA. The results of the summing are passed to a layer-master FPGA, which also receives the single electrical link for DAQ from each HGCROC in the layer. The trigger data from each layer requires a bandwidth of 2.5 Gbps. The DAQ and control data volume is much smaller (< 1 Gbps per layer), but for simplicity of firmware and software development, a separate optical link will be used.

In the ATCA crate, two DPMs will be used for trigger processing and two for DAQ and control. The trigger processing DPMs will produce cross-layer trigger sums that will be transmitted by high-speed digital links to the final trigger decision logic. The full trigger, readout, and control of the ECAL can be managed by a single COB.

4. HCAL and Target DAQ

The HCAL and target scintillators are read out using SiPM phototransducers and QIE11 ASICs developed as part of the CMS HCAL Phase 1 upgrade[16]. The QIE11 is a deadtimeless charge-integrating ADC with an internal TDC capable of 500 ps timing resolution. The data from the QIE11 is read out via an Igloo2 FPGA onto a pair of digital optical data links that operate asynchronously at 5 Gbps. The readout of the QIE cards is continuous. For every 46 MHz bunch the data links will digitize and transmit on the fiber the complete data for every channel. For the HCAL, where TDC information is not important for the analysis, the QIE11 readout cards will operate as 16-channel cards. For the trigger scintillator, where TDC information is quite important, the cards will operate as 10-channel cards. A total of 14 QIE cards will be required for the scintillator-based channels, six for the trigger scintillators and eight for the HCAL.

In the ATCA crates, each scintillator readout RCE will handle ten readout fibers, allowing two links to be used for full-speed communication to the final decision-logic RCE. As a result, one DPM (two RCEs) will be required for each of the trigger scintillator and the HCAL. Therefore, the scintillator readout can be hosted on the same blade as the final decision logic.

V. PHYSICS AND DETECTOR SIMULATION

This section describes the methods used to simulate signal and background physics reactions and the responses of various detectors to these reactions. For historical reasons, two separate simulation frameworks have been used: the “tracker simulation” implemented in SLIC comprises a full simulation of the tagger tracker and recoil tracker, with ECAL material included but not its detector response. Likewise, the “calorimeter simulation” implemented directly in Geant4 comprises a full simulation of the ECAL and HCAL, with recoil tracker material included upstream. In both cases, the target material and magnetic field map are also included. These two simulations are described in more detail below in §V A–V B 2.

While the primary simulation engine is Geant4 [17], the signal reaction (Dark Matter pair production) is modeled using an external generator based on MadGraph/MadEvent4 [18]. This generator, its validation, and the interface with Geant4 are described in §V C. All background processes are modelled directly in Geant4, with modifications for photonuclear processes as discussed in §V D.

A. Simulation of the Tagging and Recoil Trackers

The simulation of the passage of both charged and neutral particles through the tagging and recoil trackers as well as the target uses a software package, the Simulator for the Linear Collider (SLIC) [19], based on the GEANT4 toolkit. It creates realistic charge depositions in the silicon layers, simulates the APV25 readout chip amplifier chain, forms clusters and performs track finding and fitting. The simulation also contains the Ecal material, needed to understand the contribution of backscattered particles to the occupancy of the last recoil tracker layer. The simulation is used to understand track reconstruction efficiencies, purity and to evaluate the momentum and vertex resolutions. A rendering of the tagging and recoil tracker as used in the simulation is shown in Figure 19.

1. Hit Reconstruction

The energy depositions resulting from a particle traversing through the silicon planes are used in the simulation of all physical processes resulting in current signals on the readout strips. The signal present on each of the strips is used to simulate the APV25 pulse shape which is then sampled into a readout pipeline. Six of the samples from each strip, selected in such a way that they capture the signal peak, are fit using a 3-pole function in order to extract the rise and fall time, the amplitude and time of the hit, t_0 . The information extracted from the fits is used to form strip hits which, in turn, are clustered using a nearest neighbor algorithm. Finally, all clusters in adjacent pairs of layers are combined to create 3D hits used by the track finding and fitting algorithm.

2. Track Reconstruction

The track finding and fitting algorithm proceeds in steps following a specific “tracking strategy”. The strategies outline which layers are used and their roles (e.g. seed, confirm, extend), the minimum number of hits required to form a track, kinematic constraints (e.g. impact parameters) and the maximum allowable χ^2 value. The algorithm proceeds as follows:




FIG. 19: Rendering of the tagging and recoil tracker geometry used for performance studies. The figure shows a typical event where an incident electron (in blue) traverses through the tagging and recoil trackers. The white boxes represent the tracker silicon microstrip sensors while the green lines represent the strips that were hit by the incident electron.

- To begin, all possible combinations of hits on the specified seed layers are used to form 3-hit track candidates. A helix fit is then performed on each of the candidates and only those which satisfy the χ^2 requirement are kept.
- Once all seeds have been found, hits from a specified “confirm” layer are added to the seeds and a helix fit is performed once again. Those 4-hit seeds which fail to satisfy the χ^2 cut are eliminated.
- Next, tracks are “extended” using hits from the rest of the layers. After the addition of a hit, a helix fit is performed. If the track fails the χ^2 constraint, the hit is discarded. This procedure is repeated until all hits in all extend layer have been added to all seeds. However, it is possible for all hits in a layer to be discarded.
- Finally, all track candidates are merged in order to form a set of unique tracks. Tracks are allowed to share at most one hit. If a pair of tracks shares more than a single hit, the track with the most hits is chosen. If the two tracks have the same number of hits, the track with the best χ^2 is selected.

The performance of the track finding and fitting algorithm is discussed in Section VI A. The tracking strategies used in these studies required only a single confirmation hit. In the tagging

(recoil) tracker, tracks were required to have at least 5 (4) hits. Furthermore, extremely loose cuts on the track χ^2 and distance of closest approach to the beam were imposed. Unless otherwise noted, the performance studies assumed that the tagging (recoil) tracker is within a uniform dipole field of strength -1.5 T (-.75 T), while all acceptance studies used the full field map.

B. Simulation of the Calorimetry Systems

Physics event generation and detector simulation are performed with the GEANT4 simulation toolkit. For trigger and DAQ studies, this simulation was coupled with an electronics simulation that converts the energy deposits in the silicon or scintillator into simulated digital hits. The full simulation toolkit has been validated by the CMS collaboration reproducing the simulation and test-beam results of the CALICE collaboration.

1. Simulation of the Electromagnetic Calorimeter

A simulated interaction of a 4 GeV electron in LDMX is shown at left in Fig. 20. The figure displays the geometrical layout of LDMX that was used in studies reported in Sec. VI C. A typical 4 GeV electron event results in $O(800)$ MIPs being observed in the calorimeter cells.

The distribution of the energy deposition for a dark photon of mass $M_A = 0.1$ GeV is shown at right in Fig. 20. We will likely use energy deposition as our primary trigger variable because of the strong separation it provides between signal and background events.

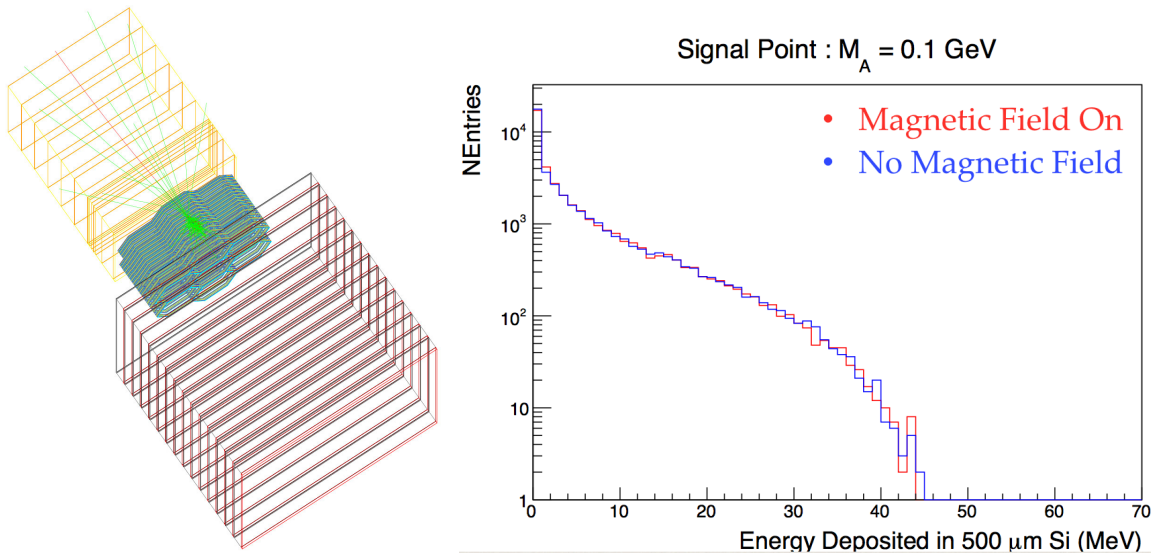


FIG. 20: An event display for a GEANT4 simulation of a 4 GeV electron in the LDMX detector (left) and the distribution of summed energy deposition in all layers of the ECAL for events in which a dark photon of mass $M_A = 0.1$ GeV is generated (right). The magnetic field (red histogram) has little impact.

2. Digitization of the Hadronic Veto System

After full GEANT4 simulation, we simulate the scintillator, optical system, photodetector and electronics responses of hadronic veto system to understand how well we can reconstruct events. Given the particle-level energy deposition per layer, we can convert that into an expected number of photo-electrons (PEs). GEANT4 simulations with 4 GeV incident muons are used to estimate the expected detector response to minimum ionizing particles (MIPs). Fig. 21 shows the distribution of energy deposited in a given layer of scintillator; on average, each MIP deposits 1.4 MeV of energy in a given layer.

For each MIP, we assume that 20 PEs are detected on average, based on CMS test beam measurements. In these tests, with a thinner scintillator than planned for LDMX, 13.5 PEs were observed for a comparable optical system and an identical readout system to that proposed for LDMX. The CMS test-beam MIP response has been scaled up to account for the differences in scintillator thickness, assuming that the light collection efficiency will remain constant.

The sum of SiPM and electronics noise is expected to be 2 PEs. Although noise is not explicitly simulated, only layers that have a minimum of 8 PEs are considered in further studies. Thus, noise effects are expected to be negligible. We are considering the possibilities for reduced-temperature operation of the SiPMs, which will reduce their noise significantly, below the 2 PE level.




FIG. 21: Deposited-energy distribution for muons to characterize MIP behavior in plastic scintillator

C. External Physics Generator for Signal Reaction (Dark Matter Production)

While background processes are modelled entirely in the Geant4 framework, the new physics of Dark Matter production is modeled using an external event generator based on MadGraph/MadEvent4 [18]. Here we describe this generator, its validation, and the interface with Geant4.

