FY22 Status Report for LDMX Dark Matter New Initiative

June 13, 2022

Sponsoring Institution: SLAC National Accelerator Laboratory

Principal Investigator:	Senior Investigator:	Senior Investigator:
Dr. Timothy K. Nelson	Prof. Philip C. Schuster	Prof. Natalia Toro
tknelson@slac.stanford.edu	schuster@slac.stanford.edu	ntoro@slac.stanford.edu
650-926-2274	650-926-5180	650-926-3625

Participating Institution: University of California at Santa Barbara

co-Principal Investigator:
Prof. Joseph Incandela
jincandela@ucsb.edu
805-893-8270

Participating Institution: California Institute of Technology

co-Principal Investigator:	Senior Investigator:
Prof. Bertrand Echenard	Prof. David Hitlin
echenard@hep.caltech.edu	hitlin@caltech.edu
626-395-6674	626-395-6694

Participating Institution: Fermi National Laboratory

co-Principal Investigator:
Dr. Nhan Tran
ntran@fnal.gov
630-840-2575

Participating Institution: University of Minnesota

co-Principal Investigator:
Prof. Jeremiah Mans
jmmans@umn.edu
612-625-8994

Participating Institution: Texas Tech University

co-Principal Investigator:
Prof. Andrew Whitbeck
Andrew.Whitbeck@ttu.edu
806-834-1205

Participating Institution: University of Virginia

co-Principal Investigator:	Senior Investigator:
Prof. R. Craig Group	Prof. E. Craig Dukes
rcg6p@virginia.edu	ecd3m@Virginia.EDU
434-243-5552	434-982-5364

Contents

1 Introduction

One of three Priority Research Directions identified in the "Dark Matter New Initiatives" (DMNI) BRN report [1] is

> **PRD 1: Create and detect dark matter particles below the proton mass and associated forces, leveraging DOE accelerators that produce beams of energetic particles**... Interactions of energetic particles recreate the conditions of dark matter production in the early universe. Small experiments using established technology can detect dark matter production with sufficient sensitivity to test compelling explanations for the origin of dark matter and explore the nature of its interactions with ordinary matter.

In discussing this Priority Research Direction, the report highlights a strong motivation for "10- to 1000-fold improvements in sensitivity over current searches" for dark matter (DM) production (Thrust 1) and discusses the unique capability of the missing-momentum technique to meet or exceed this goal over most of the MeV-GeV mass range. The report also notes a secondary emphasis on "explor[ing] the structure of the dark sector by producing and detecting unstable dark particles" (Thrust 2).

The Light Dark Matter eXperiment (LDMX) is a small experiment that realizes this missing-momentum concept. As shown in Figure 1, dark matter is produced in electron fixed-target collisions and detected through the use of tracking and calorimetry to identify events where an incoming electron lost most of its energy to DM production. Operating at high rate in a continuous-wave (CW) electron beam for only a few years, this approach can fully address Thrust 1 by achieving a 1000-fold improvement in sensitivity to dark matter production, while also broadly searching for unstable dark particles that are the objective of Thrust 2. In addition to the primary dark matter motivation, LDMX data can also provide measurements of electron-nucleon interactions at large momentum transfers that are of critical importance to interpreting the data from the flagship neutrino program at Fermilab.

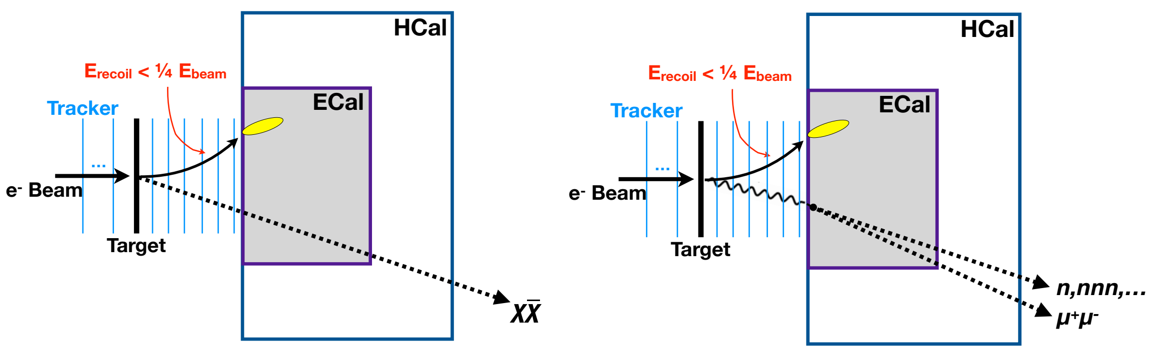

Figure 1: The conceptual layout of the LDMX apparatus demonstrating the missing-momentum technique of searching for dark matter in fixed target production (left) and key backgrounds that the detector is designed to reject (right). In signal events, an incoming beam electron loses most of its energy and experiences a hard kick in transverse momentum in the production process, producing a soft, high-angle recoil and no other detectable interaction products in strong contrast to the Standard Model scattering processes. However, in cases that are rare but relevant to a high-statistics experiment, a hard bremsstrahlung photon is produced that undergoes muon conversion or photo-nuclear reactions in the detector material that are difficult to detect. These backgrounds define the required veto performance of the detector.

While this experiment can achieve new sensitivity for sub-GeV dark matter with fewer than 10^{12} electrons on target (EOT) – only weeks of operation – achieving the full potential of the experiment requires event yields as high as 10^{16} EOT. Meanwhile, ensuring high purity for the missing momentum signature requires the ability to correctly associate all of the particles belonging to each individual event so that each incoming electron is correctly tagged to its interaction products. For large event yields, this requires a near-CW beam and granular detectors with high-rate capability and excellent time resolution. Furthermore, because the experiment is an active dump, some elements of the detector must be relatively radiation tolerant.

Most of the technologies required to meet these challenges are well established, and none of them is beyond the state of the art. High repetition rate electron beams are available within the DOE complex at both SLAC and JLab, as well as at other labs worldwide. Charged particle tracking with the granularity and rate capability required for the experiment has existed for at least 20 years, and the technologies needed for hadronic calorimetry, triggering, and data acquisition are similarly mature. Only the electromagnetic calorimeter faces rates and radiation doses that require the newest detector technologies. As a result, little fundamental development is required: the task is one of adapting technologies, designs, and hardware that already exist in a way that optimizes the performance while minimizing the technical risk, cost, and effort involved in mounting the experiment. In particular, LDMX re-purposes designs from the HPS Silicon Vertex Tracker for tracking, the CMS upgrade HGCal for the ECal, and the Mu2e Cosmic Ray veto for the HCal, and plans to use LCLS-II drive beam at SLAC in a way that is parasitic and invisible to the primary photon science program.

Commensurately, the LDMX DMNI project is a "Track 1" project, consisting of the design and prototyping required to adapt hardware and technologies developed for other experiments to the needs of LDMX and integrate them into a coherent whole, computing support for the physics studies needed to achieve a robust and efficient design, and project planning and management expertise necessary to develop a design report and project execution plan for construction and operation of the experiment. The goal of this DMNI development project is to be ready to begin construction of the LDMX apparatus in FY23 and to be ready for operations in FY25. In accordance with guidance from OHEP, we are planning towards a Design Review at the conclusion of the DMNI development project to set a baseline and begin the LDMX construction project.

The LDMX DMNI development project has received 70% of the planned funding, with most arriving at the beginning of FY22. Some tasks scheduled to begin in FY21 have been delayed as a result, while others were completed with redirected effort. The remaining work planned under the original budget has been carefully de-scoped to fit the provided funds so as to protect the goal of having a design and project execution plan ready for review by the end of FY22. The few tasks that have been pushed into the construction project plan do not introduce significant cost, schedule, or performance risks to the project.

2 Science Goals and Details

The primary goal of LDMX is a missing-momentum search for dark matter (PRD1, Thrust 1). Beyond this objective, LDMX is a multi-purpose forward experiment that can search for unstable dark-sector particles (PRD2, Thrust 2) such as axions, can make electronuclear measurements for neutrino physics, and measure or constrain rare decays of Standard Model mesons. This section summarizes the motivations and LDMX's expected capabilities in each of these domains.

Dark Matter Production: (PRD 1, Thrust 1)

The exemplar of the "compelling explanations for the origin of dark matter" emphasized in the BRN report is the idea that dark matter arose as a thermal relic from the hot early Universe. This paradigm is viable over the MeV to TeV mass range and requires a small non-gravitational interaction between dark and familiar matter. Any such interaction implies a DM production mechanism in accelerator-based experiments; in most sub-GeV realizations, electron-DM couplings are key to the thermal DM origin, and so measurements of these couplings are a priority. LDMX's missing momentum measurement directly explores this coupling, while also greatly expanding sensitivity to hadron-DM couplings via missing-energy signals of invisible decays of Standard Model mesons into dark matter [2].

Scalar, Majorana, or Pseudo-Dirac particle DM can be thermally produced through contact interactions with Standard Model leptons f (for example, $\frac{1}{\Lambda^2}\bar{\chi}\sigma^\mu\chi\bar{f}\sigma_\mu f$ for the Majorana fermion χ). All three scenarios are consistent with CMB bounds on DM annihilation [3]; the fermion models in particular are compatible with a small DM mass (i.e. technically natural) and are poorly constrained by existing terrestrial experiments. Thermal freeze-out predicts the interaction scale Λ for a given DM mass m_χ, or equivalently $y \approx 0.9 m_\chi^4/\Lambda^4$ shown in Figure 2. These predicted couplings define an important sensitivity milestone [4, 5, 1]. Most of their parameter space falls within a factor of 10 to 1000 of the interaction strengths that have been explored to date by both beam-dump experiments [6, 7] and fixed-target missing-energy searches [8]. Accelerator experiments are often capable of resolving the mediator particle responsible for the interaction, in addition to the dark matter, making the contact operator description incomplete. For this reason, we quote results (in Figures 2 and 4) using the canonical model where the interactions arise from a dark photon of mass $m_{A'}$, coupling ϵe to SM charged matter, and coupling $g_D = \sqrt{4\pi\alpha_D}$ to dark matter. In this case, the parameter predicted by thermal freeze-out abundance is typically $y = \epsilon^2 \alpha_D (m_{DM}/m_{A'})^4$, and this is commonly used to quote experimental sensitivity. The top-right panel of Fig. 2 illustrates the modest dependence of LDMX's sensitivity on the dark photon mass for fixed DM mass — a theme more thoroughly explored in [9].

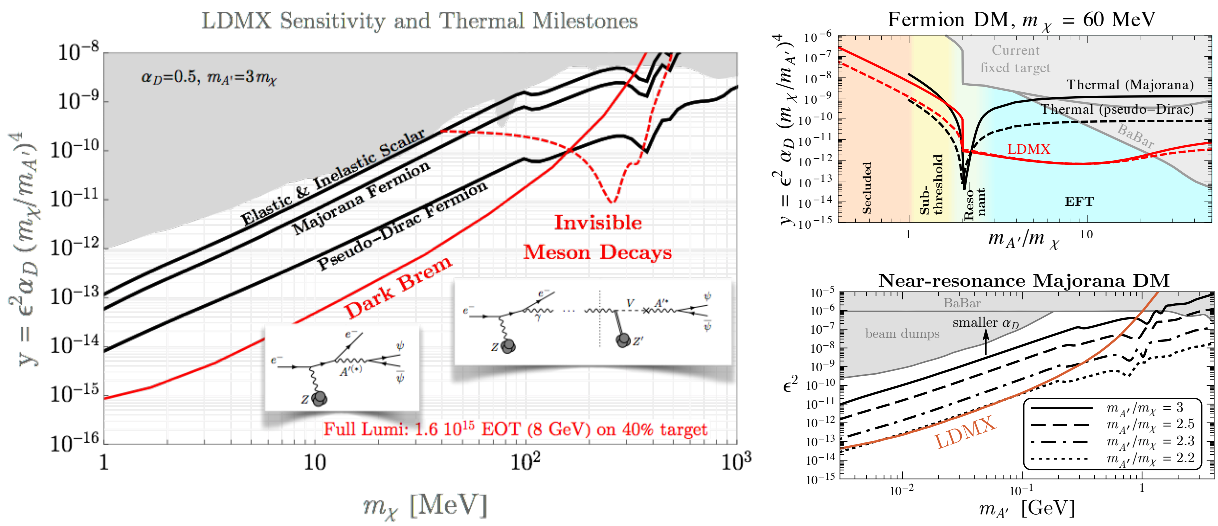

Figure 2: Left: Thermal dark matter milestones (black curves), present constraints (gray), and LDMX projected reach, in the conventions of [5]. LDMX improves over existing DM production searches by three orders of magnitude or more, which is required to robustly explore the thermal freeze-out scenarios highlighted by PRD 1. The insets show two DM production reaction observable at LDMX: dark bremsstrahlung (left inset, with sensitivity indicated by the solid red curve) and exclusive standard model meson photo-production followed by rare decay to DM (right inset, with sensitivity indicated by the dashed red curve). Right top: milestones, LDMX reach (dark bremsstrahlung only), and constraints (as of 2019) for a fixed DM mass of 60 MeV, as a function of the mediator-to-DM mass ratio, for Majorana (solid) and pseudo-Dirac (dashed) DM. Colored bands indicate different qualitative domains of DM annihilation, discussed in the text (see also [9]). Right bottom (adapted from [10]): milestones in the resonant region, plotted as in [11] for up to 10% mass tunings, LDMX dark-bremsstrahlung sensitivity and constraints (as of 2019).

The three panels of Figure 2 illustrate the power of LDMX to explore the commonly discussed thermal freeze-out scenarios, including challenging milestones such as Pseudo-Dirac dark matter [4] over a wide range of mediator masses, and thermal freeze-out with a mediator that is near-resonance [11] or below the DM pair threshold [4, 12]. Furthermore, by exploring deep in the coupling parameter space LDMX probes models such as secluded annihilation [13] of light DM into scalar mediators, motivated parameter space for DM-electron couplings through other types of mediators [12], and SIMP and ELDER models [14, 15, 16] where DM interactions with ordinary matter maintain kinetic equilibrium while DM self-interactions deplete its abundance. Recent work [2] has highlighted exclusive production of SM vector mesons, whose subsequent decay to DM leads to a missing energy/momentum signal, as a highly sensitive probe of *hadronically* coupled DM achievable at LDMX as illustrated by the red dashed line in Fig. 2(left).

Broad Dark Sector Sensitivity (PRD 1, Thrust 2)

In addition to the above-mentioned dark matter searches, LDMX makes several notable contributions to searches for unstable dark sector particles beyond DM (Thrust 2). Some of these scenarios — such as sub-MeV axions, millicharged particles, and $B - L$ gauge bosons decaying to neutrinos — are tested directly by LDMX's standard missing-momentum analysis. For example, even early missing momentum data will test much of the remaining sub-MeV "cosmological triangle" parameter space for QCD axions coupled to either photons or electrons, as considered in e.g. [17, 18, 19]. LDMX will also improve constraints on invisible ϕ and ω meson decay by 4–5 orders of magnitude [2]. These examples are both shown in Figure 3. Likewise, LDMX's excellent sensitivity to millicharged particle production has substantial overlap with the parameter

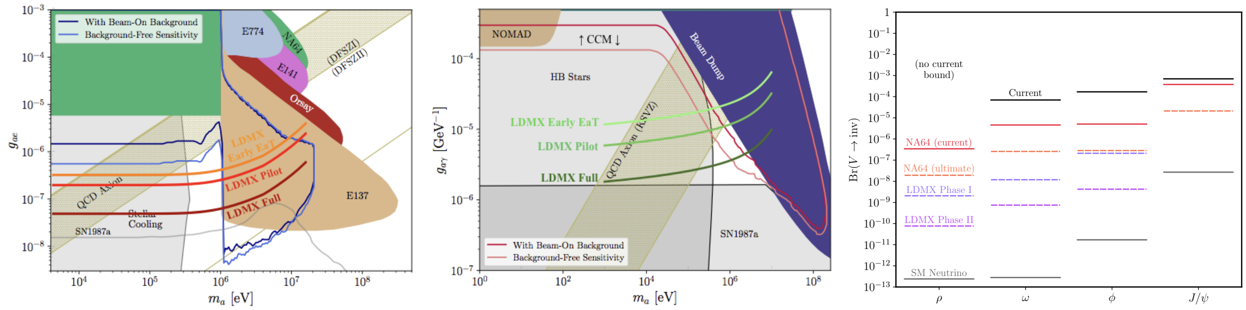

Figure 3: Three recently identified examples of the broader physics potential of LDMX (see [12] for more): **Left:** Projected sensitivity of LDMX missing-momentum analysis to QCD axions and axion-like particles via electron coupling, overlaid on figure from [19]. The three phases of LDMX running from Fig. 4(left) are shown in shades of red. Even early running EaT data could close the remaining sub-MeV QCD axion parameter space. **Center:** Projected sensitivity of LDMX missing-momentum analysis to QCD axions and axion-like particles via photon coupling, overlaid on figure from [19]. The three phases of LDMX running from Fig. 4(left) are shown in shades of green. **Right:** Figure from [2] illustrating the sensitivity of the LDMX missing-momentum analysis to fully-invisible decays of various vector mesons, relative to current bounds (black) curves, NA64 capabilities, and the expected neutrino-decay signal in the Standard Model.

space motivated by the EDGES anomaly [12], exceeds existing limits by up to an order of magnitude, and surpasses the sensitivity of dedicated millicharge detector proposals in the < 100 MeV mass range [20].

In addition, distinct analyses can search for long-lived dark sector particles decaying deep in the LDMX calorimeters [12]. Such searches are sensitive to axion-like particles and dark photons as well as large-splitting inelastic dark matter and SIMP models. The collaboration is investigating how to trigger on these late-decaying particles to maximize LDMX's sensitivity.

Intensity-Frontier Synergy: Electronuclear Measurements for Neutrino Program

Beyond dark sector physics, LDMX can make powerful measurements [21] of electron-nuclear scattering, which address key systematics for DUNE and other neutrino oscillation experiments (see [22, 23]). LDMX complements other experimental efforts in this direction (mainly at JLab [24, 25, 26, 27, 28, 29]) that can be used to improve neutrino scattering models in generators such as GENIE and GiBUU. These generators have been found to differ from *inclusive* electron-scattering cross-sections by up to $\mathcal{O}(50\%)$ [23]. Compared to the experiments and proposals above, LDMX is unique for its broad (nearly 40 degree) forward coverage, low reconstruction energy threshold in various hadronic final states, and ability to detect neutrons with high efficiency. These will allow LDMX to perform *semi-exclusive* measurements of nuclear multiplicity and kinematics within its phase space in addition to electron kinematics [21]. Moreover, due to LDMX's use of a 4–8 GeV beam, LDMX data used to search for dark matter will be taken in a range of momentum and energy transfer that closely overlaps the region most relevant to DUNE [21]. For these reasons, the LDMX collaboration has a dedicated effort to include an electronuclear trigger for data taking and is engaged in ongoing studies with neutrino physicists to refine our understanding of LDMX capabilities in this area.

2.1 LDMX On the World Stage

2.1.1 Accelerator-Based Experiments

The expected sensitivity of LDMX, compared with other accelerator-based experiments (completed, ongoing, and proposed), is illustrated in Figure 4. At low masses, LDMX is uniquely capable of 1000-fold improvements in sensitivity, with sensitivity unrivaled by other experiments. The most relevant accelerator searches to compare with LDMX include:

- *Collider missing-mass searches (Belle-II)* are most relevant to thermal DM above ~ 100 MeV produced through an on-shell mediator, complementing LDMX's sensitivity to lower-mass DM and production through off-shell mediators. A Belle II 20 fb^{-1} study [30] projects a factor of 10 sensitivity improvement over BABAR's [31] missing mass search, constraining $y \gtrsim 10^{-9}$. Neglecting background and systematic uncertainties, which dominate in the sub-GeV mass range and are known to be the limiting factor, studies suggest up to 10-100 further improvement might be possible in the upper end of the GeV mass range with the full Belle II dataset by the end of the 2020's.

- *Beam dump based searches (e.g. CCM [32], MiniBooNE-DM[7], and the COHERENT detectors at ORNL [33])* are limited by signal rate – their sensitivity scales with 4 powers of the small interaction coupling, compared to only 2 powers for missing mass/energy/momentum searches. Therefore, even with substantially increased current or geometric/kinematic acceptance, state-of-the-art proposals [34, 35, 36, 37, 38, 39, 40, 5, 33] achieve only 10-fold sensitivity increases (for clarity, only CCM and the COHERENT detectors planned for deployment in the mid-2020s are shown in Fig. 4; others are comparable). These generally test baryonic couplings, while thermal freeze-out of DM below $\sim$ 100 MeV relies on electron couplings; these couplings are of similar strength in the hidden-photon model assumed in the comparison.

- *Missing energy searches (NA64)* also compete with DMNI projects such as LDMX and CCM. The only such experiment is CERN's NA64 [41, 8], situated in the H4 beamline [42], which delivers a secondary electron beam to the front of NA64's calorimeters. NA64's published results use a sample of $2.8 \cdot 10^{11}$ EOT with a 0.5 ± 0.2 event background estimate. An August 2021 run aimed to accumulate $5 \cdot 10^{11}$ EOT, which could roughly triple their sensitivity when combined with past data [43], if backgrounds remain small with similar signal efficiency. Long-term NA64 projections reflect a goal of integrating $3 \cdot 10^{12}$ EOT before LS3 (2025), close to their setup's irreducible neutrino-background floor [44]; this projection, shown as a thin dotted line in Fig. 4(left), relies on a 10-fold improvement in background rejection over [8]. The impact of the CERN Council decisions, following the Russian invasion of Ukraine, on these operations are unclear as NA64 is a Russian-led experiment. Beyond 2025, no additional electron data is anticipated for NA64, and instead that collaboration is discussing a muon beam data taking period extending into the early 2030's after LS3, though no beam time has been approved. A muon-beam configuration could offer similar sensitivity for dark-photon models, and is interesting for testing explanations of the $(g-2)_\mu$ anomaly, but would not directly probe the electron-DM coupling relevant to thermal dark matter below O(100) MeV. Nonetheless, we show a zero background sensitivity projection (in the dark photon mediator model) for NA64 with the optimistic assumption of 2×10^{13} muons on target, which might be possible by the early 2030's. If these goals are realized before completion of LDMX, NA64 could probe much of the scalar and Majorana benchmark models outside of the resonance region. LDMX would still explore new, well-motivated parameter space including the resonance regions of these models and the pseudo-Dirac benchmark. *Moreover, a few-week LDMX pilot run could compete with NA64's most aggressive sensitivity projections, motivating accelerated preparations.*

2.1.2 Non-Accelerator-Based Experiment

Recent advances in low-threshold direct detection in both semiconductor and noble-liquid detectors [52, 53, 54, 55] offer a complementary window on sub-GeV DM. Broad comparisons are difficult to make because accelerators probe interactions of semi-relativistic dark matter while direct detection involves scattering with much lower momentum transfer. *In general, thermal DM predictions for accelerators have mild dependence on DM spin because both early-universe thermal DM production and accelerator production probe similar momentum scales. By contrast, direct detection probes very different kinematics and so thermal DM predictions for scattering cross-section span 20 orders of magnitude (with only the best-case scenario appearing on the scale of most projections).* The full range of predictions is shown in Fig. 4(right), along with model-dependent mappings of LDMX's sensitivity to the direct detection parameter space. For specific benchmark models discussed in the BRN Report [1]:

- (elastic scalar dark matter benchmark) is comparably accessible to DMNI sub-GeV direct detection and LDMX, because the dark matter scattering is velocity-independent in such models.

