LDMX Prototype Testbeams: S30XL and miniLDMX

November 2025

Rory O'Dwyer on behalf of the Stanford LDMX Group

HEPCAT Annual Meeting

Outline

Stanford’s LDMX Group

Elizabeth Berzin: Implementation of firmware draft into S30XL, integration of detector, analysis of S30XL detector, now works on second draft of full LDMX TS firmware and implementation of miniLDMX.

Majd Ghrear: detector design for mini-LDMX, oversight of light guide simulations (w/ Jackson Whitt), detector fabrication and testing prior to incorporation into ESA.

Layan Alsaraya: Performance of TS offline pile-up studies, analysis of S30XL test beam data, feasibility studies for S30XL/miniLDMX (cross-talk and confusion matrices, etc).

Mason Stobbe: Online monitoring system for miniLDMX testbeam.

Lauren Tompkins: PI of the Stanford LDMX group, involved in every task.

Me (Rory O’Dwyer): Analysis of 2022 testbeam, production of first draft of full LDMX TS firmware, now wholly works on HPS to finish thesis.

Physics at the LDMX

What Dark Matter Is LDMX Looking For?

If current DM is a freeze-out relic of the early universe (**the thermal model**), we can predict production probability as a function of DM mass.

SUSY predicts a weakly interacting thermal relic DM particle >>1 GeV (the WIMP miracle).

The parameter space for these models has been explored over decades and is very constrained [11],[20].

If we allow for a generic LDM model with a mediator particle, parameter space opens up.

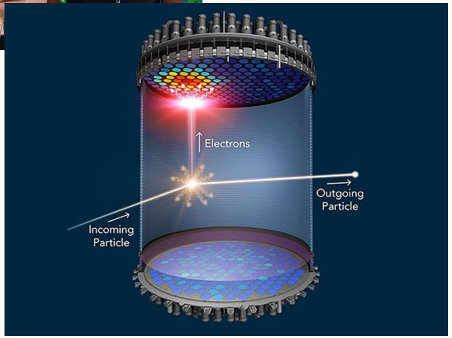

Top: The CMS detector, which (alongside 5 other CERN points) has helped tightly constrain SUSY WIMP candidates [11]. There are also many tight constraints from direct detection experiments (XENON, HESS, etc. [20])
Bottom: XENON experiment chamber looking for SUSY WIMP decays with rare nuclear interactions.

The Dark Photon And Possible Dark Sectors

The Dark Photon model (massless or otherwise) assumes another vector boson field F_{buv}.

In the rightward Lagrangian, it couples to the EM boson via a kinematic mixing term with strength ε.

It also couples to the dark sector through a traditional current term.

One of the things that the LDMX looks for is an e^- interaction on tungsten followed by A' dark photon bremsstrahlung.

$$\mathcal{L}_0 = -\frac{1}{4}F_{a\mu\nu}F_a^{\mu\nu} - \frac{1}{4}F_{b\mu\nu}F_b^{\mu\nu} - \frac{\epsilon}{2}F_{a\mu\nu}F_b^{\mu\nu}$$

Two copies of a U(1) symmetry with ε facilitating a kinematic mixing.

$$\mathcal{L} = eJ_\mu A_b^\mu + e'J'_\mu A_a^\mu$$

Interaction terms (left are dark sector, right is SM). You may have a Stueckelberg mass term for the DM photon. [12]

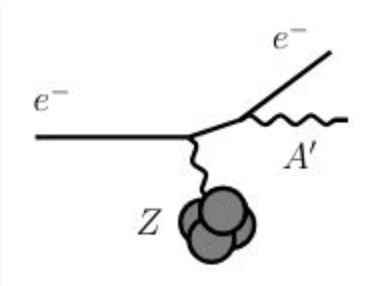

Main diagram for LDMX. Note it features a factor of ε because the actual diagram is a produced photon which mixes into the A'

Brief Look at LDMX Projected Reach [13]

This plot on the x axis shows the mass of a candidate DM particle, its interaction strength on y. The black contour are relic models, our reach is in color.