MadGraph is an automated tool for calculating the tree-level amplitudes for arbitrary physics processes, which allows users to define Feynman rules for new physics models; MadEvent is a Monte Carlo event generator based on MadGraph. MadGraph/MadEvent4 (MG/ME) was designed for the study of high-energy collider reactions, but minor modifications to the code (introducing non-zero masses for incident particles and for the electron, and an electromagnetic form factor as described in [20] for the nucleus) allow for its application to fixed-target processes. These modifications and a new-physics model that introduces a dark photon with arbitrary mass and kinetic mixing ϵ with the photon has been used for the APEX test run [21] and HPS experiment [13]. For LDMX, we have added LDM particles (either fermions or scalars) that couple to the dark photon with an arbitrary interaction strength g_D . This allows us to simulate the signal process of DM particle pair-production via either decay of an on-shell A' or off-shell A' exchange. This report focuses on the on-shell production process, though the kinematics of the two are very similar.

Within MG/ME, we generate events for the DM production process $e^- W \rightarrow e^- W A', A' \rightarrow \chi \bar{\chi}$ where χ represents the dark matter particle and $\bar{\chi}$ its antiparticle. Events are generated assuming a 4 GeV incident electron and tungsten nucleus at rest as the initial state. MG/ME computes a Monte Carlo approximation of the inclusive cross-section for this process, and generates a sample of unweighted events in the Les Houches Event (LHE) format [22]. The inclusive cross-section computed by MadGraph is stable within 1% and is consistent within $\sim 30\%$ with independent calculation of the cross-section in the Weizsacker-Williams (WW) approximation from [23]. The deviations from the WW inclusive cross-section are largest at high and low masses, and compatible with the size of errors expected in the WW approximation.

We read the four-momentum of the recoiling electron (and, for the tracking simulation, the tungsten nucleus) from the LHE file produced by MG/ME and use these to seed a Geant4 event. The electron and nucleus are assumed to originate from a common vertex at the center of the target (for the calorimeter simulation) or uniformly distributed (for the tracker simulation) over the thickness of the target and over a transverse region spanning ± 1 cm in the x direction and ± 2 cm in the y direction about the nominal center of the target.

D. Photonuclear Model

By default, Geant4 models photonuclear reactions initiated by < 3.5 GeV incident photons using the Bertini Cascade model, and those initiated by > 3 GeV photons using a high-energy string model (with a mixture of the two event generators in the intermediate region). The Bertini model is, however, believed to be accurate up to ~ 10 GeV incident energies [24], and indeed is used in Geant4 for this purpose for other primary particles. There are notable differences between the two models, in particular for photon collisions with heavy nuclei, where the Bertini cascade within the nucleus frequently results in a large number of soft hadrons that are not well modeled by the high-energy models. For this reason, all LDMX simulations have been done using a modified FTFPBERT physics list, where photonuclear processes with < 10 GeV photons are always modeled using the Bertini cascade.

The Bertini model is notably complete, including cross-sections from data up to almost 3 GeV for a large number of photon-proton interaction processes and for $2 \rightarrow N$ and $3 \rightarrow N$ processes within the nucleus (however, the final-state phase space is only taken from data for $2 \rightarrow 2$ processes, and process cross-sections above 3 GeV are extrapolated). The critical photon energy range for LDMX is 3–4 GeV, and in some cases the phase space of multi-particle final states may be peaked and therefore poorly approximated by Geant. Efforts are underway to validate the Bertini model at these energies against data for high- Z nuclei.

VI. PERFORMANCE STUDIES

At the beginning of Section IV we motivated the LDMX detector design by discussing a variety of background event types, ranging from frequent to extremely rare, that LDMX must reject. The design and optimization of the detector and of analysis strategies to provide high signal acceptance while rejecting these backgrounds are ongoing. It would be premature to claim an overall background rate based on present studies, to claim that a low-background experiment can be achieved, or to infer a limit to the experiment's ultimate sensitivity. This section aims to summarize the present status of design and performance studies and to demonstrate that the concept of each sub-system is well matched to the goals of the experiment as a whole.

Section VI A discusses the performance of the tagging tracker and recoil trackers. The focus for the tagging tracker is on ensuring an adequately low rate for low-energy particles to produce high-energy tracks in the tagging tracker, contributing to the potential mimicking a signal event. For the case of the recoil tracker we present studies of the signal acceptance and momentum resolution and some initial studies of how the recoil tracker contributes to rejecting events where photo-nuclear reactions occur in the target. Section VI C discusses the ECAL performance, focusing on the detection of photons that undergo photo-nuclear reactions in the ECAL, since these are expected to be the limiting background for the experiment. The preliminary conclusion of these studies is that the background events that are typically *not* rejectable in the ECAL alone are those where the photo-nuclear reaction produces a small number of neutrons. The rejection of these few-neutron events is the primary objective of the HCAL, studied in Section VI D. Finally Section VI E summarizes the expected performance of a physics trigger based on a maximum total energy deposition in the ECAL and outlines a possible trigger menu for the experiment.

A. Tagging Tracker Performance

The tagging tracker must identify incoming beam-energy electrons with extremely high purity, suppressing the mis-reconstruction of any incoming low-momentum charged particles as beam-energy electrons. In particular, any incoming charged particle within the recoil acceptance for signal that is wrongly reconstructed as having the beam energy in the tagging tracker would be an irreducible background. The design of the tagging tracker makes the likelihood of such errors vanishingly small, due to its good resolution for both beam-energy and off-energy incoming particles and an exceedingly low rate of mis-reconstruction for tracks within the recoil energy acceptance.

In order for an incoming low momentum particle to fake a beam-energy electron in the tagger, a number of conditions must simultaneously be met:

1. The incoming particle must reach the first tagger layer or it will not intersect with any material until it hits the wall of the vacuum chamber.
2. The particle must either scatter in each layer in order to fake a 4 GeV track or create secondaries that generate occupancy which confuses the pattern recognition in the tracker, resulting in reconstruction of a fake 4 GeV track.
3. The resulting track must have a trajectory consistent with that of a typical 4 GeV beam electron all the way through the tracker.
4. The resulting track must have an impact point at the target consistent with a reconstructed track within the signal acceptance in the recoil tracker.




FIG. 22: The longitudinal momentum reconstructed by the tagging tracker for a sample of 4 GeV beam electrons. Excellent momentum resolution allows tight selection against any off-energy component in the beam.

Taking into consideration the design of the tagging tracker, each of these requirements places a very heavy penalty on any off-energy component in the beam. First, incoming particles with momentum less than approximately 500 MeV will not hit the first layer of the tagger unless they are significantly off-trajectory as well. Furthermore, even at 500 MeV, a first scatter of more than 10° is needed in order for the incoming particle to appear to be on the correct trajectory. It is clear then that the most challenging scenario is large contamination with incoming charged particles at the top of the momentum range for signal recoils, nominally 1.2 GeV. Such particles have the highest likelihood of reaching the first layer of the tagging tracker without being bent away by the magnetic field and require much smaller scatters and/or track reconstruction errors to result in fake tags. In order for a 1.2 GeV particle to make a trajectory through the tracker consistent with a 4 GeV track, six successive scatters of approximately ten milliradians must occur, each equivalent to approximately 15σ on the multiple scattering distribution. From the tails of the Moliere scattering distribution, the likelihood of each of these scatters is smaller than one per million. Therefore, the much more likely scenario is the generation of secondaries in the material of the tagging tracker followed by mis-reconstruction of a fake 4 GeV track. Since the resulting 4 GeV track must arrive at the target on the correct trajectory, and since beam energy electrons arrive normal to the target within a one-sigma range of 250 microradians, there is very little phase space for randomly reconstructed 4 GeV fakes to have the correct trajectory. Finally, any falsely reconstructed 4 GeV track must have a common impact point in the target with a real track of matching momentum in the recoil tracker, which is unlikely for a falsely reconstructed track.

Two methods are used to benchmark the performance of the tagging tracker and better understand the constraints it provides in preventing mis-tags. First, the momentum and impact param-




FIG. 23: The x - y position at the target reconstructed by the tagging tracker for a sample of 4 GeV beam electrons. Excellent impact parameter resolution provides a strong constraint on a matching recoil track.

eter resolutions at the target are determined using an analytic model of the tagging tracker that includes the effect of intrinsic resolutions and multiple scattering in the tracker planes. Second, full simulation is used to confirm these resolutions and understand reconstruction efficiencies and susceptibility to background from both physics processes and mis-reconstruction effects.

For incoming 4 GeV electrons, the analytic model finds a longitudinal momentum resolution of approximately 1%. The corresponding full simulation results show good general agreement, as shown in Figure 22. The transverse momentum resolutions are found to be 1.0 MeV and 1.4 MeV in the horizontal and vertical directions, respectively, which are small compared to the 4 MeV smearing in transverse momentum from multiple scattering in the $10\% X_0$ target. Meanwhile, the impact parameter resolution for 4 GeV electrons is expected to be approximately $7 \mu\text{m}$ ($48 \mu\text{m}$) in the horizontal (vertical) direction. Again, the full simulation shows good general agreement, as shown in Figure 23. These results indicate that tight requirements can be made in both the energy and trajectory at the target, which serve to reject off-momentum particles that could be present in the incoming beam, as described above.

In order to test the robustness of the tagging tracker layout against mis-tagging off-energy incoming electrons, a sample of simulated 1.2 GeV electrons on a trajectory that allows them to pass through all seven layers of the tagging tracker was used to predict how often in this layout they would be mis-reconstructed as full-energy beam electrons. The results, shown in Figure 24,

confirm that such particles cannot be mistaken for 4 GeV electrons at the level of less than one part in 10^6 , limited here only by the simulation statistics.

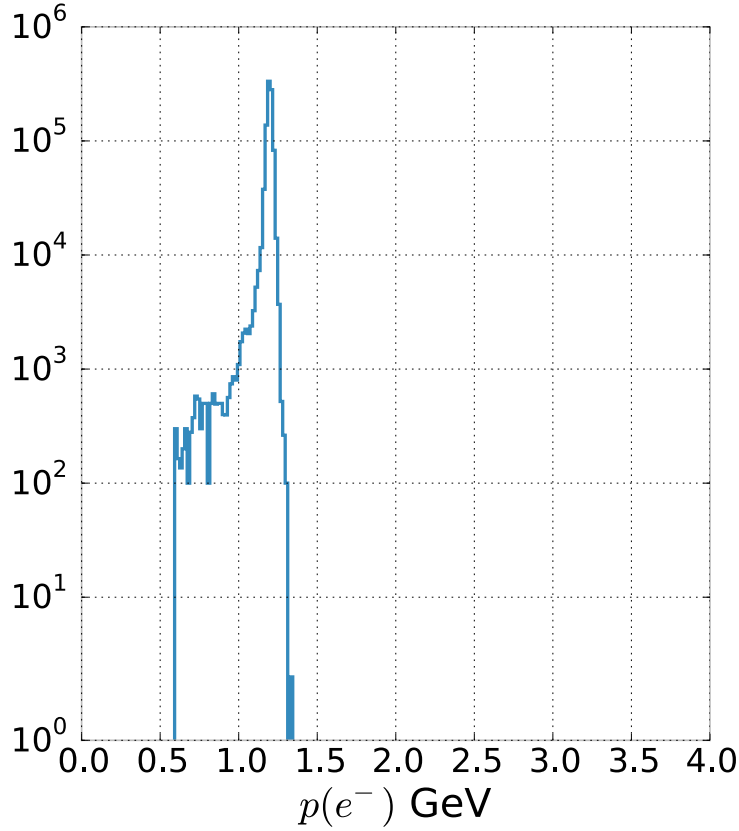


FIG. 24: Reconstructed total momentum for a sample of 1.2 GeV beam electrons in the tagging tracker. Of 10^6 incident particles, none is reconstructed with momentum greater than 1.4 GeV.