- (Majorana fermion dark matter benchmarks) is accessible to LDMX, but direct detection cross-section is suppressed by CM-frame v^2 (a $10^{-6} - 10^{-10}$ cross-section suppression).

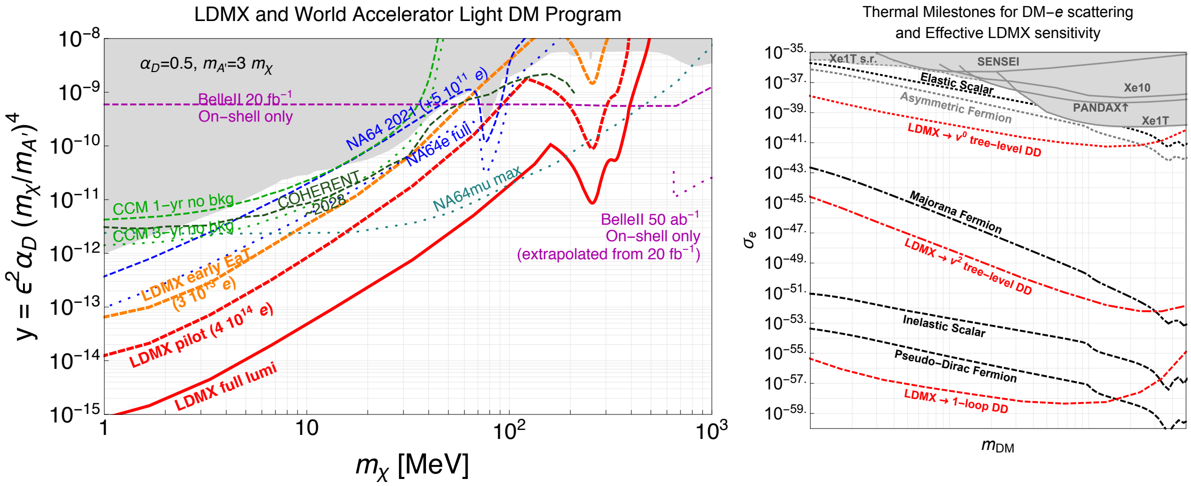

Figure 4: **Left:** The expected sensitivity of LDMX (thick red), a $4 \cdot 10^{14}$ EOT pilot run (thick dashed red), and an early running EaT analysis in $3 \cdot 10^{13}$ EOT (thick dashed orange, see Sec. 2.2) compared with current constraints (shaded) from [45, 31, 46, 8, 6, 47, 33, 48], with expected near-term analyses (dashed) and long-term projections (dotted) from ongoing experiments NA64 (blue) and Belle II (magenta), DMNI project CCM (green), and planned COHERENT detectors at Neutrino Alley (dark green). NA64 projections are scaled from [47] assuming zero background and unchanged acceptance for the maximum 3×10^{12} EOT that can be obtained before LS3. We also include a background free projection for a possible muon beam run currently under discussion at CERN for 2×10^{13} MOT (possible by early 2030's), though this does not probe the electron interactions that LDMX measures; CCM projections are scaled from [32] to 2.8 events in 1 and 3 years respectively; COHERENT projections are from [33] and Belle II curves from the Snowmass whitepaper [49] (the dotted line terminates at masses below which unknown systematics dominate and increasingly degrade sensitivity). **Right:** Milestones for thermal relic freeze-out and asymmetric fermion DM in electron-recoil direct detection, which vary over 20 orders of magnitude depending primarily on whether the structure of DM-SM interactions induces velocity-independent (dotted), v^2-suppressed (solid), or 1-loop (dashed) scattering. LDMX sensitivity can also be mapped onto this parameter space for each scenario (for LDMX, only dark brem sensitivity is shown here). Leading present direct detection constraints are taken from [50] and solar reflection constraint from [51].

- (inelastic scalar or fermion, aka pseudo-Dirac, dark matter benchmarks) is accessible to LDMX, but the leading direct-detection reaction is a one-loop diagram rather than tree-level (suppressing the cross-section by $\sim 10^{-15} - 10^{-20}$).

Thoroughly exploring thermal freeze-out for all DM spins is important, and requires a powerful accelerator-based search such as LDMX. Indeed, fermionic models may even be theoretically favored over the scalar benchmark since they are (at the effective operator level) technically natural.

The broad complementarity between low-threshold direct detection and LDMX extends beyond the thermal freeze-out paradigm. For example, IR freeze-in through an ultra-light mediator [56, 57, 58] has couplings well below LDMX's sensitivity, but can be observable in direct detection because low-momentum scattering is enhanced. By contrast, UV-dominated freeze-in [59, 60, 10] produces observable signals at LDMX without a direct detection signal. **The overarching conclusion from these model-specific comparisons is that the physical parameter spaces for dark matter detection in accelerators and direct detection are fundamentally different and highly complementary.**

2.2 Background Rejection, Early Running, and Mass Reconstruction

Performance and Design Studies To maximize its reach associated with the PRD 1, Thrust 1 physics goal, the LDMX design is driven by the requirement to have a sensitive response to Standard Model reactions associated with a multi-GeV electron interacting and showering on normal detector materials. Figure 5 depicts the types and relative rates of such reactions, and also lists the types of detector sub-systems in our concept with an efficient response. The combined detector response forms the foundation of a background veto used to identify dark matter production events. Whereas potential weak-interaction backgrounds with a recoil track are negligible at LDMX luminosities, and beam backgrounds are rejected to a negligible level by a tagging tracker, rejection of potential photon-induced instrumental backgrounds places demands on the geometric acceptances, depth, and performance of multiple detector sub-systems. These are therefore an important design driver for the experiment. While NA64 has demonstrated background rejection at a level equivalent to $\sim 10^{-12}$ per electron on target (EOT) at LDMX, LDMX aims for $100\times$ greater rejection in a 4 GeV pilot run and 1,000-10,000-fold improvement in rejection in a full-luminosity run at 8 GeV.

With these requirements in mind, LDMX has undertaken a systematic background study [61] corresponding to a pilot-run luminosity of $2 \cdot 10^{14}$ EOT at 4 GeV. Anticipating few-body photonuclear reactions and metastable charged particles as drivers of LDMX's ultimate performance, the collaboration has closely studied the modeling of these reactions in our Geant 10.2.3 Monte Carlo (MC) and implemented bug fixes in the Bertini Cascade model (subsequently incorporated in Geant4 10.5) as well as exact tree-level matrix elements for muon conversion.

The results of this study show explicitly that pilot-run backgrounds can be robustly rejected by a combination of energy deposition and pattern recognition vetoes in each sub-detector — by design, the rejection *does not use recoil electron p_T*, thereby providing a comforting margin of safety. This strategy allows LDMX to use recoil p_T as an unbiased additional handle for signal confirmation (an unexpected background excess would have lower recoil p_T than a signal) and for mass measurement as discussed below — capabilities that missing-energy experiments like NA64 cannot match.

The study also confirms our conceptual understanding of the background: the non-Gaussian tails in ECal energy deposition are dominated by rare photon-induced reactions (as illustrated in Figure 1 (right)), with processes of particular interest highlighted in Figure 5). Most such events are still rejected simultaneously by two subsystems, providing redundancy in the veto and thus the possibility to define control regions in data. More specifically, multi-body photonuclear final states are generally rejectable by **both** unusual energy deposition patterns in the ECal (e.g. MIP-like tracks, isolated hits in deep layers) *and* energy deposition in the HCal. Events with two-body final states in asymmetric kinematics, where a single hadron or muon carries most of the "missing" energy lost by the electron, have been modeled with particular care because their rejection leans strongly on either the ECal or HCal performance. Single neutrons and K_L's are rejected primarily by the HCal, while single muons and charged kaons are detected by their tracks in the ECal, with performance limited by rejection of early decays-in-flight with soft charged decay products. Armed with this understanding, the collaboration has developed higher-statistics samples of kaon decay-in-flight backgrounds, which validate the findings of earlier studies and allow for the ongoing optimization of the ECal veto on short tracks. These studies are currently being extended to a beam energy of 8 GeV as well. Studies are also underway to compare two key aspects of our detector modeling — the distribution of photonuclear final states and hadronic interactions in the HCal — to alternative particle Monte Carlo codes FLUKA, MCNP, and PHITS.

The full LDMX run luminosity will attain 50x more statistics than were used for [61], but at 8 GeV rather than 4 GeV. First studies of the leading photon-induced backgrounds show that the ECal veto *alone* rejects at least $10 - 20\times$ more background at 8 GeV than 4 GeV. HCal-only veto efficiencies also improve, while signal efficiencies are kept at the same level. This improvement is expected a priori because (a) more energetic final-state particles are easier to reject, (b) cross-sections for two-body photonuclear reactions, the most challenging to reject, fall as $1/E^3$, and (c) decay-in-flight backgrounds are further suppressed when the decaying meson is boosted. So in practice, the more difficult background situation that LDMX will encounter is actually for the planned 4 GeV pilot run.

Opportunities for Early Running with ECal as Target (EaT) LDMX will integrate its ultimate luminosity over a few-year period of running. However, competition from NA64 to explore the higher-coupling scalar and Majorana thermal relic models motivates an effort to attain the greatest possible sensitivity from a rapid analysis of early data. To this end, LDMX can exploit signal production in the ECal *in addition to* signal production in the target in an "ECal as Target" (EaT) analysis, increasing the effective luminosity by a factor of ~ 5. At high event statistics, this analysis is subject to greater backgrounds than the standard LDMX analysis, but it is a powerful tool for maximizing sensitivity in early running.

The EaT analysis uses the standard LDMX trigger but a complementary offline event selection, requiring

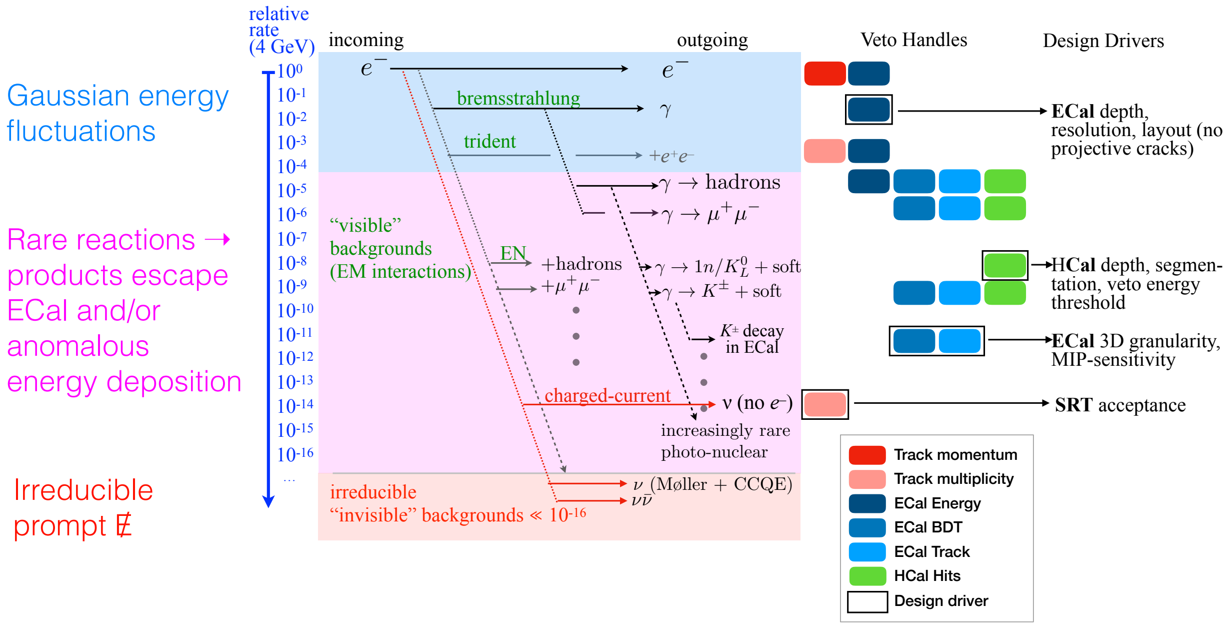

Figure 5: Standard Model reactions associated with a multi-GeV electron interacting and showering on detector materials. The LDMX design is driven by the requirement of providing a sufficiently sensitive response to these reactions to robustly veto them relative to dark matter production, in which the hard bremsstrahlung chain is replaced by dark matter that carries most of the energy of the event. Most reactions listed are redundantly vetoed by multiple detector subsystems. Reaction chains vetoed primarily by one subsystem are highlighted by black rectangles. These have been a focus of LDMX performance and design studies. Detector considerations needed to veto these reactions are highlighted to the right of the figure.

a *beam-energy* track in the recoil tracker rather than a low-energy track. A simplified EaT analysis has been developed for early running, which requires only that the ECal reconstructed energy be < 1 GeV and the summed HCal signal < 10 photoelectrons. For this very simple analysis, $\lesssim 1$ background event is expected in $3 \cdot 10^{13}$ EoT allowing a sensitivity comparable to the best-case accumulated statistics of NA64. The expected sensitivity of this analysis is shown by a dashed orange curve in Fig. 4. With modest optimizations, the EaT analysis may also enhance LDMX's sensitivity in the full $\mathcal{O}(10^{14})$ EOT pilot run.

Mass Measurement from Transverse Momenta The transverse momentum of the recoiling electron in high-missing-energy events is an important feature of LDMX's DM signal. As noted above, LDMX's strategy is to reject events *without* using this feature, reserving transverse momentum as a final cross-check and a measurement handle.

Figure 6 illustrates the power of this measurement assuming an on-shell A' decay to dark matter, with coupling corresponding to the scalar thermal target with $m_\chi = 1/3 m_{A'}$ and $\alpha_D = 0.5$. In a $4 \cdot 10^{14}$ EOT pilot run, 7–60 signal events are expected for A' masses between 3 and 300 MeV. For masses in the 10s of MeV where the signal yields are largest, the p_T distribution allows an A' mass estimate to within 50% or better. This performance is dramatically better than could be obtained using the recoil electron's energy distribution alone (e.g. at NA64 or in an LDMX EaT analysis).

A recent phenomenology paper [62] also suggests some discrimination (at higher statistics) between on- and off-shell mediators and by combining cross-section measurements at 4 and 8 GeV beam energies.

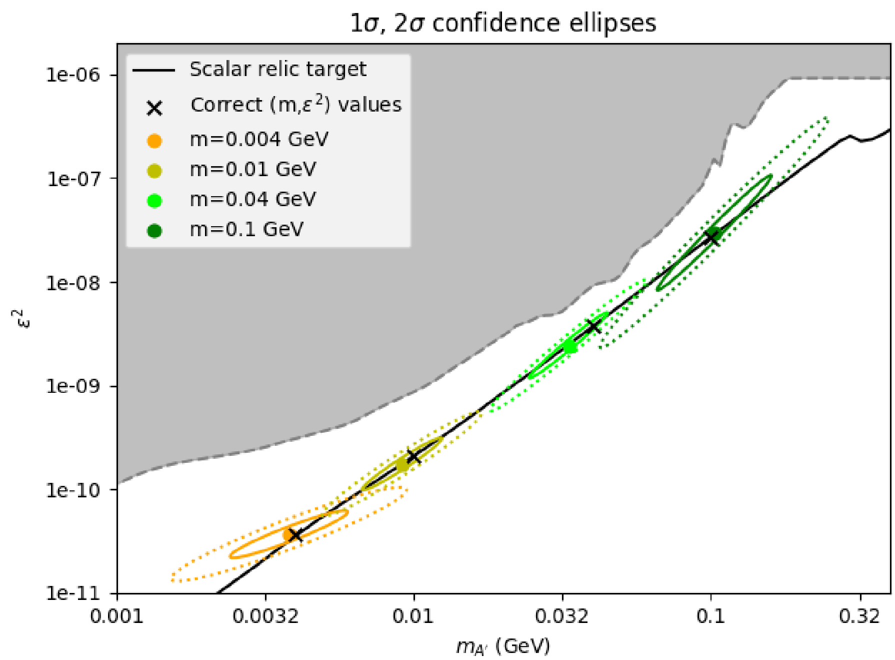

Figure 6: For several choices of A' mass, and couplings corresponding to thermal freeze-out with $m_{A'} = 3m_\chi$ and $\alpha_D = 0.5$, the solid (dashed) error ellipses indicate the 1σ (2σ) resolution on reconstructed $m_{A'}$ and ϵ^2 assuming an on-shell mediator decay interpretation of an LDMX signal in a $4 \cdot 10^{14}$ EOT pilot-run with no expected background. Parameters are inferred from the recoil electron p_T (and, to a lesser extent, energy) distribution and observed event rate. The $\times$'s denote the correct results for each reconstruction test, and the colored dots represent the best fit.

3 Description of Project

3.1 Technical

3.1.1 Description of LDMX Technical Design

The missing momentum signature exploited by LDMX to search for Dark Matter has three components:

1. substantial energy loss by the incoming beam electron, leaving the recoiling electron with a small fraction (e.g. less than 30%) of its initial energy.

2. a large transverse momentum kick of the electron, which together with the degraded energy means the recoiling electron is ejected at a large angle with respect to the incoming beam.

3. the absence of any other visible final-state particles that could carry away the significant energy lost by the electron.

These three observables, and the ability to utilize them at high rates for up to 10^{16} incoming electrons to search for only a few signal events, define the composition and layout of the apparatus.

Taken together, the first two elements of this signature require estimation of the change in vector momentum of individual electrons across a thin $(10 - 40\%\ X_0)$ target, where multiple scattering in the target determines the useful precision. Although the beam energy is known, the beam can be contaminated with off-energy electrons or other particles, so the momentum of each incoming electron must be robustly measured. This can be accomplished with a narrow, low-mass tracker upstream of the target in a magnetic field optimized for measuring beam-energy electrons. The same technology may be used for measuring downstream recoils, but the low energy and wide angles of signal recoils demand wider acceptance in a lower magnetic field. The third element of this signature requires a highly sensitive veto for additional outgoing particles, suggesting hermetic, large-acceptance calorimetry placed directly in the beamline behind the target. Because the vast majority of outgoing particles are scattered electrons or bremsstrahlung photons, the central part of this calorimeter must be optimized for electromagnetic showers (an "ECal"). Furthermore, because the ECal signal rate is of the same order as the repetition rate of the beam, the ECal must be fast and have good spatial and temporal resolution to distinguish energy deposits from different events. Indeed, the ability to resolve the ECal responses to individual electrons sets the overall ceiling on the beam repetition rate ($\lesssim$ 40 MHz) as well as limiting each bunch to a few electrons spatially separated within the beam spot. The LESA beamline (see Sec. 3.4.1) can operate at these limits, allowing LDMX to accumulate 10^{16} electrons on target in a reasonable few-year running period. Meanwhile, the most pernicious potential backgrounds involve a hard bremsstrahlung that carries away most of the electron energy, followed by a highly atypical muon conversion or photo-nuclear reaction that happens to leave little energy in the ECal (see Figure 5). Identification of these events calls for a large and highly sensitive hadronic calorimeter (HCal) surrounding the ECal to veto events with any significant in-time energy deposit.

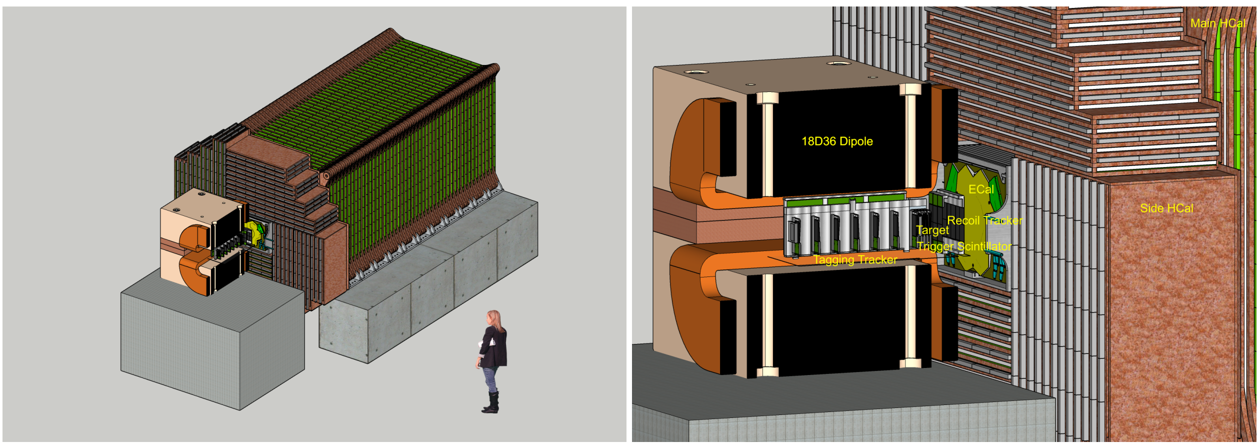

Figure 7: Left: An overview of the LDMX detector solid model. Right: A cutaway overview of the LDMX detector showing, from left to right, the trackers, trigger scintillator and target inside the spectrometer dipole, the ECal, and the Side and Main HCal.