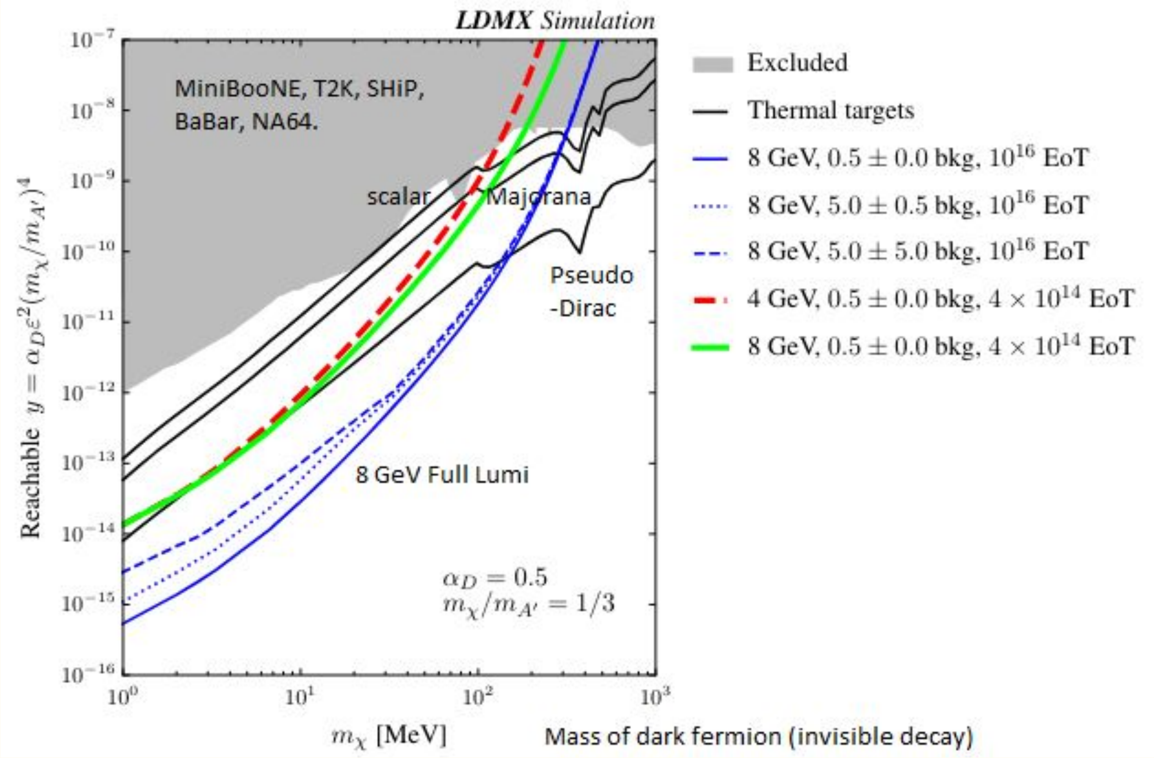

How Does The LDMX Work?

DM production is ultra-rare, and discerning individual decays needs low electron counts.

For this, we need a high energy, high repetition, low current beam electron beam: one such 8 GeV, 37 MHz, 1 pA beam is currently being serviced using LCLS II's dark current.

Every 27 ns, we need to see if an electron passed through our tracking layers, produced an A' on tungsten, and had that A' pass unseen through an ECal and HCal.

This is where our TS and its HLS firmware come in!

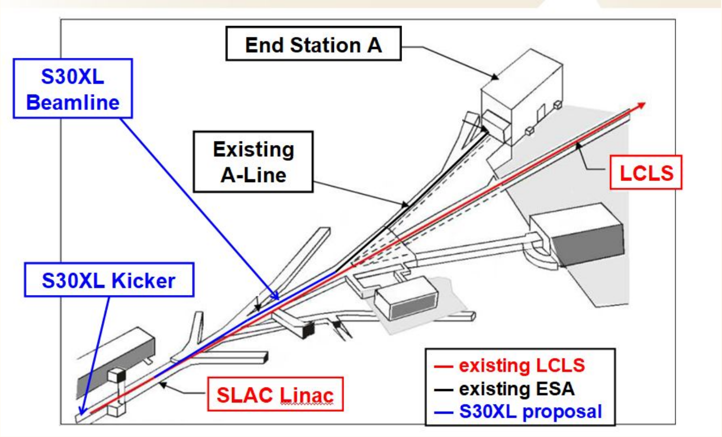

Diagram borrowed from Thomas Markiewicz (source)

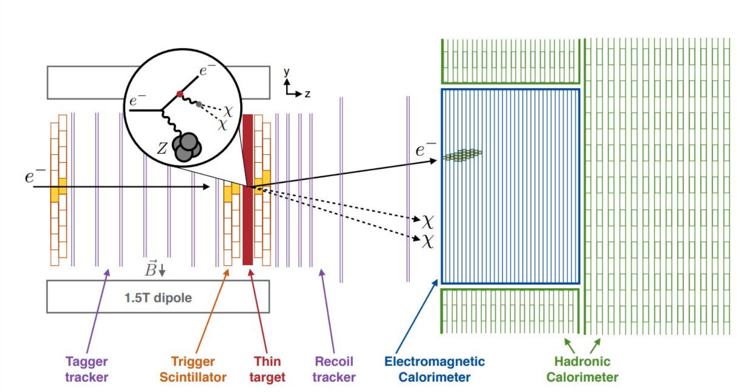

LDMX

What is the Trigger Scintillator (TS)?

The Trigger Scintillator (TS) consists of 3 arrays of 48 scintillating plastic bars.

These cover the beam spot; together (accounting for the B field) they can distinguish individual incoming electrons.

A MIP (particle with >100 MeV) deposits ~100 photoelectrons in one to form a hit.

This hits, alongside HLS firmware, are used to count e^{-}s.

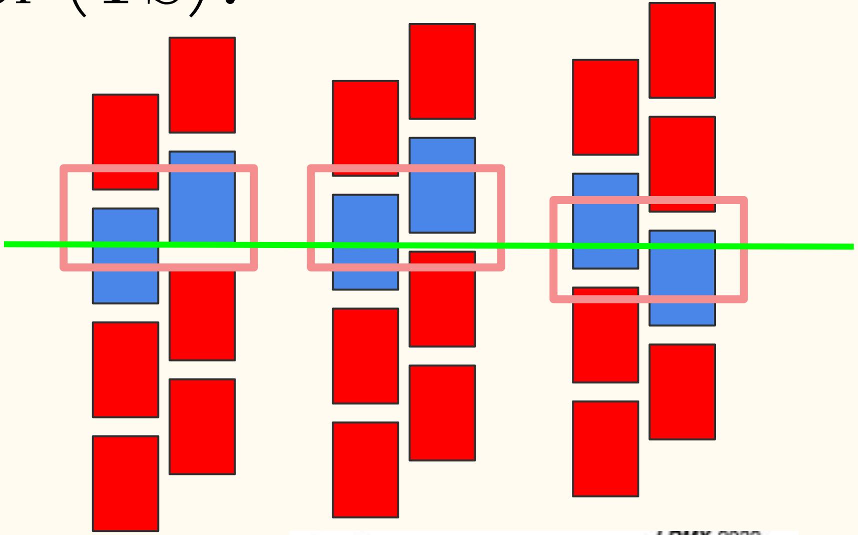

Above: Interpolated e^{-}s track
Left: PE distribution of TS bar.