In summary, the design of the tagging tracker appears robust enough to provide unambiguous tagging of incoming electrons with the nominal beam energy for Phase I of LDMX. Further study will be required to find the beam intensity limits for any future upgrade.

B. Recoil Tracker Performance

The recoil tracker must have a large acceptance for recoiling electrons characteristic of signal events with good resolution for transverse momentum and impact position at the target, both of which are critical for unambiguously associating those recoils with incoming electrons identified by the tagging tracker. While good reconstruction efficiency for signal recoils is important, it is even more important to have good efficiency for charged tracks over the largest possible acceptance to help the calorimeter veto background events with additional charged particles in the final state. In addition, the tracker should have sufficient momentum resolution that it can assist the ECAL in identifying events in which an incoming beam electron passes through the target and tracker without significant energy loss. Finally, the recoil tracker can help identify events in which a bremsstrahlung photon generated after the end of the tagging tracker undergoes a photo-nuclear reaction in the target or tracker material. As with the tagging tracker, an analytic model of the recoil

tracker is used to estimate momentum and impact parameter resolutions, while full simulation and reconstruction are used to confirm those resolutions, to estimate signal acceptances and tracking efficiencies, and to test the ability of the recoil tracker to reject backgrounds.




FIG. 25: The recoil tracker acceptance to signal recoils as a function of A' mass. The acceptance drops with A' mass as the polar angle of the recoils increases.




FIG. 26: The tracking efficiency for signal recoils as a function of the momentum for electrons that pass the acceptance criteria described in the text. Since data from all A' mass points are used, these efficiencies are averaged with respect to different recoil polar angle distributions for each sample.

Measurement of the position and transverse recoil momentum at the target requires precise determination of the track angle at the target. For the transverse momentum, a good curvature measurement is also required in order to set the overall momentum scale. At least two 3-d measurements directly downstream of the target are needed to determine the recoil angle, and at least one additional bend-plane measurement is needed for curvature. For low-momentum tracks, it is sufficient for the third measurement to be in the first four closely-spaced layers, but for high momentum tracks that are nearly straight, hits in both of the downstream axial layers are necessary for good momentum resolution. Because high-momentum signal recoils will nearly always pass through all six layers, the acceptance near the top of the energy range for signal recoils is near unity, only reduced by the small single-hit inefficiency in the last two layers. However, at low momentum a large number of tracks can escape detection. Therefore, in order to estimate the signal acceptance for well-measured recoils using simulated signal events, we require that recoiling electrons leave hits in at least three of the 3-d layers, for pattern recognition and angle estimation, and at least four hits total for reasonable purity. The resulting acceptance for the recoil tracker as a function of mediator mass is shown in Figure 25, and the average of the track finding efficiencies for accepted tracks in those samples as a function of momentum is shown in Figure 26.

The ability to distinguish signal from background using the recoil transverse momentum is obviously limited by the multiple scattering in the target. For example, multiple scattering in a 10% X_0 target results in a 4 MeV smearing in transverse momentum. Using the analytic model of the recoil tracker, the material budget and single-hit resolutions were designed so that the transverse

momentum resolution is limited by multiple scattering in the target over the momentum range for signal recoils. This has been verified in full simulation, as shown in Figure 27. Meanwhile, the impact parameter resolution, shown in Figure 28, strongly constrains the phase space for mis-reconstructed tracks to point to the same location in the target as the incoming track reconstructed in the tagging tracker.




FIG. 27: Resolutions on the components of the momentum transverse to the target for signal recoils. Simulated recoils are assumed to originate at a random depth in the target so that the average resolution is less than the 4 MeV smearing from multiple scattering in the full target thickness. Only the vertical component of the recoil momentum shows significant degradation due to detector resolution at the highest recoil momenta considered.

Although the ECAL does an excellent job distinguishing scattered full-energy electrons from potential signal recoils, the recoil tracker adds significant leverage for this critical task. A sample of incident 4 GeV electrons was used to estimate the rejection of mis-reconstructed full energy electrons in the momentum range for signal recoils, $p < 1.2$ GeV. Figure 29 shows that there is a low-momentum tail in the reconstructed tracks but that these tracks are, as expected, accompanied by bremsstrahlung photons in the final state that can be vetoed by the calorimeters.

Finally, events in which a hard bremsstrahlung occurs after the tagging tracker and undergoes a photo-nuclear reaction in the target or recoil tracker will sometimes be observable in the recoil tracker. In the case of a photo-nuclear reaction in the target or the recoil tracker, a large multiplicity of hits in the next layers of the recoil tracker can result, as shown in Figure 30. In addition, the trigger scintillator on the back side of the target generally measures a larger signal in these events, as shown in Figure 31. Use of this information will significantly reduce backgrounds due to photo-nuclear events that occur upstream of the calorimeters.




FIG. 28: Recoil tracker impact parameter resolutions. Requiring a common impact position at the target with both the tagging and recoil tracker strongly selects against associating mis-reconstructed recoils with a tagged incoming electron.

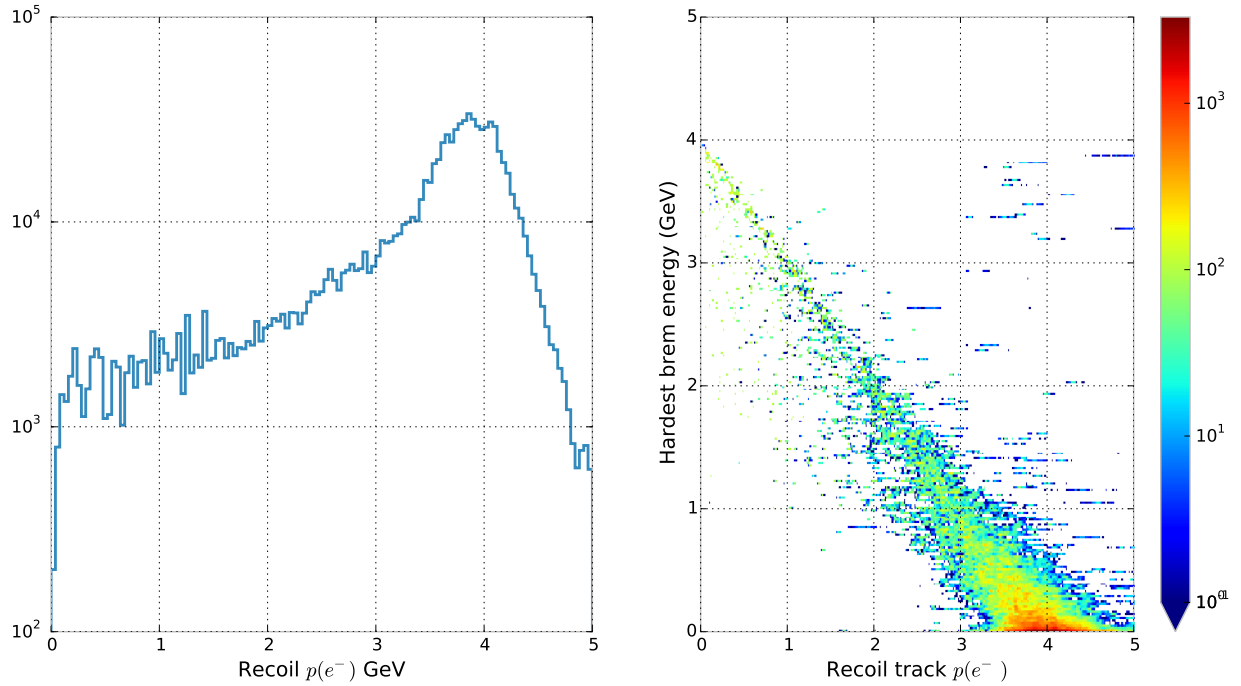


FIG. 29: The reconstructed recoil track total momentum for a sample of 4 GeV beam electrons (left) shows a low-momentum tail that extends into the signal region. However, events with low-momentum tracks are accompanied by bremsstrahlung photons that are vetoed by the calorimeters (right).




FIG. 30: The multiplicity of hits in each of the recoil tracker layers due to particles produced in photo-nuclear reactions in the target. Hits from the recoiling electron are not included.

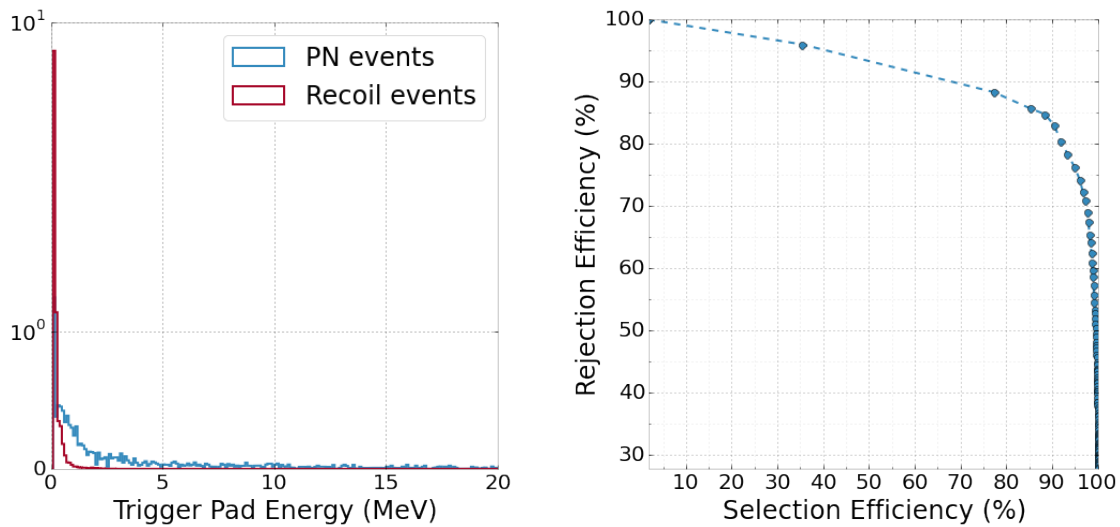


FIG. 31: On the left, the energy deposited in the trigger scintillator in events with a photo-nuclear reaction occurring in the target (blue) and signal recoils (red). On the right, the rate of rejection for photo-nuclear events as a function of the signal efficiency for a cut placed on the energy deposited in the trigger scintillator.