The LDMX apparatus, a compact realization of this concept, has been presented in detail in [63] and is shown in Figure 7. Following the beam, the detector subsystems in the magnet region are a silicon tagging tracker (STT) inside a dipole magnet and a silicon recoil tracker (SRT) in the fringe field of the magnet, with a thin tungsten target interposed between them. Behind the SRT is a compact and highly segmented Si-W electromagnetic calorimeter (ECal) with excellent MIP sensitivity that is surrounded by a large scintillator-based hadronic veto system (HCal) with low energy thresholds. Two important details follow from this picture. First, because the beam passes directly through the trackers and into the ECal, these detectors must contend with high radiation doses to enable an experiment for 10^{16} EOT. To mitigate this issue, and also reduce the peak occupancies in these devices, a large, rectangular beamspot with area on the order of 20 cm^2 is used. As a result, only the ECal has challenging requirements for radiation tolerance. Second, the rates in the detector prohibit streaming readout: a fast trigger is required. Because signal events have unusually large missing energy in the ECal, and such events are very rare, the simplest strategy is to trigger on low energy in the ECal. In order to set an appropriate energy threshold for this trigger, the number of incoming electrons in each beam bunch must be known. This can be easily accomplished with an array of small scintillator bars – a Trigger Scintillator (TS) system – placed in the path of the beam to count the number of incoming electrons in each bunch.

These detector subsystems – Beamline and Magnet, Trigger Scintillator, Tracking, ECal and HCal – along with the trigger and data acquisition electronics (TDAQ) and the software and computing environment required for simulation and analysis of the data define the scope of the technical systems for the experiment that are being developed under the DMNI project in preparation for construction. The following provides an overview of the technical details of these systems.

Beamline and Magnet Design The baseline design for LDMX assumes operation at the proposed Linac to End Station A (LESA) facility being built at SLAC. This user facility will take drive beam for LCLS-II that would otherwise go to a beam dump and divert it to an existing test beam facility in End Station A via a new transfer line, the Sector 30 Transfer Line (S30XL), which is currently under construction as a DOE Accelerator Improvement Project (AIP). Because LCLS-II fills only 1 MHz of the 186 MHz bunches accelerated by the linac and diverts them to undulator lines for producing x-ray pulses, roughly 60% of the duty cycle of the accelerator can be made available to LESA without impacting LCLS-II operations. For LDMX, LESA will produce bunches with one or a few electrons per pulse, and use collimators and quadrupole magnets to produce the large, uniform, rectangular beamspot. This leaves LDMX responsible for a small number of additional beamline elements: a final section of large-diameter beampipe terminated with a thin vacuum window in front of the apparatus, additional vacuum pumping and monitoring for this final beamline section, and basic monitoring to measure the background and radiation environment around the periphery of the beamline where it enters the detector.

The various subsystems of the detector are built in and around the spectrometer magnet for the trackers, a standard 18D36 dipole magnet with a 14" vertical gap already in hand at SLAC. The trackers, target, and trigger scintillator systems and their front-end readout electronics are installed in a support box that is inserted into the magnet bore from the upstream side, while the ECal is mounted on a support frame attached to the back side of the magnet. The magnet is supported on rails, and can be retracted upstream from the HCal for access to the ECal and the front of the HCal.

Tracking Design A tracker design similar to the Silicon Vertex Tracker of the HPS experiment [64] is well suited to the precision, timing, and acceptance requirements for LDMX. Immediately upstream of the target is the 60 cm long STT inside the 1.5 T central field of the magnet, with coverage sufficient to contain the entire beamspot. Seven evenly-spaced, low-mass double-sided layers of silicon microstrips provide the precision momentum measurement for incoming particles needed to reject low-energy beam backgrounds. The 18 cm long SRT is immediately downstream of the target, with much wider coverage to capture low energy recoils at large scattering angles. To optimize tracking for low-momentum recoils (50 MeV to few GeV), the SRT is placed in the fringe field of the magnet and includes four closely spaced stereo modules identical to those in the tagging tracker directly behind the target, and two larger axial-only layers closer to the ECal face to maximize acceptance and minimize material. The sensors are p^+-in-n type silicon with a sense (readout) pitch of 30 (60) μm, providing excellent spatial resolution at high signal-to-noise ratios with 0.3% X_0 per hit. The front end electronics are identical to those used in the HPS SVT, including readout with CMS APV25 ASICs in multi-peak mode, which enables ≈ 2 ns hit time resolution.

The strong resemblance of the LDMX trackers to the recently-built HPS SVT greatly simplifies the task of designing this system and understanding the scope, cost, and schedule for construction. In particular, the sensor modules and front-end readout electronics require only insignificant design changes to be used directly in LDMX. The main differences are in the high level mechanical interfaces with the magnet, which are much simpler than in HPS, and the Trigger Scintillator system, which requires some degree of co-design and a detailed definition of mechanical interfaces between these two systems.

Trigger Scintillator Design The Trigger Scintillator system (TS) consists of arrays of scintillator bars that are used to estimate the number of beam electrons within each beam pulse. One array of 48 bars will be placed near the target to ensure that the electron corresponding to an event has traversed the target. There will be 2 additional arrays upstream of the tagger tracker. Requiring a coincidence of hits in each array mitigates the effects of secondaries produced in the target or the tracker material. Each scintillator bar will be read out with independent Silicon Photomultipliers (SiPMs). Each bar will be 30 mm long with a cross sectional area of 2x3 mm^2, which is well suited to readily available SiPMs. Each array will be arranged in a 2x24 grid with bars stacked horizontally with the two layers separated along the direction of the beam. Layers are vertically offset by half the width of a bar to eliminate projective gaps in the detector.

The SiPM signals will be digitized by deadtimeless readout electronics developed for the phase 1 CMS HCal upgrades. These boards will produce low noise charge integrating amplitude measurements and pulse arrival time measurements with 0.5 ns precision. The front-end electronics will be controlled via a Zynq-based interface that will manage fast signals and parameter configuration of front end electronics and SiPM boards. Data will be continuously streamed to ATCA-based electronics over 5 Ghz fiber optics, where trigger primitives are computed and data pipelines are managed.

In additional to plastic scintillator arrays, we are also developing a design for an active target that would make use of the same readout. The active target will be made with 10% X_0 thick LYSO scintillators that are read out by SiPMs. Such a design is expected to help identify photonuclear reactions in the target by measuring the energy loss of by-products of these reactions. The Active Target design is being evaluated as an alternative to a passive target during the DMNI phase.

ECal Design The need for granularity, very high efficiency, radiation hardness, and speed led to the selection of high granularity silicon technology for the LDMX ECal. To this end, we adopt designs from the CMS High Granularity Calorimeter (HGC) for the CMS phase 2 upgrade [65]. The ECal, shown in black in Figure 7, is a sampling calorimeter with W absorber planes corresponding to 40 radiation lengths, interleaved with 34 Silicon layers, paired into 17 doublelayers. Motherboards carrying trigger and data signals overlay and connect the 7 modules per layer. The motherboards host data-processing and control mezzanines based on the radiation-tolerant Polarfire FPGA. Control and readout is carried on optical fibers to electronics hosted in the experiment's common ATCA DAQ crate.

The total number of modules in LDMX correspond to less than 1% of the number to be built for CMS. We will use the CMS high density (HD) HGC module design, but with thicker silicon (300to improve signal given the lower radiation environment at LDMX than CMS. The thicker sensor material is used in CMS for low density sensors and therefore the sensors simply involve using the HD mask on sensors of an existing design. The hexagonal HD sensor maximizes space available on an 8" wafer with a flat-to-flat size of ∼17 cm and is divided into 432 individual hexagonal readout pads, each with an area of 0.56 cm^2, fitting in a circle of radius ∼0.47 cm. The ECal and trigger of LDMX rely on the CMS HGCROC ASIC developed by the LLR Omega group.

An ECal based on silicon pad sensors is well-suited to identify photons and electrons with high efficiency and good energy resolution, as well as rejection of photonuclear backgrounds with modest or even very small energy depositions using shower shapes and particle tracking. The ECal energy resolution has a very small constant term and stochastic term of $\sim 20\%/\sqrt{E}$. The Moliére radius is ∼2.5-3.0 cm. The radius of containment of 68% of the energy in EM showers is less than 1 cm in the first 15 layers of the ECal, enabling discrimination of individual electrons and photons at small angular separations. It also provides efficient detection of charged hadrons that range out in a single silicon layer and tracks those that traverse multiple planes with excellent per-cell efficiency for minimum ionizing particles. The large ECal depth — $40X_0$ of W absorber — is driven by ensuring that late developing EM showers and photonuclear reactions still deposit sufficient energy in the ECal to be detected, down to the 10^{-16} level. A large depth also improves detection of muons and charged hadrons.

Each ECal doublelayer has a core cooling plane covered on both sides by W plates, followed by modules and motherboards. An additional W plane is added to one side of the doublelayer to provide the absorber layer between it and its nearest neighbor. The 7 modules per layer are arrayed in a flower configuration with a central module surrounded by a ring of six others. The sampling planes, (7 modules and 2 motherboards) are ∼6 mm thick, including the module baseplates made from a thin C-Fiber layer or simple printed circuit board (PCB) with integrated Cu shielding. The cooling layer have ∼4 mm O.D. thin-wall stainless steel cooling tubes embedded in, and covered by, C-Fiber sheets. The doublelayer thickness excluding W is ∼2.3 cm. The depth of the entire device, including W and gaps between doublelayers is ∼55 cm. The detection volume is thus about 55x 55x 55 cm^3, while the full system occupies a volume of 95x 65 cm^2 and depth of 65 cm. It is small and compact, but dense, with a mass of ∼825 kg. The support structure holds this large mass while precisely positioning doublelayers and absorber planes.

The ECal is designed to be disassembled to replace problematic components or to reconfigure plane positions should that be later deemed advantageous for rare backgrounds, systematic uncertainties, or signal sensitivity. Small transverse offsets between sensing layers will improve resolution of charged particle tracks and avoid overlap of small dead regions between modules. The designs will include manifolds for distribution of coolant and dry nitrogen. Solutions have been identified for low-voltage power using radiation-tolerant DC/DC converters and for the necessary feed-throughs and cabling harnesses.

HCal Design The HCal is a scintillator-based sampling calorimeter comprising a large number of nuclear interaction lengths of steel absorber. The main function of the HCal is, in combination with the ECal, to provide a high efficiency veto of hadrons produced in the electron-target interactions. The design is driven by the need to efficiently detect neutral hadrons - mostly neutrons produced in the target or ECal - in the energy range from hundreds of MeV to several GeV. The HCal must also measure the component of electromagnetic showers escaping the ECal, and be sensitive to MIPs, such as muons. Good efficiency for lower energy neutrons requires sampling thickness of the absorber plates less than 30% of a strong interaction length (λ_A), while ensuring that high-energy forward neutrons interact in the detector to the desired level requires a total depth of approximately 15 λ_A.

The calorimeter consists of two major sections: the Main HCal located behind the ECal, and the Side HCal, a smaller device surrounding the ECal. The Main HCal contains 96 layers of 25 mm absorber plates. GEANT simulations, including several key background processes and single particle response, are underway to optimize the dimensions and segmentation of the device. We plan to finalize the design at the end of Summer 2022. Scintillator bars similar to those used in the Mu2e Cosmic Ray Veto system are deployed in an X,Y configuration, with ambiguity resolution by relative timing at the end of each bar. The doped polystyrene bars are 20 mm thick $\times$ 50 mm wide, co-extruded with an integrated TiO_2 reflector. The extrusion has a through hole into which a 1.8 mm wavelength-shifting fiber is inserted. The scintillator response to minimum ionizing particles has been measured to be around 90 photo-electrons, providing an adequate signal for LDMX. The scintillator bars are assembled into units of four bars ("quad-counter") onto which readout electronics is mounted. The quad-counters are glued to the steel absorber plates, and fully equipped layers are fabricated as 12 modules of 8 layers each. Each module forms a rigid object that will be fastened to the floor. The Side HCal consists of 4 modules arranged in a pinwheel-like fashion around the ECal, approximating the shape of a disk. Each module contains 24 layers of 20 mm absorber with scintillator bars, read out at one end, arranged in a X,Z or Y,Z configuration to provide 3d information. A preliminary earthquake analysis of the HCal has demonstrated that it meets SLAC standards. A 384 channel prototype of this design, described in 3.1.2, has been built and operated in a CERN test beam.

The readout electronics is adapted from the Mu2e CRV system and the HL-LHC upgrade of the CMS endcap calorimeter. Each quad-counter fiber is read out at each bar end by a SiPM mounted on a Counter Mother Board (CMB) that provides bias to the SiPMs, a temperature monitor, and flasher LEDs to calibrate each bar independently. Given the mechanical constraints, the Side HCal bars are only read out at one end. The four SiPM signals are transmitted to a High Granularity Calorimeter Read Out Chip (HGCROC) board via an HDMI cable. A single HGCROC Board is designed to operate and read out the signals from 16 CMBs. Four HGCROC boards are housed on a large back plane board, together with a Mezzanine card providing the DAQ logic and the initial trigger calculations.

Trigger and Data Acquisition Design The trigger, data acquisition, and slow control (TDAQ) system of LDMX consists of a custom electronics hardware trigger that identifies events of interest, reads out data from the thousands of LDMX channels to be saved for offline analysis with off-the-shelf electronics, and provides the communication mechanism to control and monitor the detector. The main technological requirements of the TDAQ system are:

- 25 kHz DAQ rate is based on the tracker readout ASIC bandwidth including a safety factor of 2 (max 50 kHz); then we set a 10 kHz average trigger rate requirement which includes a 2.5 safety factor from the DAQ rate
- 10 Gb/s DAQ data rate based on modest modern technological capabilities
- 3.5 μs latency to deliver the trigger decision to the detector front-end controllers, based on the pipeline depth of the tracker readout ASIC; consequently we require $2 \mu s$ latency to formulate the trigger decision

The TDAQ system relies on existing electronics hardware platforms in order to reduce risk and resources. For the data acquisition path, we choose an off-the-shelf PCIe-based FPGA system being developed by the SLAC electronics group based on the Bittware XUP-VV8 card [66]. It is a modern and flexible DAQ system with processing power meeting the LDMX DAQ requirements and a 32 optical links per card. The Bittware XUP-VV8 system is a good match to perform readout, zero suppression, and control for the tracker, ECal, and HCal.

The trigger of LDMX should read out trigger primitives from the ECal, HCal, and trigger scintillator systems at full rate and reduce the event rate to 10 kHz using information from those subsystems. The primary task of the trigger system is to save events for which one of the incoming electrons lost a significant fraction of its energy. However, a wide range of secondary triggers are expected in order to collect background sideband data, calibration data, and physics data for other dark matter and nuclear measurement final states. Because the trigger system must perform a wide range of tasks under challenging latency constraints, we require a different board for triggering with a more powerful processing FPGA and more optical links per FPGA in order to aggregate all detector signals in a single processing node. For this task, we choose the ATCA-based APx (Advanced Processing) board developed for the CMS L1 Trigger upgrade.

Software and Computing Design The generation of MC and reconstruction of all data will make use of `ldmx-sw`. `ldmx-sw` is a `C++` event processing and simulation framework developed by the LDMX collaboration that implements a software bus model to facilitate communication between data processing modules. At the heart of `ldmx-sw` is the "Framework"; a library that builds the processing pipeline, manages the configuration of the modules and provides the event bus used to pass data between them. An embedded python interpreter allows for the configuration of the pipeline at runtime by using basic python commands. These features make the Framework very lightweight as only necessary modules are loaded dynamically. As discussed further in Sec. 3.1.2, the Geant4 based simulation and subsystem reconstruction are all implemented as data processing modules and used to build the data processing pipeline.

Although LDMX is a small experiment, the high statistics nature of the dark matter search dictates relatively large volumes of experimental data along with commensurately large samples of Monte Carlo (MC) simulated data. As a result, the development of a robust computing model for the experiment is fundamental to the design. The design drivers for LDMX computing resources include the storage and processing of raw data for the and the production, storage, and processing of Monte Carlo (MC) samples. The scope of the resources required at the outset of experimental operations is defined by the needs of the pilot run, as summarized in Table 1 and described below.

Sample	**Disk Storage (TB)**	**Tape Storage (TB)**	**Processing (CPU-hrs$\times 10^6$)**
Raw data	—	900	—
Reconstructed data	360	5400	13.5
Monte Carlo	—	867	5.5
Total	360	7367	14.6

Table 1: Data storage and processing requirements for the 4×10^{14} EOT pilot run, as described in the text. Data processing requirements will be met with the planned allocation of 1500 CPUs. Staging raw data and Monte Carlo to and from tape will utilize 300 TB of storage purchased with DMNI funds. These estimates include the full safety factor for DAQ operation at up to 25 kHz.

During the initial pilot run, LDMX will see 4×10^{14}, 4 GeV electrons and collect a total of 2.7×10^{11} events at a maximum trigger rate of 25 KHz. The event size is estimated using full simulation in combination with the design of the readout for each subsystem and is described in Table2. From this, we estimate storage of the unprocessed raw data from the pilot run will require $\sim$ 900 TB of space. Reconstruction of physics objects for analysis adds 10 kB/event to this total, estimated with the current digitization and full reconstruction pipeline used for physics studies with a factor 2 safety margin. Including the raw data as usable analysis objects, the reconstructed data for the pilot run is therefore expected to take $\sim$ 3.6 TB of space. All raw and reconstructed data will be transferred to tape with 10% ($\sim$ 360 TB) being kept on disk for analysis. Furthermore, we anticipate that 10% of data will be processed multiple times as improvements to the reconstruction and calibrations are made. Based on previous experience, we estimate that the 10% data could be reprocessed as many as 5 times. This will add $\sim$ 1.8 TB of tape store to our needs. Taking all of the above into account, processing of the raw data from the pilot run will require 360 TB of disk and $\sim$ 6.3 PB of tape.

In addition to the experimental data, large Monte Carlo samples are also required to compare against data and for performance studies. The size of these samples is minimized by careful use of event biasing and

Subsystem	Size (kB)	Notes
Trigger Scintillator	.1	8B for header + 20B per hit, 6 hits/event
ECal	1.2	10B/hit for TDC hits, 8B/hit for low-amplitude hits, 125 hits/event
HCal	0.4	20B/hit for channel id and data, 20 hits/event.
Tracker	1.0	20B/hit for channel id and data, 50 hits/event.
Trigger	0.5	Event accounting, ECal trigger sums, trigger counter, HCal trigger info.
Total	3.2	

Table 2: Estimated event size for each subsystem for running with an average of one electron per pulse ($\mu = 1$). All ECal hits are assumed to have TDC information.

filtering for the most challenging backgrounds, since generation and simulation of $> 10^{14}$ events is clearly impractical. Table 3 summarizes the requirements for analysis of the pilot run MC data, based on the tools and techniques developed for our physics studies. Finally, before the conclusion of the pilot run, we will need large scale Monte Carlo to develop the 8 GeV analysis at full luminosity. The requires approximately 700 TB of tape storage with more aggressive filtering and biasing of events.

Monte Carlo Sample	Total Size (TB)
1×10^{14} ECal PN sample – un-skimmed	50
4×10^{14} ECal PN sample – skimmed	21
4×10^{14} ECal PN sample multi-electron sample - skimmed	50
4×10^{14} ECal as a target sample	20
4×10^{14} Target photo/electro-nuclear sample	10
4×10^{14} Muon pair conversion sample	5
Inclusive + signal	11
Total	167

Table 3: Monte Carlo samples and expected size needed for the pilot run. The largest samples are those required to study the key photonuclear (PN) backgrounds at high statistics.

Reconstruction of the raw data dominates the processing required. The full reconstruction chain has been profiled and benchmarked at 60 ms/event using digitized Monte Carlo from physics studies on currently available CPU. Adding a factor two safety margin, 9 million CPU hours are required to fully reconstruct the pilot run data. LDMX plans for the allocation of 1500 dedicated CPU cores for LDMX, allowing the full dataset to be reconstructed in ≈ 8 months and a 10% reconstruction pass to be completed in less than a month. Meanwhile, generation, simulation, digitization and reconstruction of the Monte Carlo required for the pilot run will require approximately 1.1 million CPU hours, and can be produced in roughly one month. Monte Carlo samples required to develop the 8 GeV analysis at full luminosity will require at least 4.4 million CPU hours, even with more aggressive filtering and biasing of events. We plan to produce the samples required for these studies during the construction of the apparatus when at least 20 million CPU hours will be available if the processing nodes are allocated during the first year of construction.

In order to realize the ultimate sensitivity attainable with the LDMX apparatus, we envision a total exposure of 1.6×10^{15} electrons on a 40% X_0 Al target at 8 GeV beam energy. Naively, the requirements for data storage and processing scale roughly with EOT (4× the pilot run) and for Monte Carlo requirements like the critical backgrounds ($\sim 25\times$ the pilot run). We anticipate that stronger filtering of events may reduce the storage and processing requirements to fit within the resources required for the pilot run, where ongoing physics studies of the full luminosity analysis at 8 GeV will clarify the need for additional computing resources. In any case, additional computing resources that could be required for this second phase of operations will not be needed until after the pilot run, so we do not plan them as part of the construction project.

The data storage and processing needs described above will be hosted the SLAC Shared Scientific Data Facility (SDF). An auxiliary computing resource, the LDMX Distributed Computing System, has been developed by the collaboration, and will bolster our ability to support offsite MC production and analysis. These resources are described in Section 3.1.6.

3.1.2 Description and Status of Technical Development Plan

The LDMX DMNI project comprises the tasks required to prepare a baseline design, along with the scope, cost, and schedule for realization of the experiment. The heart of the plan is the development of the subsystem designs discussed in the last section, and their integration into a complete apparatus, to the point where the scope, cost, and schedule can be accurately presented and reviewed. The foundation of this process has been the determination of a few remaining elements of the baseline design. Examples include understanding the services required for installation and operation of the experiment in End Station A, whether an active target is a feasible mitigation for certain backgrounds, what silicon thicknesses in the ECal and absorber grading in the HCal are optimal, and how best to coat the Trigger Scintillator bars. The development of these baseline designs is monitored by the LDMX Coordination and Project Management team described in Section 3.3.