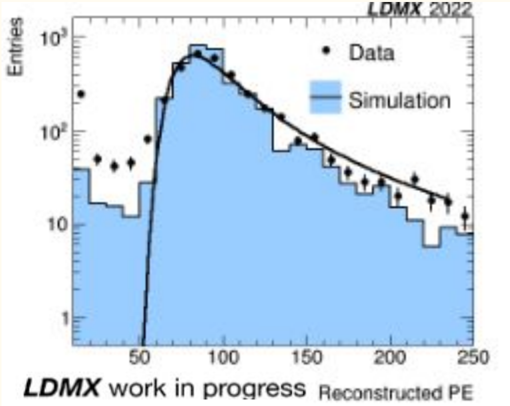

LDMX Global Trigger: An Oversimplification

Step 1: Count the number of electrons in the TS. This is down by reconstructing tracks as in Slide 5 and counting tracks.

Step 2: Count the number of electrons worth of energy in Ecal. This is done in simulation in a diagram to the right.

Step 3: f $e^-_{Ecal} < e^-_{TS}$, trigger!

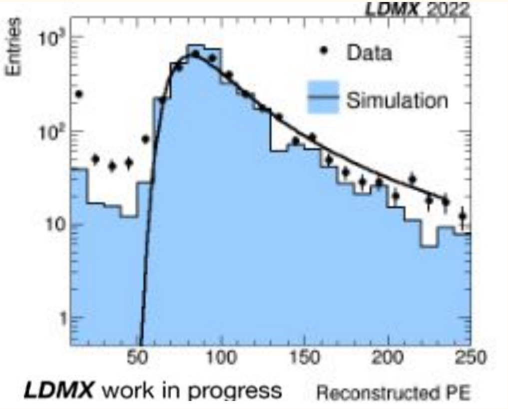

Data TS energy (PE) counts for 4 GeV electrons [4].

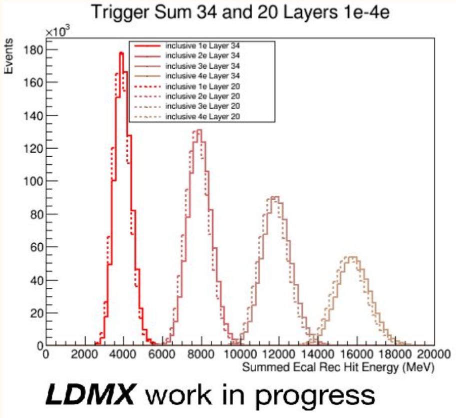

Simulated ECal energy counts for up to 4 GeV electrons. [6].