C. Forward Electromagnetic Calorimeter Performance

Reconstructed electromagnetic shower energy is the most powerful discriminant between signal and the most common, purely electromagnetic bremsstrahlung background. Our initial performance studies therefore focused on electromagnetic response and resolution for several sensor thicknesses. The mean energy deposition increases with thickness for electrons in the ranges of $0.25 \leq E \leq 4$ GeV, while the constant term obtained from a fit to the mean measured energy of the electrons is reduced as seen in the left and right plots, respectively, of Fig. 32. The distribution

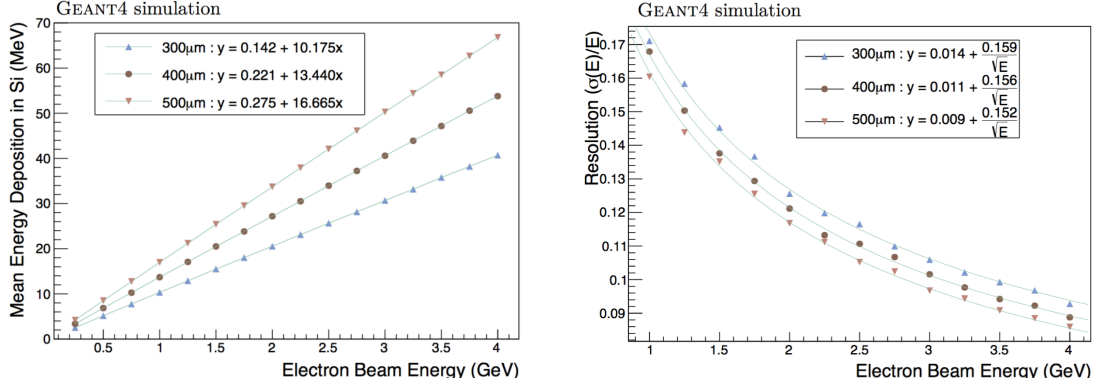


FIG. 32: Mean measured energy deposition (left) and resolution (right) in the ECAL as a function of electron beam energy as obtained from a standalone simulation of incident electrons on a 40 layer device. The study was carried out for sensor thicknesses of 500, 400 and 300 μm .

of energy deposited in a 40 X_0 ECAL with 500 μm sensors for 4 GeV incident electrons is mostly Gaussian with a mean of 69 MeV and standard deviation of 6.5 MeV. There is very rarely any escaping electromagnetic flux, defined as the summed energy of all electrons and photons that exit the bounding surface of the ECAL, as seen in the left plot of Fig. 33. This extraordinary level of containment is important for LDMX, in order to eliminate the very rare instrumental background in which a photon fails to interact while passing through many absorber layers. As seen in the right plot of Fig. 33, for instance, the escaping flux for the 23.35 X_0 ECAL configuration planned for CMS would lead to an unacceptably high background rate for LDMX.

The second class of background events occur when a photo-nuclear reaction takes place between a hard incident photon and an atomic nucleus in the target or in an absorber layer of the ECAL, converting all or most of the incident photon's energy into outgoing hadrons. The resulting final states can vary greatly because both nuclear disintegration and/or hadron production can occur. Since neutrons do not lose energy to ionization, and protons follow a very steep ionization curve, the detector response also varies greatly. To see how different photo-nuclear events are from those containing purely electromagnetic showers, we compare deposited energy distributions for the two types of events in Fig. 34. It is seen that the photo-nuclear events produce broad tails at high and low deposition energy. The low energy tail is largely made up of events in which nearly all of the incoming energy is used to liberate neutrons. Unless one or more of these neutrons interacts in the ECAL to produce an ionizing particle that can be tracked, or that generates a significant amount of charge in a single layer, we will have to rely on the HCAL to veto these events. As noted earlier in the letter, we plan to study the effect of adding materials with higher cross sections

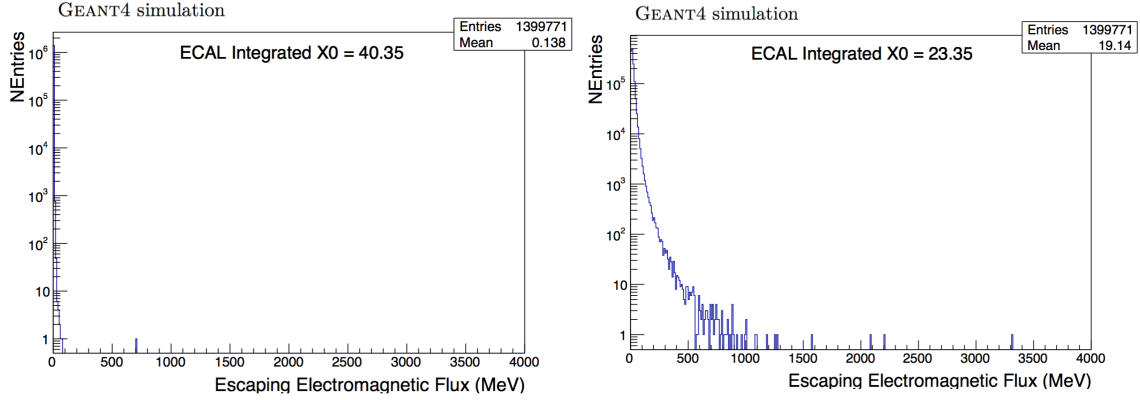


FIG. 33: Escaping electromagnetic flux for incident 4 GeV electrons for an ECAL with $40.35 X_0$ (left) and for an ECAL with $20.35 X_0$ configured similar to the CMS HGC (right).

for low energy neutrons to produce charged hadrons in various layers of the ECAL.

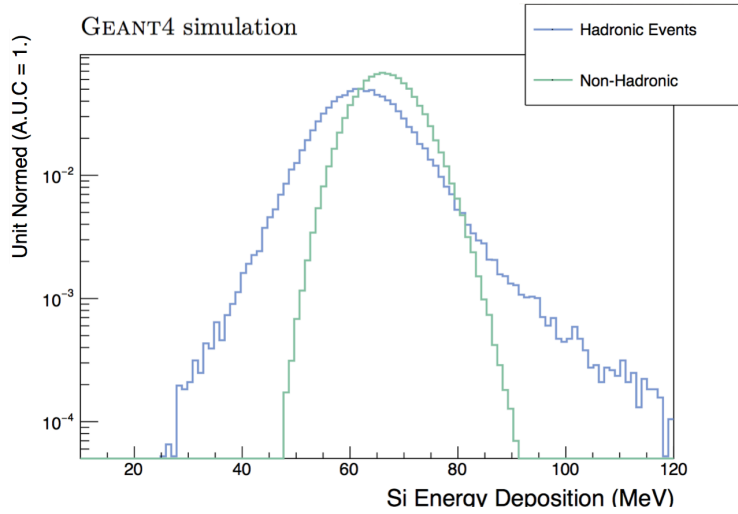


FIG. 34: Distributions of energy deposited by 4 GeV electrons. The histogram labeled “Hadronic” is produced from a sample of 25 million events in which an incident photon or electron with $500 \leq E \leq 3000$ MeV undergoes a nuclear interaction in the target or ECAL. Photo- and electro-nuclear interactions with $E \geq 3000$ MeV are not yet considered, as the physics model in `GEANT4` is currently being validated. The “Non-Hadronic” distribution is produced from a sample of 100,000 events in which no photon or electron with $E \geq 10$ MeV participates in a nuclear reaction.

Photo-nuclear reactions are a primary source of events with low energy deposition that could potentially generate the missing energy signature expected for a signal event. To study this background more carefully, we selected a large sample of photo-nuclear events arising from 3 GeV photons incident on the ECAL. These events act as a proxy for background events since our signal

model entails a soft recoiling electron accompanied by a hard bremsstrahlung photon that converts to a dark photon that decays to undetected dark matter. We expect the magnetic field to produce good transverse separation between the recoil electron and photon at the face of the ECAL and so to get a more detailed understanding of the photo-nuclear backgrounds, it is only necessary to study isolated hard photons that undergo photo-nuclear interactions in the calorimeter. To reach high statistics as efficiently as possible, we fully generate and simulate only those events in which the first interaction of the photon in the calorimeter is photo-nuclear. Such events are by far the most likely to produce a clean missing energy signature. This requirement reduces an initial sample of 10^8 incident photons to only about 75,000 fully simulated events. Note that this sample probes the backgrounds expected from $\approx 3 \cdot 10^9$ EOT since a hard bremsstrahlung ≥ 3 GeV occurs for only about 3% of incident electrons. The total energy, summed over all silicon layers, for these events is shown in Fig. 35.




FIG. 35: The full distribution of the energy deposited in the ECAL for the sub-sample of 3 GeV incident photons that undergo photo-nuclear interactions is shown at left, with a zoom in on the low end of the distribution in the middle plot. The distribution of the total energy deposited in the ECAL+HCAL for the modified ECAL design described in the text is shown at right.

Leading hadron KE [MeV]	Scattering Angle θ	N(Secondaries)	Energy Deposited
2520	10.3	4	0.000
2650	13.4	2	0.147
2020	33.5	2	0.005
1720	41.1	2	0.114
2410	24.0	2	0.03
2530	20.9	6	0.09
1980	33.5	2	0.01
1960	36.7	2	0.115
2900	1.52	2	0.05
1980	35.1	12	0.09

TABLE IV: The subset of photo-nuclear events with less than 0.15 MeV deposited. For all events in the table, the most energetic hadron is a neutron and in fact, most secondaries are neutrons in these events.

Our detailed studies of these events revealed some interesting potential sources of backgrounds. The most pernicious of these backgrounds involves events in which the incident photon interacts

with a W nucleus to produce a final state in which the photon energy has been distributed roughly uniformly over 30-40 charged and neutral hadrons. In some cases the charged secondaries did not escape the W absorber, resulting in only neutrons reaching any of the sensing layers. To help suppress this background we modified the layout of the early ECAL layers by starting with a sense layer followed by a very thin W absorber. For the next 5 alternating layers of sensor and absorber, we smoothly increased the absorber thickness to a nominal value of $1 X_0$. We also changed over to a very thin cooling layer made up of very small diameter, thin-walled cooling tubes sandwiched between thin C-fiber layers in order to avoid thick Cu cooling plates which would add to the absorber thickness. This is mainly to help keep the stack as shallow as possible to maximize the acceptance of the first $20 X_0$.

Combined performance of the ECAL + HCAL for this new layout showed significant improvement. Once again we use a reference sample of 10^8 3 GeV photons. We test a simple ‘MIP’ veto by noting that a minimum ionizing particle would deposit of order 100 KeV in a single layer of silicon in the ECAL or 8 photo-electrons in a single HCAL layer. By requiring there to be no ‘MIP layers’ in the event we reject all but 7 events in a sample of 78,739 that had photo-nuclear interactions. The characteristics of the subset of 10 events with less than 0.15 MeV of total energy deposited in the ECAL are listed in Table IV. It is worth noting that this performance would already be adequate to allow the LDMX experiment to enter into new territory in the search for low mass dark matter. We are now studying how to improve our veto capabilities with enhanced acceptance for secondary hadrons, and also by introducing some materials that could enhance nuclear interactions producing charged hadrons at some layers of the ECAL.