In addition to engineering and technical work required to establish the design, there are physics studies required by scientific personnel to optimize the design and ensure that it robustly meets the requirements of the experiment. While the effort for this task is not supported by the DMNI project, it requires computing resources that are part of the project. Finally, in order to prepare for a review of the construction project, project management effort is required to develop the project plan, along with the cost, schedule, milestones, and risk register, and to produce the documentation required to ensure that that project adheres to all standard policies and procedures. The following sections describe the elements of this plan, and their current status. We anticipate that all tasks necessary to produce a design report for review will be completed by the end of FY22.

LDMX Test Beam at CERN A focal point for development activity since the FY21 review has been the operation of Trigger Scintillator and HCal prototypes in the T9 test beam at CERN. In addition to developing and operating fully integrated subsystems, test beam data provides valuable information to tune the Monte Carlo simulation as well as optimize the design of the Trigger Scintillator, HCal, and DAQ. Tuning the simulation is particularly important for the HCal, as its calibration will partially rely on the simulation since it cannot be fully performed *in situ* with an electron beam in the shadow of the ECal.

The test beam was nominally planned for a two week period in October 2021. During that time, we commissioned the DAQ system and collected an extensive data set with the Trigger Scintillator, but due to disruptions of international shipping logistics, the HCal prototype was not delivered in time to be included in this effort. We conducted an additional data taking campaign in April 2022 with the full detector, and recorded tens of millions interactions of hadrons and electrons with incident beam energies ranging from 100 MeV to 8 GeV. We also collected a large sample of muons to calibrate the apparatus: a few events are shown in Fig. 8.

The HCal prototype consists of 19 layers of 25 mm steel absorber, sufficient to fully contain electromagnetic showers and a large fraction of hadronic showers. The scintillator bars are 2 m long whose orientation alternates between vertical and horizontal in every layer; the position of the horizontal layers could be moved to measure the response along the bars. The readout electronics, with a total of 384 channels, are based on the Counter Motherboard, HGCROC board, backplane board, and Polarfire mezzanine prototypes described below in the *Hadronic Calorimeter Development* section. The construction of the prototype was partially funded by our international collaborators, as detailed in section 3.7.

The Trigger Scintillator prototype represents roughly one twelfth of the full system planned for LDMX. A single front end readout board was used to digitize 12 SiPM signals. The scintillator module consisted of a single SiPM board and arrays of scintillators inside of a housing. The module was configurable, enabling us to test a fully plastic geometry and a hybrid plastic/LYSO geometry for the active target option. The front end electronics and the scintillator modules formed a unit that was mounted to a pair of linear stages to allow for fine adjustment of the module with respect to the beam, shown in the left image of Figure 8. Due to the schedule of our DMNI work, the APx DAQ electronics were not used for the test beam; another platform, the CAPTAN+X system, was used. This system was developed by the same group working on LDMX Trigger Scintillator DAQ development and firmware developed for the test beam is expected to be transferable to the APx system.

With the successful completion of the CERN test beam program, we plan to re-assemble the prototypes at SLAC to enable further testing and integration work. After completing the "Stage A" of the Sector 30 Transfer Line (S30XL), the Trigger Scintillator system can be temporarily installed just downstream in the beam switchyard. Once Linac to End Station A (LESA) is complete, the Trigger Scintillator and the HCal can be assembled and operated with test beam in ESA. The opportunity to operate the test beam apparatus with early LESA beam will accelerate the process of installing and commissioning the entire LDMX detector once its construction is complete. The timelines for the completion of beamline development work are discussed in Section 3.4.1.

The ongoing analysis of data collected from the CERN test beam will help develop and optimize HCal reconstruction algorithms, including 3-dimensional hit reconstruction and calibrated energy sums. For both the Trigger Scintillator and the HCal, valuable characterizations of detector noise, response functions to minimum ionizing particles, and realistic pulse shape parameterization will be derived from the data. These characterizations will be used to improve our detector simulations. We expect the bulk of this work to be done by the end of FY22 and publication of our findings to be submitted in FY23.

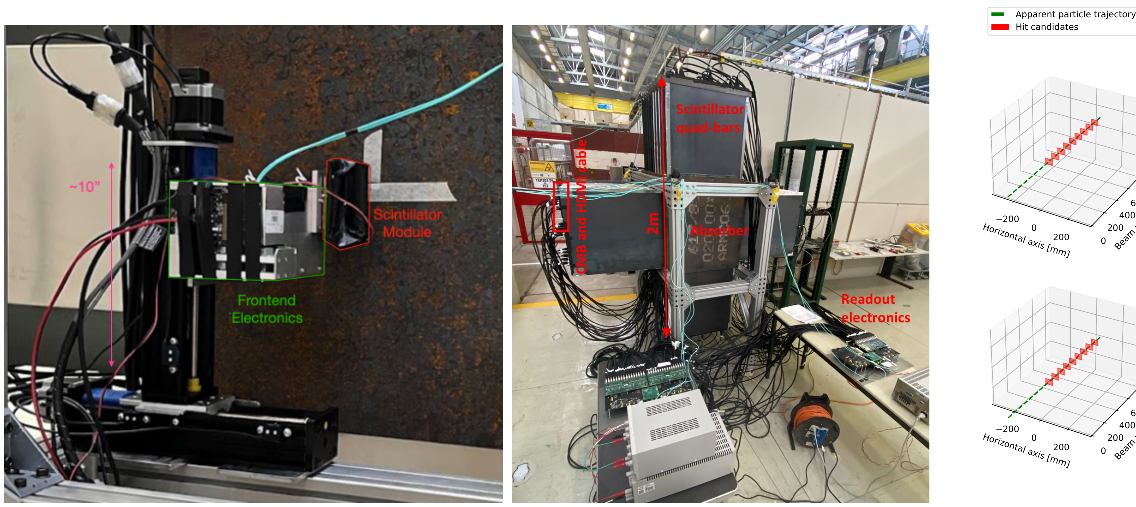

Figure 8: The prototype in the T9 test beam at CERN. **Left:** The prototype Target Scintillator front end electronics and scintillator module mounted to linear stages. **Middle:** The back of the HCal prototype with its r/o-electronics on the two tables. **Right:** Muons crossing the HCal prototype. The red rectangles represent the hits significantly above pedestal. Their location is inferred from the position of a pair of adjacent horizontal and vertical bars. The green line is an estimate of the trajectory of the muons.

Beamline and Magnet Development With most of the beamline already in place or funded as part of the LESA project, the components that LDMX is responsible for are small in scope and need no development work. These consist of a section of large diameter beampipe terminating in a thin vacuum window in front of the apparatus, additional vacuum pumping and monitoring in this region, and simple beam halo counters at the front of the detector to detect any problems with the steering and defocusing of the beam during operations.

Since LDMX is effectively a beam dump, it is important to understand any radiation protection requirements that could be imposed on the detector design. With typical beam currents of 5-30 pA and less than one Watt of beam power – it was not expected that the detector would require any shielding during normal operations. However, one must also consider the potential for brief accidental exposure to a small number (one to a few) of much larger pulses intended for LCLS-II. For the LESA project, these issues had already been studied for 10 W of beam power impinging on a target in the open area of End Station A where LDMX will be located, as shown in Figure 9. The DMNI development project has considered the Radiation Pro-

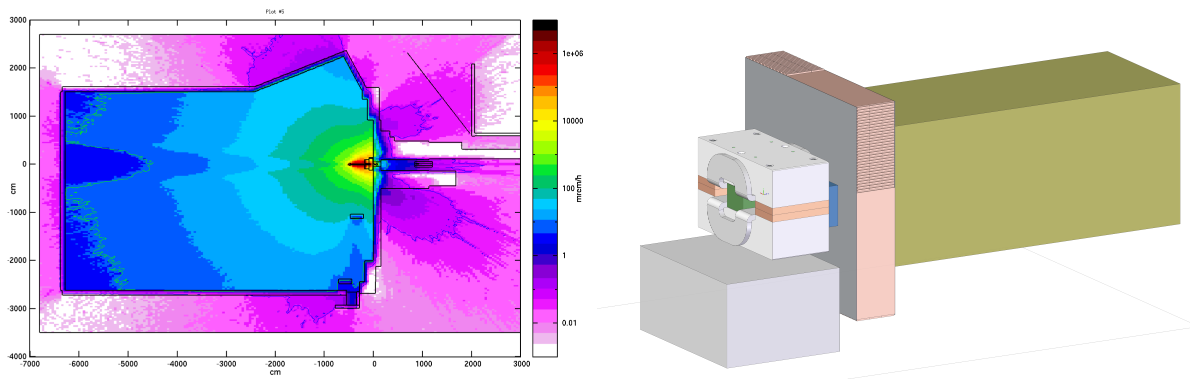

Figure 9: **Left:** The prompt dose rates (mrem/hr) in End Station A for 10 W of 14 GeV beam on a target in the open area of End Station A where LDMX will be located. **Right:** A preliminary CAD model of the detector being used to define stay clears and develop interfaces.

tection (RP) implications of placing the LDMX apparatus in this location, and defined the scope, cost, and schedule for the work that will need to be performed as part of the LDMX construction project. Because the detector will shadow existing RP monitoring devices installed on the back wall of ESA, a beam loss monitor will need to be installed in the region in front of the detector. The design calls for the relocation of an RP system that has been installed in the beam switchyard (BSY) during Stage A of the S30XL project which will be unnecessary when beam is transported all the way to End Station A. The required work has been added to the project, along with the effort required from the SLAC RP department to document and review for approval the RP plan for LDMX. In addition, the RP sections required for the LDMX Design Report are being drafted as part of the DMNI project.

The 18D36 dipole magnet LDMX plans to use, shown in Figure 9 has been in storage at SLAC for many years and requires refurbishment. Extensive test data from 1978 and experience with other magnets of the same design (e.g. used by HPS) establish the suitability of this magnet for LDMX, and an expert assessment of the magnet and the refurbishment task was used to establish a preliminary cost and effort of this undertaking. The DMNI project has funded the effort for a more complete assessment of the scope of refurbishing this magnet, as well as supporting and installing it in ESA, providing it with power and cooling, and defining its interfaces to other detector subsystems, as well as considering various risks inherent in the baseline plan.

The first task, completed early in FY22, has been to produce a 3d solid model of the magnet suitable for defining the interfaces to other subsystems, shown in 9. The development of this model has been a critical element to the technical design work for the support of the ECal, as well as defining the envelope for the tracker, target scintillator, and their associated services. The second task has been to develop a detailed plan for the refurbishment of the magnet, as well as the scope of work required to control, power, and cool it during operations. As expected, this recently completed effort has been necessary to define the full scope of the magnet project, and has identified critical work that must be undertaken on electrical infrastructure at SLAC, as well as various cost and schedule risks that must be considered. This new information has been used to update the budget, schedule, and risks for the Beamline and Magnet portion of the WBS.

Tracker Development While it would be possible to directly reuse nearly every element of the HPS tracker design without changes, the less severe operational environment and space constraints for LDMX allow for reduced complexity and greater redundancy, both of which mitigate risks and improve the robustness of the apparatus. In particular, HPS must operate in a vacuum (and fit in a small vacuum chamber), tolerate much higher rates and radiation doses, and be precisely and remotely repositioned. Defining these simplifications to allow more precise estimates of the scope, cost, and schedule for the LDMX tracker subsystem is the thrust of the DMNI development plan for the tracker. There are two elements of this development plan. First is work to understand the interfaces to the magnet and target scintillator systems to better define the space available for the components of the tracker and the routing of cables and cooling lines. With an understanding of those constraints, the second task is defining the functional arrangement and physical layout of the elements of the tracker readout system relative to the similar components for HPS.

The first task has been completed to the point of enabling the second, where the initial design of the magnet described in the previous subsection has allowed work to begin on the readout design. The remaining mechanical design work to be completed in FY22 is a clearer definition of the outer support box for the tracker and how it will be supported from the magnet, and the interface of the target scintillator and the target itself with the tracker support, for which there is only a conceptual design. These deliverables are expected soon: while not necessary to retire any risks, they will allow better cost estimates.

The second task – recasting the HPS SVT readout system for LDMX – has also made significant progress. There are three main elements of the readout chain for the HPS tracker: the APV25 hybrids that amplify signals from the silicon sensors, the Front End Boards (FEBs) that distribute power, clocking and control, and digitization of APV25 data, and the optoboards that convert LVDS to optical signals for longer-distance transmission to the back end DAQ. In this design, the hybrid and optoboard are functionally very simple, while the FEBs are extremely dense and complicated PCBs. Design work undertaken for the DMNI project focuses on defining the key simplifications relative to HPS. First, the generation of voltages required by the APV25 move from the FEBs to the hybrids, which simplifies the FEBs and reduces the cable plant connecting the FEBs to the hybrids. This change benefits the close spacing of layers in the target region and the integration with the target scintillator system. Second, individual and identical cables connect each hybrid to a FEB channel, allowing FEB channels – with a high degree of redundancy – to be treated as a pooled resource in LDMX. The FEBs can have a much larger footprint and therefore have a much less aggressive layout than the 20-layer HPS design, reducing cost and increasing yield and reliability. In addition, the FEBs can reside outside the magnet bore between the coils and be air-cooled, which makes the layout task much simpler than the FEBs for HPS, which must accommodate water cooling directly to the heat generating components on the board. The optoboards must still reside outside of the magnetic field, supported upstream of the FEBs, but they can have a much simpler layout and be much more easily serviceable and replaceable than in HPS, where they are potted into a vacuum flange to serve as the vacuum penetration for the data path. Similarly, the vacuum penetration for power in HPS can be replaced by much simpler and more standard cables and connections. Key design work for the LDMX tracker readout system is part of the DMNI project and culminates in the production of FEB prototypes as well as a design for the hybrids that includes APV25 power generation. This work is underway, but the production of prototypes may not be deliverable before the end of FY22 due to component shortages for the FEB, and FPGAs in particular. As with the mechanical work, this deliverable is not necessary to retire any significant risks, but will allow for a more precise cost and schedule estimate for the construction project.

Trigger Scintillator Development The trigger scintillator (TS) development work has three main elements: designing and prototyping scintillator modules (both plastic modules for electron counting and LYSO modules for a possible active target), designing and prototyping the front end control system, and finalizing designs for the front end readout electronics. This work will fully validate the TS design from active components through the readout chain.

For the scintillator module development work, we have procured and tested scintillators for plastic and active target designs, and both were tested with radioactive sources and cosmic rays. We have designed two SiPM boards, one to work with the plastic module geometry and the other for the LYSO module geometry. The plastic SiPM board has gone through a few iterations in FY22 to improve critical features. The mechanics that host 2x24 arrays of plastic scintillators were also improved throughout FY22 to make assembly easier. Prototype SiPM boards and a mechanical housing for the active target were developed and fabricated. Arrays of plastic and LYSO scintillator were characterized in recent test beam campaigns.

A prototype control system was fabricated in FY22. Based on experience from these prototypes, an alternative solution was identified that better leverages commercial components. We tested this alternative solution and adopted it as the new baseline design. Prototypes for this design have been fabricated and boards have been individually tested. Testing the interface between our new control system design and front end electronics (FEE) will be completed in FY22. In addition, software for monitoring and configuring the front end electronics and other peripheral elements of the control system have been developed and tested.

The design work for the FEE involved integrating the SiPM boards, the control system, and the data acquisition system with FEE boards. For current prototyping work, we have utilized the readily available

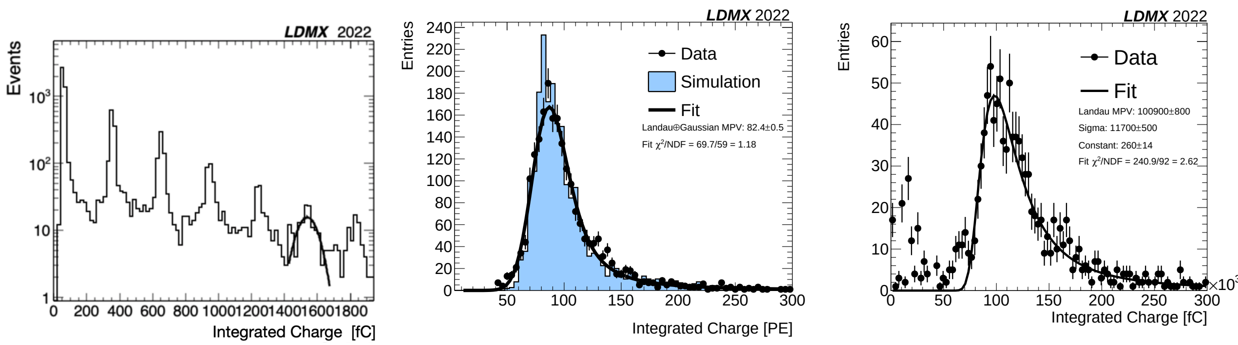

Figure 10: **Left:** Integrated charge distribution of digitized signals from a single SiPM in the Trigger Scintillator prototype. The peaks represent integer numbers of pixels firing and can be used to calibrate SiPM gains in-situ. **Middle and Right:** Plastic and LYSO scintillator responses to minimum ionizing particles are shown in the middle and right panels, respectively. Data was collected with the Trigger Scintillator prototype during the 2022 CERN test beam campaign. The blue filled histogram (black points) in the middle plot represents simulated (beam) data. The black points in the right plot is beam data; the black curve is a fitted Landau distribution. The most probable value for a MIP in the plastic (LYSO) is 82.4 (336) photoelectrons.

CAPTAN+X platform developed at FNAL as a data acquisition board. Firmware for capturing data from FEE boards over fiber optics and transmitting them to a DAQ computer was developed and used during both test beam campaigns. The final DAQ design will involve an APx board, which can receive all 24 fibers from the trigger scintillator FEE. A demonstration of the DAQ interface between the APx and the TS FEE is expected to be completed in FY22. The interface between the SiPM boards and the FEE electronics has been tested. In addition, the interface between the FEE and the control system has been partially tested. The control system involves several boards that enable signals to be appropriately distributed. These boards include the central Zynq-based clock and control module (zCCM), 24 individual FEE boards, and 12 different backplanes. Communication through the backplanes has been thoroughly tested. The complete integration of the zCCM with the FEE will be tested by the end of FY22. Finally, the development of new firmware for the FEE serializer has been deferred to the project due to reduced funding of the DMNI work, but is not required to retire any risks.

We tested several critical designs and interfaces between components during the recent CERN test beam campaigns. Using our prototype system, Figure 10 (left) shows strong separation between integer numbers of pixels firing for an individual channel, enabling in-situ calibration of the SiPM gains. The middle and right plots of Figure 10 show the response to minimum ionizing particles in plastic and LYSO, respectively. These data suggest that our plastic (LYSO) scintillator response is 88 (336) photoelectrons.

Finally, significant progress has been made to refine our studies of trigger algorithms in simulation. These studies have been improved through technical developments that include better signal modeling and the capabilities to study out-of-time pileup. In addition, large samples of 4 GeV electrons were recorded during the test beam. These samples are currently being used better to understand our modeling of the detector in simulation to help ensure that key performance metrics are achievable.

Electromagnetic Calorimeter Development Development work on the ECal subsystem focuses on the adaptation of HGC designs for the CMS phase 2 upgrade to the needs of LDMX, where the UCSB group expects to be building High Density (HD) modules for the HL-LHC CMS upgrade project that would be appropriate for LDMX by the end of 2022 and to have enough modules for a doublelayer in Winter of 2023. LDMX calls for different support and cooling for the detector planes, as well as a different layout for the readout motherboard and adaptation to the LDMX trigger and DAQ systems. The work being undertaken by the DMNI project is aimed at developing and prototyping these LDMX-specific designs to define the baseline and prepare for construction of the detector.

Advanced designs of the LDMX support structure with integrated services and the doublelayers including

Finite Element Analyses (FEA) of stresses and deflections were completed by UCSB engineering with FY20 DMNI funding. A prototype of the full support structure has been designed (Fig. 11) and is expected to be built by the end of FY22 as planned. The design will enable most of the needed studies. These include studying the double layer mounting, the routing of services, and access issues, using thin aluminum with the cutouts expected in the final structure as a mockup. One study that cannot be done with this design will be the loading/sag to be expected with the final doublelayers, since it will be made of framing materials that will not hold their weight.

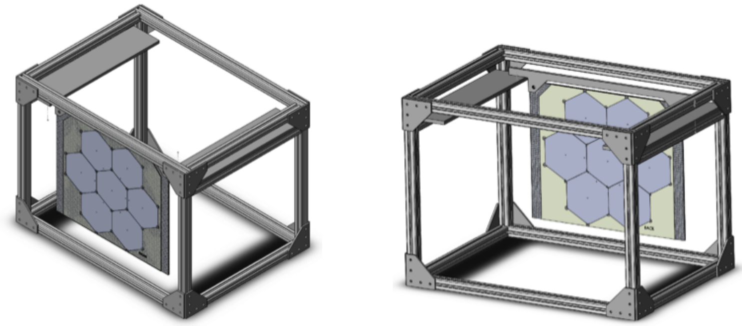

Figure 11: Front (left) and back (right) views of the ECal support structure prototype design, showing a single double layer mounted.

Another principal task of the ECal development work is designing the cooling system with the goal of ensuring that the sensors can be maintained at temperatures as low as -30°C. Radiation damage will be significant in some regions of the ECal, but FLUKA studies indicate that the worst case will be more than two orders of magnitude below the worst case anticipated in CMS. This means that it may not be necessary to operate at -30°C but the ECal will nevertheless be capable of operating this cold should that be necessary. Power use of the sensing layers is dominated by the 6 HGCROC ASICs per module that will each use about 1W of power. With other components, and assuming a worst case dark current from radiation damage, we conservatively assume 10W/module for a total of 2.4kW. Power use by other electronics within the support structure volume will be of order 2kW and so a Julabo W56 chiller with $\sim$7 kW capacity at -30°C would be adequate for the ECal. In FY22, work was done on prototyping the cooling. A prototype cooling plane for a single module was built and studied, with cooling lines laminated in carbon fiber, and resistors used to mimic the power that would be produced by the HGCROC ASICs. The cooling plane layout for a full doublelayer has been designed and will be prototyped next. A dummy doublelayer with the integrated cooling is planned to be built by the end of FY22. In coordination with this effort, the Minnesota group will be prototyping solutions for the feed-through of power, voltage, and optical data fibers between the cold volume and the external environment, while fitting within the available integration space for exit from the ECal between the magnet coils.