Firmware Validation: Simulation and Testbeam

Universal Test Vector Structure

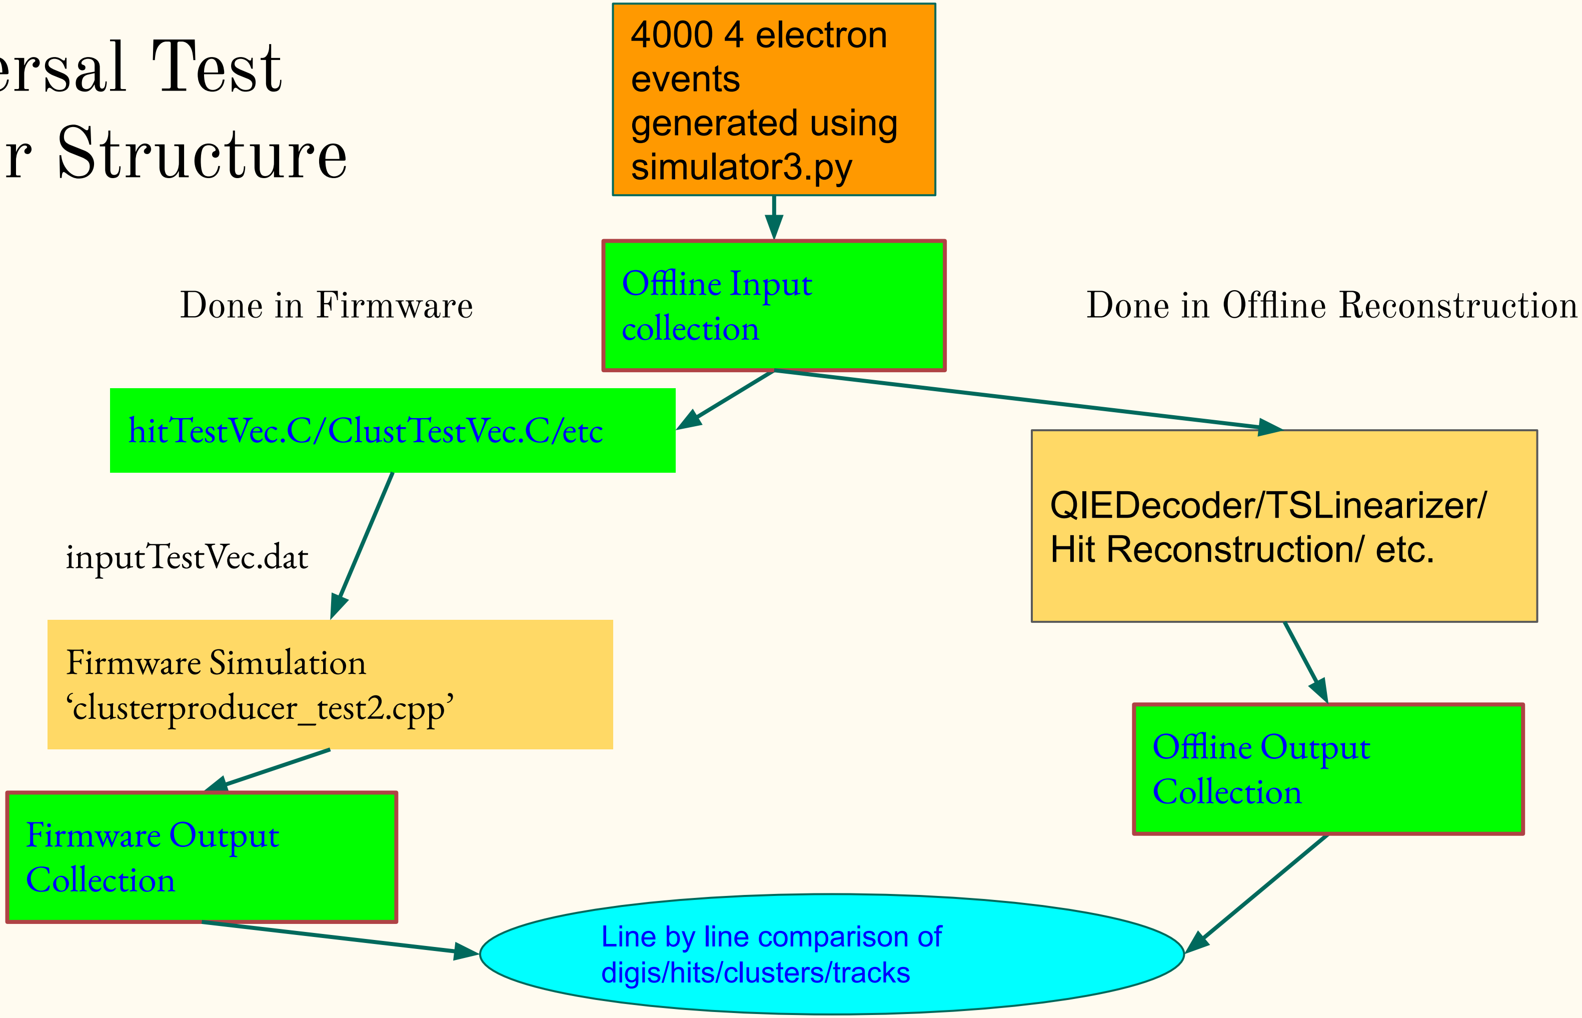

Firmware HLS Implementation and Test Vectors

The first pic shows the HLS resource usage/specs of clustering implementation.

We generate input and output objects (clusters to tracks) in MC, and feed input into firmware.

We can then event-by-event validate the firmware, checking for extra tracks or poor positions.

LDMX work in progress

Latency(cycles)	Latency(ns)	Pipelined	FF(%)	LUT(%)
5	35.000	yes	~0	7

Resource usage for clusterer as percentage (and latency), ~0 is on the order of several 1000 flip flops

```
First Firmware Cluster Centroids:
37
858
2161
2300

Now Simulated Cluster Centroids:
37
858
2161
2300
```

Hypothetical test vector pass for cluster reconstruction.

Full Feasibility Demonstration of Firmware

Here are some resource usage percentages by module:

- Hit Prod: 15% LUT, ~3 cycle interval
- Cluster Prod: 12 % LUT, ~0 interval
- Tracker: 38 % LUT, 19% FF, ~2 interval, ~127 latency

This implies that it takes 25e-9s between each input and .6e-6 s to make a decision; this is well below the 2.4e-6 s allocated to global trigger.

We use ~65% of our FPGA's resources; to get something so flat requires a lot of parallelization.

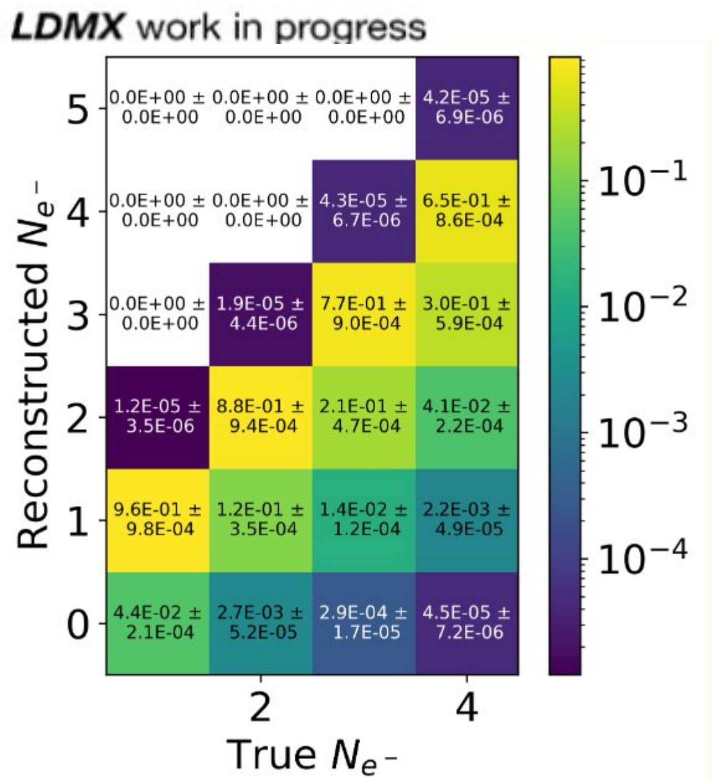

Preliminary Confusion matrix for TS system; we lose efficiency at higher input e but still have low enough fakes to so as to not overwhelm readout rate.

