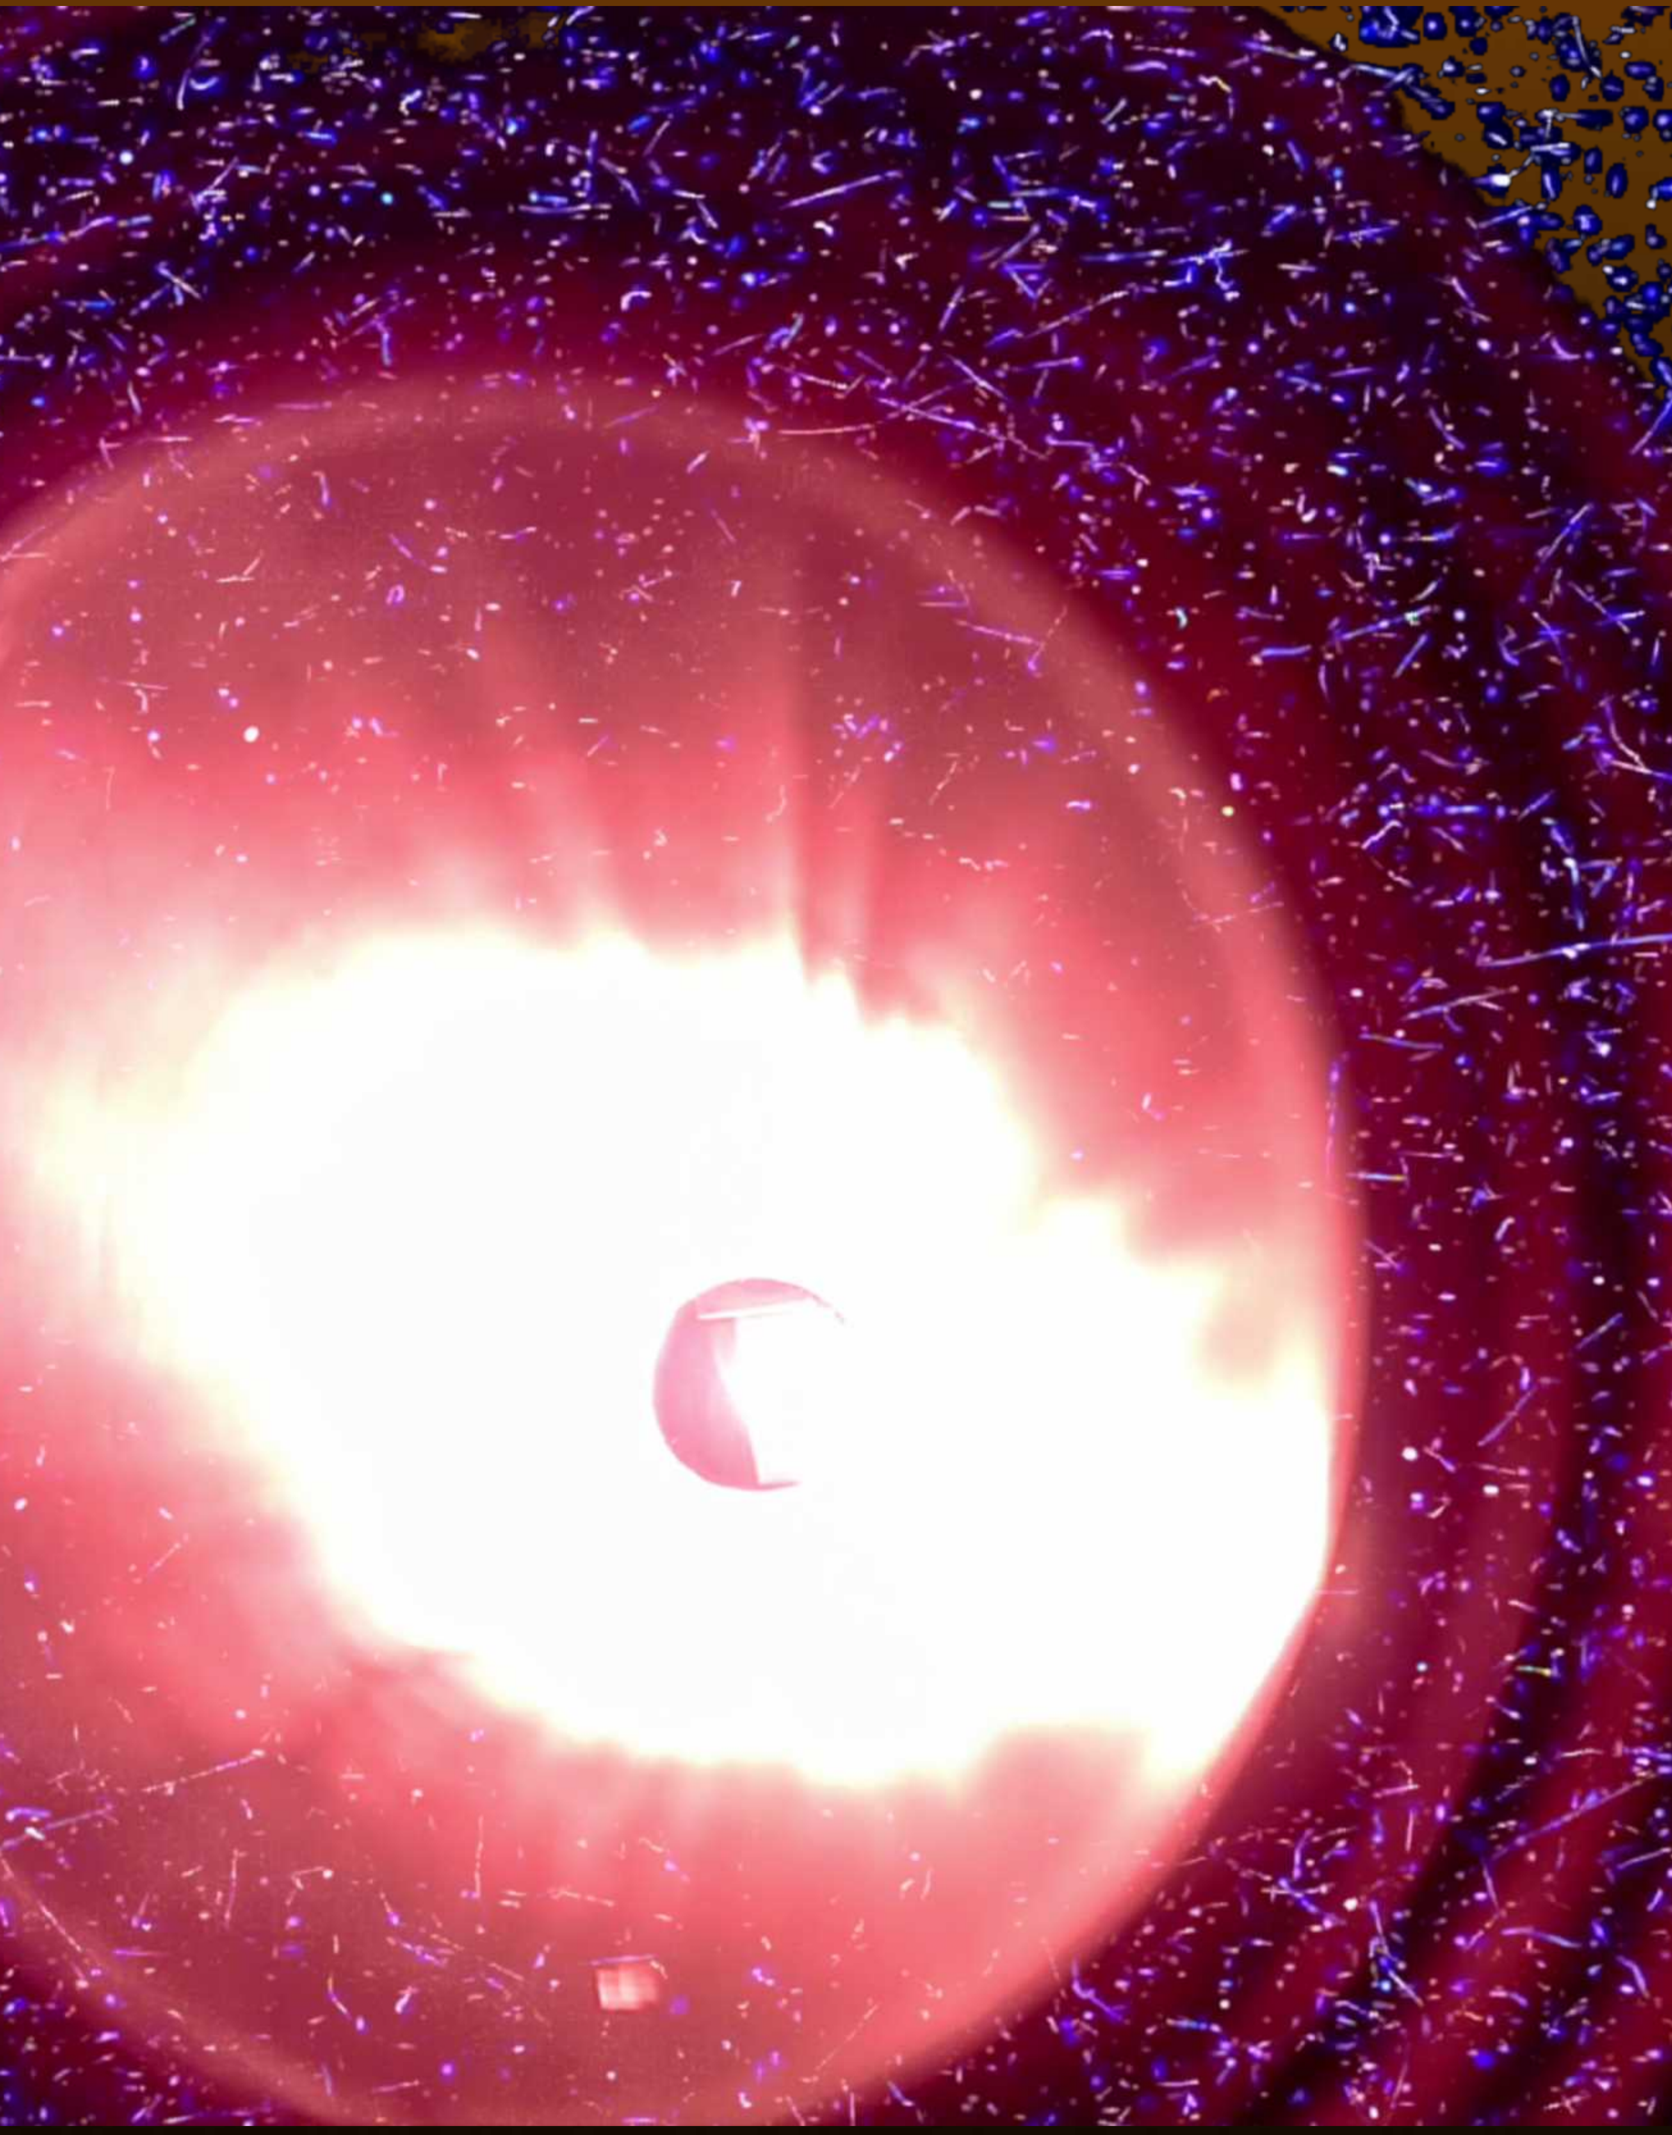

Annual Report 2025 CENTER FOR EXPERIMENTAL NUCLEAR PHYSICS AND ASTROPHYSICS University of Washington

CENPA

Center for Experimental Nuclear Physics and Astrophysics

CELEBRATING KATRIN
KARLSRUHE TRITIUM NEUTRINO EXPERIMENT
T_2
1000 DAYS

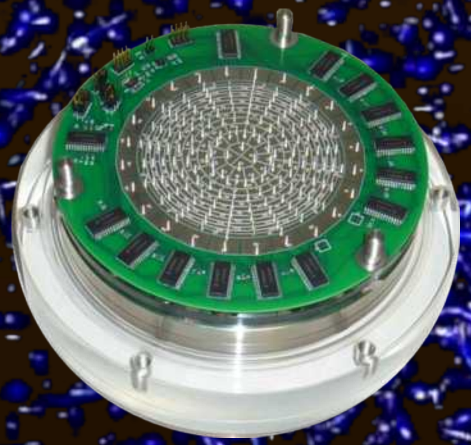

Annual Report 2025

University of Washington

CENPA ANNUAL REPORT 2025

Center for Experimental Nuclear Physics and Astrophysics
University of Washington
April 1, 2026

Sponsored in part by the United States Department of Energy
under Grant #DE-FG02-97ER41020.

Cover photos: Back and front cover overlay: CENPA's Van de Graaff accelerator is used to produce a variety of charged particle beams. In November-December 2025, the LEGEND group created a 41Ar gas target on the L30 beamline. To ensure the beam was centered, a phosphorescent screen was inserted at the target and viewed through a glass viewport via a webcam. From the video, a 20-second long exposure image was created, showing not only the bright red light from the beam's impact on the target, but also showing a large number of streaks created by radiation impacting the CCD sensor of the webcam. (Photo by Clint Wiseman). Front cover: The KATRIN experiment completed its planned 1000 days of running in December 2024. Pictured is the CENPA-built detector system with faculty member Hamish Robertson (Right photo). Graduate student Laura Bodine (Top left photo). Eric Martin (left) and Sanshiro Enomoto (Bottom left photo) assemble the vacuum and postacceleration chamber for the KATRIN detector system. The detector test assembly is at the bottom right.

CENPA Group Photo 2025

Standing: Anika Paudel, Angela Lawson, Nicholas Buzinsky, Adam Yusupov, Brooke Yommarath, Stephen Kuenstner, Eric Hsu, Edwin Solis, Yuan-Ru Lin, James Sinnis, Nicolas Oaxaca, Cyrus Goodman, Omar Beesley, Charles Hanretty, Ian Paulson, Tim Van Wechel, Samuel Bellman, David Peterson, Dana Jones, Shoshana Apple, George Marshall, Luciano Malavasi, Sanshiro Enomoto, Bradley Taylor, Kieran Enzian, Michael Huehn, Peter Kammel, Heather Harrington, Matthew Kallander, Brittney Dodson, Nathan Miedema, Mason Newman
Seated: Ali Kurmus, Meg Wynne, CJ Nave, Claudio Savarese, Alvaro Chavarria, Jeff Dillon, Jason Detwiler, Clint Wiseman, Patrick Schwendimann, David Hertzog, Hamish Robertson, Jens Gundlach, Gary Holman, Ryan Roehnelt, Paul Kolbeck, Madeleine Carhart

INTRODUCTION

The Nuclear Physics Lab at the University of Washington was established in 1948. In 1998, it evolved into CENPA, the Center for Experimental Nuclear Physics and Astrophysics. CENPA is the institutional home for a broad program of research in nuclear physics and related fields. Research activities — with an emphasis on fundamental symmetries and neutrinos — are conducted locally and at remote sites. In neutrino physics, we concentrate on the pursuit of the direct measurement of the neutrino mass. CENPA has been a lead US institution in KATRIN since its inception and is currently driving the local Project 8 experiment. Both study the tritium beta decay spectrum, whose end-point shape is sensitive to the neutrino mass. We also study neutrinoless double beta decay as a collaborating institution in the ^{76}Ge based LEGEND experiment, which derives from our extensive experience in the recently completed MAJORANA effort. Looking beyond the ton scale, we continue to explore a next-generation R&D effort with the SELENA experiment. The Muon Physics group has been a major player in the high-precision measurement of the muon's anomalous magnetic moment at Fermilab. The experiment is complete and the last results were recently published. This group's attention is now strongly focused on the PIONEER rare pion decay experiment that will measure lepton flavor universality to unprecedented levels in its first phase. Our local fundamental symmetries program also includes an experiment using our Tandem Van de Graaff (VdG) accelerator to measure the Fierz interference in ^{6}He decay, using the CRES technique developed by Project 8. This year saw considerable development of controls and operational procedures for the VdG, and the delivery of a variety of beams to PIONEER, LEGEND, and Project 8 for instrumentation measurements.

In addition to the research directly supported by DOE's Office of Nuclear Physics through the CENPA core grant, other important programs are located at CENPA, forming a broader intellectual center with valuable synergies. The "Gravity" group carries out, with both DOE and NSF support, studies of the weak and strong Equivalence Principles, fundamental precepts of General Relativity, as well as searches for non-Newtonian forces such as those predicted by theories with extra dimensions. In addition, they participate in LIGO with unique instrumentation that aids the functioning and stabilization of the interferometers. The DOE Office of High Energy Physics supports the unique ADMX axion search experiment. The NSF has supported the DAMIC experiment that looks for light dark matter.

CENPA is home to a large number of faculty, research faculty, postdoctoral scholars, graduate students, and undergraduate students. The core professional engineering and technical staff provide diverse capabilities and skills such as state-of-the-art detector development, fabrication of custom electronics, large-scale computing, and design engineering. New advancements, capabilities, and ideas are regularly shared at seminars by CENPA members and visitors alike.

Personnel updates

Faculty: Research Professor Peter Kammel obtained Emeritus Professor status following his retirement. He continues to lead many critical teams within the PIONEER collaboration and co-supervises students and postdocs in the group.

Staff: Research Scientist Erik Swanson retired after 48 years filled with exceptional scientific contributions to many programs at CENPA, including, but not limited to, the Gravity, Muon g-2, and the He-6 efforts. The ADMX group hired Research Scientist On Kim and the Gravity group promoted Michael Ross to the position of Research Scientist.

Postdocs: Brynn MacCoy (Muon g-2), Svende Braun (Muon g-2), Conner Gettings (Gravity), Heng Lin (DAMIC/Selena), Yu-Hao Sun (KATRIN), and Michelangelo Traina (DAMIC/Selena) have left the laboratory. Postdoctoral scholars at present (January 2026) include Nick Buzinsky (He-6), Patrick Schwendimann (PIONEER), Louis Varriano (LEGEND), Josh LaBounty (PIONEER), and newer arrivals include George Marshall (LEGEND), Stephen Kuenstner (Project 8).

Graduate Students: Across all projects, CENPA typically has between 18 and 20 full graduate students. Ph.D.s awarded in the last year include:

- X. Ni; TopmetalSe: Development of a CMOS Imager for Selena
- A. Hostiuc; Search for Neutrinoless Double Beta Decay Using the Full MAJORANA DEMONSTRATOR Dataset
- S. Borden; Quasi-Background-Free Neutrinoless Double-Beta Decay Searches with LEGEND: Statistical Methods and Cryogenic SiPM Characterization
- M. Guzzetti; Designing, Operating and Analyzing: The Quest for Axion Dark Matter with ADMX
- G. Song; The CAGE Scanner: Development of a Novel Surface Event Rejection Technique in High Purity Germanium Detectors

Recognitions: David Hertzog was elected to the National Academy of Sciences and the Washington State Academy of Science. Gray Rybka was elected to Fellowship in the American Physical Society. Jason Detwiler was elected Co-Spokesperson of the LEGEND experiment.