The heart of the readout and control system for the ECal and HCal is a radiation-tolerant FPGA (Polarfire) which controls and collects data from the HGCROC ASICs.The University of Minnesota group is responsible for the readout and electrical services integration for the ECal. During the last year, a first prototype mezzanine hosting the Polarfire has been designed using DMNI funds and constructed using University funds. A photo of the prototype, mounted on a test fixture, is shown in Fig. 12. Substantial firmware was developed for the DAQ and control functions through University support of an undergraduate computer engineering student. These firmware were successfully used in the HCal testbeam campaign described above. This effort demonstrated the ability of the proposed electronics to control and receive data from the HGCROC ASIC, as required in both the ECal and HCal electronics. The test used the RCE electronics which was the previous baseline DAQ electronics for LDMX.

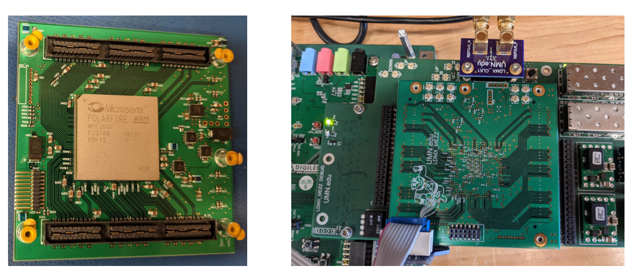

Figure 12: Prototype Polarfire mezzanine for use in the ECal and HCal readout and trigger systems (left), mounted on a test system (right).

With this successful demonstration, the project rescheduled the design of full-scale ECal motherboards into the construction project phase of LDMX. This adjustment to the schedule also matches well to the CMS development schedule, to avoid construction of a full-scale motherboard which would use a non-final connector configuration.

During the coming year, the Minnesota group will produce a few more mezzanines for extended studies for integration with the new baseline DAQ electronics (see below) and for further studies of integration with the HGCROC. The demonstrated HCal front-end will be the platform for these studies. Priorities include working on the trigger calculation processing and improving the other portions of the firmware based on the experience of testbeam. These improvements will benefit both calorimeters.

Hadronic Calorimeter Development The hadronic calorimeter is modeled on the Mu2e Cosmic Ray Veto (CRV) system and the HL-LHC upgrade of the CMS end cap calorimeter, greatly simplifying the design effort. Development activities were only required to (a) adapt the fabrication procedure and the front-end electronics to the LDMX environment, (b) engineer the HCal structure, (c) develop the readout electronics, and (d) construct a prototype, described and showed in section *LDMX Test Beam at CERN*, to validate the full hardware chain and benchmark the Monte Carlo simulations.

During FY21, we adapted the Mu2e CRV fabrication process to the LDMX quad-counter geometry. In FY22, we fabricated 50 quad-counter prototypes for the test beam, and developed procedures to mount them onto steel absorber plates (see Fig. 13). In the course of this activity, we identified several optimization opportunities, and are currently revising the design of a few components to facilitate the assembly of the full HCal. The front-end electronics was already re-designed for the quad-counter geometry in the previous funding period. In FY22, 110 pre-production boards for the HCal prototype were fabricated, together with a QC station to test 16 CMBs simultaneously. The board performance was found to be adequate during the test beam, but we have identified a few avenues to further reduce the noise level. We plan to produce new prototypes and to finalize the CMB design at the beginning of the project phase.

The second task – engineering the HCal structure – has made significant progress. A preliminary design of the Main and Side HCal structures has been developed in FY22, including a finite element analysis to evaluate the response of the detector to seismic motion. Due to reduced FY22 funding, we have postponed the development of a fully engineered mechanical design to the beginning of the project phase. Rescheduling this task will have a marginal impact on the overall project schedule.

The initial design of the readout electronics, including the HGCROC board and back plane board was completed in FY21. During this funding cycle, we have produced ten prototypes of the HGCROC board, four backplane boards and three mezzanine cards to test the full readout electronics chain (Fig. 13). The HGCROC boards were fabricated with the HGCROCv2 ASIC, although the final design will use the HGCROCv3 version produced by the LLR Omega group. Part of this effort is driven by our international collaborators (see section 3.7 for details). The mezzanine card was developed for the ECal electronics and has been readily integrated into the HCal readout. The readout electronics have been extensively evaluated during the test beam. The parameters of the HGCROC circuit have been tuned to optimize the performance, and several reconstruction algorithms have been developed to calibrate the detector response. Efforts to further characterize the system are underway. Minor revisions of the HGCROC and backplane designs that will further improve the overall stability and performance of the system were identified during the test beam campaign. These revisions were originally planned for FY22, but due to reductions in funding, they have been rescheduled for the beginning of the project phase; the resulting impact on the project schedule is minimal.

This effort has greatly benefited from the construction of the prototype, described in section *LDMX Test Beam at CERN*, to evaluate the performance of the hardware chain under realistic conditions. During the test beam in April 2022, we collected several million interactions of hadrons and electrons with incident beam energies ranging from 100 MeV to 8 GeV, together with an additional muon sample to calibrate the apparatus. We are currently analyzing these data to assess the detector performance, such as the noise level, timing resolution, energy resolution, shower profile or e/h response. The distribution of the sum of the ADC counts for a single channel from a run with a 4 GeV muon beam is shown in Fig. 13 as example. The results will be integrated into the Monte Carlo simulation and reconstruction algorithms. In parallel, we pursue studies to optimize and validate the detector geometry, refine the estimation of the various backgrounds and assess the overall HCal performance (displaced vertices reconstruction efficiency, photon/neutron separation,

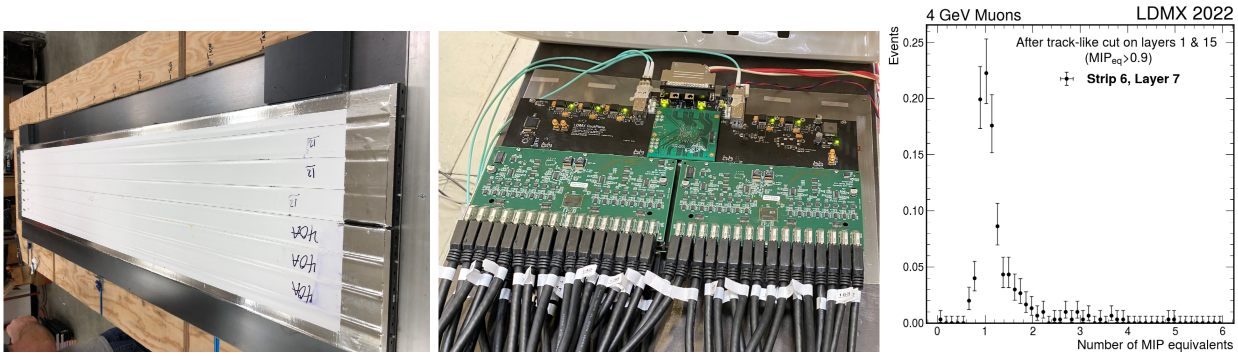

Figure 13: **Left:** A two quad-bar HCal prototype layer constructed at Caltech with bars fabricated by UVA. **Middle:** From the test beam set-up: Four HGCROC boards developed by Lund University mounted (two on each side) on a Backplane board developed by Fermilab, that holds a Mezzanine board developed by University of Minnesota, from which the serialized output is read through optical fibers. **Right:** The distribution of ADC counts for a single channel (expressed as number of MIP equivalent) from a 4 GeV muon beam run. Muons are required to leave a minimum of 0.9 MIP equivalent in counters in the front and back layers of the HCal.

impact of slow neutrons on the signal efficiency, ...).

Trigger and Data Acquisition Development The main technological goals for this proposal are system demonstrations in three primary areas:

- data readout paths for each of the detector subsystems
- trigger critical path for the primary trigger demonstrating the decision within the allotted 2 μs latency
- software and firmware interfaces for timing, fast control, and clock distribution

The activities of the TDAQ the work plan are focused towards the latter part of the development period as development of each detector subsystem progresses. Much of the initial development for the detector readout was done in concert with the test beam effort described above as data acquisition for the HCal and TS prototypes. This development work has produced a first version (V0) of the data readout firmware for these subsystems, where firmware for the HCal forms a basis for the ECal also, which uses the same FPGA mezzanine cards. For the TS readout, the CAPTAN+X firmware will be ported to the APx platform, where both use AMD/Xilinx FPGAs. Recent and on-going work has been on software and firmware development for the timing and fast control system. We have defined a specification for the LCLS-II bunch train timing messages and fast control and L1 trigger messages, and are currently developing the firmware prototype using the SLAC PGP protocol employed by the DAQ.

The remaining milestones for the technical development plan focus on two goals. The first is the implementation and demonstration of the LDMX timing and fast control firmware using a prototype of the FPGA-based PCIe DAQ platform developed as a test stand for HPS tracker readout at SLAC. The second is demonstration of the trigger data flow from back-end TS, ECal, HCal electronics to the trigger system APx ATCA boards. The former is important for synchronizing the entire experiment and defining communication with detector front-ends. The latter will demonstrate communication on high bandwidth trigger data links and define the latency constraints for the trigger algorithms within the total 2 μs budget. This will be demonstrated on APx boards available at FNAL.

Software and Computing Development The sole element of the DMNI development project in Software and Computing is the purchase of computing resources for the SLAC SDF. This investment enables scientific effort to make critical contributions to software development as well as physics studies that are required to optimize the design of the apparatus and develop the physics program of the LDMX experiment.

As discussed in Section 3.1.1, LDMX has developed the `ldmx-sw` framework for the simulation and processing of all data. The effort has been focused on ensuring timely generation of samples and processing of data including:

- optimizing the algorithms and filtering used by the simulation
- development of the reconstruction pipeline that can handle MC and test beam data

- creation of the tools to facilitate multi-site deployment of the LDMX software stack.

Each of the `ldmx-sw` modules contains sub-detector specific algorithms used to digitize and reconstruct data, apply filtering based on event conditions and implement vetoes. The simulation is also implemented as a module and wraps a custom version of the Geant4 toolkit (10.02.p03) that includes bug-fixes to the Bertini Cascade model as well as to the tree-level matrix elements for muon conversion[63].Other enhancements to the simulation include, dedicated process and cross-section biasing,ROOT based persistence, custom dark-brem physics list and a GDML based geometry system.

Generating the samples used to perform the studies described in Section 2.2 for a (1-4) $\times 10^{14}$, 4 GeV pilot run was achieved by using a combination of bremsstrahlung pre-selection and the Geant4 occurrence biasing toolkit. Updates to these algorithms have resulted in a factor of 5 improvement in the performance of the simulation. However, even with these improvements, generating a full luminosity sample will require $\sim$ 50M CPU hours and several PB of disk space. Ongoing work aims to improve the filtering of background events which can be trivially vetoed by the LDMX detector. This will allow LDMX to focus its computing power on rare, difficult to veto events allowing for timely generation of full luminosity backgrounds while also reducing the storage footprint of such a sample.

Development of the emulation of the front end electronics for all sub-detectors has been a collaboration wide effort and significant progress has been made to put in place a baseline design. Currently, the digitization and reconstruction of hits in both the ECal and HCal is complete. The robustness of both of these pipelines is currently being validated using test beam data. The digitization of both the tracker and trigger scintillator is in an advanced stage and is expected to be complete in early FY23. Futhermore, reconstruction of tracks in both the recoil and tagger trackers has been demonstrated using ACTS. [67]

All software is open source and hosted on GitHub to facilitate collaboration. Currently, the branching flow being used is based on GitHub Flow. Features are developed on branches which are either merged into trunk or a release branch. GitHub Actions are leveraged to automate unit testing of new features before being merged into trunk. Once a release is made, the packaging of `ldmx-sw` and its dependencies into a Docker container is done automatically. This facilitates the deployment of the software stack to multiple-sites without having to recreate the dependency environment. Meanwhile, the capability to run analysis jobs has been added to LDCS, which will allow collaborators to more transparently access and analyze data stored at various sites and write output to SDF at SLAC, where all LDMX collaborators have access.

Subsystem Integration Development and Project Planning Planning the high-level integration of the detector subsystems is a key task in developing a design that is ready for construction. This begins with the process of defining key interfaces between subsystems and developing the necessary tools to ensure their smooth integration into a coherent design with no missing pieces. This task requires the assignment of an experienced engineer as the Project Engineer. Beginning in FY21, a SLAC Lead Mechanical Engineer has taken on the role of Project Engineer for LDMX. This has enabled critical work to begin on understanding integration between the detector subsystems, and in particular between the Magnet and the ECal, HCal, and Trackers.

A vital component in preparing for construction of the experiment is the Project Management task required to develop a baseline scope, cost, and schedule that can be reviewed for funding of the project. The SLAC Project Engineer is also taking on the formal project management role, and together with the Technical Coordination of the LDMX Collaboration (see Section 3.2) and the PIs of the collaborating institutions, is developing the project plan. With the addition of a formal project manager as part of the DMNI project, the collaboration has initiated a formal project management process, including the development of a complete set of Basis of Estimate (BOE) documents for construction of the experiment, an accompanying resource loaded schedule, and a risk register. The cost and schedule information being developed for the Design Report is the key deliverable of the Project Management component of the DMNI project. A current snapshot of this process is presented in Section 3.6 and the accompanying documentation for the current status review.

3.1.3 Use of DOE Lab Infrastructure and Capabilities

The use of DOE Laboratory infrastructure and technical capabilities is an integral and critical part of the LDMX project. First and foremost, the entire experimental concept depends upon DOE accelerators to deliver the electron beam required for the experiment, either with the LESA facility at SLAC, or by using the CEBAF at JLab. Without a CW or high-rep-rate multi-GeV electron beam, the experiment cannot be realized. Related to this are other technical capabilities at the labs that are needed to develop and mount an accelerator-based experiment: expertise in beamline components and vacuum systems, special expertise required to refurbish and characterize the dipole magnet for the experiment, and the heavy fabrication and conventional facilities that support these systems.

Another area where LDMX leverages DOE Laboratory expertise are the electronics required for the readout, data acquisition, and control systems for the experiment. The design and implementation of these systems for LDMX rely upon specialized hardware developed at SLAC and FNAL for other experiments, where key experts at the labs are needed to update these designs and adapt them to the needs of the LDMX experiment.

Finally, the extensive experience at SLAC in developing and managing projects at all scales is critical to the stewardship of the LDMX project. This includes formal project planning and coordination, project scheduling, budgeting and financial oversight and reporting, procurement, quality assurance, and management of ES&H required. Coordination of these elements is provided by Project Manager Lange with oversight by the SLAC Project Management Assurance Group.

3.1.4 Investigation of Alternatives

Before the start of the DMNI effort, the collaboration considered a number of design alternatives, including the use of a crystal ECal and alternative absorber materials in the HCal. The technologies identified herein were determined to be the best choices with respect to performance and cost to meet the physics goals of the experiment.

One alternative being investigated at the time of the FY21 status report was the possibility of utilizing different hardware platforms for the data acquisition or trigger systems of the experiment. As a result of this investigation, the previous baseline platform for the back-end DAQ, the Reconfigurable Cluster Element (RCE) system developed at SLAC, has been replaced in the baseline plan by commercially available FPGA-based PCIe hardware, which has revolutionized computing for other specialized tasks in recent years. While the RCE platform has been used by HPS, it is aimed primarily at larger experiments such as LSST, and its custom nature implies relatively large costs for construction and maintenance for a small experiment. This design and related development work are described in Sections 3.1.1 and 3.1.2 and leverage work taking place independently at SLAC to develop this platform as a general purpose DAQ platform for smaller experiments, as described in Section 3.4.2.

Some alternatives and options in the final design are still under study. We are studying the opportunity of using an LYSO active target, as discussed in Section 3.1.2. The active target could offer opportunities to further suppress photo-nuclear interactions in the target. Additionally, as a risk mitigation, we include in the baseline plan design work to allow the use of an alternative readout ASIC for the tracker.

3.1.5 ES&H Planning

The LDMX Project has the duty to follow the guiding principles and core functions of the Integrated Safety and Environmental Management System to conduct all work safely, effectively, and efficiently to ensure the protection of workers, the public, and the environment. The ES&H plan covers all phases of the project including design, development, fabrication, assembly, handling, transportation, storage, integration, test, and operation. Personnel safety will include site-specific training as deemed appropriate by the project's SLAC safety representative, Norm Picker. ES&H deliverables for the design report will include:

- Integrated Safety Management Plan
- Hazard List
- Preliminary Hazard Analysis Report (later to become a HAR)
- Complete National Environmental Policy Act (NEPA) strategy, as required by DOE O 451.1B
- Prepare Environmental Compliance Strategy
- Prepare Construction Safety and Health Plan

- Prepare Radiation Protection Documentation

Hazards associated with the project are identified and evaluated at both the subsystem level and for the integrated experiment. The risks associated with all identified hazards are controlled to acceptable levels and documented in standard hazard assessment matrices.

3.1.6 Computing and Data Management Plan

The design of the computing infrastructure for the experiment is described in detail in Section 3.1.1, and meets the needs of the experiment with the deployment of 360 TB of disk space, 6.3 PB of tape storage, and 1500 cores of CPU. These resources will be hosted at the SLAC Shared Scientific Data Facility (SDF), which is being developed as the next-generation pooled resource for scientific computing at SLAC. SDF is aimed at serving the major storage and data processing needs of LCLS-II and the Vera Rubin Observatory, as well as smaller experiments such as LDMX. SDF currently consists of 88 dual 64-core Rome servers (11264 cores and $\sim$ 400 TFLOPS) and 2 DataDirect Networks ES18K storage appliances with a capacity of 17 PB interconnected via high speed networking fabric. As shown in Figure 14, the total computing and storage capacity is expected to increase dramatically by FY25, when LDMX could begin the pilot run. The figure also shows the expected LDMX needs, which amount to a tiny fraction of the total projected capacity. The current SDF compute model grants all SLAC users access to baseline computing capabilities including 25 GB of storage per user and access to the shared CPU node partitions. However, higher priority goes to groups that have contributed hardware to SDF. The cost for LDMX's storage and CPU needs assume that LDMX will contribute to SDF the full needs of the experiment, which will ensure dedicated access to the required storage and CPU. However, as a pooled (batch) processing resource, the availability of CPU can often be larger.

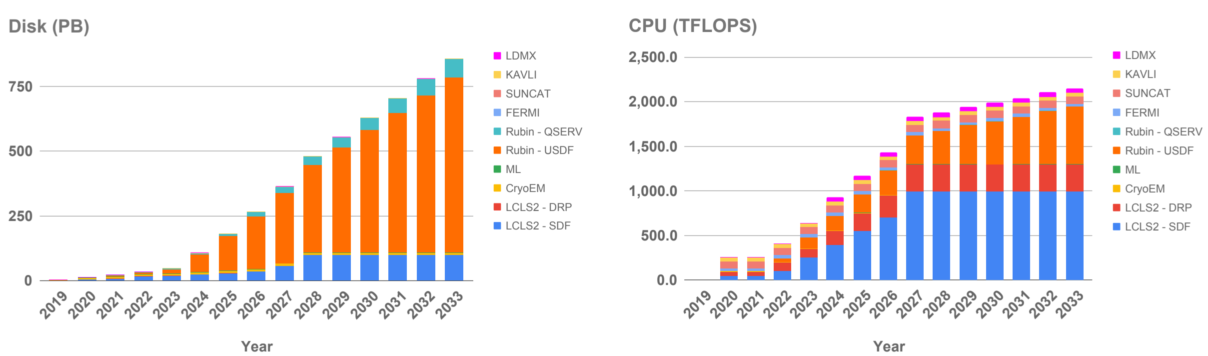

Figure 14: Projected SDF computing and storage capacity. The LDMX needs amounts to a small fraction of the total capacity.

To augment the production of MC and offline analysis, LDMX will continue to develop the LDMX Distributed Computing System (LDCS)[68]; a distributed computing system currently consisting of 4 sites: Caltech, Lund, SLAC and UCSB. All sites have access to local storage along with 100 TB of storage at SLAC accessible via GridFTP. All simulation jobs are run at all sites and the resulting files are catalogued using Rucio. LDCS has played a key role in generating the large scale MC samples needed for design studies and is expected to continue contributing in an auxiliary capacity.

The data distribution and access policy is set by the Collaboration Board on behalf of the LDMX collaboration. The policy can be revised by the Collaboration Board at any time after a review process with input from the collaboration.

Data Description & Processing of Products LDMX will produce data from the following sources

- Raw data from testing and calibration of detector prototypes at collaborating institutions
- Monte Carlo data generated using SDF and LDCS as described in Section 3.1.1

- Raw and reconstructed data from the two different periods of operation described in Sections 3.8 and 3.1.1 and the test beam.

As discussed in Sec.3.1.1, `ldmx-sw` will be used to build the reconstruction and analysis pipelines needed to process and persist the data. All data is persisted to a ROOT based data model. All data will be centrally stored at SLAC and be made available to all members of the LDMX collaboration.

Plan for Serving Data to the Collaboration and Community Before being released to the LDMX collaboration, data is tagged using the framework version used to produced it. These tagged releases will serve as the standard data sets that will be used for analysis and publication. Dissemination of the data beyond collaborators will be cost prohibitive.

Plan for Making Data Used in Publications Available In all cases of publications, data in the plots, charts and figures, and Digital Object Identifiers will be made available in accordance with policy at the time of publication by using mechanisms provided by the publisher, hosting by a collaborating institution, or services provide by INSPIRE. This includes publications resulting from research data from experiments, simulation, and research and development projects such as detector prototype data.

Responsiveness to Office of Science Statement on Digital Data Management The data management plan fully adheres to the recently implemented policy of the
DOE Office of Science: http://sciences.energy.gov/funding-opportunities/digital-data-management/.

3.1.7 QA Planning

The goal of the LDMX quality assurance program is to provide mechanisms for controlling activities that affect product quality, or protect the environment and health and safety of both the public and personnel involved with the project. These mechanisms are intended to establish a graded approach to quality assurance, invoked to the extent consistent with the importance of the activity. Not all items, processes, activities, and services have the same effect on health and safety, reliability, environmental protection, or program objectives. Therefore, such a graded approach acknowledges the importance in establishing the applicability of aspects of quality assurance to specific activities and to the degree to which they need to be applied. Considerations include:

- The relative importance to safety, safeguards, and security
- Compliance with SLAC and other institutional Policies and Regulations
- LDMX project mission and programmatic impact

The objective of this graded approach is to ensure that activities affecting quality are managed through adequate systems and procedures that are commensurate with the complexity and hazards of the work being performed. LDMX project management and collaboration members are responsible for identifying the activities that are subject to these requirements, and for carrying out an analysis to justify the degree of rigor to be applied.