S30XL Testbeam

With our firmware validated in simulation, it was time to test it (and the full DAQ) at the newly commissioned S30XL beam.

This beam, originally, could have been up to .2 nA of 8 GeV, 37 MHz, beam.

We got lucky, and they installed a new cathode; we obtained a beam precisely like the one we needed (1 pA).

This gave us the opportunity to test the performance of our hit making firmware on real beam conditions.

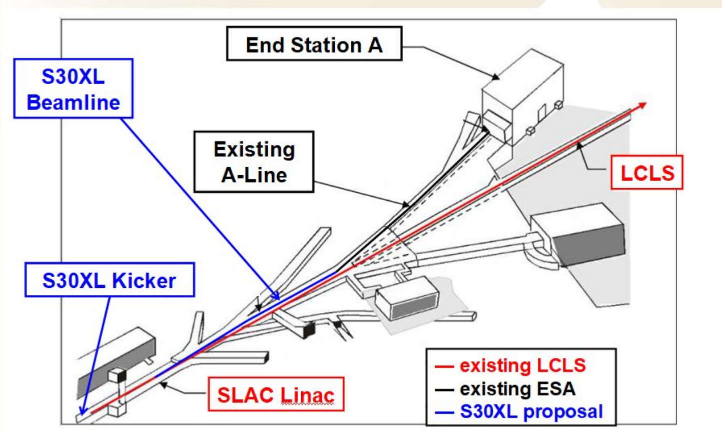

Diagram borrowed from Thomas Markiewicz (source)

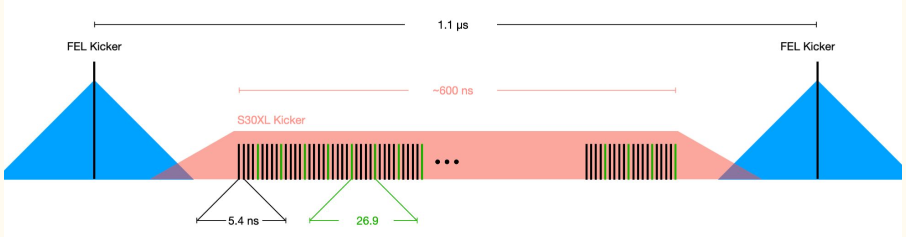

Timing profile for S30XL Kicker

Testbeam SetUp

The S30XL Teststand consists of two elements.

The upstairs DAQ system in SLAC's klystron gallery

- APx board (board w/ trigger FPGA fast control)
- zCCM (aligns our clock w/ LCLS II clock)
- Server

Downstairs (in the beam tunnel) we have a prototype TS w/ 12 serviced channels and readout electronics.

DAQ rack w/ FPGA. Blue fibers are optical; one optical fiber feeds through S30 port down

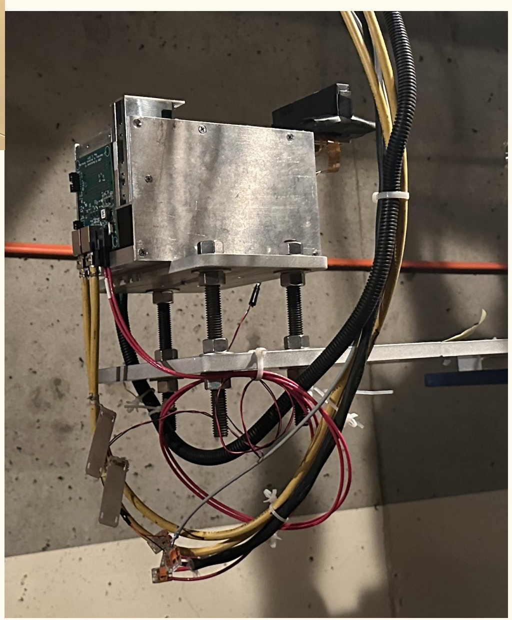

TS prototype. The 20 by 80 mm TS array is in black on top. The wires are HV, LV, optical fast control and ethernet.

Measured Charge

Plots courtesy of Elizabeth Berzin.

Plots show total charge integrated over full kicker flat-top window.

Peaks corresponding to individual electrons are visible in channels where we see the most charge.

- MIP response is ~100 PEs.

Rough dark current estimate:

- ~6 e- per 810 ns window: ~ 1 pA
- May be an underestimate if a portion of the beam is missing the instrumented area.