Recent Accomplishments

- KATRIN has successfully achieved its goal of 1,000 days of beta-spectrum measurements. The detector system provided by CENPA has performed reliably throughout the entire measurement period, and the data-processing chain developed at CENPA has processed all KATRIN raw data at CENPA and produced fitting-ready spectra for the collaboration.
- Based on the analysis of 259 days of data, KATRIN has published a result in Science setting an upper limit of 0.45 eV on the neutrino mass.
- Using the same data set, KATRIN has also conducted a search for sterile neutrinos. This analysis excludes the majority of previously reported positive claims, including those associated with the recently discussed gallium anomaly, and has been published in Nature.

- Project 8 has developed and published the novel dynamic evaporative cooling concept for production of atomic tritium at mK temperatures.
- The Project 8 Cavity CRES Apparatus (CCA) for demonstrating Cyclotron Radiation Emission Spectroscopy in a resonant 26-GHz microwave cavity has been completed and installed in our 1-T ex-MRI magnet. Commissioning is in progress.
- Project 8 has completed the design and construction of an ECR (electron cyclotron resonance) source for the production of atomic hydrogen. The prototype now assembled at Indiana University is undergoing commissioning.
- The LEGEND-200 detector published its first results, providing leading limits on NDBD in 76Ge, including numerous contributions from our CENPA team. The detector array was redeployed in 2025 and is currently taking data with higher mass and lower background.
- The LEGEND-1000 project was awarded NSF design funds in 2025, with CENPA in charge of the argon lock systems and the analysis software framework, both of which are well underway and slated for preliminary and final design reviews in the coming year. LEGEND-1000 also completed a successful DOE CD1 review in November.
- The Muon $g-2$ Collaboration published its final results with a precision on the measurement of the muon's anomalous magnet moment at 127 ppb. The result is fully consistent with our previous publications, but the SM theory has shifted and become less certain at this time.
- The PIONEER Experiment design has evolved considerably, with leadership by UW in many areas of hardware and simulation/physics software. The CENPA Van de Graaff was used on numerous occasions to evaluate LGAD sensor strips. A highly successful test-beam run at PSI demonstrated the excellent performance of full-sized, tapered LYSO crystals, which have become the baseline choice for PIONEER.
- The He6-CRES experiment successfully implemented a number of improvements and completed its first ^{19}Ne data-collection campaign.
- The Selena R&D program demonstrated excellent charge response of the CMOS sensors designed at CENPA. Hybrid amorphous selenium/CMOS devices fabricated with these sensors demonstrated a remarkable spatial resolution of 15 micrometers for MeV beta tracks, an important step toward the project's technical goals.
- The DAMIC-M program published world-leading results in the search for sub-GeV dark matter with data from two prototype CCD modules operated in the Modane Underground Laboratory (LSM). These results rule out well-motivated hidden-sector models of dark matter for the first time.
- The DAMIC-M team at CENPA completed the production of the final 28 CCD modules for the experiment. The modules were delivered to the Modane Underground Laboratory, where testing confirmed their excellent performance. The modules now await installation in the cryostat to complete the detector in 2026.
- ADMX has been operating in the range from 1.1 to 1.3 GHz, thereby improving sensitivity to DFSZ axions.

- The Eöt-Wash group completed a science campaign with its upgraded rotating torsion balance. The results set leading limits on long-range interactions between normal matter and dark matter, B-L coupled ultra-light dark matter, and equivalence principle violations towards the sun.
- LIGO has procured and began testing six ultra-precise inertial rotation sensors developed by the Eöt-Wash group which will improve the seismic isolation system of LIGO by an order of magnitude and its low frequency noise significantly.
- The Van de Graaff accelerator team delivered beams for the He6-CRES beta decay program (production of ^{19}Ne and ^{79}Kr), PIONEER LGAD detector development (Rutherford backscattering), and LEGEND background studies (^{41}Ar and ^{56}Co), all while continuing to perform major maintenance and upgrades to the old control system.

As always, we encourage outside applications for the use of our facilities. As a reference for potential users, the table on the following page lists the capabilities of our accelerator. For further information, please contact Gary Holman, Associate Director (holman@uw.edu) or Brittney Dodson, Research Engineer (wynn8@uw.edu) CENPA, Box 354290, University of Washington, Seattle, WA 98195; (206) 543 4080. Further information is also available on our web page: `http://www.npl.washington.edu`. We close this introduction with a reminder that the articles in this report describe work in progress and are not to be regarded as publications nor to be quoted without permission of the authors. In each article the names of the investigators are listed alphabetically, with the primary author underlined in the case of multiple authors, to whom inquiries should be addressed.

David Hertzog, Director
Gary Holman, Associate Director
Clint Wiseman and Jason Detwiler, Editors

TANDEM VAN DE GRAAFF ACCELERATOR

Our tandem accelerator facility is centered around a High Voltage Engineering Corporation Model FN purchased in 1966 with NSF funds, with operation funded primarily by the U.S. Department of Energy. See W. G. Weitkamp and F. H. Schmidt, "The University of Washington Three Stage Van de Graaff Accelerator," *Nucl. Instrum. Methods* **122**, 65 (1974). The tandem runs the majority of its time using the Direct Extraction Ion Source (DEIS) to produce up to 9 MeV proton and 18 MeV deuterium beams for the He6-CRES experiment. The Model 860 Sputter Ion Source is available for heavier ion beams.

Some Available Energy Analyzed Beams

Ion	Max. Current (particle μA)	Max. Energy (MeV)	Ion Source
^{1}H or ^{2}H	50	18	DEIS or 860
^{3}He or ^{4}He	2	27	Double Charge-Exchange Source
^{3}He or ^{4}He	30	7.5	Tandem Terminal Source
^{6}Li or ^{7}Li	1	36	860
^{11}B	5	54	860
^{12}C or ^{13}C	10	63	860
*^{14}N	1	63	DEIS or 860
^{16}O or ^{18}O	10	72	DEIS or 860
F	10	72	DEIS or 860
* Ca	0.5	99	860
Ni	0.2	99	860
I	0.001	108	860

*Negative ion is the hydride, dihydride, or trihydride.

Several additional ion species are available including the following: Mg, Al, Si, P, S, Cl, Fe, Cu, Ge, Se, Br and Ag. Less common isotopes may be generated from enriched material. We have recently produced rare-isotope calibration sources such as ^{56}Co and ^{79}Kr, and provide proton beams at variable energies as low as 1 MeV. We maintain capabilities to create pulsed ("chopped") ion beams for experimental background rejection, including a separated beam of 15-MeV ^{8}B at 6 particles/second. We maintain hardware for a terminal ion source with its non-inclined beam tube, which enable high intensity beams of hydrogen and helium isotopes from 100 keV to 7.5 MeV. We are capable of production of positive ion beams of the noble gases He, Ne, Ar, and Kr at ion source energies from 10 keV to 100 keV for implantation, in particular the rare isotopes ^{21}Ne and ^{36}Ar.

Contents

1 Neutrino Research

LEGEND

1.1 Commissioning and status of the LEGEND search for $0\nu\beta\beta$ in Ge-76

S. Bellman, S. Borden, J. Detwiler, Y.-R. Lin, C. J. Nave, D. A. Peterson, R. Roehnelt, C. Savarese, G. R. Song, L. Varriano, T. D. Van Wechel, and C. Wiseman

LEGEND is a next-generation neutrinoless double-beta ($0\nu\beta\beta$) decay search in ^{76}Ge. By combining the technological expertise and experience from the MAJORANA DEMONSTRATOR and GERDA experiments, LEGEND is expected to reach a design sensitivity two orders of magnitude greater than its predecessors.

The phased program advances from the intermediate-scale LEGEND-200 experiment to the future LEGEND-1000 experiment. LEGEND-200 is currently collecting physics data with ∼140 kg of enriched HPGe detectors deployed at the Laboratori Nazionali del Gran Sasso (LNGS) in Italy. The first physics results from LEGEND-200 have been presented in Neutrino 2024 and have been accepted for publication by *Physical Review Letters* (`https://doi.org/10.1103/25tk-nctn`) on 29 September 2025.

LEGEND-1000 is a leading candidate for a future US-led tonne-scale $0\nu\beta\beta$ experiment. It uses advances from the ^{76}Ge program, through LEGEND-200, to mitigate risks by demonstrating the scalability of the array and the achievement of background goals. It will yield an additional order of magnitude sensitivity to the $0\nu\beta\beta$ half-life, covering the neutrino mass parameter space allowed by the inverted ordering of neutrino masses.

Members of the LEGEND group at CENPA play a leading role in the development of the analysis framework for the project, which is actively tested on the data collected with our local test stands. Professor Detwiler is a member of the LEGEND Steering Committee, an Analysis Co-Coordinator, and the lead of the Analysis Framework task group. The UW group is responsible for the development of the Lock system for LEGEND-1000, a critical component for the staged deployment of the germanium detector array in the next experiment. Locally, we are pursuing R&D efforts to understand the surface response of HPGe detectors and to characterize the light detection properties of SiPM detectors used in the liquid argon veto. These efforts are described in the following sections.

1.2 LEGEND-1000 Lock System

S. Bellman, J. Detwiler, M. Huehn, R. Roehnelt, C. Savarese, and L. Varriano

The experience of the GERDA, MAJORANA and LEGEND-200 experiments has guided the design and development of the lock system. Previously, in LEGEND-200, the entire detector array was able to be raised and lowered within a single lock. However, since LEGEND-1000 is a much larger and heavier stage, this is not possible. LEGEND-1000 will deploy individual strings of eight detectors that will be suspended from their own flange and sealed on their own port on the cryostat lid, then submerged in a bath of underground liquid argon UGAr. The underground liquid argon shall maintain purity levels of less than 0.1 ppm oxygen and 1 ppm nitrogen and shall not be mixed with the outer bath of atmospheric argon. The LEGEND-1000 lock system enables long strings of germanium detectors to be installed and removed from the cryostat containing UGAr. This design avoids coiling and uncoiling of cables and other mechanical supports but requires a large headspace above the main cryostat to fit the tubes that allows for string deployment. The lock system also minimizes the lock deployment cycle time, reducing the deployment time to be no more than a 24- or 48-hour period.

Figure 1.2-1. *(left:)* Close up view of the lock table with gate valve. *(right:)*: View of the tube and lock system with full length of the lock tube loaded with detector string.

One of the tasks the LEGEND team at UW has taken on is the design and prototyping of the Lock apparatus for LEGEND-1000. The primary components of the system include vacuum pumps and UGAr gas handling systems, controls and interlock systems to monitor operation, sealing actuator, lock tube, internal crane system, and a system to move and align the flange and the cable. The design allows the strings to be assembled in the LEGEND-

1000 glovebox, attached to a flange, and secured inside an 8m long tube, then purged and pressurized with argon before being moved to the cryostat port for installation. The sealing actuators, gate valves, and interlock systems will then allow the string and flange to be installed on the cryostat lid without contaminating the UGAr.

Figure 1.2-2. *(left:)* Assembly of stage one prototype in the high bay of CENPA. *(right:)* CAD model of stage one prototype in the high bay, with linear-plus-rotary feedthroughs, DPRF, and vacuum chamber with port and mock string flange (inside chamber).

At UW, we have several stages of prototyping underway. The first stage has been constructed in the high bay of CENPA, where we have begun prototyping the sealing actuator, which consists of the linear plus rotary feedthroughs, chamber, string-flange port, and pumps. This prototype has successfully demonstrated flange sealing and unsealing under vacuum. The second stage is being constructed in Cave 1 of CENPA, where we will prototype the overhead crane movement, lock interface with crysotat, raising and lowering of mock strings, insertion and removal of cryogen, and gasket exchange in the glovebox.

1.3 Characterizing low-energy surface backgrounds with KrSTC

J. A. Detwiler, N. Fuad*, N. Miedema, E. Miller, C. J. Nave, D. A. Peterson,
W. Pettus*, T. D. Van Wechel, L. Varriano, and C. Wiseman

Understanding the energy spectrum produced by high-purity germanium (HPGe) detectors used in the LEGEND experiment is a critical check against our understanding of the components making up the background for the experiment. This is of particular importance near the $Q_{\beta\beta}$ value of 2039 keV in ^{76}Ge, the energy at which we expect neutrinoless double beta decay ($0\nu\beta\beta$) to occur. Due to the extremely slow (if existing) decay rate of $0\nu\beta\beta$, the region near $Q_{\beta\beta}$ must be practically background-free to claim a discovery. There are a number of backgrounds that can contribute to the measured LEGEND energy spectrum. One background of interest for the UW LEGEND group arises from surface events. Ionizing radiation that does not penetrate far (< 1 mm) into an HPGe detector can produce signals that do not fully reflect the total energy deposited by said radiation, degrading the reconstructed energy. CAGE studies this effect with surface alphas, which is relevant for $0\nu\beta\beta$ analysis, while the Krypton String Test Cryostat (KrSTC) uses low-energy electrons, gammas, and x-rays to study this phenomenon at low energy. These energies are well below $Q_{\beta\beta}$, but this region is worth studying due to LEGEND's sensitivity to potential beyond the Standard Model (BSM) physics.

A picture of the experimental setup of KrSTC is shown in Fig. 1.3-1. Inside the cryostat, we install a p+ point contact (PPC) ^{76}Ge detector connected to the electronics inside and outside the cryostat. All of these electronics stages are designed and built by engineers in CENPA, and are optimized for our detector and data-taking needs. The detector is also connected to a high voltage supply, which allows us to bias the detector up to its maximum operating voltage of 2700 V. During operation, the cryostat is evacuated by a two-stage vacuum pump, which achieves cryostat pressures below 10^{-6} mbar when warm. There is also a copper cold finger that connects the detector assembly and a liquid nitrogen dewar to cool the detector to ~ 90 K. These cold surfaces provide us with cryopumping, allowing our cryostat to achieve pressures below 10^{-7} mbar.

The cryostat is also connected to a ^{83}Rb source that produces radioactive $^{83\mathrm{m}}$Kr gas. This krypton gas emits monoenergetic photons and electrons below 33 keV. The fact that these photons and electrons are monoenergetic makes it, in theory, easier to determine the nature of the energy degradation of the krypton decay products.

---

*Indiana University, Bloomington, IN.

Figure 1.3-1. The KrSTC hardware setup. The main cryostat body contains a PPC ^{76}Ge detector and is connected to a vacuum pump and our $^{83\mathrm{m}}$Kr source. Above the cryostat are the electronics used to extract signals, which are then sent to a digitizer for data collection.

Data were taken across about a year (February 2023 - April 2024), typically starting with background runs before introducing krypton gas into the cryostat. The most basic analysis is producing a background-subtracted krypton spectrum. Three examples of these spectra are shown in Fig. 1.3-2. The most interesting part of this study is that the krypton spectra drastically change depending on data-taking conditions, most heavily correlated with the cryostat pressure at which we began cooling with liquid nitrogen. One theory is that, when cooling at higher pressures, more water vapor and other residual gases can condense on cold surfaces, changing the surface charge configuration on the insulating passivated surface of the germanium detector.

If free charges collect on the insulating passivated surface (made of amorphous germanium), those charges can influence the drift of charge carriers within the Ge crystal. In particular, a negative surface charge can pull holes towards the passivated surface and cause incomplete charge collection (the holes get “stuck” on the passivated surface and are released on time scales longer than our waveform collection time). Since our signals are hole-dominated, this would generate a large region of the detector near the passivated surface that is not fully active. An activeness map for a surface charge of $\sigma = -4.5 \times 10^9$ e/cm^2 is shown in Fig. 1.3-3.

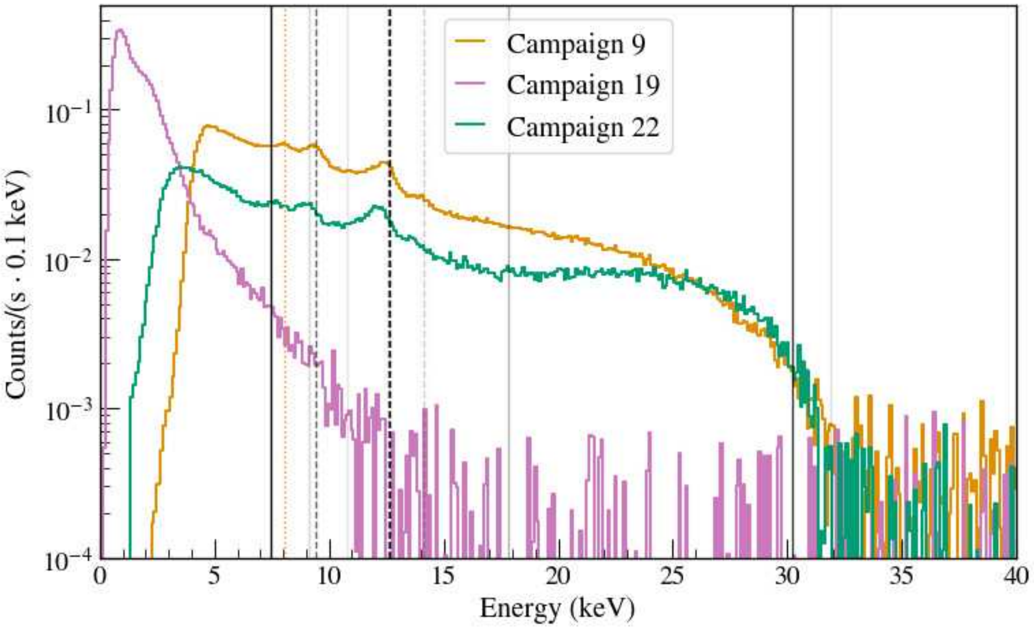

Figure 1.3-2. Three background-subtracted krypton energy spectra from different data-taking campaigns are shown above. The spectra labeled “Campaign 9” and “Campaign 22” were taken from campaigns with higher pressure-at-cooldown values, with the “Campaign 19” spectrum showing a spectrum with a lower pressure-at-cooldown. There are distinct differences in each energy spectrum, suggesting a strong dependence on this pressure-at-cooldown value. Expected krypton decay products are depicted as vertical lines with opacity scaling as expected intensity.