Work planning control (WPC) is the formal process by which the project will identify and mitigate risks and hazards. This formal process occurs at the earliest stage of the project during the initial design and planning and continues through releasing and performing the actual work. Norm Picker (SLAC) has been identified as our project safety officer and will assist in WPC for the LDMX project.

3.2 The LDMX Collaboration

The LDMX collaboration is centered around the SLAC team which provides core leadership scientifically, technically, and managerially for the collaboration. In total, the collaboration consists of eight institutions. The collaboration developed informally through the period 2016-2018 and took formal shape through the adoption of bylaws and selection of officers including co-spokespersons in 2019. The organization chart for the

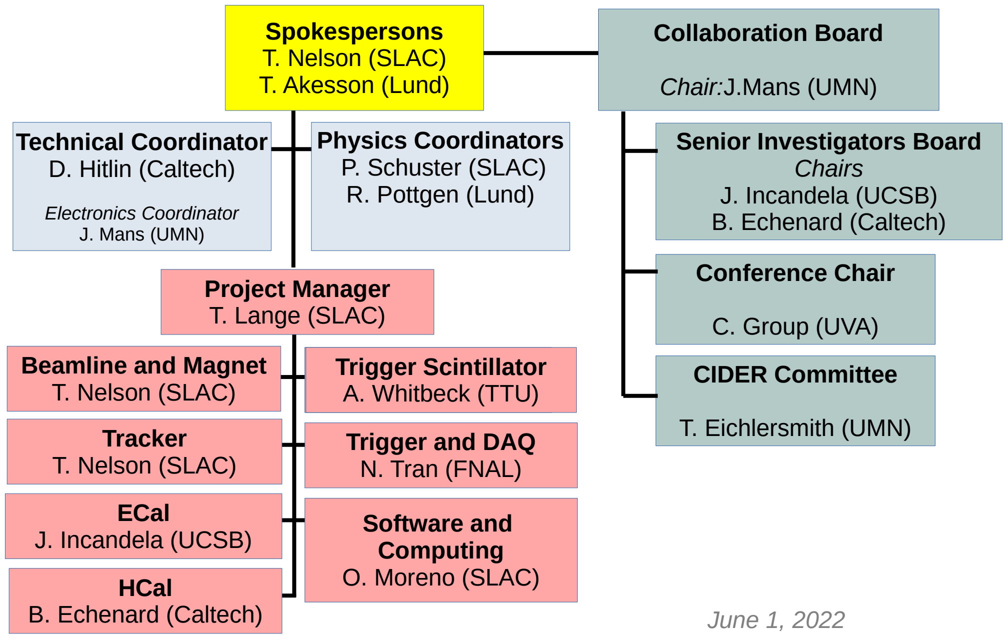

Figure 15: Organization chart of the LDMX Collaboration.

collaboration is shown in Fig. 15. The CIDER committee is responsible for collaboration climate, diversity, and outreach activities.

Each of the institutions in the collaboration has particular roles in the development and planned construction of the experiment and brings specific and important knowledge and capabilities to the collaboration. The list of institutions below is given in alphabetical order.

- *California Institute of Technology* (PI: B. Echenard) – The Caltech group is responsible for overall design and integration of the HCal detector. B. Echenard is the HCal coordinator and co-chair of the Senior Investigators Board. D. Hitlin serves as the Technical Coordinator for LDMX. The Caltech group also leads the exploration of the potential of a LYSO-based active target.

- *Fermilab* (PI: N. Tran) – The FNAL group is responsible for off-detector trigger electronics for the ECal and HCal systems (based on work underway for the HL-LHC CMS upgrade), as well as the motherboard for the HCal readout/trigger electronics. The FNAL group is also active in simulation tasks for TDAQ. N. Tran is coordinator for Trigger and DAQ.

- *Lund University (Sweden)* (PI: T. Akesson) – The Lund University group is responsible for the digitizer board for the HCal, based on the HGCROC ASIC. The group is also very active in computing, including supporting the LDCS distributed computing system [68]. T. Akesson is a co-spokesperson for LDMX and R. Pottgen is a co-physics coordinator.

- *SLAC/Stanford* (PI: T. Nelson) – The SLAC group provides core management and leadership for the collaboration as well as holding responsibility for specific subsystems. The group is responsible for the preparation of the beamline and interfaces with LESA, the refurbishment of the experimental magnet, the construction and integration of the tracker, and the readout of the trigger scintillator and its integration in the trigger. The group is also responsible for central DAQ, offline software, and computing activities. The tracker design builds extensively on the design from HPS, the DAQ is based on general-purpose DAQ technologies developed at SLAC, and many aspects make heavy use of DOE laboratory expertise and unique capabilities. T. Nelson is a co-spokesperson and coordinator for the tracker and beamline, P. Schuster is a co-physics coordinator, and O. Moreno is the software and computing coordinator.

- *Texas Tech University* (PI: A. Whitbeck) – The TTU group is responsible for the construction of the trigger scintillator system and its integration with electronics developed for the CMS Phase 1 HCAL Upgrade. A. Whitbeck is the Trigger Scintillator coordinator.
- *University of California, Santa Barbara* (PI: J. Incandela) – The UCSB group is responsible for the overall ECal design and the construction of the ECal modules and ECal absorber/cooling/support structure. The module construction leverages expertise with the same task for the CMS HL-LHC upgrade, while the cooling and other engineering aspects leverage experience with construction of several trackers. J. Incandela is the ECal coordinator and co-chair of the Senior Investigators Board.
- *University of Minnesota* (PI: J. Mans) – The UMN group is responsible for the readout, on-detector trigger calculations, and services of the ECal detector. A common UMN-designed mezzanine is used to interface between the HGCROC ASICs and the off-detector electronics in both the ECal and HCal, leveraging experience from the CMS HL-LHC upgrade. The Minnesota group is also highly active in core offline-software tasks. J. Mans is the collaboration board chair and electronics coordinator.
- *University of Virginia* (PI: C. Group) – The UVA group is responsible for the construction and test of the scintillator units for the HCal detector, an activity which directly leverages experience from the Mu2e experiment construction. C. Group is the LDMX conference chair.

The project plan anticipates the participation of a modest number of additional groups in the construction, operation, and analysis of data from the experiment. The total scale of the required scientific effort for these different phases of the experiment is discussed in 3.6 and Section 3.8.

3.3 Organization and Management

The organization and management of the current DMNI development project and the planned construction project for LDMX are closely tied to the formal LDMX Collaboration described in Section 3.2. The organization of the technical arm of the collaboration, shown in pink in Figure 15, includes a set of Subsystem Coordinators, an Electronics Coordinator, and a Technical Coordinator. The PIs of the DMNI development project are the corresponding Subsystem Coordinators. The Electronics Coordinator is a co-PI and responsible for developing standards, interfaces, and an integration plan for electronic systems of the experiment, including the DAQ, trigger, online computing, detector monitoring, interlocks, and slow controls. The Technical Coordinator is a Senior Investigator on the DMNI project responsible for overall coordination of technical work across subsystems and developing an overall subsystem integration plan for the apparatus, including the mechanics, cooling and other environmental systems, and utilities. The Technical, Electronics, and Subsystem Coordinators all report directly to the spokespeople, one of whom is the lead PI of the DMNI project.

As discussed in Section 3.1.2, an element of the DMNI project is support for Lead Mechanical Engineer Travis Lange at SLAC, who has experience from the LSST project, to take on the role of Project Engineer. The DMNI project plan calls for 25% of an FTE for his time beginning in FY21. In the absence of full funding for the project, SLAC has provided internal support for this effort. In this role, Lange takes on leadership of formal project management for the DMNI project and a corresponding position in the organization of the collaboration. Addition of this effort has allowed project planning to proceed to a more formal process for developing and documenting the budget and schedule for both the DMNI project and the construction project, beginning with the Basis of Estimate (BOE) documentation, Resource Loaded Schedules, and Risk Register presented with this report. With a transition to the construction project for the experiment, this role is expected to grow from 25% to roughly a full FTE throughout the construction and installation of the apparatus.

Management of Budget, Schedule, and Milestones Subsequent to the DMNI award, the project was descoped to the awarded amount and internal milestones were established to track progress. These milestones are presented in Section 3.5 and are tracked at bi-weekly Technical Coordination meetings organized by the Technical Coordinator, and attended by the Project Manager, the Electronics Coordinator and all of the Subsystem Coordinators. Missed milestones are promptly rescheduled so that minimal risk is introduced to achieving the project goals.

The Project Management effort at SLAC has enabled the collaboration to develop a formal means of tracking the budget and schedule for the DMNI project and to develop plans for the construction of the experiment. Toward this end, we have adopted Microsoft Project in the Cloud as a collaborative tool. This choice appears to be appropriate for a project of the scale of LDMX. Project is used in conjunction with our Basis-of-Estimate (BOE) documents, in the form of Excel spreadsheets in a common format which are kept in the Fermilab doc-db database, that are used to apply rule-based contingencies on an item-by-item basis. In the future, this system will be able to level resources, and track earned value. The program is substantially less expensive than Primavera P6, requires less training, and will require a smaller group of people to administer. The WBS and resource-loaded schedule is contained in linked Project files driven from a Master file that contains common labor resources and information on indirect charges and escalation rates at each institution. The details of individual tasks are contained in the BOE documents, which break down the labor and M&S content of individual WBS tasks and provide engineering estimates, RFPs, invoices, *etc.* to support the estimates. These tools have been used to produce the budget and schedule information presented in this report.

Management of Risks The project follows a formal risk management plan which involves the Technical Coordination, Project Engineer, and all subsystem coordinators. A risk register has been developed for the project following common best practices and utilizing formalism from the LSST project. Each entry in the risk register has a clearly responsible individual for developing and tracking the risk. As part of the technical coordination activities of the project, the risk register is will be reviewed and updated on a quarterly basis. For each risk, mitigations are considered and tracked as appropriate, particularly for any planned mitigations which have not yet been implemented into the project plan.

Each identified risk is assessed for cost impact (potentially across a range of cost impacts), delay (also across a range), performance impact, and probability. These impacts are combined with the probability to determine the overall impact on the project. The overall exposure level is categorized into one of five levels (Insignificant, Minor, Moderate, Major, Critical). Currently, the project has 11 Minor risks and 7 Insignificant risks. The project does not have any Moderate, Major or Critical risks. The highest-exposure risks currently include supply chain risks, which would impact the schedule, as well as cost-impact risks related to lack of scientific personnel, departure of a key engineer, or the requirement to make a minor design change or addition to the detector design. The full risk register is available from the review website.

In addition, a Monte Carlo simulation is performed using the information in the risk register to determine the necessary risk-based contingency for the project. At the 95% confidence level, the risk Monte Carlo predicts a requirement for $1.332M of risk-based contingency. The delay impacts in the risk register are not carried through a Monte Carlo process and are preliminary and conservative pending further analysis of the project critical path.

3.4 Dependencies on Outside Resources Planned

3.4.1 Linac to End Station A (LESA) Beamline

The unique capability of a missing momentum experiment hinges on a CW multi-GeV electron beam to achieve high statistics. The proposed Linac to End Station A (LESA) beamline at SLAC is well suited to deliver such a beam and will support long-term occupancy for LDMX. LESA beamline installation and commissioning is planned for FY23-24. LESA will leverage the LCLS-II superconducting linac and of the Sector 30 Transfer Line (S30XL) AIP, both of which are currently under construction. S30XL, an important precursor to LESA, will parasitically extract trains of "empty" (dark current) bunches from the LCLS-II linac's dump line, with a duty factor of roughly 60% (one 600 ns macro-pulse every 1.1μs). The window duration is limited by the time required to ramp LCLS-II and S30XL kickers before and after each pulse, so when LCLS-II operates at lower rate, longer macro-pulses and hence higher duty factors will be possible. LESA will build on S30XL by connecting this transfer line to the existing End Station A beamline and in FY24 or 25 a low-power gun laser will be added to seed dark current bunches within the S30XL macro-pulse at 37 to 46 MHz (the 5th or 4th sub-harmonic of the gun frequency). LESA will be completely parasitic to LCLS-II, and has been reviewed extensively for non-interference.

The End Station A line provides the capability to tune the beam current to a level of one electron per pulse, with a wide beam spot as required by LDMX — capabilities used extensively in the End Station Test Beam program. The LDMX experiment is the flagship LESA user, and is anticipated to receive the majority of LESA's available beam. LDMX will be installed near the down-beam end of End Station A for a multi-year term of running. Other comparably-low-current experiments are also anticipated to run at LESA for relatively short periods, such as test beam users and dedicated electronuclear scattering experiments. These will be installed up-beam of LDMX, with a stopper preventing beam from reaching the LDMX detector. The impact of such experiments on LDMX in terms of lost beam-time and radiation are minor. Access to the detector will, however, be impossible whenever another experiment is receiving beam in End Station A. LESA completion on the timescale of FY24 will allow for early commissioning of LDMX in End Station A with a low-current 4 GeV CW electron beam, and — depending on the LDMX project profile — possibly initial physics running. Beam availability will continue until the LCLS-II linac will be upgraded to 8 GeV energy as part of the LCLS-II-HE project, currently planned for the 2026-28 timeframe. Upon completion of this upgrade, LESA will deliver 8 GeV beam to LDMX.

Higher beam energies are advantageous to LDMX, as they sharply decrease the rates of challenging few-particle backgrounds and improve the detector's single-particle rejection capabilities. However, LDMX studies have shown that 4 GeV beam is sufficient for a low-background search at the pilot-run LDMX statistics of $1 - 4 \times 10^{14}$ electrons on target. Thus, LDMX can proceed with this run and achieve powerful physics sensitivity even if the LCLS-II-HE project is delayed.

Although LESA is a logical and efficient approach to achieving LDMX's beam requirements, and is on track for completion on the same timescale as LDMX, alternatives exist and will be pursued aggressively by the LDMX collaboration if LESA is not completed. One possibility is to run at JLab's CEBAF, which has demonstrated delivery of pA-scale currents for short periods, for example for QWeak detector commissioning, and can deliver beam at energies up to 11 GeV to each of its four halls. Foreign facilities such as the proposed eSPS [69] at CERN could also be explored.

3.4.2 External Resources for Detector Design / Construction

LDMX leverages a number of existing HEP technologies and facilities to develop and fabricate the detector. A complete list of facilities and equipment from the collaborating institutions is presented in Appendix A.4 and Appendix A.5 of the LDMX DMNI proposal. There we include descriptions of available laboratory space at collaborating institutions, the availability of engineering resources, and specialized facilities for detector fabrication. Examples of these dedicated facilities and equipment include clean rooms for silicon detector construction at SLAC and UCSB, machine shops and detector assembly facilities at TTU and UVA, the Fermilab scintillator fabrication facility, calorimeter testing equipment at Caltech, and electronics assembly and testing equipment at the University of Minnesota.

In addition to available laboratory and university facility and equipment, LDMX is leveraging technology currently being designed for other projects. LDMX is planning to use sensors designed for the CMS HGCal which are currently in pre-series production. The ECal readout electronics are also being developed for the CMS HGCal, in particular the HGCROC readout ASIC, which will be submitted for its engineering run by the end of 2022.

LDMX plans to use ATCA-based trigger cards being developed for the CMS L1 Trigger upgrade. These cards are currently in their pre-production phase and would already be usable by LDMX. Finally, LDMX is leveraging the independent development of a new DAQ platform using commercial FPGA PCIe hardware taking place at SLAC using KA-25 R&D funds. A test stand based on similar hardware, assembled for the testing of HPS detector modules and readout electronics, is being used as a prototype of the LDMX DAQ as part of this KA-25 project. Where appropriate, dependency of the LDMX project and schedule on these external designs have been discussed as risks in Section 3.3. Scientific effort for design and prototyping comes from external institution resources and includes activities such as physics simulation, performance, and sensitivity. This is described in detail in Section 3.7.

3.5 Budget and Schedule for DMNI Project Phase

The LDMX DMNI proposal for Track 1 submitted in May 2019 included a budget and schedule for a two-year project to complete a technical design and produce a baseline cost, scope and schedule for construction of the experiment. This budget totaled $1.96M and is summarized in Table 4. The proposal was awarded on

DMNI Phase Budget Summary				
	Year 1	**Year 2**		**Total**
DMNI Proposed (K$)	890	1071		**1961**
	FY20	**FY21**	**FY22**	**Total**
DMNI Awarded (K$)	150	675	675	**1500**
Funds Received (K$)	150	400	500	**1050**

Table 4: High-level summary of the proposed budget for the LDMX DMNI project, the awarded budget and actual funding received. Initial funding was received in Q3 of FY20 and the latter two have been stretched to three years.

February 28, 2020, at the level of $1.5M and with the planned profile of $150K in FY20, $675K in FY21, and $675K in FY22. Before submitting the lab FWP for this work, the project was descoped to the award amount and planned profile in order to preserve the goal of preparing the project for a Design Review in the latter half of FY22 and construction in FY23 with minimal additional risk. A schedule for the work plan, including internal milestones, was developed together with this budget, in order to track progress on the project. These milestones and their current status are shown in Figure 16.

FY20 funding was made available in Q3, and work began on the project according to this work plan. Funding in FY21 was delayed, where redirected effort and outside resources enabled critical work to continue into FY21 to meet some milestones and maintain progress towards the planned test beam activity at CERN. With $400k arriving in Q4 FY21, we were able to successfully mount the CERN test beam amidst complications from the pandemic and restart work on critical tasks. In FY22, the project was funded at $500K, for a total of $1050K over the three year period. Upon careful examination of the work plans and the key deliverables, the project was descoped to fit this total, new milestones for the DMNI project were established, and some work planned for the DMNI period was pushed into the plan for the construction project. While these changes do not present significant risks to the construction project, they add to the project cost and schedule. Table 4, shows the budget for the LDMX DMNI development project presented in this report and Figure 17 is a Gantt chart showing a Level 3 roll-up the DMNI resource-loaded schedule.

3.6 Budget and Schedule Estimate for the Small Project

The LDMX DMNI proposal included a preliminary estimate of the project cost for construction of the experiment. This initial attempt at a complete bottom up estimate did not include a full accounting of overhead rates at some institutions or the pass-through overhead that applies for funding the project through SLAC. The computing resources required to record, process, and analyze the data were also not included, and significant scope in a number of subsystem concepts had not been identified. Contingency was included at the level of 25% on the detector subsystems and 50% on the Installation and Management tasks. This estimate assumed full funding of the DMNI proposal at the amount of $1.96M. With these assumptions, the total estimated cost without (with) contingency was $6.4M ($8.4M).

As discussed in Section 3.3, we have since adopted a more formal process for budgeting and scheduling of the LDMX Project, with the development of a formal Work Breakdown Structure (WBS), a complete set of Basis of Estimate (BOE) documents for all labor and M&S, and a corresponding Resource-Loaded Schedule (RLS), which can be used to track budgets and schedules in detail. A complete set of preliminary BOEs and a RLS for the construction project were produced for the FY21 review, and used to generate a more complete cost estimate for the project, shown in Table 5. Based upon this exercise, the total cost of the LDMX construction project was $11.65M without contingency, not including contributions through Lund University discussed in Section 3.7 valued (M&S) at $547K. While the BOE documents defined item-by-item contingencies according to established practice, these were not yet integrated into the RLS. Given the level of maturity of the estimate, a 45% overall contingency rate was applied, resulting in a cost estimate with contingency of $16.9M. The differences in scope, costs, and overheads from the estimate in the proposal

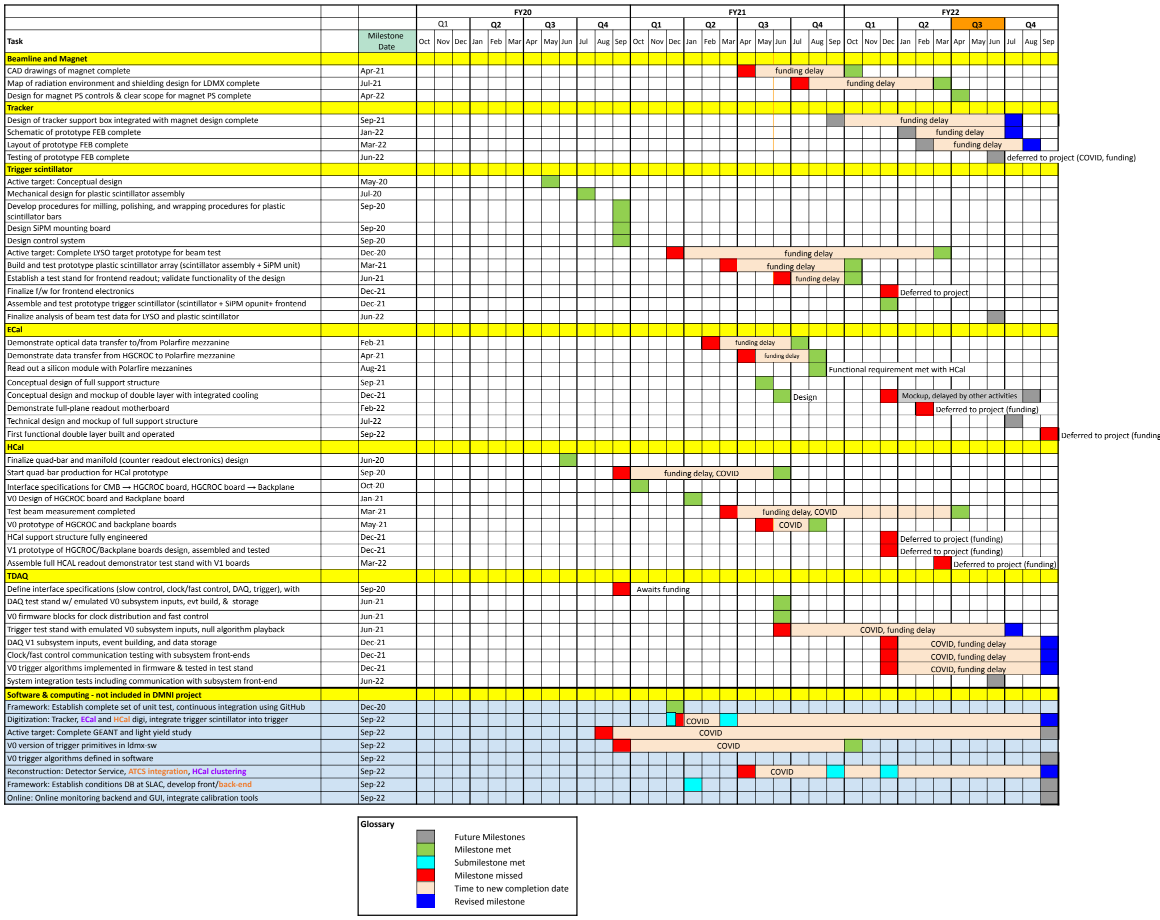

Figure 16: Milestones for the DMNI phase. We have used the color code at the bottom of the figure, together with some annotations, to compactly summarize milestones that have been met, those that have been missed due to COVID delays and the delay of FY21 funds, and those that have been replanned.

were discussed in detail in the FY21 Status Report. Note that no contingency is applied on international contributions as there is already a 30% contingency applied to those funds.