D. Hadronic Veto System

The hadronic veto system is designed for detecting rare processes and therefore studies of its performance focus on “bottom up” vetoes of photo-nuclear reactions. In the most challenging case, the rare processes described in Section IV produce a low multiplicity of neutrons, so we perform studies to compute the efficiency of detecting a single neutron in the hadronic veto system. This will give us an idea of the efficiency with which we can veto any type of photo-nuclear reaction based on the hadronic system only. Lower inefficiencies are expected for multi-neutron events but these have not yet been studied in detail.

We benchmark the hadronic veto system by considering neutrons generated at the face of the hadron calorimeter, with various:

- total energies: 0.95–3.5 GeV
- number of calorimeter layers: 5, 10, 15, 20, 25
- incident angles: 0, 10, 20, 30 degrees

For each neutron phase space point, we generate 5×10^4 events. This discussion focuses on the detector performance for forward (0 degree) neutrons, focusing first on the propagation of neutrons through the HCAL and then on the efficiency of a veto including a simple model of digitization.

To characterize the propagation of neutrons through the HCAL using `GEANT`, we study the “kinetic energy flux”, defined as the amount of kinetic energy passing through the front face of a given HCAL layer. In Fig. 36, we show the neutron kinetic energy flux as a function of HCAL layer for neutrons produced at an incident angle of 0° and with a total energy of 3.5 GeV (kinetic energy of 2.56 GeV). We note that the “shower max” for hadronic showers in this system is typically

around layers 6/7. The horizontal line at 2.56 GeV in the figure indicates the neutrons that have passed through a portion of the calorimeter without interacting. From this, we conclude that roughly $100/50000 \sim 0.3\%$ of neutrons do not interact in the first 15 layers of the system. This number drops to roughly 0.06% (0.02%) for a 20 (25) layer system.

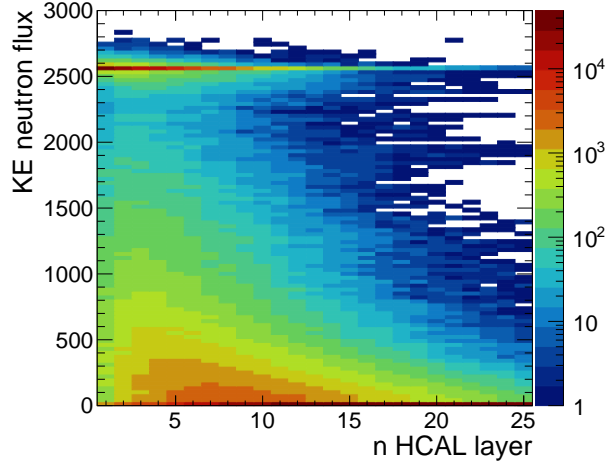


FIG. 36: Neutron kinetic energy flux through the HCAL as a function of layer for a neutron produced with initial total energy 3.5 GeV (kinetic energy of 2.56 GeV) at an incident angle of 0.0° at the front face of the calorimeter.

For neutrons that do interact in the hadronic veto system, we then compute the efficiency of detecting them based on a simple model of the digitization process described in Section V B 2. As a reminder, we sum all sim hits in a given layer and only consider those layers that produce at least 4 times the number of PEs expected from noise. This corresponds to 0.4 MIPs and will be referred to as a *layer-hit*. A veto criterion is defined in which any event that produces ≥ 1 layer-hits is considered to be consistent with having at least one hadron passing through the hadronic veto system.

The number of layer-hits is plotted in Fig. 37 for incident neutrons with total energy of 2.5 GeV (left) or 1.3 GeV (right) for a system of 5, 10, 15, 20, or 25 layers; or in other words, counting only layer-hits in the first 5, 10, 15, 20, or 25 layers, respectively. The number of layer-hits depends strongly on the neutron energy. Furthermore, for lower energy neutrons, increasing the depth of the calorimeter beyond 10–15 layers has little effect on the distribution (and in particular, the veto inefficiency indicated by the leftmost bin) because low energy neutrons typically only interact in the first 5–10 layers, while higher energy neutrons still deposit significant energy in later layers.

In Fig. 38, we count the number of events in which 0 layer-hits were found — these events would *not* be vetoed. The fraction of missed neutrons depends strongly on the number of layers in the system, with the greatest differences for neutrons with 1.5 GeV kinetic energy and above. This suggests, for higher energy neutrons, that the inefficiency is dominated by neutrons which do not interact in the system or are not contained. As discussed above, for lower energy neutrons the efficiency is less dependent on the total depth. Instead, the large inefficiency for low energy neutrons is largely due to a low sampling fraction in the early layers where most of the interactions occur. Work is ongoing to understand how to balance the effects of high sampling fraction and large total interaction lengths for optimizing sensitivity to the expected hadron energy spectrum from photo-nuclear interactions.




FIG. 37: Layer-hit multiplicity for 2.5 GeV, left, and 1.3 GeV, right, neutrons produced with an incident angle of 0° at the front face of the calorimeter for systems with 5, 10, 15, 20, or 25 layers.

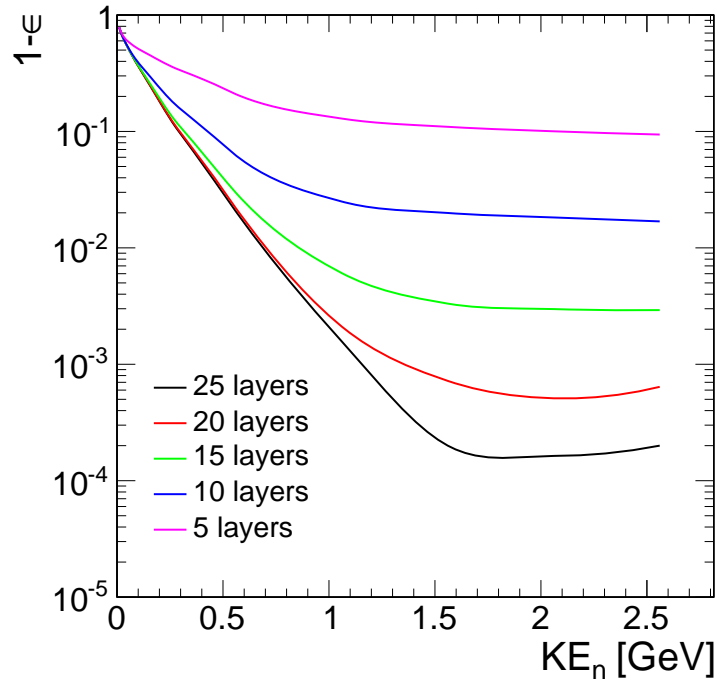


FIG. 38: Fractions of single neutron events which failed the ≥ 1 layer-hit cut, i.e. 1 minus the efficiency of tagging a hadron, versus the neutron's energy. A number of geometries are considered where the number of layers are varied.

TABLE V: Draft trigger menu for LDMX, showing the primary contributions to the trigger budget for a 46 MHz beam rate, with no corrections for any overlaps.

Trigger	Prescale factor	Rate (Hz)
<i>Physics Trigger</i>		
E(ECAL) < 1.8 GeV	1	4000
<i>Background-Measurement Triggers</i>		
E(ECAL) > 3 GeV	3000	100
E(ECAL) > 2 GeV	25	100
Multi-MIP event in trigger scintillator		100
HCAL single layer-hit trigger	2500	100
HCAL multilayer trigger	75	100
Multi-electron trigger (E(ECAL) > 6 GeV)	100	100
<i>Detector-Monitoring Triggers</i>		
Beam-arrival (trigger scintillator)	2.3×10^5	200
Empty-detector (trigger scintillator veto)		200
Total Trigger Budget		5000

From single neutron studies, we conclude that we are able to veto 0° neutrons at the level of one event per $10^1 - 10^4$ with kinetic energies ranging from 400 MeV to 2.5 GeV, depending on number of layers in the system. Low-energy neutrons at greater incident angles have similar veto inefficiencies, while wider-angle high energy neutrons are more easily rejected because they pass through more material (e.g. inefficiencies down to 10^{-5} for the highest-energy neutrons at 30° incidence). In the next Section, we present some results on the veto capabilities of the entire calorimeter system, ECAL + HCAL, for rare photo-nuclear processes.

E. Trigger Performance

As described above, the primary physics trigger is based on the total energy observed in the calorimeter, combined with a requirement of a single incoming electron observed in the trigger scintillator. Figure 39 shows the simulated performance of the primary physics trigger for signal and background.

Besides the primary physics trigger, the LDMX trigger system will also allow the selection of events for calibration, alignment, and background studies. Each event will be marked with the set of triggers that fired. The trigger will include input from the scintillator calorimeter to allow selection of events with hadrons or muons. Based on simulation studies, 0.5% of beam electrons will produce at least one layer-hit in HCAL and one electron in 6000 will produce at least three layer-hits.

Pre-scaled samples of events will also be acquired with only a requirement on energy deposition in the trigger scintillator and with a veto on energy in the trigger scintillator.

An initial draft trigger menu is shown in Table V. Given the projected performance of the primary physics trigger and the design capability of the DAQ, there is sufficient capacity to add additional monitoring triggers, to increase the bandwidth assigned to these triggers, or to define new physics triggers, e.g. triggers appropriate for nuclear physics studies.




FIG. 39: Performance of the primary physics trigger for LDMX. The efficiency for signal electrons of differing energy and the trigger rate for all backgrounds induced by beam electrons are shown as a function of the trigger threshold in MIP units.

VII. CONCLUSION

As this note described, LDMX aims to decisively explore thermal dark matter in the sub-GeV mass range using an electron fixed-target missing momentum approach. This exploration is a priority for the broader physics community, and a timely goal in light of recent advances in calorimeter technology for HL-LHC applications. A key enabler for LDMX is the use of a low-current (\sim pA) but high-repetition electron beam with multi-GeV energy. The recently-proposed DASEL beam line at SLAC would provide such a beam, and offers an excellent opportunity to run LDMX in the early 2020's, thereby enabling LDMX to compete with overseas efforts at CERN. As was described above, initial design studies show considerable promise for the missing momentum approach employed by LDMX. Future work will focus on expanding the detector rejection power for rare photo-nuclear backgrounds, but initial results in this direction are encouraging. Future revisions of this note will detail progress with the detector design, performance studies, and cost estimation, culminating in a letter of intent and conceptual design report in 2017.

Appendices

Appendix A LIGHT DARK MATTER THEORY

The overwhelming evidence for the existence of dark matter (DM) arises from multiple independent sources. Observations of galactic rotation curves, the power spectrum of the Cosmic Microwave Background (CMB), the cosmological matter power spectrum, light element yields from Big Bang Nucleosynthesis (BBN), gravitational lensing, and galaxy cluster collisions all require roughly 85% of the matter in our universe to be cold and non-baryonic. Together, these data constitute smoking gun evidence for physics beyond the Standard Model BSM, which contains no viable DM candidate. For a comprehensive review, see [25].