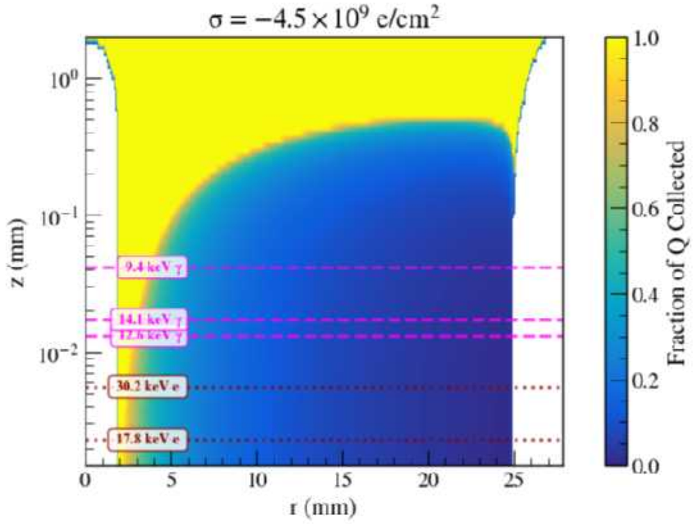

Figure 1.3-3. Depiction of “elephant ear” dead layer near the passivated surface of the OPPI detector. Surface charge density of $\sigma = -4.5 \times 10^9$ e/cm^2 applied to passivated surface ($z = 0$). Based on simulations from the `siggen` package. Deadness along the p+ point contact surface at small r and n+ bevel at large r is due to interpolation along physical boundaries. Penetration depths of most intense photons and electrons are shown.

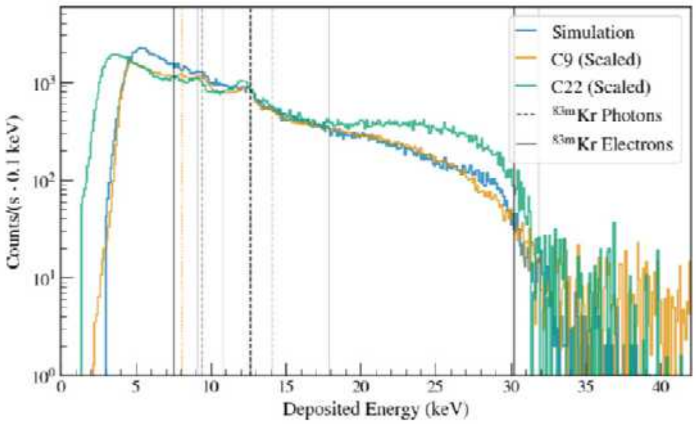

Figure 1.3-4. Comparison between background-subtracted krypton energy spectrum from data and the resultant simulated spectrum. The two spectra are scaled to match at the high-energy rise at ~ 31 keV and the expected 12.6 keV x-ray peak. There is disagreement in the continua above this 12.6 keV peak as well as below it and studies are ongoing to rectify these differences.

Applying this "elephant ear" activeness map to a GEANT4 simulation of ^{83m}Kr decays results in a simulated energy spectrum, shown in Fig. 1.3-4. Many of the features in the data spectra (especially the Campaign 9 spectrum) are recreated by simulation. We see a steep event rate increase in both data and simulation below ~ 32 keV due to degradation of the higher-energy electron lines. No electron peaks are visible, but some photon peaks are. The continua above and below the expected 12.6 keV x-ray peak do not agree perfectly across all campaigns (with a number of degenerate parameters in our model), but this is overall a promising explanation of what we observe in the KrSTC data. Campaign 19 data cannot yet be explained by our simulations, but this phenomenon has been seen before, namely in the MAJORANA DEMONSTRATOR low energy spectrum[1].

Another way to further study this surface charge-induced energy degradation is to solely use photons of different energies incident on the passivated surface. An x-ray fluorescence source[2] has been developed, and the hardware required to deploy it has been fully designed and built. Initial data-taking has begun and will continue when noise issues with our system are solved. Additionally, we have plans to install an electrode above the passivated surface and will attempt to "tune" the surface charge and observe any difference in the low energy spectrum. There is also interesting analysis work to be done on the KrSTC data set. We see a clear difference in the risetimes of krypton waveforms compared to background events, and a proper pulse shape study could help tag these surface events in KrSTC and in LEGEND data.

---

[1] I. J. Arnquist, *et al*, Phys. Rev. Lett. 132, 041001 (2024).

[2] S. R. Elliott, B. Dodson, C. Wiseman, *et al*, arXiv:2211.05691 [physics.ins-det] (2022).

1.4 CryoPDE: SiPM characterization at cryogenic temperatures

S. Borden, J. A. Detwiler, D. A. Peterson, C. Savarese, I. Wang, T. D. Van Wechel, and L. Varriano

The current iteration of the cryogenic silicon photomultiplier (SiPM) test stand at CENPA has completed its main physics data collection, the results of which are presented in Borden *et al*[1]. SiPMs are solid-state light detectors capable of a single-photon resolution that LEGEND uses to detect light from the liquid argon detector. The overall detection efficiency of the liquid argon system is directly proportional to the photon detection efficiency (PDE) of the SiPMs devices LEGEND uses; however, it was not known *a priori* if the PDE of the SiPMs in LEGEND's liquid argon detector, which is at cryogenic temperature, would match the manufacturer's value reported at room temperature. The results of the test stand built at CENPA have shown that the cryogenic PDE is roughly 20% smaller than the measured room temperature PDE at the green wavelength of light (562 nm) LEGEND uses in its liquid argon scintillation detection scheme (Fig. 1.4-1). More optical modeling studies are needed to quantify the impact of this decrease on the projected background index for LEGEND-1000.

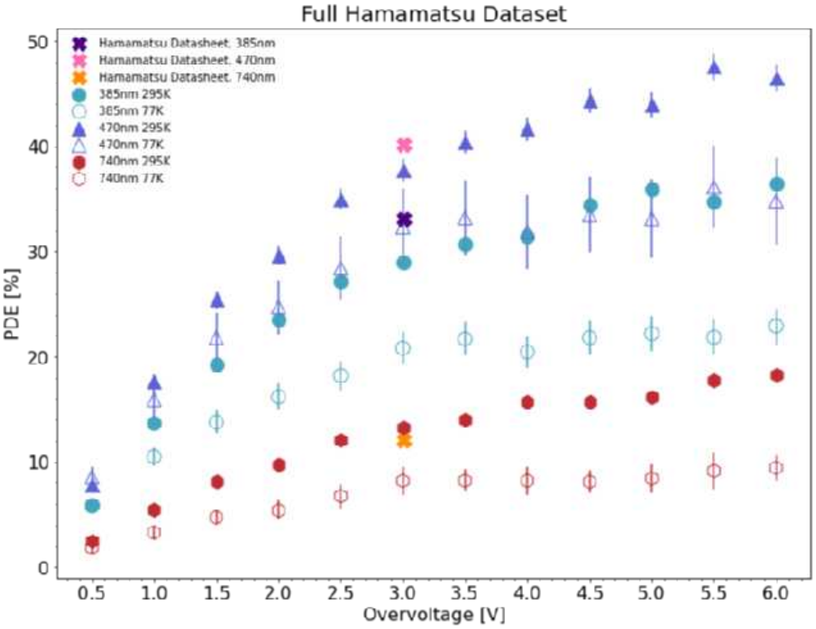

Figure 1.4-1. Cryogenic PDEs for the Hamamatsu S13360-3050CS SiPM at three different wavelengths. All three wavelengths show a decrease in the PDE at cryogenic temperature.

Significant work has been put into building a new SiPM test stand at CENPA capable of achieving cryogenic PDE measurements in vacuum. The original test stand required the use of optical fibers to prevent the reference SiPM from being cooled by liquid nitrogen boil-off, and the measurement of the optical fibers' transmission coefficient introduced measurement complexity and systematic uncertainties; directly cooling the SiPM under test with a cold

---

[1] S. J. Borden, J. A. Detwiler, W. Pettus, and N. W. Ruof, JINST **19**, P12014 (2024).

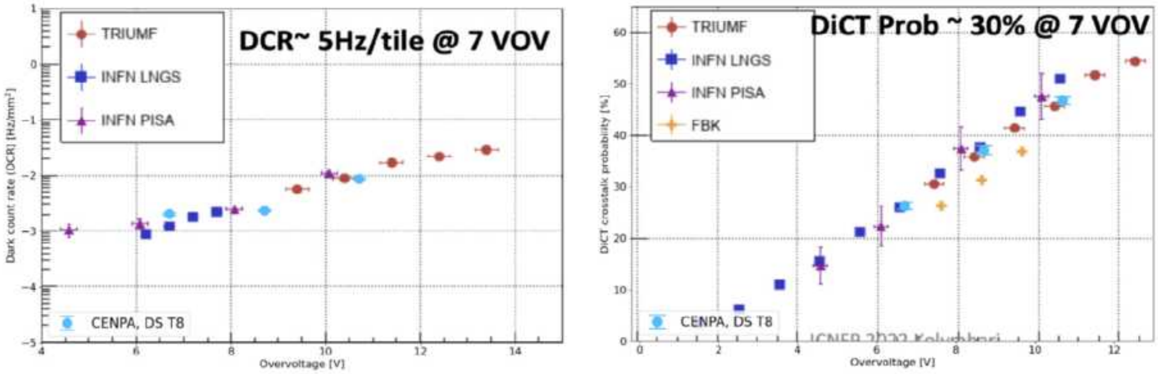

Figure 1.4-2. Dark Count Rate (left) and Direct CrossTalk Probability (right) for one of the DarkSide-20k SiPMs tested at CENPA. The measurements are compatible with results obtained in other facilities.

finger in a vacuum environment eliminates these issues. There were several attempts to engineer a vacuum-compatible cold finger feed-through, but many of these did not have enough cooling power to reach cryogenic temperatures. We settled on a feed-through based on a part provided by Professor Garcia. A copper cold plate was attached to the vacuum side of the feed-through to provide a larger surface area for cooling.

While the final holder of the new test stand was being designed and machined, we decided to perform the characterization of a few SiPMs produced for the DarkSide-20k experiment, which will use a dual-phase argon time projection chamber to search for dark matter. These SiPMs are optimized for operation in liquid argon and will also be deployed in LEGEND-1000 for its veto system. The samples tested at CENPA showed noise properties (*i.e.* dark count rate and direct cross-talk probability) fully compatible with results obtained in other setups.

At the end of the calendar year, the integrating sphere and mounting parts for the PDE measurement in vacuum have been installed in the cryostat and the temperature performance has been verified. The wiring and final assembly have just been completed. Data will be taken soon.

1.5 Alpha event characterization with the CAGE Scanner

J.A. Detwiler, N. Miedema, D. A. Peterson, G. Song, T. D. Van Wechel, and
C. Wiseman

The CAGE Scanner is designed to study the response of the HPGe detector to potential background sources that penetrate the passivated surface. Radiation that enters the active volume through the passivated surface can become energy degraded, causing normally high-energy radiation to be reconstructed in our region of interest. In particular, alpha radiation from the ^{238}U and ^{232}Th decay chains is emitted in the 4-6 MeV range, but can be reconstructed with energies near the 2039 keV $0\nu\beta\beta$ decay Q-value. CAGE uses a ^{241}Am source to direct a collimated beam of 5 MeV alpha particles onto the passivated surface at specific radii, rotary angles, and incidence angles. This year we continued to perform exploratory runs with new hardware to improve gain and pulse shape discrimination (PSD) stability. We also continued the development of the new PSD parameter, early charge (EQ), and quantified its ability to discriminate alphas in CAGE compared to other traditional PSD parameters.

Figure 1.5-1. Copper can installed around steel blank of ICPC detector. The HV contact with the N+ surface can be seen on the PEEK supporting structure above the detector.

Previously, we modified the high voltage connection with the ICPC detector to utilize the aluminum pad bonded to the N+ contact, rather than the copper HV ring the detector sits on. This connection is shown in Fig. 1.5-1. The continued analysis of the data taken last year with the new HV connection to the aluminum pad showed that while stability was not affected, the noise in the waveforms improved. Noise was measured as the standard deviation of the baseline. The variation of this noise over time is shown in Fig. 1.5-2. Data taken with the aluminum pad contact was found to be a factor of 2 smaller, with less variation, compared to data taken earlier with the copper HV ring at multiple bias voltages. Many factors, such as temperature fluctuations and vacuum imperfections, contribute to the overall noise level, so the exact contribution from the new HV connection cannot necessarily be quantified.

In addition to the new HV contact, a copper can was installed around the detector, shown in Fig. 1.5-1, to investigate gain shifts observed after the source motion. It is believed that motion of the copper IR shield is interfering with the internal electric fields of the detector, and the copper can in between the detector and IR shield should shield from those effects. A comparison of the uncalibrated location of the ^{40}K 1460 keV peak across six movements of the source is shown in Fig. 1.5-3. Fewer discontinuities over the same number of motor movements show promise in improving system stability. However, copper can limit the motion of the system, so further studies will need to be performed to see if these trends are maintained over longer distances.

On the analysis side, the performance of our new PSD parameter, EQ, was compared for a full scan at multiple rotary positions with traditional PSD parameters. The traditional

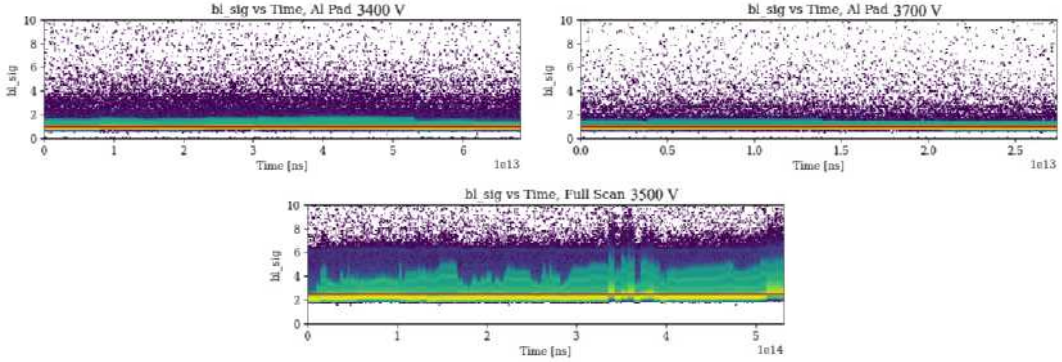

Figure 1.5-2. Baseline noise over time for data taken with the aluminum pad at 3400 V *(top left)* and 3700 V *(top right)* vs. data taken with copper HV ring at 3500 V *(bottom)*. Average values are shown as red horizontal lines.

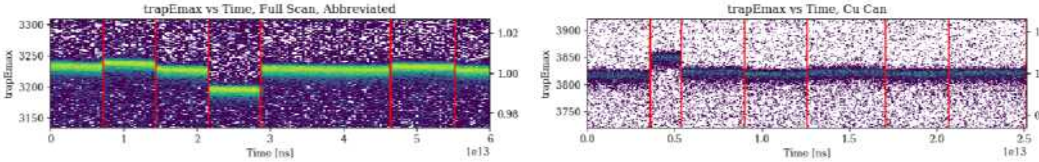

Figure 1.5-3. The uncalibrated position of the 1460 keV peak is shown as a function of time. Red vertical lines mark times when datataking was stopped to move the source to a new position. On the left, data without the copper can is shown in 7 different positions. On the right, data with the copper can is shown to have fewer discontinuities over the same number of source movements.

PSD parameters used were DCR, a measure of the slope of the tail of the waveform, T/E a measure of the slope of the rising edge of the waveform, and `tp_20-tp_02`, the time it takes for the waveform to rise from 2% of its maximum to 20% of its maximum.

The overall alpha rate was found by performing a background subtraction on the alpha data taken at each location and summing the resulting energy spectrum over 1000 keV, avoiding prominent known gamma lines. A similar procedure is performed on the events cut by each parameter, with the cuts set at a survival fraction of 90% of the background events. The absolute difference in the overall alpha rate and the alphas identified by each parameter is shown as a function of the radial position, averaged over the rotary position, in Fig. 1.5-4. EQ can be seen to most accurately reproduce the overall alpha rate.

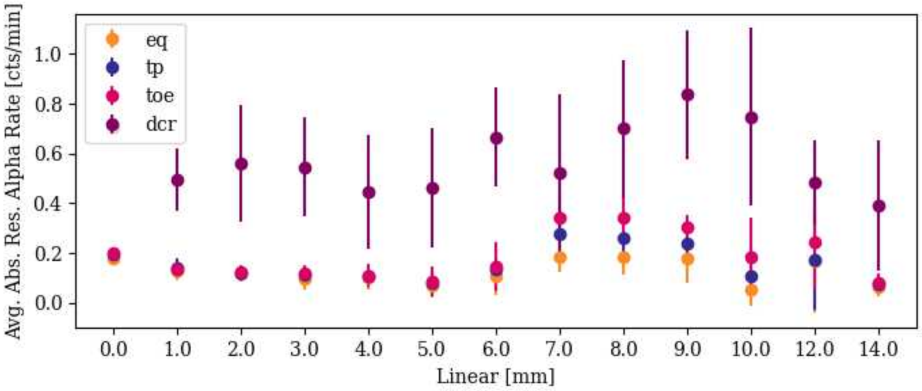

Figure 1.5-4. At each linear position, the absolute difference between overall alpha rate and identified alpha rate by each parameter is shown. Each data point is averaged over each of the 3 rotary positions.

1.6 Production of ^{42}Ar for LEGEND

S. Bellman, S. Borden, S. Enomoto, J. Detwiler, D. Peterson, C. Savarese,
L. Varriano, and Y. Zhang

One of the most important background for LEGEND comes from ^{42}Ar, a cosmogenically activated isotope with a half-life of 32.9(11) years. Although ^{42}Ar has a Q_β of only 599 keV, it undergoes β-decay to ^{42}K, which subsequently β-decays with a Q_β=3.5 MeV and a half-life of 12.359(4) hours. Energy-degraded β particles from this decay form a background at the ^{76}Ge $Q_{\beta\beta}$; in LEGEND-200, this is the second largest contribution to the background before the analysis cuts. Furthermore, the ^{42}K is ionized or forms ions, and can migrate in the liquid argon towards the germanium detector surfaces through some not-well-understood combination of convection and drift in the weak electric fields in the vicinity of the biased detectors. Presently, LEGEND-200 uses nylon tubes around the detectors as a physical barrier to the ion drift, preventing a large fraction of ^{42}K from reaching the detector surface, but ions still reach the detectors in sufficiently high numbers to form the second largest background before analysis cuts.