We have recently repeated this exercise, including a number of significant changes. First, scope that had been removed from the DMNI development plan was added to the construction project. Second, some components of the DMNI development plan have been completed that were specifically aimed at defining the scope of the project, and have added to the project plan. Third, all labor, M&S, and overhead rates have been updated. Fourth, the item-by-item contingencies defined in the BOE documents have been implemented in MS Project, allowing for more precise picture of the contingency required. Fifth, the risk assessment process has defined some mitigations to be added to the baseline, as well as providing an estimate of the risk-based contingency to be added to the project. The updated cost estimate for the project is shown in Table 6, and totals \$15.3M(\$21.3M) without(with) all contingency, where the major differences are discussed below. We have also begun tracking scientific effort required during construction, in the form of graduate students and postdocs that are assumed to be supported on research budgets. An estimate of this effort is summarized in Table 7 and the potential for a shortfall in research funding is included in the risk register.

The most significant changes in scope were identified as part of the DMNI project to develop the beamline and magnet project. In particular, LDMX will be responsible for significant upgrades to electrical infrastructure at SLAC to provide power to the magnet. The other significant additions come from moving \$450K

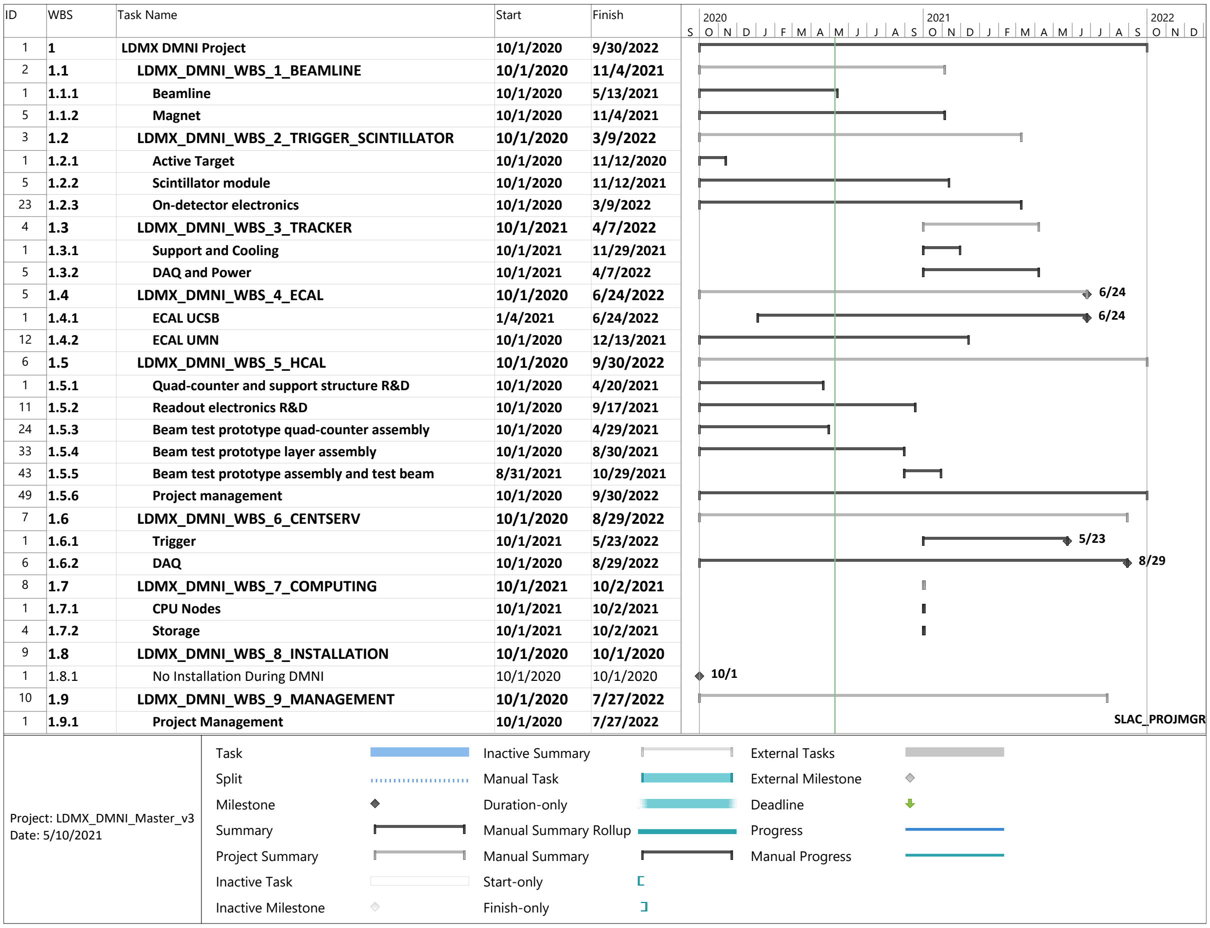

Figure 17: Summary Gantt chart at Level 3 for the DMNI phase

in scope from the DMNI project to the LDMX project, from better understanding the installation and commissioning efforts, and adding scope to the baseline to mitigate risks identified in the Risk Register. These changes add \$1.3M to the baseline cost. An even larger impact comes from updating the anticipated labor, M&S, and overhead rates. In particular, the currently anticipated escalation of labor costs is much higher than expected last year, which compounds with rates for FY22 that have already increased substantially at many institutions. The same is true of M&S across the project, where some cost drivers like steel for the HCal have seen increases that approach a factor of 2. Finally, the addition of risk-based contingency to the project adds \$1.332M.

Fig. 18 is a Level 3 roll-up the resource-loaded schedule for the LDMX construction project, including all work required prior to commissioning the experiment with LESA beam. Most additions to the project scope do not lie on the critical path, so that the schedule is only moderately impacted, and shows LDMX ready to take beam approximately 2.75 years after project start.

This fully integrated project plan will be further developed for the design report, including the full budget and schedule at an increased level of detail, resource-leveling, and the identification of critical path items and project milestones. In addition to managing and tracking the use of DOE resources and infrastructure to build the experiment, the management will also coordinate the resources contributed by international collaborators and be adapted to deal with the operations phase of the experiment.

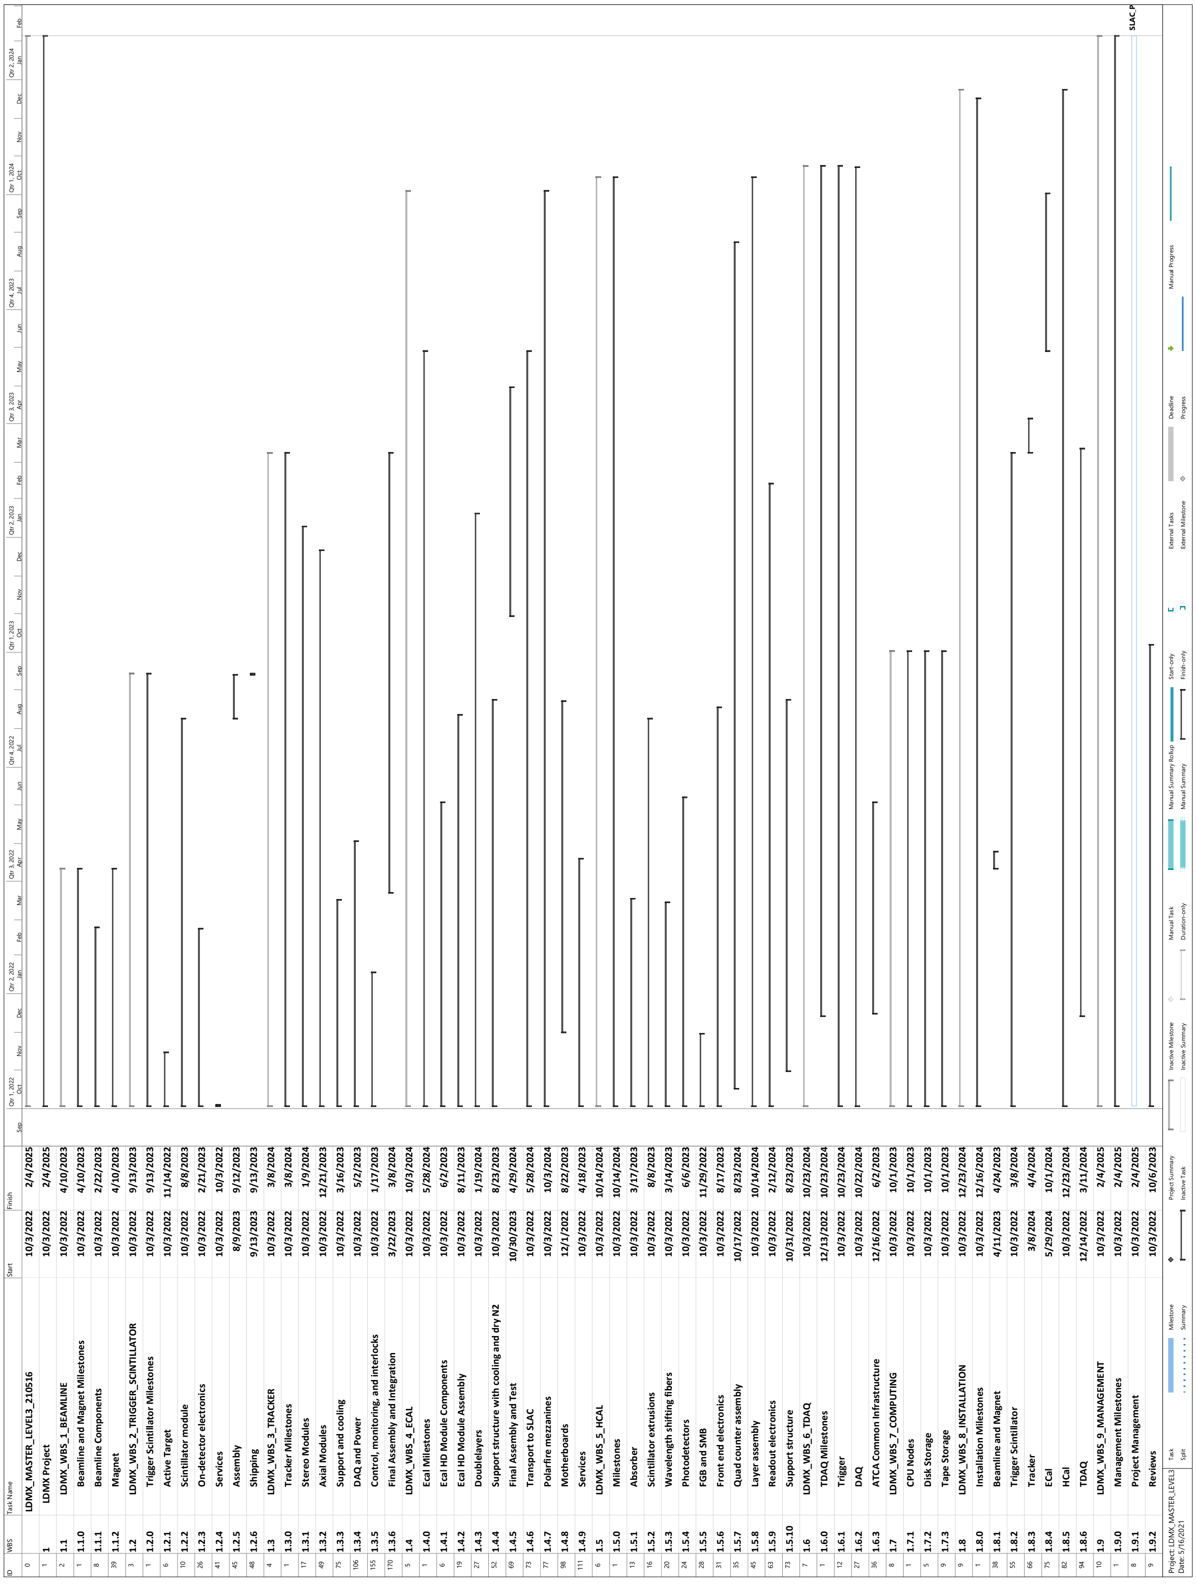

Figure 18: Summary Gantt chart at Level 3 for the LDMX Small Project. No attempt at resource-leveling has as yet been made.

FY21 Cost Estimate for the LDMX Project						
WBS	**Item**	**M&S Total (K$)**	**Labor Total (K$)**	**Total (K$)**	**Cont. (K$)**	**Total w/ Cont. (K$)**
1	LDMX Detector	4,214	7,435	11,649	5,242	16,890
1.1	Beamline	113	445	558	251	809
1.2	Trigger Scintillator	103	109	212	95	307
1.3	Tracker	499	1,555	2,054	924	2,978
1.4	ECal	1,353	973	2,327	1,047	3,374
1.5	HCal	1,085	1,109	2,194	987	3,181
1.6	Trigger/DAQ	482	1,594	2,076	934	3010
1.7	Computing	442	0	442	199	641
1.8	Installation	109	434	542	243	785
1.9	Management	28	1,217	1,246	561	1,807

Table 5: The FY21 estimate of the cost of the LDMX project broken down by WBS, assuming funding of the DMNI project at the planned amount of $1.5M. Given the preliminary nature of this estimate, an overall contingency rate of 45% was assumed (excluding international contributions).

Current Cost Estimate for the LDMX Project						
WBS	**Item**	**M&S Total (K$)**	**Labor Total (K$)**	**Total (K$)**	**Cont. (K$)**	**Total w/ Cont. (K$)**
1	LDMX Detector	5,179	10,070	15,250	6,096	21,346
1.1	Beamline	192	1,186	1,378	676	2,054
1.2	Trigger Scintillator	208	99	307	87	395
1.3	Tracker	541	2,438	2,979	982	3,962
1.4	ECal	1,655	1,155	2,809	718	3,528
1.5	HCal	1,499	942	2,441	656	3,097
1.6	Trigger/DAQ	449	2,454	2,903	932	3,835
1.7	Computing	481	0	481	169	650
1.8	Installation	123	613	736	327	1,063
1.9	Management	30	1,184	1,214	216	1,430
	Risk Contingency				1,332	1,332

Table 6: A current (June 2022) estimate of the cost of the LDMX project broken down by WBS. Total contingency, including risk-based contingency, is 40% of expected project cost.

3.7 Other Contributions

LDMX is a Collaboration of eight institutions of which five are US universities, two are national laboratories, and one is a Swedish university. This section outlines other contributions than the DMNI funding used (and for Lund University also to be used) by those institutions. When the project becomes fully funded, we expect these other contributions to grow commensurately with the level of effort required to support the project.

Scientific Labor for LDMX Project	
Effort type	Total (hours)
Postdoc	32,562
Graduate Student	11,936

Table 7: An estimate of total scientific effort planned to be available to the LDMX construction project and funded by research budgets.

3.7.1 Lund University, Sweden

Lund University[1] participates with two faculty members, one postdoctoral researcher, one engineer, graduate students and undergraduate students.

The support was obtained for the DMNI-phase through the Royal Physiographic Society[2] (M&S $ 27,000), the Knut & Alice Wallenberg Foundation[3] (M&S $ 20,000), and the Crafoord Foundation[4] (M&S $ 60,000). The latter two also included the travel support, and the latter some postdoctoral support. These grants covered the majority of the M&S costs for the HCal prototype, and some M&S costs for the Trigger Scintillator prototype. All this, roughly M&S $ 107,000 and 1 200 hours of engineering have been spent by the end of May this year.

The Knut & Alice Wallenberg Foundation[5] approved in October 2019 a project entitled *Light Dark Matter*, that started 1 July 2020. This project has four work packages of which the largest is the Lund University participation in the LDMX Collaboration. It includes $ 696,000 M&S support for the HCal (of which $ 20,000 was spent on the testbeam prototype), $ 51,000 for LDMX data storage at Lund University, and labor support for 2 400 hours of electronics engineering. This grant covers in addition postdoctoral and PhD-student support, and the required travel and subsistence for a proportionate presence at SLAC.

All amounts above in $ are based on the 31 May 2022 exchange rate (which has become less favorable since the FY21 DMNI report), while the budgets are granted in SEK.

The faculty member participation is supported by the University research budget, the Swedish Research Council[6], and by the Knut & Alice Wallenberg Foundation.

LDMX is using the LUNARC Center for Scientific Computing[7] with 3M CPU hours allocated by Lund University and by the Swedish National Infrastructure for Computing (SNIC)[8]. For LDMX, 107 TB storage was purchased in 2018 and 125 TB dedicated to LDCS (see next paragraph) in eof 2019 for in total $ 25,000 university infrastructure funding, and another 250 TB in 2022.

The group has designed the first version of the so-called HGCROC board for the readout of the HCal. The group has also initiated and coordinated that a distributed computing system (LDCS) [68] was developed and deployed, allowing LDMX to use computing resources from several of the participating institutions, currently including at Caltech, Lund University, UCSB and SLAC.

3.7.2 Caltech

PI B. Echenard is the system coordinator for the HCal system. D. Hitlin is the Technical Coordinator of LDMX. Visitor J. Oyang, who is not supported by DOE, is making major contributions to the development effort on the HCal and the active target option. In FY21-22 we built a nineteen-layer HCal prototype (mechanics, quad-bar planes and counter mother boards) and an active target LYSO prototype and successfully tested them at the CERN T9 test beamline.

[1] https://www.lunduniversity.lu.se/
[2] https://www.fysiografen.se/en/
[3] https://kaw.wallenberg.org/en
[4] https://www.crafoord.se/en/
[5] https://kaw.wallenberg.org/en
[6] https://www.vr.se/english.html
[7] https://www.lunarc.lu.se/
[8] https://www.snic.se/

3.7.3 Fermilab

The PI N. Tran, along with other Fermilab scientists, was awarded a Fermilab LDRD award in 2020 for advancing accelerator-based dark matter experimental concepts. A fraction of the award, roughly $\frac{1}{4}$, was allocated to advance the LDMX detector concept and doing simulation studies to further the physics concept. The award is providing research fraction to N. Tran (0.1 FTE) and a postdoctoral fellow at 0.25 FTE time. Prior to the LDRD effort PI N. Tran worked on LDMX at the level of 0.1 FTE.

3.7.4 University of Minnesota

During 2020 $2,500 of University funds was invested in hardware and in undergraduate engineering labor to develop firmware. In addition $15,000 of University funding was invested in summer research assistant support for several graduate students since 2017. In 2021, $6,000 in University funding was invested in hardware to maintain schedule in prototyping.

The PI J. Mans has worked 10% of his time and graduate student T. Eichlersmith 50% of his time, on LDMX. T. Eichlersmith is supported by university funds (Teaching Assistantship).

LDMX is making use of computing and data storage resources made possible by a combination of ARRA funding and university funds, supported operationally by university funds.

3.7.5 University of California Santa Barbara

The UCSB work has been supported by university and donor funds as well as the synergy from DOE supported upgrade work for CMS. This has resulted in $ 50,000 per year on support for technical personnel and students since 2015[9].

LDMX is using the California Nano-Science Institute (CNSI) POD computing cluster at UCSB, supported by the university and CNSI.

Postdoctor V. Dutta is 15% on LDMX and is mainly funded by the DOE Energy Frontier, but about 10% of her salary is UCSB funds.

The graduate students A. Li and P. Masterson have been at 25% on LDMX and funded by UCSB.

Many undergraduate students have have worked for free and have made significant contributions to LDMX.

3.7.6 SLAC

The PI T. Nelson is working 25% on LDMX and the Project Scientist O. Moreno is working 75% on LDMX, funded by the DOE Intensity Frontier.

The Mechanical Engineer T. Lange is working 25% on LDMX as Project Engineer/Project Manager, funded internally at SLAC.

SLAC is using $ 125,000 of KA-25 funding towards developing generic DAQ infrastructure for LESA that will provide event building and run control for LDMX.

LDMX is utilizing computing resources funded by the DOE Intensity Frontier.

3.7.7 Stanford University

The PI L. Tompkins works one summer salary month from university funds for LDMX.

The group had 80% FTE of an electrical engineer from June 2020 through Feb 2021 funded by L. Tompkins start-up funds.

L. K. Bryngemark, postdoctoral scholar, leads the development and operation of LDMX distributed computing system LDCS. They started in the group as a Wallenberg Stanford Postdoctoral Fellowship from June 2019 to June 2021, on a grant from the Knut & Alice Wallenberg Foundation, and is now employed by Stanford University.

The group has had three short term (3 months) students funded by Stanford departmental funds.

[9]We acknowledgement this support from University of California Santa Barbara and The Pat and Joe Yzurdiaga Chair in Experimental Science

3.7.8 Texas Tech University

N. Gogate, research assistant, has been supported by the PI A. Whitbeck's start-up funds.

Several Texas Tech undergraduate students have been contributing thanks to support from the Honors College at the university.

3.7.9 University of Virginia

The group has two faculty members C. Group and C. Dukes each working 10% on LDMX. In addition they have a post-doc M. Solt working ∼35% and a research scientist R. Ehrlich working ∼10% on LDMX. In addition two graduate students have been contributing at the level of about ∼30%. All funded by the DOE Intensity Frontier. UVA also has ∼300k CPU hours of high-performances computing resources available per year that could be used for LDMX efforts.