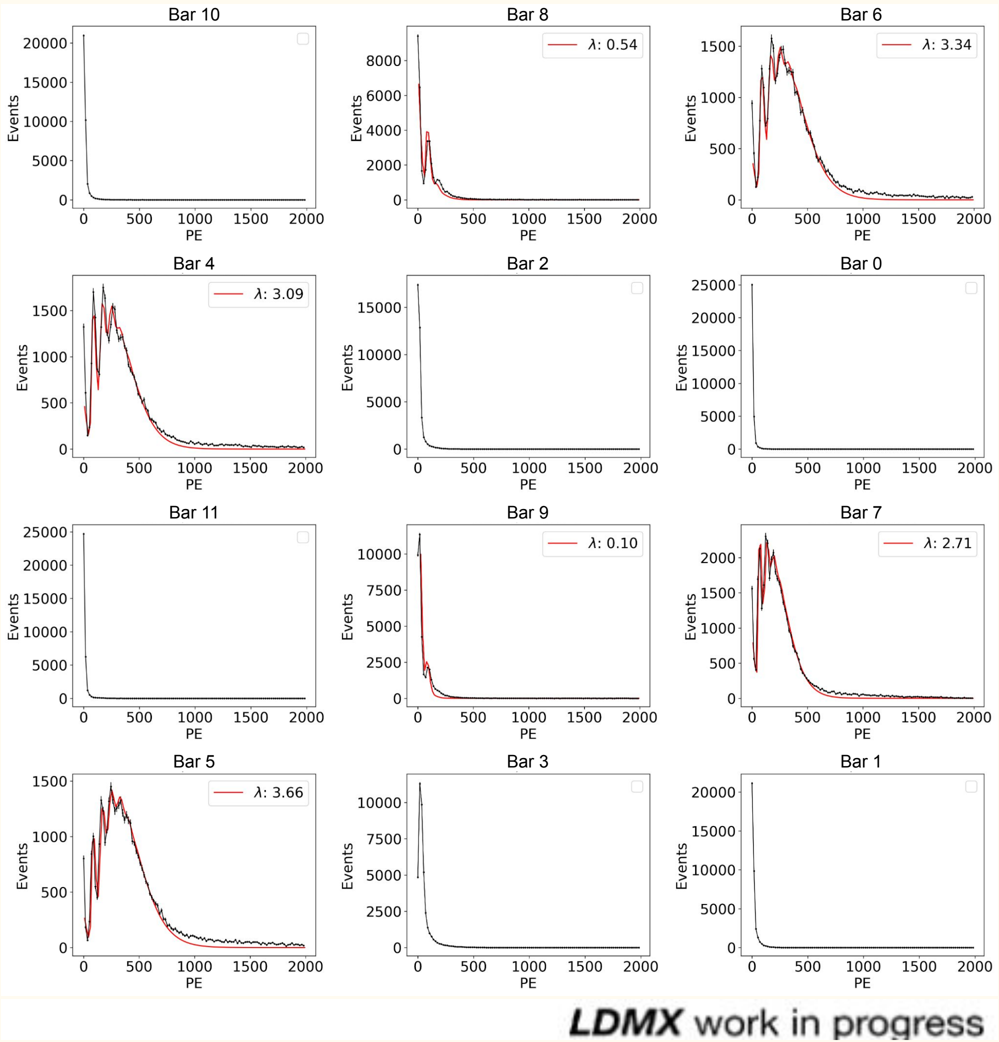

Event Visualization

Plots courtesy of Elizabeth Berzin.

The kicker window is ~25 clock cycles or 625 ns.

You can see here we get roughly 1 every few clock cycles.

The decay width is also ~1 clock cycle.

The beam is fit extremely well by a poisson distribution around 6 electrons per kicker window (1 pA).

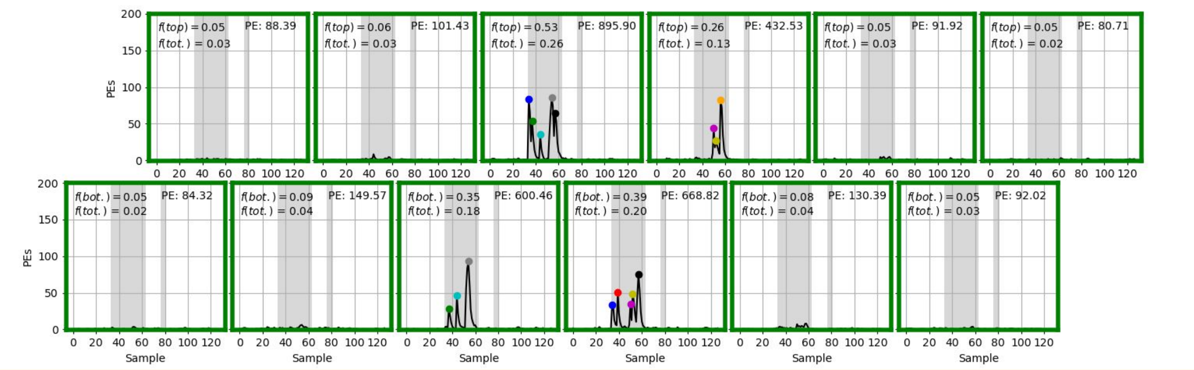

LDMX work in progress

Top: Electron pulse profiles, expressed in time. Shaded region is kicker pulse.
Right: Histogram of number of electrons in kicker window, alongside fit..

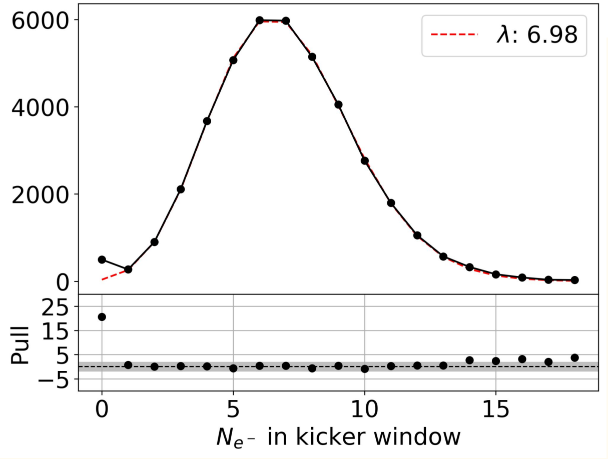

LDMX work in progress

miniLDMX: Design, Implementation, and Testing

What is miniLDMX?

miniLDMX is our collaboration's first implementation of every prototype detector subsystem in ESA with realistic DAQ infrastructure.