For the next-generation experiment LEGEND-1000, the background after analysis cuts must be improved by at least a factor > 20 compared to LEGEND-200, meaning that the current ^{42}K shielding strategy is inadequate. To mitigate the ^{42}K background for LEGEND-1000, the collaboration plans to use underground-sourced argon, which is depleted in ^{42}Ar compared to atmospheric argon by a factor $\gtrsim$ 1400. However, given that the procurement of this underground argon remains a risk for the experiment, alternate mitigation strategies for shielding and pulse-shape discrimination-based cuts are also being pursued in parallel. A promising technique is the use of tighter-fitting scintillating plastic enclosures to physically block ion drift, as well as provide additional scintillation light for veto, coupled with advances in machine-learning-based pulse shape analysis. LEGEND-1000 is currently undergoing various funding reviews with indicated strong support from DOE, NSF, and international agencies; assuming that funding schedules proceed as planned, LEGEND-1000 expects to begin to operate in ∼2030, making the need for a more detailed study of mitigation alternatives quite urgent.

Although the ^{42}Ar background is extremely important for LEGEND, it is quite difficult to study in above-ground laboratories due to the low natural activity of ^{42}Ar in atmospheric argon of ∼ 100 µBq/kg. To further develop these mitigation strategies in the lab, it is necessary to have an artificial source of ^{42}Ar with sufficiently high activity. An activity of closer to ∼ 40 Bq/kg would allow for the study of new shielding and discrimination techniques; in a typical lab-scale test stand, this could be achieved with a 40 kBq (1 µCi) ^{42}Ar source.

We have developed an argon gas target (Figure 1) at CENPA for an ^{42}Ar production using high-pressure argon gas for use with α and tritium beams, which we are currently prototyping. To extract the irradiated argon, we will use a gas handling system that we are presently prototyping at CENPA. A vacuum pump is used to evacuate the cell before argon is introduced into the cell. A pressure gauge monitors the pressure of the cell. After irradiation, a pneumatic valve is remotely actuated to open the cell to the sample bottle, which is kept at 77 K by a liquid nitrogen bath. The argon gas leaving the irradiation cell condenses and freezes inside the sample bottle. To increase the collected yield, due to the phase diagram of argon, we will introduce a new argon to flush the system several times. A getter can be used to trap reactive species, though we do not expect many by-products with a tritium beam. A programmable logic controller operates the system.

The freezing of argon has been tested and works well; we expect a yield of nearly 100% of produced ^{42}Ar after ∼20 flushes, which can be accommodated by the pressure rating of the sample bottle. The sample bottle would be shipped back from the irradiation facility to CENPA after a cooling-down period of a few days or weeks.

Figure 1.6-1. Annotated design of the gas cell for production of ^{42}Ar.

1.7 Solid Argon as a ^{42}Ar mitigation strategy

J. Buetow[*], S. Borden, J. A. Detwiler, Y.-R. Lin, C. Savarese, and L. Varriano

As already mentioned in Sec. 1.6, ^{42}Ar is one of the main backgrounds for LEGEND-1000.

---

[*] Pacific Lutheran University, Parkland, WA.

The baseline strategy of the LEGEND Collaboration to minimize the presence of this isotope in the LAr veto bath consists of procuring approximately 20 tonnes of underground-sourced argon (UAr), which is expected to be depleted in ^{42}Ar compared to atmospheric argon by a factor $\gtrapprox 1400$. Although the extraction and refinement infrastructures, already developed by the DarkSide-20k Experiment, are currently under construction, the UAr procurement remains one of the items with the highest associated risks for the LEGEND Project. Therefore, the development of alternative ^{42}Ar mitigation strategies is a priority for LEGEND-1000.

In addition to the production of an ^{42}Ar source as a way to explore alternative mitigation strategies for LEGEND-1000 (Sec. 1.6), our group decided to pursue the development of a solid argon (SAr) veto system. In liquid argon (LAr), the decay product of ^{42}Ar, ^{42}K, is observed to migrate and drift towards the germanium detectors, thus exacerbating the impact of this isotope on the background budget of the experiment. Conversely, we expect ^{42}K ions generated in argon ice to remain trapped at their point of origin, lifting by orders of magnitude the requirements on the isotopic abundance of ^{42}Ar for the argon that will be used in the LEGEND-1000 veto system.

First, we developed a Geant4-based simulation to estimate the effectiveness of such an idea. All simulations and analysis jobs have been run on the Perlmutter system at the National Energy Research Scientific Computing Center (NERSC). For this study, we instantiated a simple geometry consisting of an inverted-coaxial point contact (ICPC) detector contained within a large cylinder of liquid or solid argon. For the solid argon data, ^{42}K decay events were generated at random positions evenly distributed throughout the surrounding argon environment. For liquid argon data, the locations of the events were constrained to a cylindrical surface 1 μm outside the detector volume, as a consequence of the experimentally observed ^{42}K drift to the detectors. For each decay event, all the generated particles were recorded in their energy depositions at each step, including the position of the initial ^{42}K decay vertex and whether the resulting energy deposits happened in the germanium detector or in the surrounding argon. The energy depositions within the germanium crystal were summed to generate decay spectra (see the top panels of Fig. 1.7-1, which were then normalized with the appropriate activities. The spectrum for solid argon simulations shows that the vast majority of energy-depositing events come from the 1524.6 keV γ line, associated with $\sim 18\%$ of the β decays, and its Compton scattering continuum. At energies above that γ line, the energy deposits in the detector are dominated by β particles. β radiation has a much lower penetration power than a γ's and only reaches the detector if generated within a few millimeters of the detector surface. The number of events falling in the $0\nu\beta\beta$ analysis region of interest (ROI - highlighted in grey in Fig. Fig. 1.7-1) is two orders of magnitude lower than the γ peak, because most events are generated far enough away from the detector that only γ photons reach the surface. In contrast, the spectrum resulting from the liquid-argon setup is dominated by β depositions, as all ^{42}K decays are generated at the detector surface.

For the purpose of comparing the effectiveness of a solid argon veto, we normalized the obtained spectra, producing four scenarios (see bottom panel of Fig. 1.7-1 for the LEGEND-1000 experiment: (i) use of atmospheric argon liquefied, (ii) use of atmospheric argon liquefied, combined with the use of nylon mini-shrouds (NMS) mounted on the detector strings, as currently done in LEGEND-200, (iii) use of underground argon liquefied, and (iv) use of

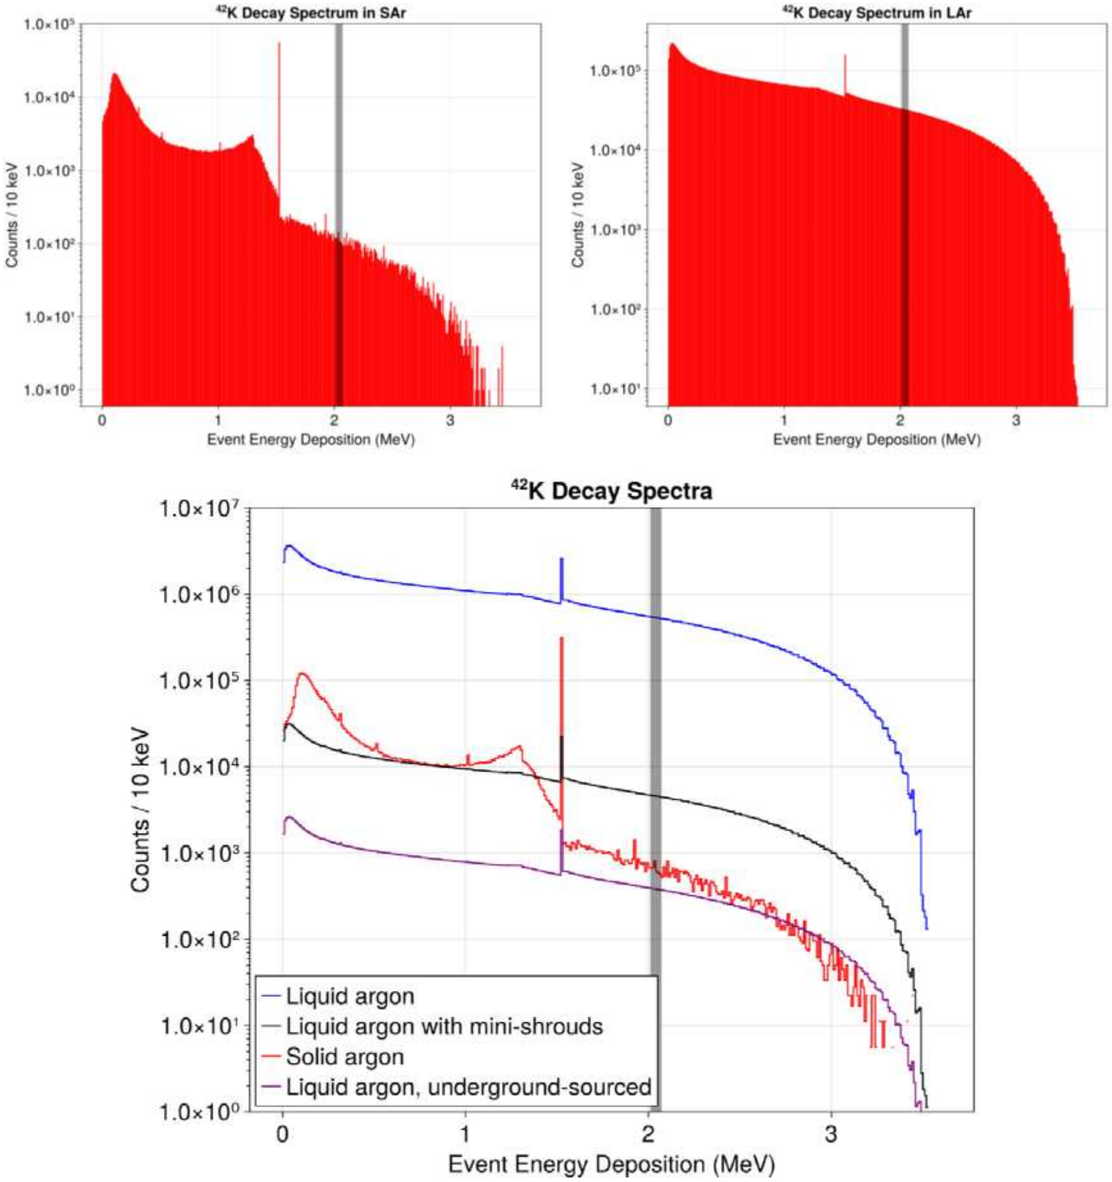

Figure 1.7-1. Energy spectrum from the beta decay of ^{42}K as detected by the ICPC germanium detector. *(Top Left:)* not normalized spectrum for the solid argon scenario, ^{42}K uniformly distributed in the argon ice volume. *(Top Right:)* not normalized spectrum for the liquid argon scenario, ^{42}K uniformly distributed on the surface of the detector. *(Bottom:)* comparison of ^{42}K decay spectra normalized for four selected scenarios, as described in the text.

atmospheric argon frozen. These simulations show that the ^{42}K contribution to the LEGEND-1000 background in the analysis window is the lowest when using UAr. At the same time, the use of solid atmospheric argon grants a ^{42}K-induced background only a factor ~ 2 above the UAr.

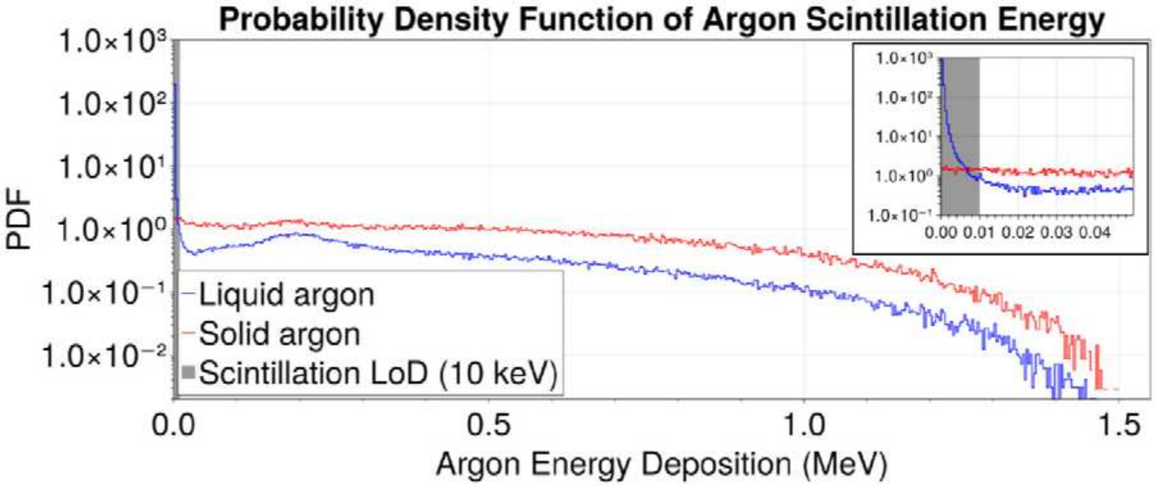

Figure 1.7-2. Spectra of energy deposits from ^{42}K decays in LAr (blue) and SAr (red), not normalized.

However, this is true before veto cuts are applied. When a β produced by a ^{42}K decay deposits part of its energy in the argon surrounding the detector, it stimulates the production of scintillation light. This light, suitably wavelength-shifted, collected, and recorded, can be used to veto such events by placing an anti-coincidence cut. To understand the relative impact of this cut on the solid and liquid argon veto systems, the energy deposited in the argon was recorded as part of the Geant-4 simulation. By selecting events falling in the $0\nu\beta\beta$ analysis ROI and plotting their energy deposits in the argon, we could predict the efficacy of the argon anti-coincidence (AC) cut in the solid and liquid scenarios, as shown in Fig. 1.7-2.

Upon a closer examination of Fig. 1.7-2, it can be seen that for the liquid argon scenario, there is a sharp peak in the probability of depositing near 0 keV. In contrast, this is absent from solid argon. Ultimately, this is a consequence of the spatial distribution of ^{42}K decays around the detector. Assuming a simple energy threshold of 10 keV for the veto system, where the detection efficiency is null for any event lower than that and 100% for all events higher, we obtained a first rough estimate for the probability of a ^{42}K β decay falling in the $0\nu\beta\beta$ analysis ROI surviving the argon AC cut by integrating the spectrum of Fig. 1.7-2 in the energy window that goes from 0 to 10 keV. Note that this extrapolation could be refined in the future by adopting a more realistic model of the veto efficiency. The results of this simple exercise show that ^{42}K events in a liquid argon veto have a $\sim$ 60% AC cut survival probability, while in solid argon this population decreases to $\sim$ 1.5%. This is an improvement of a factor of $\sim$ 40 on the efficiency of the argon AC cut, making the adoption of a solid argon veto a promising background reduction method.

In conclusion, Monte Carlo simulations of the ^{42}K background indicate that the use of a frozen argon scintillator could substantially enhance the sensitivity of the experiment. After accounting for the argon veto AC cuts, this idea promises to be the most powerful background reduction strategy considered. However, it is worth noting the current limitations and simplifying assumptions present in these results. First, when developing the simulations, the distribution of ^{42}K was assumed to be perfectly uniform in solid argon and was exactly

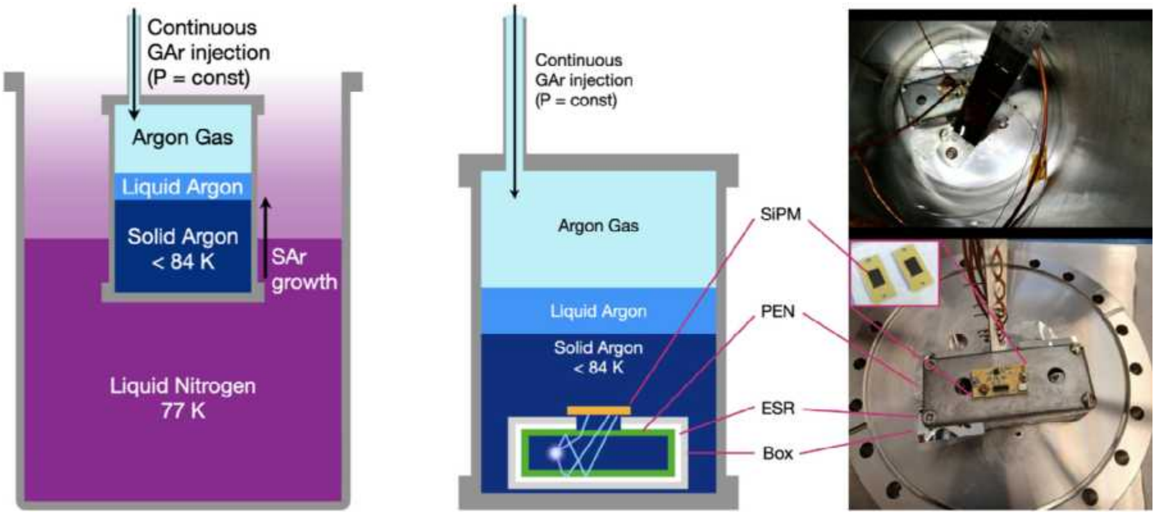

Figure 1.7-3. *(Left:)* schematic representation of the argon freezing setup. *(Center:)* detail of the argon freezing setup showing a section of the box to study the argon scintillation light. *(Right:)* pictures of the inside of the freezing cell during argon liquefaction and of the SiPM light detector mounted on the box.

constrained to a cylindrical surface that envelopes the detector at a distance of 1 μm in liquid argon. Although this model does not seem unreasonable for the solid case, the actual distribution of ^{42}K in the liquid argon scintillator is not yet well understood. Second, the dead layer arising from the diffusion of lithium across the surface of the n+ contact in the ICPC detector has not been modeled. Adding this feature to the simulations may shift the spectra to lower energies, potentially affecting the final results. Third, the normalization factors used to compute the relative backgrounds for each method are not known with high precision and may contain sizable errors. Of these, the specific activity of ^{42}Ar in UAr is particularly important but is not well known. If the activity is significantly lower than the assumed value of 46 nBq/kg (1400 times lower than the value for atmosphere-sourced argon), UAr could easily outperform solid argon, making this development moot. Lastly, the differences in scintillation behavior between solid and liquid argon were not rigorously taken into account in this study. The veto efficiency and the energy threshold might change significantly in the two considered veto scenarios.