3.8 Planning for Operations and Analysis

The LDMX experiment is planning for two different periods of operation, where we assume a full year of operation can provide 250 days of beam. During a pilot run of 6 months to a year, we expect to collect $2-4\times10^{14}$ electrons on target at a beam energy of 4 GeV. Soon thereafter, LCLS-II will be upgraded to an 8 GeV drive beam (LCLS-II-HE) which mitigates key backgrounds allowing for LDMX to more easily achieve its full sensitivity. During this "Phase II" of operations, LDMX would run for 2-4 more years to achieve full sensitivity, depending upon how robust the design is to higher occupancies required to compress operations with higher beam intensities.

While the LESA facility at SLAC is imagined as a multi-user facility, the mechanism by which LESA operations are funded has not yet been determined. For the time being, we assume that the LDMX experiment will be responsible for maintaining all of equipment and services that are part of the construction project, and that LESA will otherwise be operated as a user facility, responsible for operation and maintenance of LESA to bring beam to ESA and the basic infrastructure in End Station A required to support all users. This assumption excludes from LDMX responsibility the maintenance and operation of S30XL and the A-line, operation of the LESA laser at the injector, and maintenance and administration of facilities and safety systems in End Station A.

With those assumptions, the major operations costs for the experiment consist of electrical power, engineering and technician support for the LDMX apparatus, and travel costs to support shift work. There are some additional small costs anticipated for consumables and the upkeep of infrastructure specific to LDMX in End Station A. The required labor and M&S for one year of operation have been compiled in a BOE document for operations and are summarized in Table 8.

In order to control and monitor the experiment during operations, we plan to have a small control room located on the SLAC campus rather than in End Station A: we find no justification for maintaining ESA as a habitable workplace in order to operate the experiment. Furthermore, we anticipate setting up the control and monitoring to allow collaborators to take routine experimental shifts remotely. Implementing remote shifts for routine control and monitoring has been successfully demonstrated by the HPS experiment in 2021 and reduces the travel costs associated with operations be roughly a factor of two. However, it becomes even more critical then to have a team of experts staffing local operations at SLAC who together are capable of doing hands-on work on all of the detector systems. Complete coverage for all of the detector systems and management of local operations requires a set of six experts at SLAC, three of whom must travel to SLAC from remote institutions. Assuming travel and per diem for two-week expert shifts with overlaps, one can estimate the cost of staffing operations. These annual costs are also captured in the Operations BOE document and summarized in Table 8.

In addition to cost that must be carried on the Operations budget, scientific effort is also required to operate the experiment. The Operations BOE document contains a proxy estimate for this effort in the form of postdocs for expert duties and graduate students for remote shifts, and is summarized in Table 9. Finally, once the experiment produces data, there is the task of calibration, reconstruction, and analysis of the data and the associated work on simulation and software. Based on our experience producing published results from physics studies as described in Section 2, analysis teams of 3 postdocs and 4-5 students are optimal for the complexity of the apparatus and analyses. Meanwhile, we envision approximately three simultaneous

Annual LDMX Operations Cost in FY26			
Item	Cost (FY26 K$)	Contingency (FY26 K$)	Total (FY26 K$)
Electricity	404	121	525
ESA services and consumables	97	29	126
Travel	651	195	847
Technical support	427	139	567
Total Cost	1607	493	2100

Table 8: A preliminary estimate of the annual cost of operating LDMX at SLAC in FY26.

analysis thrusts: the primary dark matter search, the visible dark sector search, and the electron-nucleon measurement program. This implies a need for 9 postdocs and 12-15 students for data analysis, also shown in Table 9.

Annual Research Effort for Operations and Analysis			
Effort type	Operations (hours)	Analysis (hours)	Total (hours)
Postdoc	6000	15912	21912
Graduate Student	6000	23868	29868

Table 9: An estimate of scientific effort required annually in order to staff operations and analyze LDMX data.

4 Response to the Previous Review

We would like to thank last year's reviewers for thoughtful comments and recommendations aimed at helping us realize the LDMX experiment. In this section, we address the recommendations made in the final report, as well as key comments of the review team.

Recommendations

A. *Consider a staged approach for building the detector and reaching the desired sensitivity [by the next review].*

Prior to the FY21 DMNI status review, the collaboration held a full-day "Early Running Workshop" with the aim of understanding the physics that could be done by a smaller scale detector and the cost and schedule reductions that could be achieved within each subsystem for an apparatus of reduced size. While the major takeaways from this exercise were discussed last year, we understand the desire to examine this possibility more closely and aim to more fully illuminate the issues here.

As shown at the FY21 status review, the early physics goal for LDMX, achievable with only weeks of data and the simplest apparatus, is the electron as target (EaT) missing energy analysis described in Section 2. The EaT analysis requires only order 10^{13} EOT rather than the order 10^{15} EOT of the full experiment, which enables smaller calorimeters. In addition, the EaT analysis does not require the recoil tracker, and the DAQ and computing requirements are smaller.

Due to the exponential nature of interactions in the calorimeters, the depth required to use them as veto devices naively scales logarithmically with the number of electrons on target: to sufficiently contain electromagnetic showers, the ECal can be roughly 36 X_0 thick instead of 40 X_0 thick. Full simulation verifies this naive expectation for the main HCal also, showing that it can be at most approximately 20% shallower while still providing an effective veto for the EaT analysis. The time and cost of design and engineering dominates the schedule and budget for these devices, while the per-unit labor and M&S cost of producing the individual calorimeter planes and their attendant readout electronics is rather small once production has been established. As a result, the cost savings of reducing the depth of the calorimeters to enable only the EaT experiment with early running are approximately 2% and 2.5% of the total project cost for the ECal and HCal respectively.

While most of the depth of the ECal and main HCal must be maintained to enable the EaT analysis, studies indicate that the smaller side HCal may not be necessary. Delaying the side HCal for deployment after initial operations would reduce the total project cost by approximately 4.5%.

The potential savings from eliminating the recoil tracker are also easily estimated. First, one would eliminate production of the larger axial sensor modules, used only in the recoil tracker, and their support structure. Second, one would produce roughly half of the readout electronics relative to the full tracking system. Much like the calorimeters, the costs and effort for all of these components are dominated by design and engineering, where much of the engineering for the recoil tracker must be completed to allow it to be added later without incurring major design risks. As a result, the cost savings of eliminating the recoil tracker are on the order of 4% of the total project cost.

A knock-on effect of reducing the scale of the main detector subsystems is a reduction in DAQ hardware required. DAQ costs are overwhelmingly dominated by engineering labor, and the cost savings here are smaller, on the order of 1% or less.

The potential to reduce computing costs are more difficult to estimate, but could be significant relative to the requirement for the full experiment. Assuming an order of magnitude reduction in computing would represent a 3% reduction in the total project cost.

Taken all together, we estimate that a downsized LDMX apparatus aimed only at the EaT analysis with a few months of early operation would reduce the cost of construction by roughly 17%. It should be noted that an apparatus built with this sole intent will not be capable of addressing any other physics, most notably the electron-nucleon cross section program, which requires (at least) the recoil tracker. The potential reduction in time to completion of a minimal apparatus is smaller, perhaps a month or two, because most of work is not on the critical path. Since the staged approach requires full disassembly, establishment of production and testing for multiple subsystems, and re-assembly, installation, and re-commissioning of the detector, the time and cost of subsequently implementing an apparatus capable of collecting larger datasets would be increased by more than the initial time and cost reductions by a large factor. However, some elements such as the HCal, DAQ, and Computing can be more easily staged than the ECal and Trackers, where the process of completing these systems is more invasive and involves significant risks.

B. *Develop a detailed plan for allocation of computing resources [by the next review].*

While an estimate of the computing resources required was present for the FY21 review, we agree that it was lacking clarity, detail, and context, both with regards to the assumptions being made and the scale of LDMX computing needs relative to the scientific computing facilities to be hosted at SLAC during the relevant period.

To remedy this, the LDMX Computing Coordinator has worked with SLAC Scientific Computing personnel to more clearly define these computing needs, to develop a detailed plan to allocate the computing resources required for the experiment, and to collect collect the information required to put these needs in context. Using conservative assumptions for the trigger rate and event size, the data volumes required are easily calculated, leading to clear estimates for the storage requirements of the experiment. Meanwhile, large statistics physics studies using full simulation provide a clear understanding of the CPU and storage requirements for processing the data, as well as the requirements for generating and processing Monte Carlo samples for analysis. These estimates have been used to plan, in collaboration with Scientific Computing personnel at SLAC, the required allocation of resources for the experiment, as described in Sections 3.1.1 and 3.1.6. The required processing and storage are a very small addition to the scale of these facilities at SLAC during the proposed time period for operations and analysis, and SLAC Scientific Computing has provided the quotes necessary to budget the computing resources in the LDMX construction project and has included the purchase and support of these resources in their future plans.

As described in Section 3.1.6, the collaboration plans to augment the baseline capabilities at SLAC with a distributed computing farm, the LDMX Distributed Computing System (LDCS). This tool, which has already been used by the collaboration for large scale generation of Monte Carlo, provides useful flexibility for offsite Monte Carlo production and analysis tasks during the analysis phase of the experiment.

Finally, we have included in the project the risk that the data volumes could be significantly higher, presumably because a higher trigger rate could be required to enable the broader physics program that is still under development, including visible dark sector searches and electron-nucleon cross section measurements. This risk is incorporated in the risk register and costed in the project.

Comments

- *The risk assessment plan is only a sketch at this point. For the final design and execution plan, full risk assessments need to be provided.*

  We agree that our risk assessment as shown at the last review was only a brief sketch of some of the more obvious risks, and that a more robust framework, including a risk register, simulated outcomes, and risk based contingencies are needed in order to baseline the project. Subsequent to the FY21 review, we have adopted a risk assessment and tracking tool developed for the LSST camera project at SLAC, as described in described in Section 3.3, and have begun the process of entering and formally tracking risks to the construction project. This includes a first estimate of cost, schedule, and performance impacts, a set of mitigations included in the baseline project, and a post mitigation model of potential costs to the project, which are added as contingency to the cost estimate presented in Section 3.6.

References

[1] Basic Research Needs for Dark Matter Small Projects New Initiatives. In *Summary of the High Energy Physics Workshop on Basic Research Needs for Dark Matter Small Projects New Initiatives, October 15-18, 2018*, 2018.

[2] Philip Schuster, Natalia Toro, and Kevin Zhou. Probing invisible vector meson decays with the NA64 and LDMX experiments. *Phys. Rev. D*, 105(3):035036, 2022.

[3] P. A. R. Ade et al. Planck 2015 results. XIII. Cosmological parameters. *Astron. Astrophys.*, 594:A13, 2016.

[4] E. Izaguirre, G. Krnjaic, P. Schuster, and N. Toro. Analyzing the Discovery Potential for Light Dark Matter. *Phys. Rev. Lett.*, 115(25):251301, 2015.

[5] M. Battaglieri et al. US Cosmic Visions: New Ideas in Dark Matter 2017: Community Report. 2017.

[6] B. Batell, R. Essig, and Z. Surujon. Strong Constraints on Sub-GeV Dark Sectors from SLAC Beam Dump E137. *Phys. Rev. Lett.*, 113(17):171802, 2014.

[7] A. A. Aguilar-Arevalo et al. Dark Matter Search in Nucleon, Pion, and Electron Channels from a Proton Beam Dump with MiniBooNE. *Phys. Rev.*, D98(11):112004, 2018.

[8] D. Banerjee et al. Dark matter search in missing energy events with NA64. *Phys. Rev. Lett.*, 123(12):121801, 2019.

[9] Asher Berlin, Patrick deNiverville, Adam Ritz, Philip Schuster, and Natalia Toro. Sub-GeV dark matter production at fixed-target experiments. *Phys. Rev. D*, 102(9):095011, 2020.

[10] A. Berlin, N. Blinov, G. Krnjaic, P. Schuster, and N. Toro. Dark Matter, Millicharges, Axion and Scalar Particles, Gauge Bosons, and Other New Physics with LDMX. *Phys. Rev.*, D99(7):075001, 2019.

[11] J. L. Feng and J. Smolinsky. Impact of a resonance on thermal targets for invisible dark photon searches. *Phys. Rev.*, D96(9):095022, 2017.

[12] A. Berlin, N. Blinov, S. Gori, P. Schuster, and N. Toro. Cosmology and Accelerator Tests of Strongly Interacting Dark Matter. *Phys. Rev.*, D97(5):055033, 2018.

[13] M. Pospelov, A. Ritz, and M. B. Voloshin. Secluded WIMP Dark Matter. *Phys. Lett.*, B662:53–61, 2008.

[14] Y. Hochberg, E. Kuflik, T. Volansky, and J. G. Wacker. Mechanism for Thermal Relic Dark Matter of Strongly Interacting Massive Particles. *Phys. Rev. Lett.*, 113:171301, 2014.

[15] E. Kuflik, M. Perelstein, N. R.-L. Lorier, and Y.-D. Tsai. Elastically Decoupling Dark Matter. *Phys. Rev. Lett.*, 116(22):221302, 2016.

[16] E. Kuflik, M. Perelstein, N. R.-L. Lorier, and Y.-D. Tsai. Phenomenology of ELDER Dark Matter. *JHEP*, 08:078, 2017.

[17] Paul Frederik Depta, Marco Hufnagel, and Kai Schmidt-Hoberg. Robust cosmological constraints on axion-like particles. *JCAP*, 05:009, 2020.

[18] Pierluca Carenza, Oscar Straniero, Babette Döbrich, Maurizio Giannotti, Giuseppe Lucente, and Alessandro Mirizzi. Constraints on the coupling with photons of heavy axion-like-particles from Globular Clusters. *Phys. Lett. B*, 809:135709, 2020.

[19] A. A. Aguilar-Arevalo et al. Axion-Like Particles at Coherent CAPTAIN-Mills. 12 2021.

[20] Yu-Dai Tsai, Jonathan Assadi, Matthew Citron, Albert De Roeck, Saeid Foroughi-Abari, Gianluca Petrillo, Yun-Tse Tsai, and Jaehoon Yu. Snowmass 2021 letter of interest: Accelerator probes of millicharged particles and dark matter. 2020.

[21] Artur M. Ankowski, Alexander Friedland, Shirley Weishi Li, Omar Moreno, Philip Schuster, Natalia Toro, and Nhan Tran. Lepton-Nucleus Cross Section Measurements for DUNE with the LDMX Detector. *Phys. Rev. D*, 101(5):053004, 2020.

[22] SLAC National Accelerator Laboratory. *Electron-Nuclear Scattering Prospects at S30XL*, March 2019.

[23] Artur M. Ankowski and Alexander Friedland. Assessing the accuracy of the GENIE event generator with electron-scattering data. *Phys. Rev. D*, 102(5):053001, 2020.

[24] H. Dai et al. First measurement of the Ar$(e, e')X$ cross section at Jefferson Laboratory. *Phys. Rev. C*, 99(5):054608, 2019.

[25] H. Dai et al. First Measurement of the Ti(e, e')X Cross Section at Jefferson Lab. *Phys. Rev. C*, 98(1):014617, 2018.

[26] M. Murphy et al. Measurement of the cross sections for inclusive electron scattering in the E12-14-012 experiment at Jefferson Lab. *Phys. Rev. C*, 100(5):054606, 2019.

[27] L. Gu et al. Measurement of the Ar(e,e′ p) and Ti(e,e′ p) cross sections in Jefferson Lab Hall A. *Phys. Rev. C*, 103(3):034604, 2021.

[28] F. Hauenstein et al. Electrons for neutrinos: Addressing critical neutrino-nucleus issues. a proposal to jefferson lab pac 45, 2017.

[29] A. Ashkenazy et al. Electrons for neutrinos: Addressing critical neutrino-nucleus issues. a run group proposal resubmission to jefferson lab pac 46, 2018.

[30] W. Altmannshofer et al. The Belle II Physics Book. 2018.

[31] J. P. Lees et al. Search for Invisible Decays of a Dark Photon Produced in e^+e^- Collisions at BaBar. *Phys. Rev. Lett.*, 119(13):131804, 2017.

[32] S. Biedron and R. Van de Water. "coherent captain-mills (ccm): Lujan psr short pulse upgrade". In *"Snowmass NF09 Workshop"*, 12 2020.

[33] D. Akimov et al. First Probe of Sub-GeV Dark Matter Beyond the Cosmological Expectation with the COHERENT CsI Detector at the SNS. 10 2021.

[34] E. Izaguirre, G. Krnjaic, P. Schuster, and N. Toro. New Electron Beam-Dump Experiments to Search for MeV to few-GeV Dark Matter. *Phys. Rev.*, D88:114015, 2013.

[35] M. Battaglieri et al. Dark matter search in a Beam-Dump eXperiment (BDX) at Jefferson Lab. 2014.

[36] M. Battaglieri et al. Dark Matter Search in a Beam-Dump eXperiment (BDX) at Jefferson Lab. 2016.

[37] B. Batell, M. Pospelov, and A. Ritz. Exploring Portals to a Hidden Sector Through Fixed Targets. *Phys. Rev.*, D80:095024, 2009.

[38] P. deNiverville, M. Pospelov, and A. Ritz. Observing a light dark matter beam with neutrino experiments. *Phys. Rev.*, D84:075020, 2011.

[39] P. deNiverville, D. McKeen, and A. Ritz. Signatures of sub-GeV dark matter beams at neutrino experiments. *Phys. Rev.*, D86:035022, 2012.

[40] P. deNiverville and C. Frugiuele. Hunting sub-GeV dark matter with NOνA near detector. 2018.

[41] D. Banerjee et al. Search for vector mediator of Dark Matter production in invisible decay mode. *Phys. Rev.*, D97(7):072002, 2018.

[42] Laura Molina. Dark sector searches with na64 experiment @ cern. In *A Rainbow of Dark Sectors (Aspen Winter Conference)*, March 2021.

[43] Sergei Gninenko. Na64 status report 2020 to cern spsc. Technical report, August 2020.

[44] E. Izaguirre, G. Krnjaic, P. Schuster, and N. Toro. Testing GeV-Scale Dark Matter with Fixed-Target Missing Momentum Experiments. *Phys. Rev.*, D91(9):094026, 2015.

[45] P. deNiverville, C.-Y. Chen, M. Pospelov, and A. Ritz. Light dark matter in neutrino beams: production modelling and scattering signatures at MiniBooNE, T2K and SHiP. 2016.

[46] A. A. Aguilar-Arevalo et al. Dark Matter Search in a Proton Beam Dump with MiniBooNE. *Phys. Rev. Lett.*, 118(22):221803, 2017.

[47] Yu. M. Andreev et al. Improved exclusion limit for light dark matter from e+e- annihilation in NA64. *Phys. Rev. D*, 104(9):L091701, 2021.

[48] A. A. Aguilar-Arevalo et al. First Dark Matter Search Results From Coherent CAPTAIN-Mills. 5 2021.

[49] Belle II Collaboration. Snowmass white paper: Belle ii physics reach and plans for the next decade and beyond. Technical report, April 2022.

[50] Rouven Essig, Graham K. Giovanetti, Noah Kurinsky, Dan McKinsey, Karthik Ramanathan, Kelly Stifter, and Tien-Tien Yu. Snowmass2021 Cosmic Frontier: The landscape of low-threshold dark matter direct detection in the next decade. In *2022 Snowmass Summer Study*, 3 2022.

[51] Timon Emken. Solar reflection of light dark matter with heavy mediators. *Phys. Rev. D*, 105(6):063020, 2022.

[52] O. Abramoff et al. SENSEI: Direct-Detection Constraints on Sub-GeV Dark Matter from a Shallow Underground Run Using a Prototype Skipper-CCD. *Phys. Rev. Lett.*, 122(16):161801, 2019.

[53] R. Agnese et al. First Dark Matter Constraints from a SuperCDMS Single-Charge Sensitive Detector. *Phys. Rev. Lett.*, 121(5):051301, 2018. [erratum: Phys. Rev. Lett.122,no.6,069901(2019)].

[54] R. Essig, T. Volansky, and T.-T. Yu. New Constraints and Prospects for sub-GeV Dark Matter Scattering off Electrons in Xenon. *Phys. Rev.*, D96(4):043017, 2017.

[55] P. Agnes et al. Constraints on Sub-GeV Dark-Matter–Electron Scattering from the DarkSide-50 Experiment. *Phys. Rev. Lett.*, 121(11):111303, 2018.

[56] R. Essig, J. Mardon, and T. Volansky. Direct Detection of Sub-GeV Dark Matter. *Phys. Rev.*, D85:076007, 2012.

[57] S. Knapen, T. Lin, and K. M. Zurek. Light Dark Matter: Models and Constraints. *Phys. Rev.*, D96(11):115021, 2017.

[58] C. Dvorkin, T. Lin, and K. Schutz. Making dark matter out of light: freeze-in from plasma effects. 2019. arXiv: 1902.08623.

[59] L. J. Hall, K. Jedamzik, J. March-Russell, and S. M. West. Freeze-In Production of FIMP Dark Matter. *JHEP*, 03:080, 2010.

[60] M. A. G. Garcia and M. A. Amin. Prethermalization production of dark matter. *Phys. Rev.*, D98(10):103504, 2018.

[61] T. Åkesson et al. A High Efficiency Photon Veto for the Light Dark Matter eXperiment. *JHEP*, 04:003, 2020.

[62] Nikita Blinov, Gordan Krnjaic, and Douglas Tuckler. Characterizing Dark Matter Signals with Missing Momentum Experiments. *Phys. Rev. D*, 103(3):035030, 2021.

[63] T. Åkesson et al. Light Dark Matter eXperiment (LDMX). 2018. arXiv:1808.05219.

[64] P. Hansson A. et al. Heavy Photon Search Experiment at Jefferson Laboratory: proposal for 2014-2015 run, 2013.

[65] D. Contardo et al. The Phase-2 Upgrade of the CMS Endcap Calorimeter. Technical Report CERN-LHCC-2017-023, CMS-TDR-019, 2017.

[66] Bittware XUP-VV8 PCIe Card with Xilinx FPGA. `https://www.bittware.com/fpga/xup-vv8/`.

[67] Xiaocong Ai et al. A Common Tracking Software Project. *Comput. Softw. Big Sci.*, 6(1):8, 2022.

[68] L. K. Bryngemark et al. Building a Distributed Computing System for LDMX: Challenges of creating and operating a lightweight e-infrastructure for small-to-medium size accelerator experiments. 2021. arXiv: 2105.02977.

[69] T. Åkesson et al. A primary electron beam facility at CERN. 2019. arXiv:1905.07657.