We aim to incorporate every subsystem into DAQ firmware, as well as most into trigger.

The full 3 layer TS will be tested (with its LUT based tracking).

This all must occur prior to the December shutdown of LCLSII and the subsequent year long beam upgrades.

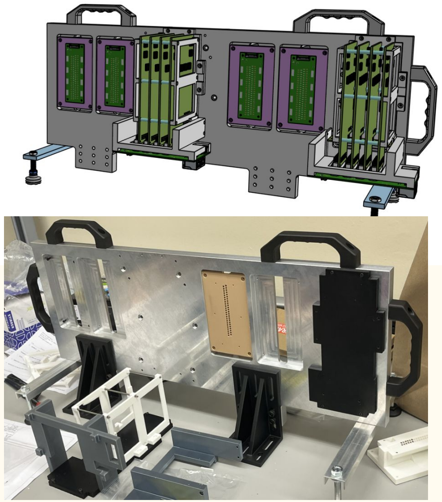

TS and tracker manifold in CAD and realized (thanks to Majd for design and pic)

1T magnet located in ESA.

TS Light Guides Simulation and Manufacture

Our TS requires light guides to reduce the activated material to roughly the target's size.

Fabricating these with optical glue is finicky, we can only tolerate 80% reduction.

At the moment, we are still trying to improve the light yield fraction at the interface of light guide to scintillator.

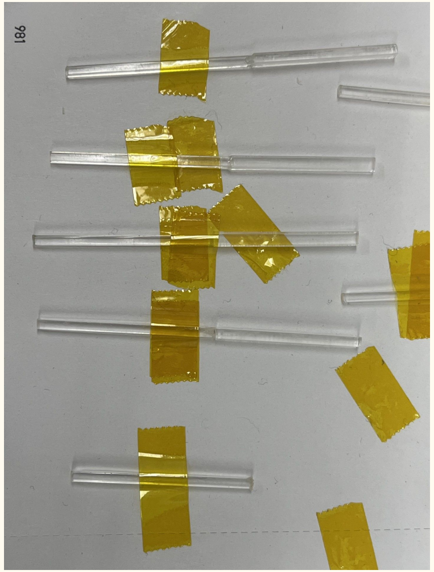

Glued acrylic pieces; the thinner piece corresponds to a a light pipe. The thicker piece will be substituted with transparent scintillator.

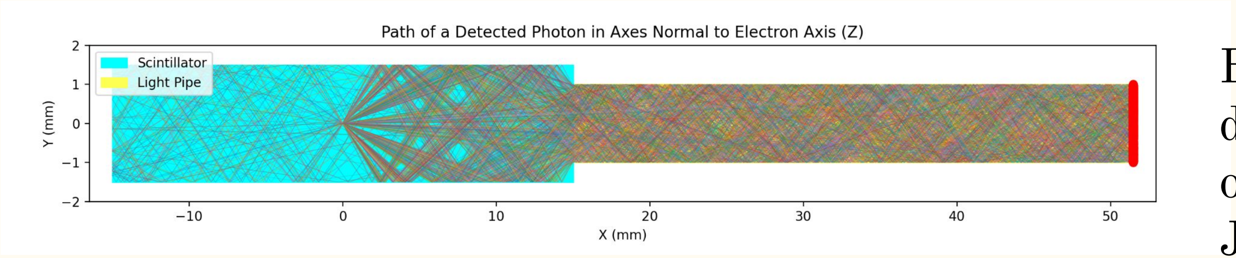

Ray tracing studies to determine light yield with optical gluing. Thanks to Jackson Whitt.

Online Monitoring of Sub-detector Systems

We have to develop tools to detect things like gradual voltage creep due to ionizing radiation damage, increasing temperatures, etc.

These tools have to communicate directly with the DAQ wires.

Mason developed a number of running histograms.

This one logs the PE counts from runs passing a trigger (possibly random) to allow us to detect dead channels and/or tune hit thresholds.

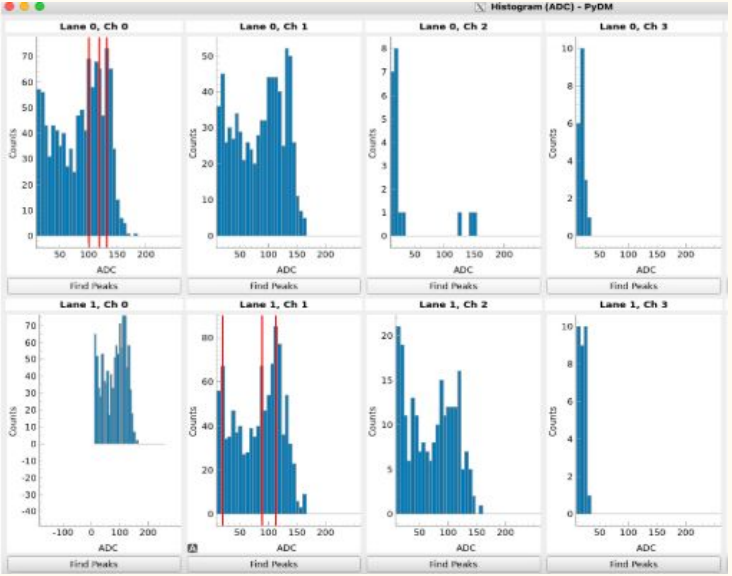

Online monitoring plot histogramming bar PE counts from test vector data taken at S30XL. It is built into the same pyrogue language as the DAQ and Trigger infrastructure at ESA

Testing 2nd Draft HLS Algorithms.