In addition to the simulation work, we started an early investigation of the practical feasibility of a solid argon veto. Our experimental efforts focused on: (i) the development of a protocol to reliably freeze argon into an optically transparent block and (ii) the study of solid argon scintillation properties, both for pure argon and after the injection of nitrogen contamination. The freezing apparatus and procedure are inspired by the work described in[1]. A representation of the experimental setup can be found in the left panel of Fig. 1.7-3. A DN160CF stainless steel nipple is used as the freezing cell and the temperature at the bottom of the cell is monitored with a platinum temperature sensor (PT1000). The height of liquid argon inside the cell is monitored with a small camera-on-PCB and an LED light mounted inside the cell close to the top flange. A vertically mounted stainless steel ruler in the cell allows reading the liquid level with the camera (top right panel of Fig. 1.7-

---

[1] Zhang L. et al., Crystals 12, no.10:1416 (2022).

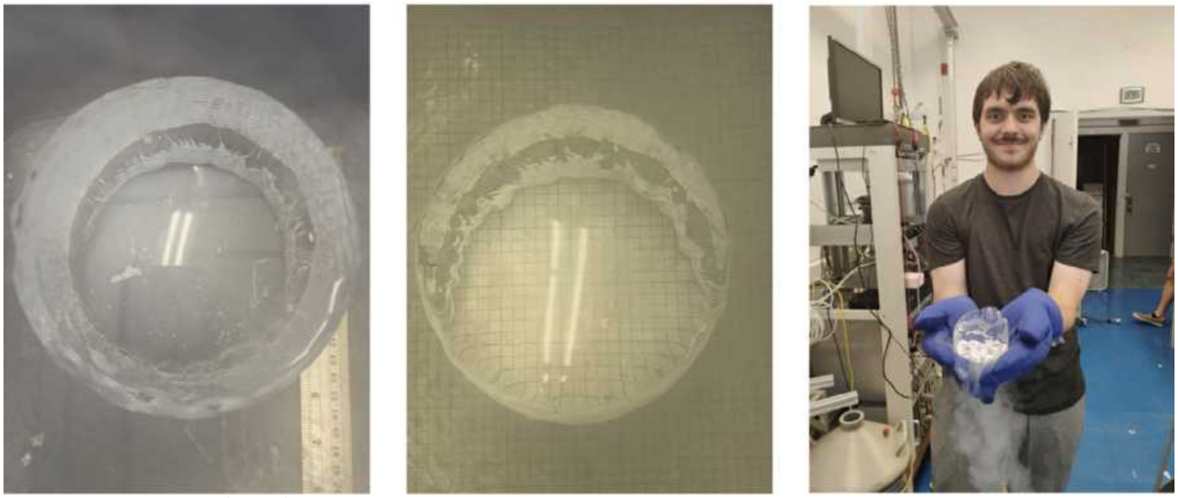

Figure 1.7-4. *(Left:)* picture of a crystal of solid argon minutes after its extraction from the freezing cell. *(Center:)* same block if solid argon lying on a sheet of graph paper with 5 mm pitch. *(Right:)* REU student J. Buetow holding the rapidly melting block of argon ice.

3). Thermocouples are secured at regular intervals on the ruler to provide a more granular temperature reading inside the cell. The flow of gaseous argon into the cell is controlled with a pressure regulator mounted on the argon cylinder, and the pressure in the chamber is monitored with a piezoelectric gauge. To start the freezing process, the cell filled with gas is lowered into a large open-neck cryostat containing liquid nitrogen that laps the bottom of the nipple. The resulting temperature decrease triggers the liquefaction of gas, which in turn lowers the pressure in the cell and causes the regulator to open, infusing more gas in a continuous process. At the bottom of the cell, a layer of ice starts to form, gradually growing in thickness as more argon accumulates in the cell. The fraction of cell surface immersed in the liquid nitrogen bath sets the speed of liquefaction and freezing. As solidification progresses, the cell can gradually be immersed deeper into the thermal bath to keep the ice growth rate constant. To preserve the optical transparency of the argon ice, the pressure is maintained between 1.1 atm and 1.3 atm for the duration of the freeze. With this protocol, we found that we can reliably produce several kilograms of SAr in as little as 5 hours. An example of solid argon obtained with this experimental setup is shown in Fig. 1.7-4.

In order to compare the scintillation properties of solid and liquid argon, we deployed a silicon photomultiplier (SiPM) in the freezing apparatus. This SiPM model has been developed by the DarkSide-20k Collaboration, in partnership with Fondazione Bruno Kessler, to operate optimally in LAr[1]. The SiPM has an active area of 7.9×11.7 mm^2 and is mounted on a PCB that hosts a custom-designed transimpedance amplifier (TIA) to read out the signals[2]. The TIA output is transmitted outside the cell via a $50\,\Omega$ coaxial cable and subsequently digitized using a CAEN DT5730 tabletop digitizer (8 channels, $2\,\mathrm{V_{pp}}$, 14-bit, 500 MS/s). The SiPM+TIA board is mounted on a steel box with cutouts to allow any light produced in the box to reach the SiPM, and LAr to flow in during the filling (bottom right panel of Fig. 1.7-3). The inner walls of the box are lined with an enhanced specular reflector

---

[1] F. Acerbi et al., IEEE TED **64(2)**, p. 521-526 (2017).

[2] M. D'Incecco et al., IEEE TNS, **65(4)**, p. 1005-1011 (2018).

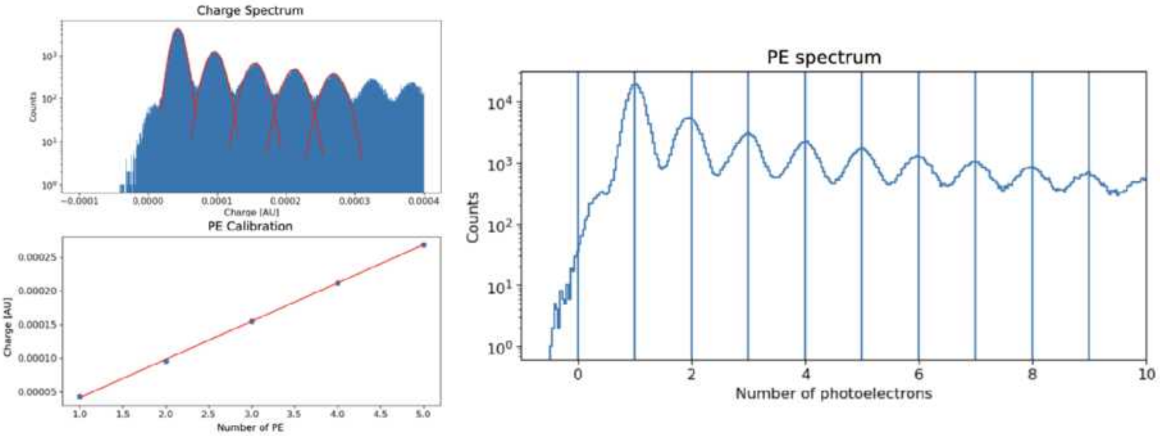

Figure 1.7-5. *(Top Left:)* raw charge spectrum from the SiPM. *(Bottom Left:)* SiPM gain extrapolation. *(Right:)* calibrated spectrum in PE.

(ESR) film and a 25 μm thick PEN film (central panel of Fig. 1.7-3). The latter is needed to wavelength-shift the argon VUV scintillation pulses from 128 nm to visible light ($\sim$ 450 nm). In this way, we maximize the amount of photons that the SiPM can detect for energy depositions that occur within the box. At the same time, the box's reduced dimensions limit the event rate, thus preventing excessive pile-up and dead time during digitization.

Scintillation pulse data was collected for both liquid and solid argon, and the data were obtained using a gas cylinder of unknown purity. For this reason, all of the following results are considered preliminary. Data were collected using external gamma calibration sources (^{137}Cs, ^{133}Ba, ^{60}Co) as well as without (background runs). We developed a simple analysis pipeline to extract and compare the system light yield and triplet lifetime for LAr and SAr runs, which are the most interesting observables for extrapolating the hypothetical performance of a future argon ice veto system. By taking data with the same argon before and after freezing, and with the same photosensor, we minimize most of the systematic uncertainties. The dominant remaining systematic arises from the temperature dependence of the SiPM's signal in response to a single photon, called a photoelectron (PE). While LAr data is acquired at 85 K, the SAr temperature can drop as low as 77 K. This difference slightly alters the duration of the 1 PE signal, as well as the bandwidth and noise performance of the TIA. To correct for this effect, we perform an in-situ calibration of the SiPM for each run by Gaussian fitting the first five peaks of the charge spectra, corresponding to events with a varying number of photons from 1 to 5. A linear fit of the mean values of each Gaussian curve returns the SiPM gain. We can then rescale all the spectra, transforming the variable on the x-axis from integrated charge to number of PEs (Fig. 1.7-5)). This is directly proportional to the energy deposited in the box.

After the PE calibration, we can compare the spectra of a gamma source (or the background) obtained with LAr and SAr. An example is shown in Fig. 1.7-6), which shows the background spectra during a cycle of freezing, melting, and re-freezing. Given that for this particular data taking we do not know the exact purity of the argon (which might have changed during phase transitions due to the opening of a pressure relief valve), we can not

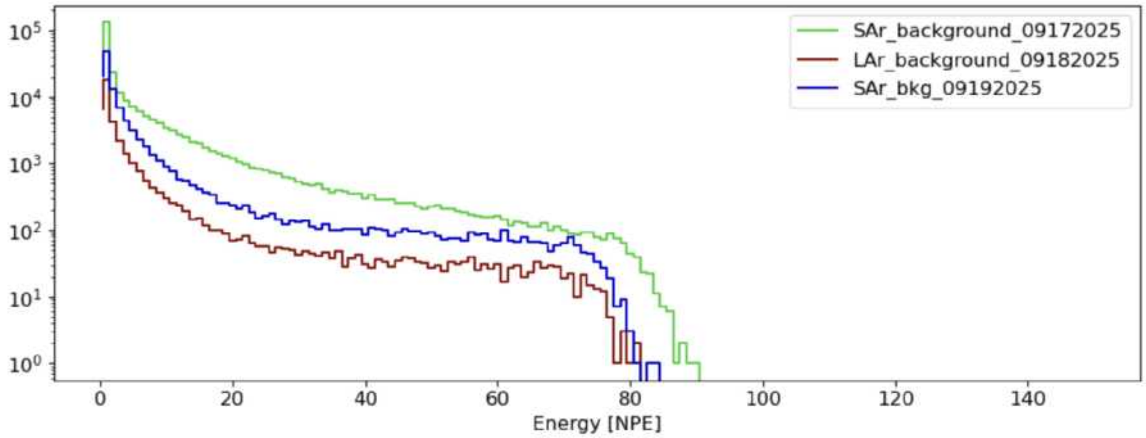

Figure 1.7-6. Energy spectra obtained without gamma sources from LAr (red) and SAr (green, blue). The spectra are normalized to the same number of events.

yet draw conclusive results on the SAr scintillation properties. With the analysis pipeline in place, the acquisition of new data will yield more quantitative results.

Another important figure that characterizes every scintillator is its scintillation time profile. Liquid argon is known to have a composite pulse made of two exponential decays with remarkably different time constants ($\tau_s \sim 7$ ns and $\tau_t \sim 1400$ ns). This is because excited argon forms metastable dimers whose spin state defines their decay time: singlet dimers are short lived (τ_s), while triplet dimers are more stable (τ_t). During our tests, we want to confirm the observation of such a pulse shape in LAr and to see if this time profile is preserved also in SAr. Moreover, the presence of nitrogen (or oxygen) in argon strongly quenches its scintillation. This quenching acts more efficiently on the long lived triplet dimers. Therefore, the measure of the argon triplet lifetime can be used as a reliable proxy for its purity. We developed the machinery to perform such measurement. During the single photo-electron calibration, we also extract the 1 PE template (or super-pulse). For our setup, we found that the 1 PE shape can be accurately described with a single exponential. After fitting 1 PE, we can use this function as a deconvolution core that acts on the average scintillation waveform. The analysis process and results are sketched in Fig. 1.7-7). The resulting pulses clearly show a fast peak (singlet), followed by a longer exponential decay (triplet). This analysis process does not account for the time response of the PEN wavelength shifter, but this can be ignored in the first approximation due to its relatively short time constant ($\tau_{PEN} \sim 20$ ns) with respect to the argon triplet. Once again, we do not have conclusive results on the LAr/SAr comparison, but the analysis infrastructure is ready for the next data taking.

During 2025, we successfully developed a test stand capable of reliably producing optically transparent SAr. We demonstrated the reproducibility of the freezing process as well as its non-destructiveness. A wire-bonded SiPM has survived several full cycles of argon liquefaction, freezing, melting, and vaporization. During the coming year, we will concentrate on extracting quantitative measurements of SAr scintillation properties, directly comparing the results with LAr. The availability of data for LAr and SAr from the same setup and batch

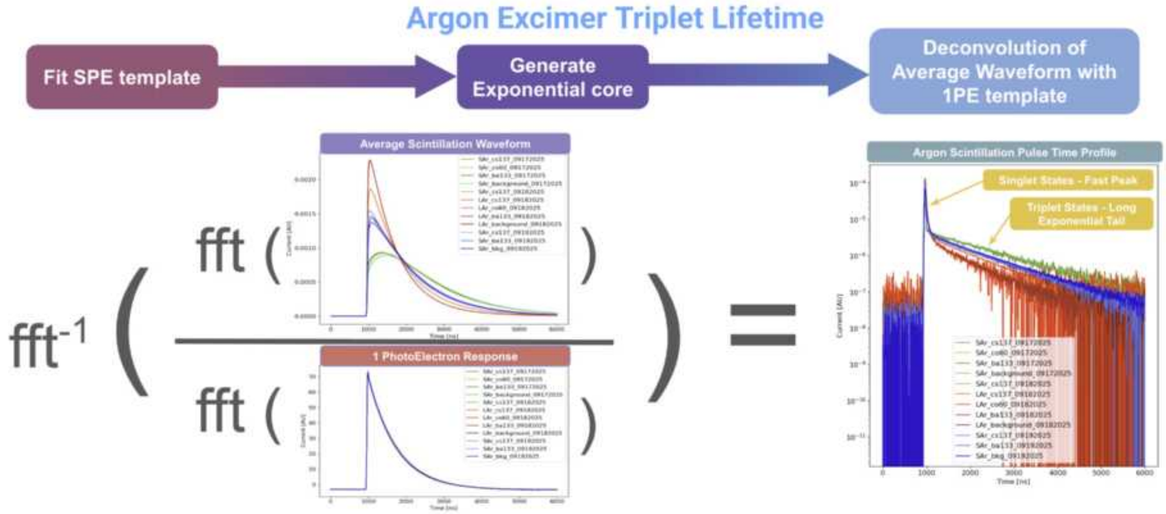

Figure 1.7-7. Sketch of the deconvolution process to extract the true argon scintillation time profile from the average waveforms. The deconvolution process removes the shape of the SiPM 1 PE response.

of argon minimizes systematic effects.

1.8 Silicon detector α, β pulse shape discrimination

S. Borden, J. A. Detwiler, A. Garcia, P. Kammel, I. Kunen, C. J. Nave, G. R. Song, F. Tschirhart, and L. Varriano

Motivated by recent measurements of the ^{8}B neutrino spectrum[1] and ^{8}B and ^{8}Li angular correlation measurements to search for a tensor contribution to the weak interaction [2], we are interested in applying some of the pulse shape discrimination techniques developed by LEGEND to silicon detectors. In these recent measurements, ^{8}B or ^{8}Li ions are trapped between four 1 mm thick silicon detectors that detect the α and β decay products < 16 MeV. These decay products are differentiated by their deposited energy alone, making it nearly impossible to distinguish < 1 MeV α from the minimum-ionizing β. However, due to the different topology of the electron-hole pair ionization tracks in the detector from α and β, it may be possible to discriminate between α and β by their pulse shape.

To investigate LEGEND-style pulse shape discrimination techniques, F. Tschirhart and P. Kammel developed an experimental setup to quickly switch between α and β sources, allowing for rapid testing of legacy silicon detectors to determine which were still operable. F. Tschirhart also developed and ran simulations to predict the expected pulse shape difference and identify appropriate β sources to use. For this work, a ^{207}Bi source was used to provide minimum-ionizing β and a ^{244}Cm source was used to provide 5.8 MeV α. These α were then degraded by an aluminum foil to around 1 MeV. Using logistic regression, I. Kunen has developed a classification algorithm with good precision in distinguishing between α and β

---

[1] B. Longfellow et al., Phys. Rev. C 107, L032801 (2023).

[2] B. Longfellow et al., Phys. Rev. Lett. 132, 142502 (2024).

particles in the energy range of interest. Features based on several key parts of the waveform, such as current, rise, and fall time, are generated. During training, the analysis routines learn which features are most important for distinguishing between the labeled data sets and use these in the classification. C. J. Nave performed Geant4 simulations to compare expected and observed spectra, which indicate that only a fraction of particles are MIP-like, depending on the energy. These different event topologies are expected to influence the discrimination power, with more MIP-like events easier to distinguish from α events.

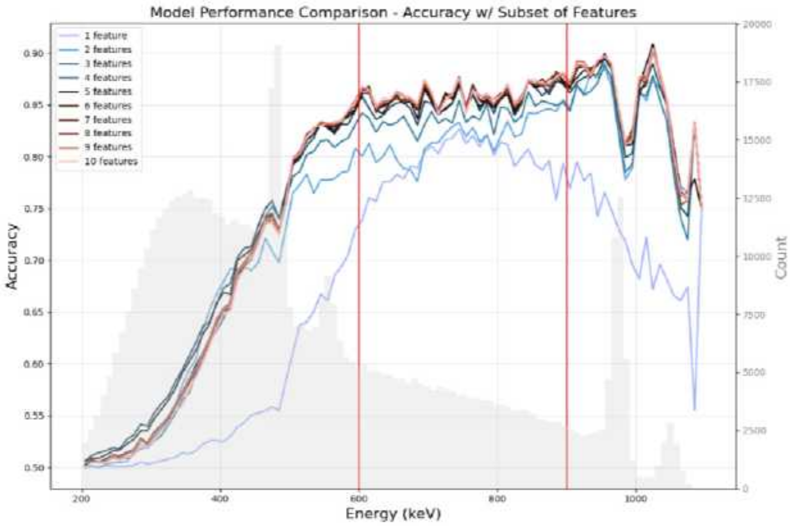

Figure 1.8-1. Overall model performance as a function of the number of included training features. Using equal numbers of α and β events in 10 keV bins, the classification accuracy is shown. The training region was 600-900 keV (red lines). The β energy spectrum is shown in gray for comparison.