However, despite this impressive body of evidence, the particle nature of DM remains completely unknown; all indications of its existence derive ultimately from its gravitational influence on visible matter in different contexts. Without making any further assumptions about these hypothetical interactions or DM’s cosmological history, its viable mass range is wildly unconstrained: 10^{-22} eV – $10M_{\odot}$. If DM is lighter than $\sim 10^{-22}$ eV, its Compton wavelength does not fit inside the smallest known DM dominated objects [26] and if it’s heavier than $\sim 10M_{\odot}$ it would have distorted the CMB power spectrum at early times [27]. Given such an enormous range of viable masses, a realistic DM discovery effort requires a well motivated organizing principle to manageably and systematically test this vast space of possibilities without being narrowly tailored to specific models.

1 Thermal Origin

A compelling, well motivated organizing principle for the DM search effort is the hypothesis that it is produced through its interactions with the SM in thermal equilibrium during the early universe. This criterion selects a broad class of popular DM models (including Supersymmetry) and is defined by the requirement that the DM-SM interaction rate exceed the Hubble expansion rate at early times. Once achieved, thermal equilibrium offers some generically attractive features that do not depend sensitively on particular model details:

- **Simply and Generically Realized:** Equilibrium is hard to avoid even with tiny coupling constants between dark and visible matter; thus, most *discoverable* models of DM with non-gravitational interactions fall into this category. Alternative mechanisms (e.g. axions produced via misalignment [28] or feebly coupled DM produced via freeze in [29]) require such small couplings that comprehensive experimental probes are prohibitively difficult in most regions of their parameter space – see [30].
- **Minimum Annihilation Rate:** Equilibrium at temperature T populates the DM with a thermal number density $n_{\text{DM}} \propto T^3$, which is many orders of magnitude larger than the observed abundance. Thus, viable thermal DM must be depleted with an annihilation cross section $\sigma v \geq 3 \times 10^{-26}$ cm³/s via the process of thermal “freeze-out.” This inequality must be saturated (=) if DM is particle/antiparticle symmetric and exceeded (>) if there is an additional primordial DM asymmetry that contributes to the total abundance in addition to the freeze-out component. In addition, any particle that makes up a *fraction* of the DM must exceed this minimum annihilation cross-section.




FIG. 40: The allowed mass range over which DM can thermalize with the SM in the early universe and yield the observed relic abundance via annihilation. For masses below $\lesssim 10$ keV, DM is too hot to form the observed structure of the universe on large scales [7] and for masses above $\gtrsim 10$ TeV, a perturbative annihilation rate cannot achieve the correct relic abundance in simple models [8].

- **Sharply Defined Mass Window:** The viable DM mass range becomes much narrower – see Fig. 40. Thermal DM below $\lesssim 10$ keV erases small scale structure in conflict with observation; thermal DM above $\gtrsim 10$ TeV requires nonperturbative and/or non-unitary couplings to realize a sufficiently large annihilation rate.

2 “WIMP” Dark Matter

If thermal DM is realized in the upper half of the thermal mass window $\sim \text{GeV} - 10 \text{ TeV}$, it can be a Weakly Interacting Massive Particle (WIMP) charged under the familiar electroweak force and achieve the observed relic abundance by annihilating through via SM gauge interactions. This class of models exploits the numerical coincidence between electroweak-sized cross sections and the requisite annihilation cross section for thermal freeze-out $\sigma \sim \alpha^2/m_Z^2 \rightarrow \sigma v \sim 3 \times 10^{-26} \text{ cm}^3/\text{s}$ (see [31] for a pedagogical treatment). This so-called “WIMP miracle” also naturally arises in popular extensions of the SM that address the electroweak hierarchy problem (e.g. Supersymmetry and extra dimensions) and may even be realized with fifth forces beyond the SM so long as the appropriate coupling-to-mass ratio is of electroweak size, but the latter possibility is not obligatory for masses in the WIMP range.

The discovery techniques available for WIMP and WIMP-like DM are well known and firmly established in the experimental community:

- **Direct Detection:** Terrestrial searches for non relativistic WIMP-nucleon scattering in a shielded underground detector [32]. This technique is powerful for DM masses near $\sim 10 \text{ GeV} - 10 \text{ TeV}$, but loses sensitivity in the neighborhood of 1 GeV where typical nuclear recoils are $\lesssim \text{keV}$; below detection thresholds. Moreover, the sensitivity of this approach depends critically on the presently unknown DM velocity distribution in the terrestrial neighborhood and is therefore subject to potentially large systematic uncertainties.
- **Indirect Detection:** Typically space-based searches of DM annihilation in regions of high DM density (galactic center, dwarf galaxies etc.) [33]. This technique is approaching sensitivity to thermal DM annihilation rates for a variety of scenarios, but poorly constrains DM below the few-GeV scale. As with direct detection, the signal strength for indirect detection depends on an unknown DM phase space profile in regions of high density, so the systematic uncertainties of this approach may also be very large.

- **Collider Production:** Laboratory based searches focus on DM produced in association with visible final states in SM particle collisions [34]. This technique is powerful and not limited by halo uncertainties or the limitations of non-relativistic scattering off DM particles in the Earth’s neighborhood. However, for DM below the GeV scale, the transverse missing energy in a typical production event is too low to impart sufficient recoil P_T to the other visible object(s) in the final state, so sensitivities remain weak for the lower half of the thermal mass window [3].

3 Light Dark Matter (LDM)

Although the thermal freeze-out production mechanism can apply equally well to LDM in the keV-GeV range, achieving the required annihilation rate is no longer possible with SM forces. Annihilation to SM particles via virtual electroweak gauge boson exchange scales as $\sigma v \sim \alpha^2 m_{\text{DM}}^2 / m_Z^4 \ll 3 \times 10^{-26} \text{ cm}^3/\text{s}$, which is insufficient for freeze-out if $m_{\text{DM}} \ll m_Z$. Thus, for light thermal DM to be viable, the following are necessary:

- **Light New Forces:** There have to be comparably light force carriers to mediate an efficient annihilation rate for thermal freeze-out.
- **Portals:** Both the DM and the mediator must be singlets under the full SM gauge group; otherwise they would have been produced and detected at LEP. Thus, in order for the DM to efficiently annihilate away its abundance, the mediator particle must have an unsuppressed tree level coupling [67] to SM particles through a mass-dimension < 4 singlet “portal” operator built out of SM fields.

The second bullet point sharply constrains the options for LDM mediator particles since the only renormalizable SM operators that satisfy this requirement are

$$\hat{\mathcal{O}}_{\text{portal}} = H^\dagger H \ , \ HL \ , \ B_{\mu\nu} \ , \quad (4)$$

and are known respectively as the Higgs, lepton, and vector portals (for a discussion, see [35]). Here H is the SM Higgs doublet, L is the SM lepton doublet, and $B_{\mu\nu} = \partial_\mu B_\nu - \partial_\nu B_\mu$ is the $U(1)_Y$ hypercharge field strength tensor, which is independently gauge invariant. Each portal corresponds to a different choice of mediator spin – only a scalar mediator can couple to the $H^\dagger H$ bilinear, only a fermionic mediator can couple to the lepton portal operator LH , and only a spin-1 mediator can couple to the $B_{\mu\nu}$ vector portal. However, for the most predictive model variations, the scalar mediator is strongly disfavored by a variety of experimental constraints [36] and the lepton portal interaction is generically proportional to factors of neutrino masses $m_\nu \lesssim 0.1 \text{ eV}$, so it is difficult to sustain thermal contact between DM and SM sectors. This leaves the vector portal, which we focus on for the rest of this discussion, for which the “dark photon” A' is the representative mediator of LDM interactions.

For completeness we also note that there are other viable vector mediators in addition to the dark photon A' , which mixes with the vector portal with a kinetic mixing parameter ϵ . Instead of this mixing (and eventually the photon after electroweak symmetry breaking), a new vector interaction between dark and visible matter can arise by directly gauging a linear combination of global SM quantum numbers. For example, SM fermions can be charged under one or many of the following new gauge groups: $U(1)_{B-L}$, $U(1)_{\ell_i-\ell_j}$, $U(1)_{3B-\ell_i}$, where B/L represent SM baryon and lepton number, while ℓ_i represents a particular lepton flavor number. All of these




FIG. 41: Schematic representation of the different DM annihilation modes (top row) and A' decay modes for $m_\chi/m_{A'}$ ratios. a) Secluded annihilation scenario with a visibly decaying mediator. Since the annihilation rate is independent of the A' SM coupling, this scenario has no thermal target and cannot be presented on the y vs. m_χ plane. However, there is an active and growing program to probe dark photons in this regime by observing their visible decay products (see [1, 30] for more details). b) Compressed region with direct annihilation, but a visibly decaying mediator. Since the annihilation rate in this regime depends on ϵ , there is a testable thermal target; probing sufficiently small values of ϵ can decisively test this scenario. c) Direct annihilation and invisibly decaying mediator particle. This regime is the primary focus of this document.

groups are anomaly free and their gauge bosons can serve as mediators to the dark sector if the DM candidate also carries charge under the new group(s). However, most of these groups feature couplings between the new vector mediator and SM electrons, so for the purpose of this document, we can focus the discussion by considering only the kinetically mixed dark photon, whose phenomenology is qualitatively very similar to the examples listed above. Typically the only difference involves reinterpreting experimental bounds and projections with the replacement $\epsilon e \rightarrow g_{\text{new}}$ where g_{new} is the gauge coupling of the new $U(1)$ group [68]. Thus, for the remainder of this document, we can consider a single dark photon without loss of generality.

4 Predictive LDM Targets

We define the LDM particle to be χ and the mediator to be a “dark photon” A' with lagrangian

$$\mathcal{L} = i\bar{\chi}\not{\partial}\chi + m_\chi\bar{\chi}\chi + g_D A'_\mu \bar{\chi}\gamma^\mu\chi - \frac{1}{4}F'_{\mu\nu}F'^{\mu\nu} + \frac{\epsilon}{2}F'_{\mu\nu}F^{\mu\nu} + \frac{m_{A'}^2}{2}A'_\mu A'^\mu, \quad (5)$$

where $F'_{\mu\nu} = \partial_\mu A'_\nu - \partial_\nu A'_\mu$, $\epsilon \ll 1$ is the kinetic mixing parameter, which controls A' mixing with the SM photon, and $g_D \equiv \sqrt{4\pi\alpha_D}$ is the A' coupling to the DM. Although we've assumed the DM

is a fermion, the qualitative features of this discussion are insensitive to its spin and compatible with scalar candidates as well.

After diagonalizing the kinetic mixing interaction, the dark photon A' acquires a coupling to the SM electromagnetic current

$$\mathcal{L} \rightarrow i\bar{\chi}\not{\partial}\chi + m_\chi\bar{\chi}\chi + A'_\mu\left(g_D\bar{\chi}\gamma^\mu\chi + \epsilon e\sum_f Q_f\bar{f}\gamma^\mu f\right) - \frac{1}{4}F'_{\mu\nu}F'^{\mu\nu} + \frac{m_{A'}^2}{2}A'_\mu A'^\mu, \quad (6)$$

where f is a SM fermion and Q_f is its electromagnetic charge.