As previous talks mentioned, we use test vectors generated from S30XL and offline reconstruction to evaluate our HLS firmware.

We see very good agreement between HLS and offline, with an efficiency loss in higher luminosities.

This loss is consistent with what we've seen from earlier studies; its better than overcounting (which could saturate our trigger with fake rate).

We are working on better algorithms in HLS to reduce the efficiency hit from pile-up.

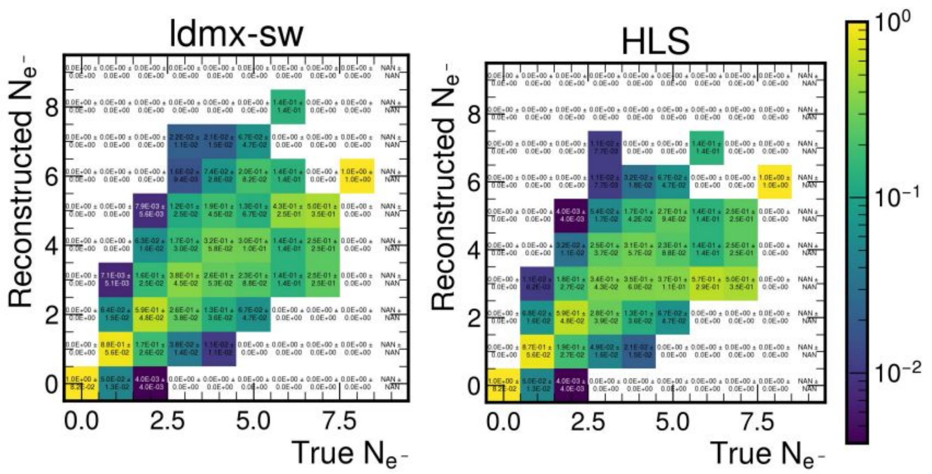

Offline reconstruction and HLS confusion matrix comparison (thanks to Elizabeth).

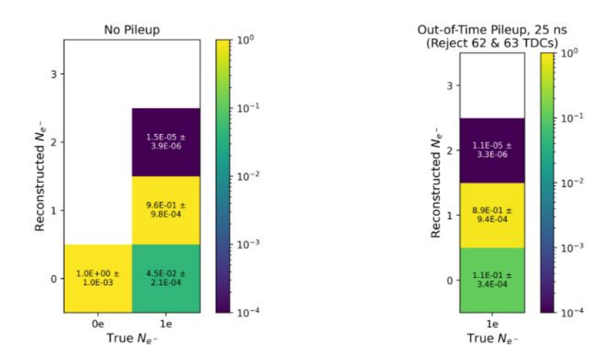

7% loss in efficiency from pile-up. (thanks Layan)

Conclusion

Conclusion

In intensity frontier HEP experiments, like the LDMX, one often must make a trigger decision at a rate of several MHz.

There are really only one general purpose option for this; building your trigger using firmware on an FPGA and similarly designing your detector to communicate with a larger DAQ.

Previous talks went over the FPGA aspects, this talk focused on aspects of detector construction/integration, testing, online monitoring etc.

The miniLDMX hopes to take some data in late December. This is one of the first steps to taking physics runs with LDMX in the coming years!

Works Cited

[1].https://arxiv.org/abs/2203.08324

[2].https://arxiv.org/abs/1808.05219

[3].https://accelconf.web.cern.ch/e04/papers/TUPLT164.pdf

[4].https://indico.slac.stanford.edu/event/7373/contributions/4969/attachments/2296/6497/collab_meeting_simps_11722.pdf (Alic Two Body SIMP)

[5].https://indico.slac.stanford.edu/event/8165/contributions/8145/attachments/3869/10488/HPS_3_Body_SIMP_Update.pdf (Tom IDM)

[6]. Courtesy of Sean Littleton

[7].https://arxiv.org/abs/2211.00764

[8].https://link.springer.com/article/10.1007/s41781-021-00065-z

[9].https://indico.cern.ch/event/1103637/contributions/4821838/attachments/2453795/4205231/CTD_31052022.pdf

[10].https://indico.fnal.gov/event/62016/contributions/278833/subcontributions/9132/attachments/174151/236024/Visibles.pdf

[11].https://www.frontiersin.org/articles/10.3389/fphy.2019.00075/full

[12].https://arxiv.org/pdf/2005.01515

[13].https://arxiv.org/pdf/2308.15173

[14].https://journals.aps.org/prd/abstract/10.1103/PhysRevD.108.012015

[15]. Courtesy LK of the LDMX collaboration

[16]. https://arxiv.org/abs/1707.08966

[17].https://indico.jlab.org/event/411/contributions/6727/attachments/5544/6882/ColabMeeting_Nov2020_ResSearch.pdf

[18].https://journals.aps.org/prd/pdf/10.1103/PhysRevD.108.012015

[19].https://arxiv.org/pdf/1801.05805

[20].https://arxiv.org/abs/1609.0673

[21].https://indico.cern.ch/event/1188759/contributions/5230349/attachments/2621850/4533261/bravoUCLADM23.pdf

[22].https://confluence.slac.stanford.edu/display/hpsg/Publications+and+Presentations