KATRIN

1.9 The KATRIN experiment

P. J. Doe, S. Enomoto, Q. Li, R. G. H. Robertson, and Y. H. Sun

This year, the KATRIN experiment reached an historic milestone: 1000 live days of data on the beta decay of tritium. Its planned campaign of data acquisition on the electron (anti)neutrino is now complete. KATRIN, located at the Tritiumlabor Karlsruhe in Germany, is the world's most sensitive experiment to measure the mass of the neutrino by a direct, model-independent method.

The KATRIN collaboration formed in 2001 with UW as one of the founding institutional members. SNO, at that time our primary neutrino project, had just discovered that solar neutrinos changed flavor on their journey to earth, proving that electron neutrinos were mixtures of massive eigenstates. The phenomenon of neutrino oscillations shows that neutrinos have mass (in contradiction to the prediction of the minimal standard model), but does not reveal the mass itself. It was essential to turn to other approaches, such as KATRIN. The design of KATRIN was based on the experience of previous experiments; the windowless gaseous molecular tritium source was based on the Los Alamos scheme and the MAC-E filter spectrometer was a greatly expanded version of the devices built at Mainz and Troitsk. The latter experiments had reached a limit of about 2 eV on the effective mass m_β, and KATRIN's goal was an order of magnitude improvement. To give that some substance, the statistical precision scales as the fourth root of the number of events, i.e. one needs 10,000 times more events with a commensurate improvement in systematic uncertainties. If background is present, the scaling is even less favorable.

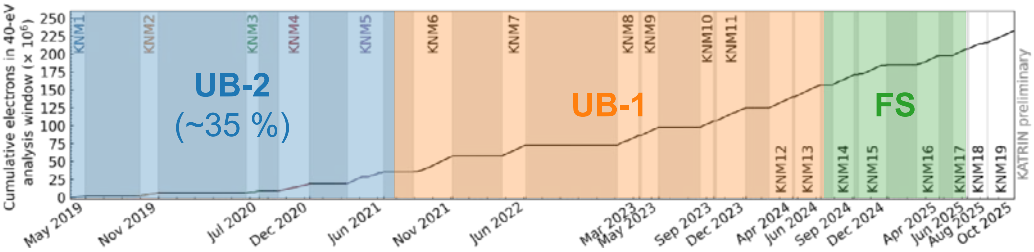

Figure 1.9-1. KATRIN's accumulated luminosity (the number of beta-decay events recorded). The shaded intervals are upgrade, maintenance, and calibration activities. The colored regions are explained in the text.

Fig. 1.9-1 shows the number of beta events accumulated over time. The data are initially processed at UW to ensure that the files are uniform and error-free, and are then turned over to the analysis group. The data are doubly blinded. Three different analysis codes are used to provide cross-check on the results. The first step is called 'figure skating' (FS) in analogy with the grading of that sport – the data are first simulated for the same statistical accuracy and calibration inputs, and then analyzed. Any deviations between codes must be tracked down. In the first unblinding (UB-1) the actual data are turned over to the analysts. The neutrino mass is still hidden via an undisclosed modification of the final-state distribution, which is strongly correlated with the derived neutrino mass. The final unblinding step (UB-2) replaces that modified FSD with the true one, and the neutrino mass is extracted.

Data from KNM1 through KNM5 runs have been fully analyzed and published in April (Science **388**, 180 [2025]), as shown in Fig. 1.9-2. The neutrino mass is found to be less than 0.45 eV at 90% CL.

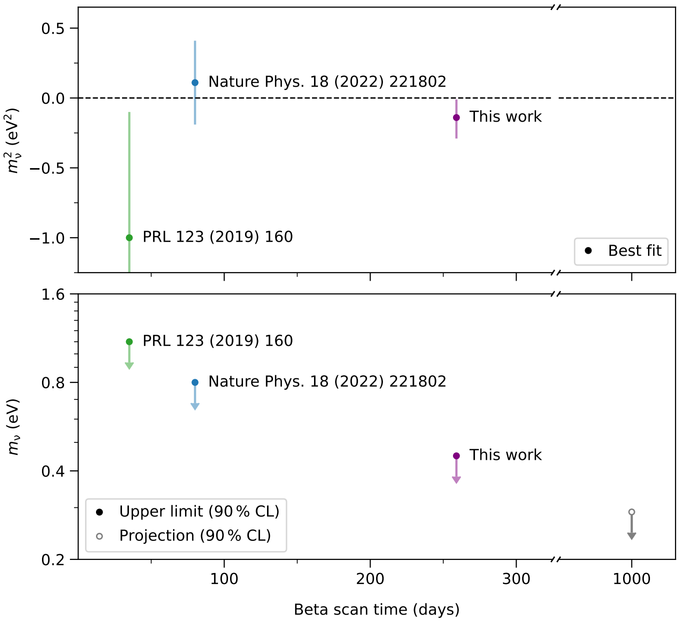

Figure 1.9-2. The evolution of KATRIN's results to date. The point at lower right indicates the anticipated limit with all the data and a neutrino mass of 0.

The results have attracted great attention because of developments in cosmology. Until recently, it has been thought that cosmology would provide a value for the sum of the neutrino mass eigenvalues. However, new instruments and data lead to 'tensions' when the measurements are interpreted in the standard ΛCDM model that has been so successful so far. Most troubling, the Hubble appears to depend on redshift, and the summed neutrino mass best fit is significantly negative. These issues can be ameliorated with additions to the model, but there exists a welter of possibilities. The importance of a terrestrial neutrino mass measurement is now fully understood.

The same data set has also been analyzed for the presence of light sterile neutrinos. Several experimental results in neutrino oscillation searches on short baselines can be individually interpreted as evidence for a sterile neutrino in the 1-10 eV mass range. Particularly striking is the 'Gallium Anomaly', a $\sim 20\%$ shortfall in the capture rate in Ga of neutrinos from an intense ^{51}Cr source. KATRIN's data (Nature **648**, 73 [2025]) exclude this interpretation for the most concordant region of the existing results, as shown in Fig. 1.9-3.

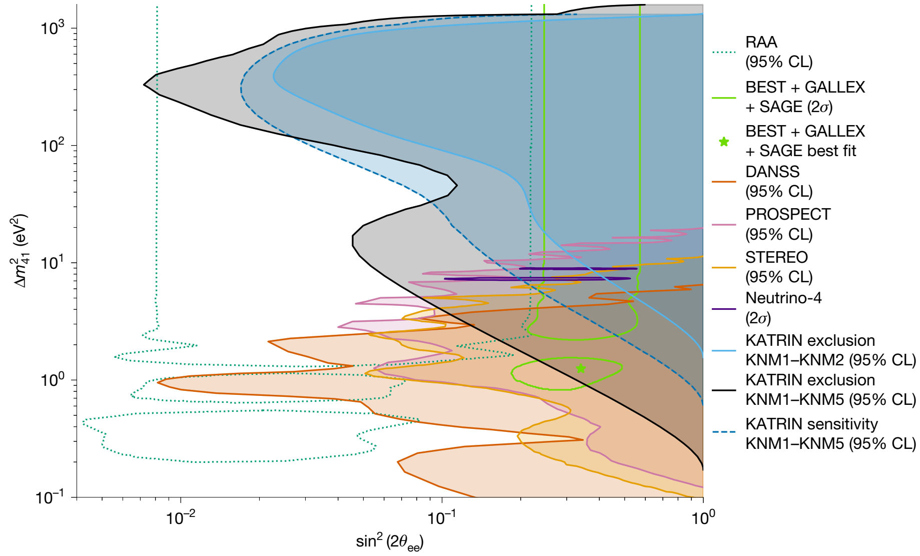

Figure 1.9-3. Limit at 95% CL (solid black line) on admixture of a light sterile neutrino from KATRIN's data sets KNM1 through KNM5.

We began our work on KATRIN immediately upon joining, designing, and building the inner electrode system for the pre-spectrometer and designing the focal-plane detector (FPD) system for the main spectrometer. A proposal was written to DOE and capital equipment funding began in 2007. The FPD consists of a 148-pixel Si detector, two 6-T superconducting magnets, a high-voltage post-acceleration electrode, a calibration device (PULCINELLA) to produce monoenergetic electrons across the detector face, a source insertion mechanism for an ^{241}Am source, and two separate vacuum systems at high and ultra-high vacuum. The entire system was completed and operationally ready in 2012 (Fig. 1.9-4), and was shipped to Germany for installation. We also wrote the Object-oriented Real-time Computer

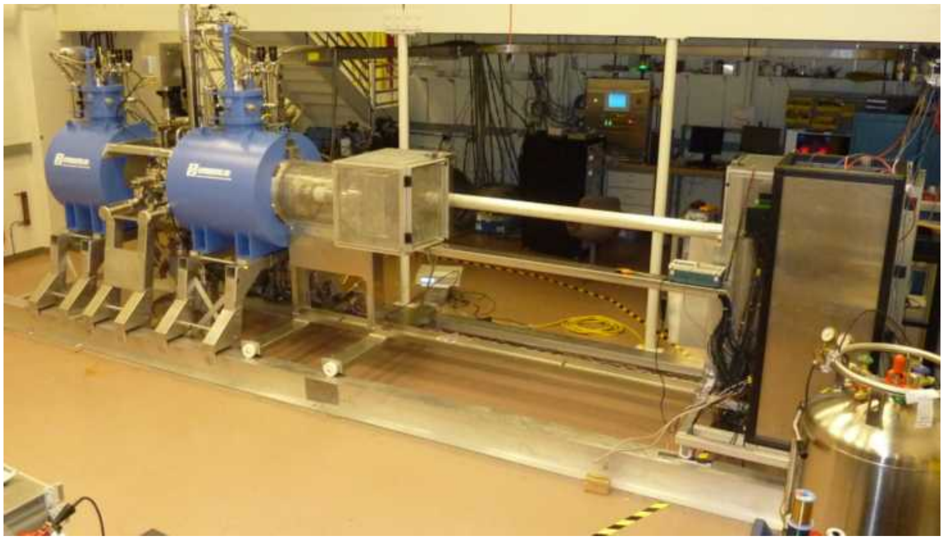

Figure 1.9-4. Completed detector system in Room PAB B037 at UW, about 2011.

Acquisition software ORCA, which is now in use at many installations. The detector was cooled by a single-stage cryocooler coupled by a heavy copper rod to the detector mount. This arrangement did not produce enough cooling, and the copper was replaced with a liquid-nitrogen heat pipe of our design that worked well. Later we designed and delivered a new cosmic-ray veto system and a special isolation valve that went inside the beamline within the first magnet, to allow isolation of the detector system from the main spectrometer. (A design challenge, this valve was a source of significant headaches for us but eventually was satisfactory.) The post-acceleration electrode operated reliably up to 10 kV, sufficient for reducing background and systematics, but not up to the originally planned 30 kV. Studies here and in Karlsruhe have led to a new design using a commercial NEC acceleration column that is expected to function at least to 20 kV, as needed for TRISTAN. This is based on our successful TRIMS system that operates at 60 kV.

The detector has functioned well over the years, with satisfactory resolution, stability, and background. There was one incident in which possibly a discharge occurred in a spectrometer, requiring replacement of the wafer. The temperature stability was adequate, and its small remaining uncertainty contribution was reduced by a detailed analysis and correction algorithm. This year is the last year for the Si *pin*-diode array. It will be replaced by a new highly pixellated array of Si drift detectors for the TRISTAN project, a search for sterile neutrinos of keV mass and very small mixing.

We are continuing our significant role in KATRIN during the final analysis of the active neutrino data sets. Our analysis work has largely been in support of detector studies, evaluating and reducing the systematic uncertainties, but we have contributed across the board. We identified a systematic effect that would have been serious, the time dependence of background as an intrinsic Penning trap between the pre-spectrometer and main spectrometer filled. This systematic was then eliminated for subsequent runs by changing the data-taking procedures.

The publication of the final results of the KATRIN active-neutrino mass measurement is expected within two years with a sensitivity of ~ 0.3 eV.

Project 8

1.10 Project 8: Overview

E. Novitski and R. G. H. Robertson

In the CRES technique[1] pioneered by the Project 8 collaboration[2,3,4,5], radioactive gas is confined within a uniform magnetic field. When a charged particle is emitted in decay, it is trapped axially in a weak magnetic bottle trap, undergoes cyclotron motion, and emits electromagnetic radiation with a frequency f_{cyc}, given by

$$f_{\mathrm{cyc}} = \frac{1}{2\pi}\frac{|e|B}{m + E_{\mathrm{kin}}/c^2}, \tag{1}$$

where B is the magnetic field, m is the mass of the particle, c is the speed of light, and E_{kin} is the kinetic energy of the particle. With the charged particle in a magnetic bottle to increase the observation time, the ∼aW to ∼fW of cyclotron radiation (depending on the magnitude of B) can be detected directly with an RF receiver, amplified, and digitized, generating data as shown in Fig. 1.10-1.

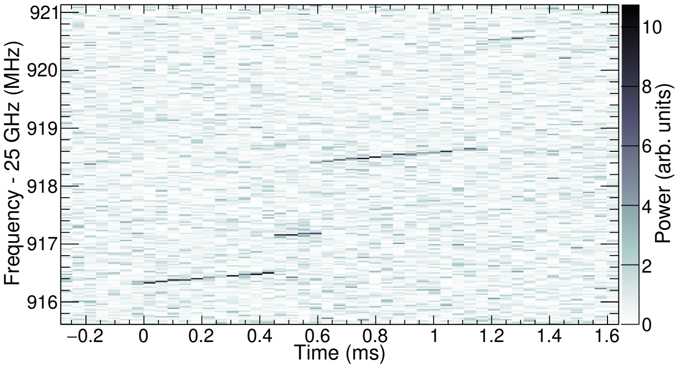

Figure 1.10-1. Spectrogram of the first CRES event detected from a tritium beta decay electron.

The energy of the charged particle at the time of its creation can be determined from f_{cyc} at the beginning of the signal. CRES has several advantages for tritium β^- spectroscopy over previous methods. Electron energies are determined via frequency, which can be measured with extremely high precision. As a differential spectroscopy technique, it continuously builds the entire endpoint spectrum, which increases statistical efficiency and protects against time-dependent systematic effects. The lack of a need to extract electrons from the source makes the CRES technique scalable by volume and compatible with the use of an atomic tritium

---

[1] Monreal, B. & Formaggio, J. A. Phys. Rev. D 80, 051301 (2009).

[2] Asner, D. M. et al. Phys. Rev. Lett. 114, 162501 (2015).

[3] Esfahani, A. A. et al. Phys. Rev. C 99, 055501 (2019).

[4] Ashtari Esfahani, A. et al. Phys. Rev. Lett. 131, 102502 (2023).

[5] Ashtari Esfahani, A. et al. Phys. Rev. C 109, 035503 (2024).

source, which must be magnetically isolated from the walls to prevent recombination. Finally, CRES is an extremely low-background technique. For example, electrons from radioactive decay on the walls collide again with the wall after a single cyclotron period and do not produce a detectable signal. Electrons that are outside the energy range of interest or that are not trapped within the magnetic bottle also do not contribute to CRES spectrum. The dominant background, thermal RF noise, can be precisely characterized and excluded given sufficient SNR.

With these advantages, an experiment using CRES is potentially sensitive to neutrino masses significantly smaller than KATRIN's anticipated 0.3 eV/c^2 limit. The neutrino mass sensitivity for Project 8 has been estimated both analytically [1] and in a Bayesian framework [2], with agreement in their conclusions. The statistical sensitivity for an observation interval ϵ below the Q-value is: $\sigma^{\mathrm{stat}}_{m^2_\nu} = \frac{2}{3rt}\sqrt{rt\epsilon + \frac{bt}{\epsilon}}$, where r is the rate in the last 1 eV of the $m_\nu = 0$ spectrum, t the observation time, and b the background rate per eV.[3] Fig. 1.10-2 shows the dependence of neutrino mass sensitivity on energy resolution and on the precision to which the width of the energy resolution (detector response function) is known for planned cavity-based CRES experiments using an atomic tritium source. For a molecular tritium source the final state spectrum of the daughter He^{3}T$^+$ ion (Fig. 1.10-2) would set the ultimate limit at around 100 meV. For an atomic source at 1 mK the sensitivity can be further improved, approaching 40 meV in a very large instrument.

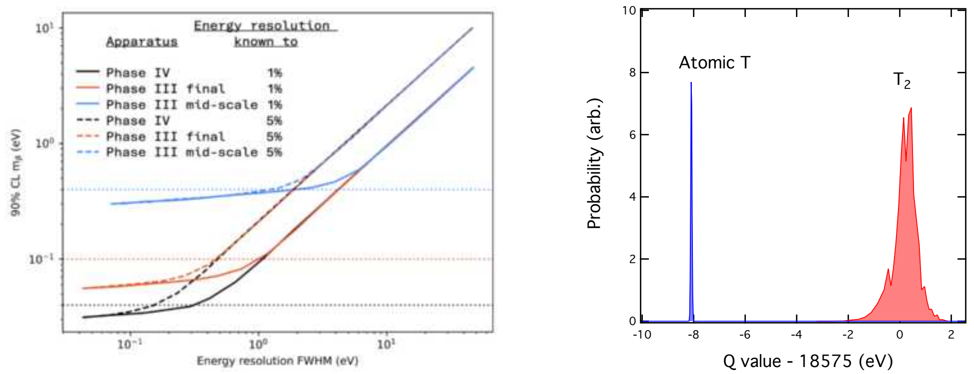

Figure 1.10-2. *Left*: Predicted neutrino mass sensitivity vs. FWHM energy resolution for the LFA (here listed as "mid-scale Phase III") and Pilot-scale (here listed as "final Phase III") experiments and the Phase IV experiment. The solid lines represent the energy resolution needed if the uncertainty on the resolution width is characterized to 1% and the dashed lines if it is characterized to 5%. The dotted lines represent the target neutrino mass sensitivities for each of these experiments. *Right*: Line widths for T and T$_2$ decay to electronic ground states.

---

[1] Doe, P. J. et al. arXiv:1309.7093 (2013).

[2] Ashtari Esfahani, A. et al., Phys. Rev. C 103, 065501 (2021).

[3] Contributions from molecular daughter ion final states, magnetic field non-uniformity, and thermal Doppler and collisional broadening (*i.e.*, finite mean free lifetime) are also considered (not shown). For concreteness, an assumption is made that the systematic contributions are each known to 1% unless otherwise stated.

Phases of Project 8

A neutrino mass measurement sensitive to m_β of 40 meV/c^2 would detect neutrino mass or disfavor the entire parameter space of the inverted ordering. Project 8 is progressing toward this full sensitivity in four phases (Fig. 1.10-3) by systematically addressing each of the factors that contribute to the improvement in sensitivity.