We distinguish between two distinct annihilation regimes depicted schematically in Fig. 41

- **Secluded Annihilation:** For $m_{A'} < m_\chi$, DM annihilation will predominantly proceed through $\chi\chi \rightarrow A'A'$, followed by $A' \rightarrow ff$ decays to SM fermions. However, the annihilation rate in this regime is independent of the SM- A' coupling ϵ and therefore difficult to test since thermal freeze-out can proceed even for tiny values of ϵ . This regime is depicted on the leftmost column of Fig. 41
- **Direct Annihilation:** For $m_{A'} > m_\chi$, the mediator decays predominantly to DM and annihilation proceeds via $\chi\chi \rightarrow A'^* \rightarrow ff$ to SM fermions f through a virtual mediator. This regime is depicted in the middle and rightmost column of Fig. 41. The middle column is the compressed region for which $m_\chi < m_{A'} < 2m_\chi$ and the annihilation rate depends on ϵ but the mediator decay to DM is kinematically forbidden.

Since the cross section for direct annihilation is proportional to all the parameters in the DM lagrangian, it is convenient to define the dimensionless interaction strength y as

$$\sigma v(\chi\chi \rightarrow A'^* \rightarrow ff) \propto \epsilon^2 \alpha_D \frac{m_\chi^2}{m_{A'}^4} = \frac{y}{m_\chi^2}, \quad y \equiv \epsilon^2 \alpha_D \left(\frac{m_\chi}{m_{A'}}\right)^4 \quad (7)$$

thus, for each choice of m_χ there is a unique value of y compatible with thermal freeze-out independently of the individual values of α_D , ϵ and $m_\chi/m_{A'}$. Reaching experimental sensitivity to this benchmark for masses between 10 keV – GeV would provide decisive coverage of these scenarios.

5 Current Bounds on LDM

Although LDM represents half of the viable thermal DM mass range, there has never been a dedicated search for this class of models. Most constraints that apply to the scenarios considered here are reinterpretations, by theorists, of experimental results collected for other purposes and fall into the following categories:

- **CMB Power Spectrum:** LDM can annihilate out of equilibrium to SM particles during the time of recombination and ionize the newly formed hydrogen, thereby modifying the CMB power spectrum in conflict with observations from PLANCK [37]. If this annihilation is s -wave and the DM is particle symmetric, the CMB rules out LDM below ~ 10 GeV. However if the annihilation is p -wave, as would easily be achieved for a scalar DM candidate coupled to a dark photon mediator, or if the DM population is different at the time of CMB (e.g. asymmetric or Majorana-inelastic, see [3] for a discussion) this constraint no longer applies.

- **CMB/BBN N_{eff} :** LDM is also constrained by the observed number of relativistic species present during BBN and recombination. Since thermal DM freezes out near temperatures $T \sim m_\chi/20$, for LDM this process can occur after SM neutrinos have decoupled from SM radiation, so DM annihilation increases the temperature of photons relative to neutrinos, which introduces a deficit in N_{eff} in conflict with observation [38]. However, a deficit in N_{eff} can be compensated by additional sources of hidden sector radiation or sterile neutrinos and such additional contributions may alleviate the recent tension between late and early time extractions of the Hubble rate [39].
- **Colliders:** Although high energy colliders (LHC, Tevatron, LEP) generally have poor sensitivity to LDM, high intensity B-factories can perform analogous searches for $e^+e^- \rightarrow \gamma A'$ production and set impressive constraints in the 100 MeV - few GeV mass range. The best current constraint of this form arises from re-interpreting [12, 40] a BaBar search for $\Upsilon(3S) \rightarrow \gamma + a$ search where a is an invisibly decaying scalar particle [41]. While this search has some of the best sensitivity for $O(\text{GeV})$ A' masses, its sensitivity below a GeV is limited by a peaking background from un-vetoed photons that fakes the DM signal.
- **Beam Dumps:** The LSND proton beam dump experiment [42] is sensitive to LDM production and scattering in their measurement of the electron-neutrino neutral current cross section. At LSND, DM can be produced in pion decays $\pi^0 \rightarrow \gamma A' \rightarrow \gamma \chi \chi$ and subsequently scatter in a downstream detector. The E137 electron beam dump experiment to search for axion-like particles is sensitive to similar processes in which DM is relativistically produced via the kinetic mixing interaction in the beam dump, passes through a downstream detector and deposits electromagnetic energy by scattering off detector targets [43].
- **Rare Kaon Decays:** LDM can also be produced in rare Kaon decays $K^+ \rightarrow \pi^+ A' \rightarrow \pi^+ \chi \chi$ that can contribute to the signal region of E787/E949 [44, 45] that's used to measure the $K^+ \rightarrow \pi^+ \bar{\nu} \nu$ branching ratio.
- **Electron Direct Detection:** The results of XENON10 S2-only study of electron recoil signals can be used to constrain LDM that scatters elastically off detector electrons [46]. Although the backgrounds are not well understood, a conservative extraction of the DM scattering limit can be used to constrain this parameter space. This bound is competitive with E137 and E787/E949 in Fig. 2, but slightly weaker for the benchmarks presented.

These constraints are collected in the y vs m_χ parameter space depicted in Fig. 2 alongside projections for LDMX and for other proposed experiments, described below. Also shown are the thermal targets for fermion and scalar LDM candidates, which are invariant in this parameter space regardless of what assumptions about ϵ , α_D , and $m_\chi/m_{A'}$ we choose. Direct detection constraints are also most naturally expressed as functions of y and m_χ , but accelerator-based constraints are not — for example, collider production bounds depend on $m_{A'}$ and ϵ . However, within the predictive framework of direct DM annihilation, both α_D and $m_\chi/m_{A'}$ are bounded from above, accelerator-based constraints typically become *stronger* as either α_D or $m_\chi/m_{A'}$ is decreased (specifically, lowering α_D expands the constrained regions downward in Figure 2 while lowering $m_\chi/m_{A'}$ moves them downward and to the left). Therefore, in the lower panel we plot conservative versions of these constraints with α_D and $m_\chi/m_{A'}$ near their upper limits to reveal all gaps in the DM parameter space consistent with the assumption of direct annihilation.

The upper panel of Figure 2 includes only missing mass/momentum/energy experiments, whose sensitivity can be directly compared to LDMX irrespective of α_D and $m_\chi/m_{A'}$. In par-

ticular, these are the only experiments sensitive to LDM in the limit of small α_D . This plot illustrates the large, orders of magnitude gaps in coverage between existing and projected constraints, and the thermal relic contours for LDM. In order to decisively cover thermal LDM in the direct annihilation regime, a dedicated effort is required. To this end, we also show the projections for LDMX at DASEL, which is the only proposed effort to probe the thermal target for both scalars and fermions down to the MeV range.

6 Future Bounds on LDM

As of the Dark Sectors 2016 meeting at SLAC, the community undertook an effort to establish projections for dark force and light dark matter sensitivity for experiments into the mid 2020's. Below we summarize some of the key advances expected in light dark matter searches.

- **Colliders:** The Belle II experiment at SuperKEKB is expected to start taking data in 2018 using an e^+e^- asymmetric collider at the Υ resonance near ≈ 10.3 GeV. Expected to achieve integrated luminosities of 50/ab (~ 100 times greater than BaBar's), Belle II is sensitive to DM through $e^+e^- \rightarrow \gamma(A' \rightarrow \chi\bar{\chi})$ production and the observational signature is a monochromatic photon produced in association with missing energy. See Ref. [40] for more details. Projections are not shown for dark photon masses below 1 GeV because they are expected to be limited by instrumental effects (especially calorimeter hermiticity), and may not improve over BaBar limits despite the higher statistics.
- **Electron/Positron Fixed Targets:** Proposed fixed-target experiments will both improve on the beam dump searches discussed above as well as pursuing several distinct approaches to search for DM, elaborated on in Section IID. The proposed BDX experiment at JLab plans to utilize the upgraded 11 GeV CEBAF beam to search for dark matter. In this setup, dark matter is produced in the electron beam dump, passes through shielding material (to range out SM particles), and scatters by depositing electromagnetic energy in a detector placed behind the beam dump. Nominally, BDX will dump $\sim 10^{22}$ electrons on target (EOT) onto an Al beam dump; see Ref. [12, 47–49] for more details. The DarkLight experiment [50] is also sensitive to DM production $ep \rightarrow epA'$ followed by the invisible decay $A' \rightarrow \chi\chi$ which appears as a resonance in the reconstructed ep invariant mass spectrum [51]. DarkLight is currently being set up at the 100 MeV Low Energy Recirculating Facility (LERF), formerly the Free Electron Laser, at Jefferson Lab [64] and involves a 5 mA current of electrons impinging on a windowless gas target of molecular hydrogen. Finally, MMAPS uses a 6 GeV positron beam incident on a Be target at the Wilson Lab at Cornell University [52]. The DM signal of interest is the A' invariant mass reconstructed from $e^+e^- \rightarrow \gamma A'$ annihilation. A similar e^+e^- annihilation search with a 500 MeV positron beam is planned at VEPP-3 at the Budker Institute for Nuclear Physics in Novosibirsk [53].
- **Proton Fixed Targets:** The MiniBooNE experimental setup in beam dump mode can be sensitive to LDM. The main DM production modes are $\pi^0/\eta\eta' \rightarrow \gamma(A' \rightarrow \chi\bar{\chi})$ and $q\bar{q} \rightarrow A' \rightarrow \chi\bar{\chi}$ and detection proceeds via $\chi e \rightarrow \chi e$, or $\chi N \rightarrow \chi N$ elastic scattering, or via inelastic such as $\chi N \rightarrow \chi(\Delta \rightarrow N\pi^0)$. MiniBooNE has completed running in off-target mode with 1.86×10^{20} POT and the analysis is currently in progress. See Ref. [42, 54–62] for more details.

- **Electron Direct Detection:** A promising technique for observing halo DM involves directly observing electron recoils induced by incident DM particles from the halo. This technique has already been used to constrain the DM-electron cross section using XENON10 calibration data (see Fig 2), but this technique remains background-limited and is likely to improve only modestly with upcoming results from XENON100 and other direct detection experiments [46, 63, 64].
- **Low-Threshold Direct Detection** A strong thrust of the next-generation direct detection experiments has been lowering their recoil energy thresholds to become sensitive to O(GeV) dark matter. Among this cohort of proposed/planned experiments, SuperCDMS-SNOLAB is promising greatly improved coupling sensitivity [65] down to ~ 500 MeV dark matter masses, improving on CRESST-II's current limit by 5 orders of magnitude, while NEWS-SNO [66] aims to explore dark matter as light as 100 MeV using hydrogen nuclear recoils in a gas target. Projections for both experiments are shown in the lower panel of Figure 2. However, an important caveat to these experiments' sensitivity is that they assume dark matter has elastic (mass-diagonal) interactions, which is not generic for LDM. If the interactions are inelastic, then direct detection is only sensitive to one-loop scattering processes and therefore explores far less parameter space in these models. Thus, while SuperCDMS-SNOLAB tests $y \sim 10^{-15}$ for elastically coupled DM, the analogous constraint on inelastic DM is only at the level of 10^{-6} (weaker than BaBar) [3].