Phase I was the first demonstration of CRES. Phase II recently set the first CRES-based neutrino mass limit. Phase III is a multi-pronged program that will demonstrate the technologies needed to reach full scale, combine them to do a mid-scale experiment with m_β sensitivity comparable to KATRIN, and culminate in a “Pilot-scale” experiment with world-leading 0.1 eV/c^2 sensitivity. With the CRES-enabled ultra-low background, it will set limits on sterile neutrino mass without assumptions about active neutrino masses. Phase IV will be the full-scale experiment that reaches a sensitivity of 40 meV/c^2 sensitivity.

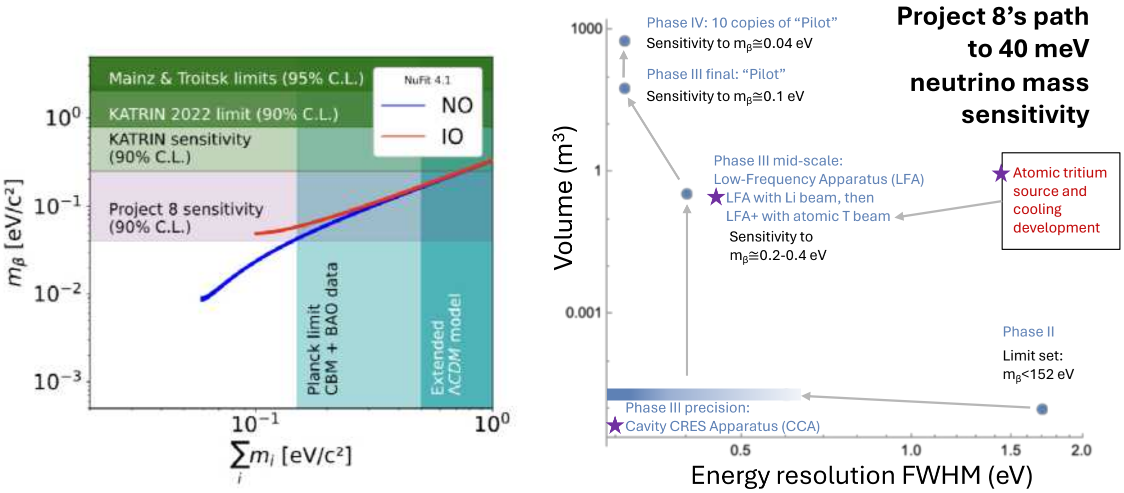

Figure 1.10-3. *Left*: Neutrino mass parameter space excluded by or accessible to direct kinematic measurements, and excluded by cosmological observations under two sets of model assumptions. *Right*: Volumes and resolutions of Project 8’s major CRES apparatuses.

1.11 The Cavity CRES Apparatus (CCA)

M. Huehn, M. Kallander, E. Novitski, R. G. H. Robertson, M. Wynne, P. Kolbeck, S. Kuenstner, A. Kurmus, and Y. Yin

The Cavity CRES Apparatus (CCA), in the final stage of commissioning at UW and pictured in Fig. 1.11-1, is a platform for developing sub-eV-precision CRES techniques. It will enable the first demonstration of CRES in a cavity, a detector geometry that scales to the larger volumes and lower frequencies that will eventually be necessary for improved neutrino mass sensitivity. It will also enable key improvements over Phase II in energy resolution (expected $\sim$ 0.3 eV FWHM) and effective volume. It is a testbed for calibration tools, including an electron gun.

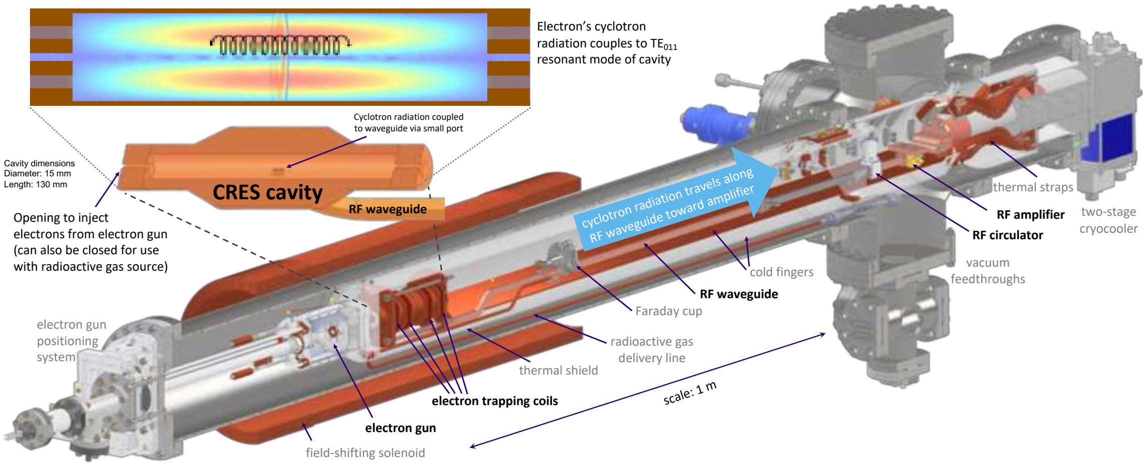

Figure 1.11-1. A schematic view of the Cavity CRES Apparatus (CCA). The TE_{011} cavity is maintained at 85 K to keep $^{83\mathrm{m}}$Kr in the gas phase; in an alternative running mode, removable endcaps allow the electron beam a path into the cavity. Electron trap coils around the cavity create small distortions in the magnetic field to confine electrons axially. Cyclotron radiation travels to the right in the RF waveguide toward the amplifier, held near 4 K. Not pictured is the medical MRI magnet in which this sits, which provides the high-homogeneity 1-T background field.

We have designed the CCA with improvements to make sub-eV resolution with CRES possible for the first time. The resonant cavity has one readout channel corresponding to the TE_{011} resonant mode, and it can be swapped out for multi-mode or NMR-compatible cavities (Sec. 1.14) in the future. This cavity-based detector architecture takes advantage of cyclotron radiation enhancement on resonance due to coupling with the cavity[1,2]. The cavity architecture is far less susceptible to the Doppler effect than the Phase II detector design was, increasing useful signal power by preventing excessive shunting of the RF main-carrier power into RF sidebands[3]. Sideband data, which was not available in Phase II, will be present in the CCA and will permit axial-amplitude-dependent field-shift corrections to improve resolution. The electron trap coils were carefully designed to provide a magnetic trap shape that enables magnetic field information to be optimally imprinted in these sidebands.

The CRES frequency in the CCA is the same as in Phase II, allowing for reuse of the receiver and digital DAQ from Phase II. A repurposed medical MRI magnet provides a 0.04 m^3 volume with 1-ppm magnetic field homogeneity, and an array of custom NMR probes to monitor field stability and homogeneity (Sec. 1.13). A new gas-handling system is capable of controlling and analyzing conditions more precisely than in Phase II. The cavity has removable vacuum endcaps, allowing for operation either with an $^{83\mathrm{m}}$Kr source or with an electron gun that will inject known-energy electrons via Mott scattering with helium gas (Sec. 1.12). With its expected 0.3-eV linewidth, it will allow for a precise characterization of detector response. Data-taking with both types of electron source, and a detailed response

---

[1] Purcell, E. M. Phys. Rev. 69, 681. (11-12 1946).

[2] Brown, L. S., Gabrielse, G., Helmerson, K. and Tan, J. Phys. Rev. A 32, 3204–3218 (1985).

[3] Ashtari Esfahani, A. et al. Phys. Rev. C 109, 035503 (2024).

characterization, is planned with the baseline CCA hardware.

This year has seen many milestones for the CCA. The fabrication, assembly, installation, and testing of the cavity, trap coils, cryogenic platform, DC wiring, RF detection chain, and the gas-handling system have been completed. This has been led by the local UW team, with close collaboration with MIT for the cavity and Yale for the trap coils (Fig. 1.11-2). The CCA has undergone multiple cooldowns in which the temperature-control system, wiring, and vacuum conditions were tested and improved. The electron gun design has been iterated, its achievable energy increased, and it has been integrated with the detector. New instruments were integrated into an upgraded slow-controls system based around the Dripline and SlowDash tools, which are now also being used by other experiments. The local computer network was upgraded and the software infrastructure—both a local online data quality monitoring system in collaboration with MIT and a long-term storage system at PNNL—were built. UW postdoc S. Kuenstner and MIT postdoc H. Binney are operations coordinators; they have begun a system of hybrid remote and local operations shifts bringing in collaborators from across Project 8. With all of this in place, the CRES data is expected imminently.

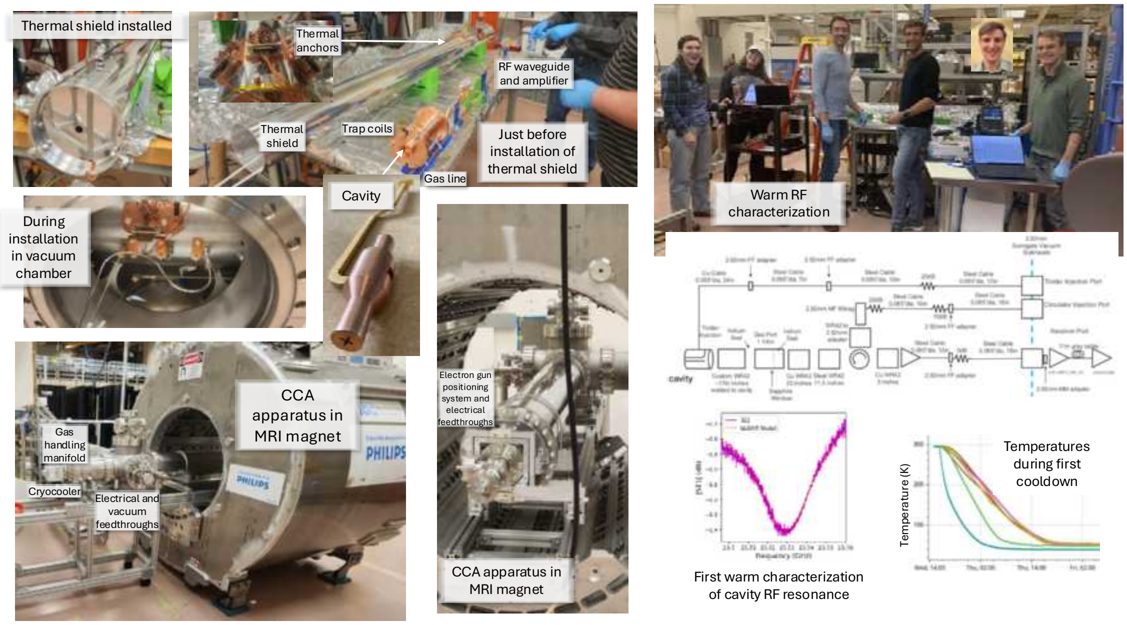

Figure 1.11-2. CCA and electron gun during assembly and testing; installed in the MRI magnet; RF detection chain; early warm RF characterization and cooldown temperatures.

1.12 Electron gun for CRES calibration

M. Kallander, A. Kurmus, and R. G. H. Robertson

The electron gun (egun) is a direct calibration tool for Project 8 as a controlled electron source. It lets us inject electrons of known and nearly continuously variable energy from 0 to 30keV into the cavity volume. With this capability, we can perform absolute energy calibration at the ~ 1 eV level, test magnetic-field uniformity, verify frequency-energy conversion, and exercise the full reconstruction pipeline.

Our design uses a Y_2O_3 cathode held at a high negative voltage. We heat the cathode via an optical link using an LED array that powers a small solar panel at source potential. Electrons emitted from the cathode accelerate toward the ground and enter the 0.959 T magnetic field. The electrons leave the egun head parallel (pitch angle $\theta \approx 0$) to the magnetic field, and their subsequent motion is guided almost entirely by the magnetic field. Designing an electron source in which the beam is guided primarily by magnetic fields rather than electrostatic optics is a distinctive feature of the Project 8 egun effort.

In normal operation, most electrons pass straight through the cavity. However, a small fraction undergoes elastic Mott scattering off helium gas that we introduce into the cavity. Because the scattering is elastic, aside from small recoil-energy losses, the scattered electrons retain the same kinetic energy but acquire trappable pitch angles. For a helium pressure of 10^{-6} mbar and a 1 μA beam, we expect to observe approximately one trapped electron per second.

Over the past year, we built a movable rail system that carries both the high-voltage (HV) box— which contains the LED array, solar-panel filament supply, and HV bias circuitry—and the egun head itself. This rail system allows us to insert the gun reproducibly into the vacuum flange and achieve coarse alignment with the cavity. Once the gun is inside the vacuum system, two piezoelectric motors drive a gimbal mechanism with a balanced vacuum-bellows linkage. This arrangement provides precise steering and allows us to inject electrons at different radial positions within the cavity. Fig. 1.12-1 shows the current setup before insertion into the vacuum chamber.

Figure 1.12-1. The current version of the electron gun with the updated rail system and high-voltage box, before insertion to the vacuum chamber. Matt Kallander is establishing the final connections.

A major challenge throughout the last year was achieving stable high-voltage operation in the collinear magnetic field. We addressed the possibility of breakdowns associated with secondary-emission-enhanced avalanche (SEEA) processes in the initial design by avoiding insulating surfaces parallel to the fields. Insufficient vacuum quality was found to be a problem, and the vacuum was improved by cryopumping. We then identified a third mechanism: the emitted electrons dislodge positive ions after striking grounded surfaces. The ions then accelerate back toward the negatively biased cathode, releasing additional electrons when they impact it and triggering breakdown.

To suppress this ion-feedback mechanism, we redesigned the front end of the gun and added a cylindrical electrode, biased at +24 V, which we call the "ion curtain." This electrode creates an approximately 10 V potential barrier along the axis and prevents ions from propagating back toward the cathode region. We validated both our diagnosis and the effectiveness of the ion curtain: without it, we observed breakdowns at around 9 kV; with the ion curtain installed, we saw no comparable breakdowns under the same conditions up to 17 kV.

With these improvements, we achieved the first full integration of the egun with the CCA insert. We transported a stable 0.8 µA beam at 16 kV through the cavity and all the way to the Faraday cup. Sixteen kilovolts is the minimum energy required to produce detectable cyclotron radiation in the geometry of the current cavity. Therefore, the electron gun is already capable of serving as a CRES electron source for system-level testing.

Our next steps include testing higher-energy beams and demonstrating long-term stability. Based on recent results, we believe that the remaining limitations at higher voltages are related in part to conditioning. Further improvements to the acceleration structure and vacuum are also planned. We are making progress in achieving the operation of the electron gun in the full 0–30 kV range.

1.13 NMR sensors adapted from the muon $g-2$ experiment

P. Kolbeck and S. E. Kuenstner

A sub-ppm understanding of the background magnetic field inside the CRES cavity is necessary to reach the target energy resolution of the Cavity CRES Apparatus (CCA). In the CCA, this field is provided by a repurposed hospital MRI magnet shimmed to sub-ppm homogeneity in an ellipsoidal region 40 cm in diameter centered on the bore of the magnet. However, superconducting magnets in neighboring labs, the residual para- and dimagnetic materials in the magnet bore, drift of MRI currents, and drift of the earth's field can perturb the background field. Further, the field will be actively adjusted on the few-percent level using a large normal conducting "field-shifting solenoid (FSS)" within the superconducting magnet bore to bring the emission lines of $^{83\mathrm{m}}$Kr into resonance with the TE_{011} mode of the cavity. Managing these changes will require active field monitoring.

This active monitoring system is currently in development and uses materials and design adapted from the muon experiment $g-2$ experiment. To monitor the field in their storage ring, the muon $g - 2$ collaboration developed a custom NMR probe[1,2] with a target field of 1.45 T and a precision of 240 ppb to be distributed to 378 locations throughout the ring (Fig. 1.13-1 (a)).

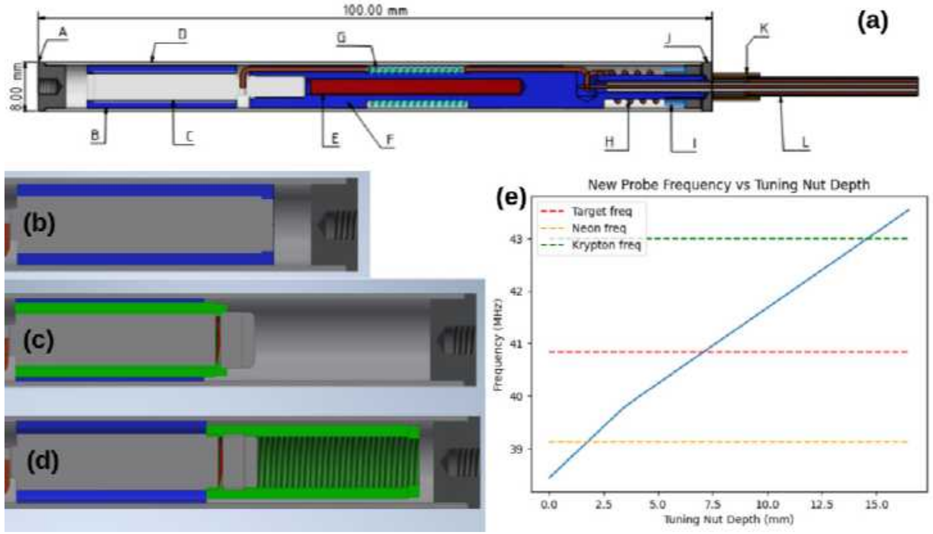

Figure 1.13-1. (a) Design of original muon $g - 2$ probe. (b) Previous design of the coaxial capacitor in the LC circuit. The gray block at the center is the threaded aluminum inner conductor, the blue is a threaded teflon nut, and on the outside is the aluminum sleeve outer conductor. (c) New design of the coaxial capacitor, tuned to the minimum frequency. There is now an aluminum nut (green) around the threaded rod, straddling the new cut in the rod that contains a phosphor bronze leaf spring (brown), which is tightened down to ensure a good electric connection between the nut and the rod. A 0.4-mm-thick layer of HDPE (blue) is epoxied to the inside of the sleeve. (d) New design of the coaxial capacitor, tuned to the maximum frequency. (e) Tuning range of the new probe design, with the frequency related to the depth of the aluminum nut.

With the completion of the muon $g - 2$ experiment, excess probe components, assembly tools, electronics, and software were made available to the Project 8 collaboration for use in the CCA.