-
- [1] J. Alexander et al. (2016), 1608.08632, URL <http://inspirehep.net/record/1484628/files/arXiv:1608.08632.pdf>.
- [2] E. Izaguirre, G. Krnjaic, P. Schuster, and N. Toro, Phys. Rev. **D91**, 094026 (2015), 1411.1404.
- [3] E. Izaguirre, G. Krnjaic, P. Schuster, and N. Toro, Phys. Rev. Lett. **115**, 251301 (2015), 1505.00011.
- [4] S. N. Gninenko, N. V. Krasnikov, M. M. Kirsanov, and D. V. Kirpichnikov (2016), 1604.08432.
- [5] DASEL Design Team, Tech. Rep. (2016), URL <https://confluence.slac.stanford.edu/download/attachments/213891196/DASEL%20LOI.pdf?version=2&modificationDate=1473892060000&api=v2>.
- [6] S. Ritz et al. (HEPAP Subcommittee) (2014).
- [7] M. Viel, G. D. Becker, J. S. Bolton, and M. G. Haehnelt, Phys. Rev. **D88**, 043502 (2013), 1306.2314.
- [8] K. Griest and M. Kamionkowski, Phys. Rev. Lett. **64**, 615 (1990).
- [9] K. A. Ulmer (ATLAS, CMS), in *3rd Large Hadron Collider Physics Conference (LHCP 2015) St. Petersburg, Russia, August 31-September 5, 2015* (2016), 1601.03774, URL <https://inspirehep.net/record/1415576/files/arXiv:1601.03774.pdf>.
- [10] B. W. Lee and S. Weinberg, Phys. Rev. Lett. **39**, 165 (1977).
- [11] M. S. Madhavacheril, N. Sehgal, and T. R. Slatyer, Phys. Rev. **D89**, 103508 (2014), 1310.3815.
- [12] E. Izaguirre, G. Krnjaic, P. Schuster, and N. Toro, Phys. Rev. **D88**, 114015 (2013), 1307.6554.
- [13] P. Hansson Adrian et al. (HPS Collaboration), *Heavy Photon Search Experiment at Jefferson Laboratory: proposal for 2014-2015 run* (2013), URL https://confluence.slac.stanford.edu/download/attachments/86676777/hps_2014.pdf.
- [14] CMS Collaboration, Tech. Rep. CERN-LHCC-2015-010. LHCC-P-008. CMS-TDR-15-02 (2015), URL <https://cds.cern.ch/record/2020886>.

- [15] G. Battistoni, S. Muraro, P. R. Sala, F. Cerutti, A. Ferrari, S. Roesler, A. Fasso, and J. Ranft, AIP Conf. Proc. **896**, 31 (2007), [,31(2007)].
- [16] J. Mans et al. (CMS), Tech. Rep. CERN-LHCC-2012-015, CMS-TDR-010 (2012), URL <http://cds.cern.ch/record/1481837>.
- [17] S. Agostinelli et al. (GEANT4), Nucl. Instrum. Meth. **A506**, 250 (2003).
- [18] J. Alwall, P. Demin, S. de Visscher, R. Frederix, M. Herquet, F. Maltoni, T. Plehn, D. L. Rainwater, and T. Stelzer, JHEP **09**, 028 (2007), 0706.2334.
- [19] J. McCormick, eConf **C050318**, 1003 (2005).
- [20] Y.-S. Tsai, Phys. Rev. **D34**, 1326 (1986).
- [21] S. Abrahamyan et al. (APEX), Phys. Rev. Lett. **107**, 191804 (2011), 1108.2750.
- [22] J. Alwall et al., Comput. Phys. Commun. **176**, 300 (2007), hep-ph/0609017.
- [23] S. Andreas, C. Niebuhr, and A. Ringwald, Phys. Rev. **D86**, 095019 (2012), 1209.6083.
- [24] D. Wright, private communication (2015).
- [25] G. Bertone, D. Hooper, and J. Silk, Phys. Rept. **405**, 279 (2005), hep-ph/0404175.
- [26] J. F. Navarro, C. S. Frenk, and S. D. M. White, Astrophys. J. **462**, 563 (1996), astro-ph/9508025.
- [27] S. Bird, I. Cholis, J. B. Muñoz, Y. Ali-Haïmoud, M. Kamionkowski, E. D. Kovetz, A. Raccanelli, and A. G. Riess, Phys. Rev. Lett. **116**, 201301 (2016), 1603.00464.
- [28] L. Visinelli and P. Gondolo, Phys. Rev. **D80**, 035024 (2009), 0903.4377.
- [29] L. J. Hall, K. Jedamzik, J. March-Russell, and S. M. West, JHEP **03**, 080 (2010), 0911.1120.
- [30] R. Essig et al., in *Proceedings, Community Summer Study 2013: Snowmass on the Mississippi (CSS2013): Minneapolis, MN, USA, July 29-August 6, 2013* (2013), 1311.0029, URL <http://inspirehep.net/record/1263039/files/arXiv:1311.0029.pdf>.
- [31] E. W. Kolb and M. S. Turner, Front. Phys. **69**, 1 (1990).
- [32] T. Marrodán Undagoitia and L. Rauch, J. Phys. **G43**, 013001 (2016), 1509.08767.
- [33] J. Conrad, in *Interplay between Particle and Astroparticle physics (IPA2014) London, United Kingdom, August 18-22, 2014* (2014), 1411.1925, URL <https://inspirehep.net/record/1326617/files/arXiv:1411.1925.pdf>.
- [34] A. Askew, S. Chauhan, B. Penning, W. Shepherd, and M. Tripathi, Int. J. Mod. Phys. **A29**, 1430041 (2014), 1406.5662.
- [35] M. Pospelov, Phys. Rev. **D80**, 095002 (2009), 0811.1030.
- [36] G. Krnjaic, Submitted to: Phys. Rev. D (2015), 1512.04119.
- [37] P. A. R. Ade et al. (Planck), Astron. Astrophys. **594**, A13 (2016), 1502.01589.
- [38] K. M. Nollett and G. Steigman, Phys. Rev. **D89**, 083508 (2014), 1312.5725.
- [39] A. G. Riess et al., Astrophys. J. **826**, 56 (2016), 1604.01424.
- [40] R. Essig, J. Mardon, M. Papucci, T. Volansky, and Y.-M. Zhong, JHEP **11**, 167 (2013), 1309.5084.
- [41] B. Aubert et al. (BaBar), in *Proceedings, 34th International Conference on High Energy Physics (ICHEP 2008): Philadelphia, Pennsylvania, July 30-August 5, 2008* (2008), 0808.0017, URL <http://www-public.slac.stanford.edu/sciDoc/docMeta.aspx?slacPubNumber=slac-pub-13328>.
- [42] P. deNiverville, M. Pospelov, and A. Ritz, Phys. Rev. **D84**, 075020 (2011), 1107.4580.
- [43] B. Batell, R. Essig, and Z. Surujon, Phys. Rev. Lett. **113**, 171802 (2014), 1406.2698.
- [44] S. Adler et al. (E787), Phys. Rev. Lett. **79**, 2204 (1997), hep-ex/9708031.
- [45] A. V. Artamonov et al. (E949), Phys. Rev. Lett. **101**, 191802 (2008), 0808.2459.
- [46] R. Essig, A. Manalaysay, J. Mardon, P. Sorensen, and T. Volansky, Phys. Rev. Lett. **109**, 021301 (2012), 1206.2644.
- [47] M. Battaglieri et al. (BDX) (2014), 1406.3028.

- [48] E. Izaguirre, G. Krnjaic, and M. Pospelov, *Phys. Rev.* **D92**, 095014 (2015), 1507.02681.
- [49] M. Battaglieri et al. (BDX) (2016), 1607.01390.
- [50] J. Balewski et al. (2014), 1412.4717, URL <http://inspirehep.net/record/1334361/files/arXiv:1412.4717.pdf>.
- [51] Y. Kahn and J. Thaler, *Phys. Rev.* **D86**, 115012 (2012), 1209.0777.
- [52] J. Alexander (2016), URL <https://indico.cern.ch/event/507783/contributions/2150181/attachments/1266367/1874844/SLAC-MMAPS-alexander.pdf>.
- [53] B. Wojtsekhowski, D. Nikolenko, and I. Rachek (2012), 1207.5089.
- [54] B. Batell, M. Pospelov, and A. Ritz, *Phys. Rev.* **D80**, 095024 (2009), 0906.5614.
- [55] P. deNiverville, D. McKeen, and A. Ritz, *Phys. Rev.* **D86**, 035022 (2012), 1205.3499.
- [56] R. Dharmapalan et al. (MiniBooNE) (2012), 1211.2258.
- [57] B. Batell, P. deNiverville, D. McKeen, M. Pospelov, and A. Ritz, *Phys. Rev.* **D90**, 115014 (2014), 1405.7049.
- [58] D. E. Morrissey and A. P. Spray, *JHEP* **06**, 083 (2014), 1402.4817.
- [59] Y. Kahn, G. Krnjaic, J. Thaler, and M. Toups (2014), 1411.1055.
- [60] D. Gorbunov, A. Makarov, and I. Timiryasov, *Phys. Rev.* **D91**, 035027 (2015), 1411.4007.
- [61] J. Blumlein and J. Brunner, *Phys. Lett.* **B731**, 320 (2014), 1311.3870.
- [62] P. deNiverville, C.-Y. Chen, M. Pospelov, and A. Ritz (2016), 1609.01770.
- [63] R. Essig, J. Mardon, and T. Volansky, *Phys. Rev.* **D85**, 076007 (2012), 1108.5383.
- [64] R. Essig, M. Fernandez-Serra, J. Mardon, A. Soto, T. Volansky, and T.-T. Yu, *JHEP* **05**, 046 (2016), 1509.01598.
- [65] P. Cushman et al., in *Proceedings, Community Summer Study 2013: Snowmass on the Mississippi (CSS2013): Minneapolis, MN, USA, July 29-August 6, 2013* (2013), 1310.8327, URL <https://inspirehep.net/record/1262767/files/arXiv:1310.8327.pdf>.
- [66] URL <https://www.snolab.ca/news-projects/index.html>.
- [67] If the operator were nonrenormalizable, there would have to be additional, sub-electroweak states integrated out to generate such an interaction since electroweak sized suppression scales in higher dimension operators would reintroduce the overproduction problem. However, the states that UV complete such an interaction would be electroweak charged and ruled out by LEP searches for light, electroweak charged matter.
- [68] The only exceptions to this general feature are groups for which the mediator couples predominantly to the second and third generations, e.g. $U(1)_{\mu-\tau}$ for which the sensitivity of electron beam based experiments are generally diminished.