To work in the CCA, the probes required modifications. Their resonant frequency had to be adjusted from 61.54 MHz to 40.84 MHz due to the CCA's field of 0.959 T, and their design had to admit frequency adjustments of up to 5%. P. Kolbeck, in consultation with engineer R. Roehnelt, went through several iterations of design and prototyping, resulting in the adjustable design shown in Fig. 1.13-1 (b)-(d). A prototype of this probe was built with new parts machined by N. Miedema, and tested in the lab using a vector network analyzer to confirm the resonant frequency range as in Fig. 1.13-1 (e), and using partially modified muon $g - 2$ electronics to acquire an FID and measure the field in MRI, as in Fig. 1.13-2.

We developed a nonmagnetic scaffolding that could hold the probes around the vacuum

---

[1] Swanson, E. et al. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 1075, 170338 (2025).

[2] Corrodi, S. et al. Journal Of Instrumentation 15, 11008 (2020).

jacket in a limited space between it and the FSS, see Fig. 1.13-3. Through 3D printing, we developed a scaffolding unit that can be assembled into a larger system cheaply and quickly. We 3D-printed a set of these scaffolding units and confirmed that they could secure probes around the vacuum jacket.

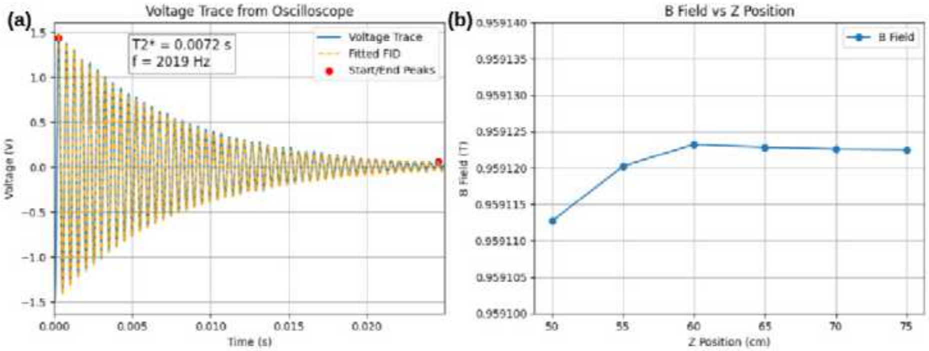

Figure 1.13-2. (a) FID signal gathered from a probe in the MRI field. $T_2{}^*$ is within the 1-10 ms range found by muon $g-2$, indicating good field homogeneity and sample conditions. (b) An uncalibrated measurement of the B-field within the MRI using the prototype probe showing the sub-ppm variation of the B-field in the bore and the end of the shimmed region.

Figure 1.13-3. (a) Scaffolding unit holding the prototype probe. (b) CAD model showing how the scaffolding units can be combined into a larger array. The fasteners are nonmagnetic, being a combination of aluminum, brass, and plastic

1.14 Hyperpolarized helium-3 magnetometry

P. Kolbeck, S. E. Kuenstner, and A. Kurmus

Achieving a good energy resolution with the CRES technique requires a precise measurement of f_{cyc} and B: in order to achieve an energy resolution of 0.1 eV near the tritium endpoint, the field experienced by the electron must be characterized to 0.2 parts per million (ppm) or better. This characterization is complicated by the magnetic bottle that confines the electrons axially, because each CRES electron experiences a slightly different spatially-averaged magnetic field. If these bottle-induced frequency effects are not measured and accounted for, they worsen the energy resolution of the event reconstruction. The current phase of the Project 8 experiment, known as the Cavity CRES Apparatus (CCA), requires a high-performance magnetometer capable of reconstructing the magnetic field experienced by each CRES electron to reach its energy resolution targets.

The magnetometer must simultaneously meet several challenging specifications, including sub-ppm accuracy and precision, and *in situ* operation at the low-pressure cryogenic conditions in the CCA. We selected a nuclear magnetic resonance (NMR) magnetometer based on hyperpolarized helium-3 vapor, which is capable of meeting all requirements for the CCA. Because the helium-3 gyromagnetic ratio is known to high accuracy, helium-3 magnetometers can achieve absolute accuracies of around 30 parts per billion. At the same time, the high nuclear polarization enables high-SNR-frequency measurements and good precision within short measurement times. Finally, the helium-3 vapor remains gaseous up to the operating temperature of the CCA, which helps maintain long spin coherence times and good SNR. The helium-3 magnetometer for the CCA will be made up of 2 key subsystems: the room-temperature optical polarizer and the cryogenic NMR probe circuit. CENPA researchers have made progress on both subsystems this year.

The optical polarizer will function via metastability exchange optical pumping (MEOP), using a high-power infrared laser to generate a high degree of electronic polarization in metastable helium-3 atoms. This angular momentum can be transferred to the nuclear spins in ground-state helium-3 atoms via a process known as metastability exchange. The resulting highly-polarized ensemble of nuclear spins can be introduced to the CCA, allowing precision *in situ* magnetometry via NMR. This year, we designed the polarizer system and began procuring the major components needed for the polarizer, including a 2 Watt 1083nm Keopsys laser, other optical components, the helium-3 gas, and various pieces of the gas-handling manifold. The system needs to maintain an ultra-pure helium-3 sample in the polarizer and recover the gas once the experiments are complete.

The second key component of the magnetometry setup is the cryogenic NMR probe circuit, which must allow the spin precession to be induced and detected while preserving the resonant microwave mode within the cavity that allows for readout of the CRES signal. The scheme for combined NMR and CRES is shown in Fig. 1.14-1, and is inspired by the “birdcage coil” design common in medical magnetic resonance imagers. The design is adapted so that the birdcage structure supports two resonant modes: the microwave TE_{011} standing wave along the axis at ∼26 GHz for CRES detection and a lumped-element RF resonance at ∼30 MHz for detecting the helium-3 spin precession.

We designed, fabricated, and tested a prototype of the birdcage circuit, as shown in Fig. 1.14-1. Nate Miedema fabricated the copper cavity halves at the CENPA machine shop. The full setup made extensive use of other prototyping facilities at CENPA, including the 3D printer to fabricate various brackets and a pulley system and laser-cut copper foil leads for the capacitors. The electronics originally designed for NMR magnetometers in the muon g-2 experiment allowed us to verify that the birdcage circuit can perform NMR measurements, as shown in Fig. 1.14-2.

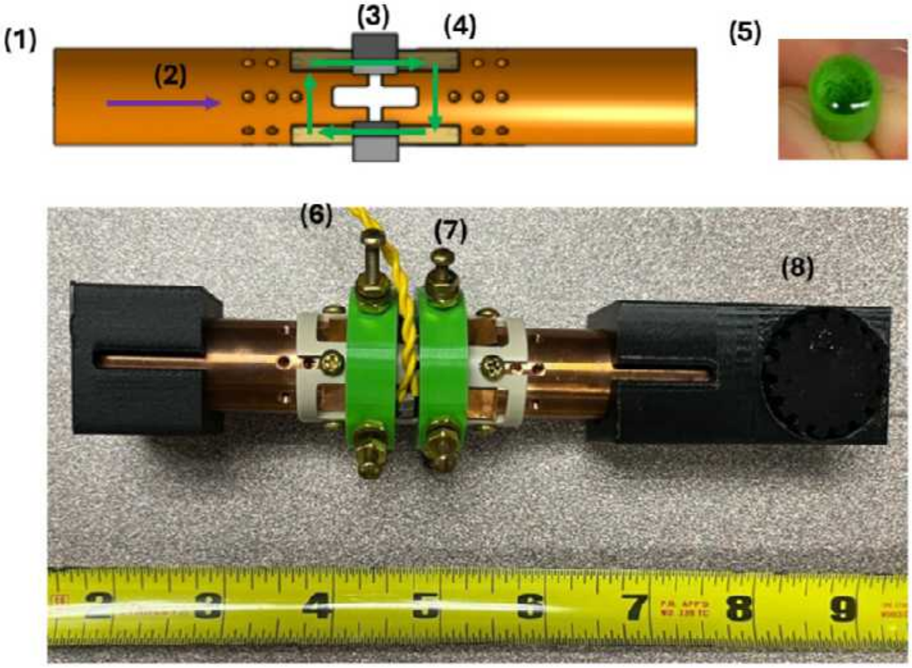

Figure 1.14-1. A prototype of the combined NMR probe circuit and CRES cavity. (1) A CAD model of the birdcage cavity, showing: (2, in purple) the orientation of the main $\sim$ 1 Tesla magnetic field along the axis of the cylindrical cavity, (3) the lumped-element capacitors that set the resonance frequency of the NMR probe circuit, (4, in green) the current path of the circulating currents in the NMR probe circuit, which generate an oscillating transverse field through the windows at the cavity midplane and allow the spins to be tipped. (5) a 3D-printed sample holder filled with petroleum jelly, which provides a dense ensemble of hydrogen nuclear spins, which was used to verify the birdcage design and capture the NMR signals shown in Fig. 1.14-2. (6) The twisted pair transmission line used to tip the spins and read out their precession, (7) brass screw terminals used to make electrical contact between the capacitor leads and the birdcage cavity, and (8) a 3D-printed pulley used to adjust the sample position along the axis of the cavity, which allows the magnetic field distribution of the birdcage circuit to be inferred, as shown in Fig. 1.14-2.

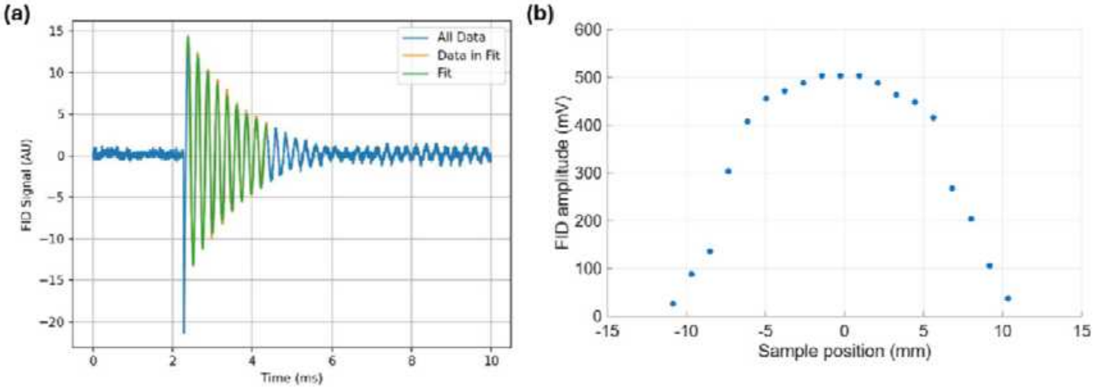

Figure 1.14-2. (a) A free induction decay (FID) NMR signal from hydrogen in the birdcage system, showing high-SNR detection of the petroleum jelly sample. Because the hydrogen and helium-3 gyromagnetic ratios are similar, petroleum jelly is a useful proxy sample for developing and validating the birdcage circuit. The FID signal is amplified and mixed down using electronics adapted from the muon g-2 proton magnetometry system. (b) The amplitude of the FID signal as the petroleum jelly sample is moved along the axis of the birdcage axis, showing that the transverse field is localized to the center of the cavity, near the windows.

1.15 The UW atomic source test stand

P. Doe, S. Enomoto, and M. Wynne

A near-term responsibility of the UW in developing the Project 8 atomic source is to provide an apparatus that will be used to:

- Characterize the electron cyclotron resonance (ECR) source currently under joint development between UW and Indiana University,
- Provide a beam of $\sim 10^{19}$ atoms of an atomic hydrogen isotope at $\sim$ 40K that will be used to develop a magnetic evaporative cooling beamline (MECB) to feed $\sim$ 1 mK atoms into the detector cavity.

The apparatus currently uses a commercial dissociator, the Hydrogen Atomic Beam Source (HABS), to provide the atomic beam. An accommodator, which is a small cold aluminum tube, cools the beam. The chopper is a rotating stainless steel disc with two ten-degree slits. The chopper wheel interrupts the atomic beam and generates timing signals that can be used in conjunction with the Hiden mass spectrometer to measure time of flight and thereby derive the temperature of the beam. The spectrometer can be traversed across the beam to determine its profile. A wire scanner based on one developed at Johannes Gutenburg University in Mainz has been installed downstream in the fourth vacuum chamber to independently record the beam profile and flux. The apparatus is shown in Fig. 1.15-1.

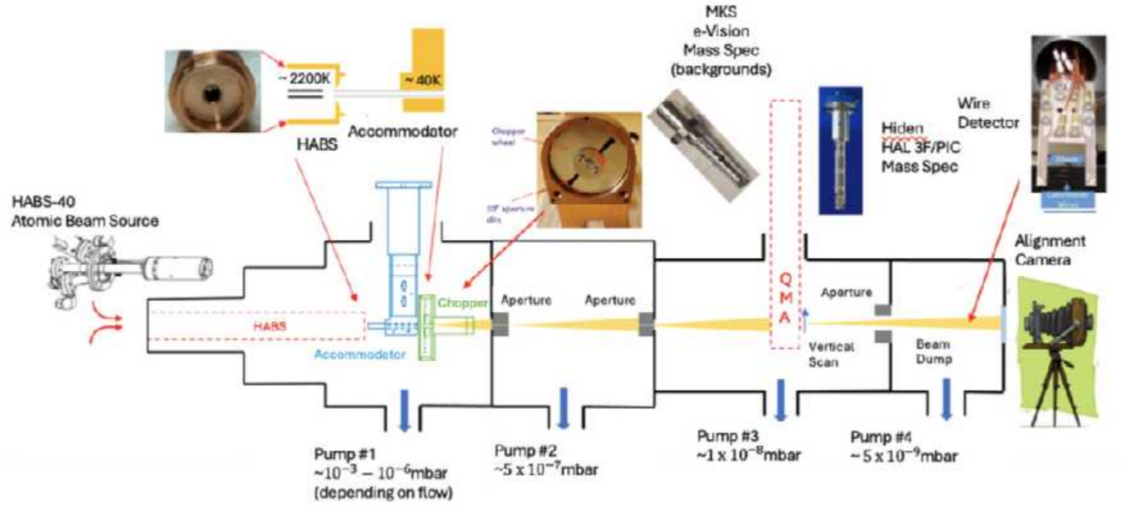

Figure 1.15-1. Atomic Tritium Demonstrator Test Stand

Efforts to cool the atomic hydrogen beam and measure the flux have been the focus of the last year. We have produced a new design for an accommodator as well as installing and commissioning the wire detector.

The purpose of the accommodator is to collimate and cool the hot hydrogen atoms coming from the source while minimizing recombination. Without the accommodator present, atoms from the HABS dissociator have an angular distribution with a sharp minimum on the axis at low flow rates. The result is a weak beam of hot atoms instead of a collimated beam of cold atoms. The present design for the accommodator includes an aluminum transfer tube that sits 1 mm away from the HABS and a barrier in the cylindrical body. The transfer tube enhances the capture of hot atoms and the barrier forces these hot atoms to bounce on a cold aluminum surface before exiting the accommodator. The size and location of the barrier has been selected based on the velocity distribution and bounce numbers generated with MolFlow simulation software. Fig. 1.15-2 shows the cross section of the current version of the accommodator.

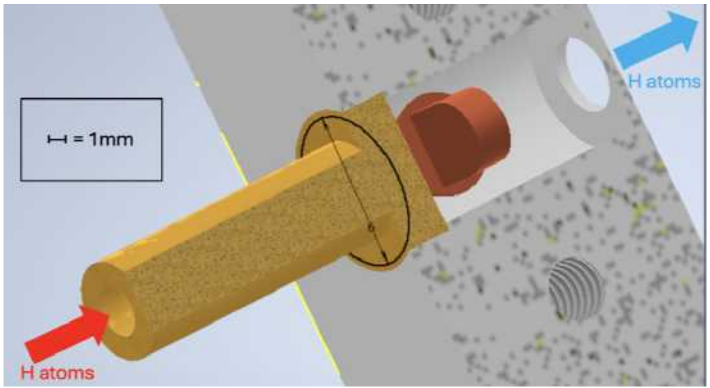

Figure 1.15-2. CAD cross section of new accommodator. The aluminum transfer tube (gold) guides atoms from the HABS to the cylindrical body of the accommodator (grey). An aluminum barrier (red) ensures that atoms cannot exit the accommodator without bouncing on its cold surface.

Currently, the Hiden mass spectrometer is the only instrument in the beamline that measures atomic flux, among its many other applications. Once the MECB is designed and commissioned, the resulting magnetic fields on the beamline will render the QMA inoperable. Collaborators at JGU Mainz have developed a wire detector that measures the atomic flux using the heat produced when atoms combine into molecules on the surface of an extremely thin platinum wire. When two atoms recombine, 2.23eV of energy is deposited on the wire. The estimated atomic flux on the wire is 5×10^9 atoms per second, which produces nanowatts of power, heating the wire and resulting in a tiny but measurable change in its resistance. Using a reference signal synchronized with the beam chopper, shown in Fig. 1.15-1, a lock-in amplifier is used to detect this small signal.

Initial tests found that the lock-in amplifier's $1/f$ input noise was obscuring the signal. The time constant of the wire's thermal response required a frequency of 1 Hz, but at this frequency the noise was much stronger than the signal. Thus, a pre-amplifier, designed by CENPA engineers, was recently installed to boost the signal before reaching the lock-in amplifier.

1.16 ECR source for atomic tritium

M. Kallander, M. Oueslati*, A. Reine*, W. C. Pettus*, and R. G. H. Robertson

We are building an electron-cyclotron-resonance source for the production of intense currents of neutral atomic hydrogen and deuterium, and eventually tritium. Such a source is an alternative approach to the method of dissociation on hot tungsten for making the atomic tritium required in a more sensitive neutrino mass experiment.

Work is being carried out in collaboration with our Project 8 partner Indiana University, where construction of the prototype source is now in its final stages. A physical model of the

---

*Indiana University.

conditions in the plasma needed to optimize the atomic fraction has been developed. The 2.45-GHz cavity design, initially TE011, is now in the TE111 mode to reduce the volume. In our model, this reduces the required microwave power about fourfold. Calculations by Tianhuan Luo (LBNL) in the CST software have enabled the design of the sapphire window for bringing microwaves into the cavity and provided an estimate for the resonant-frequency shift as the plasma forms.

The microwave power is provided by a tunable 1-kW solid-state generator coupled via a cable and a WR-284 transition into the cavity through the sapphire window. The magnetic field of 0.0875 T is produced by two of the magnet coils used in the TRIMS experiment. The cavity is cooled by water and pumped by a turbomolecular pump. An aluminum accommodator coupled to the aluminum cavity cools the beam to room temperature and (in the future) to below when a cryocooler is installed. A minimal gas system with a capacitance manometer for measuring the cavity pressure supplies hydrogen to the cavity. The apparatus is illustrated in Fig. 1.16-1.

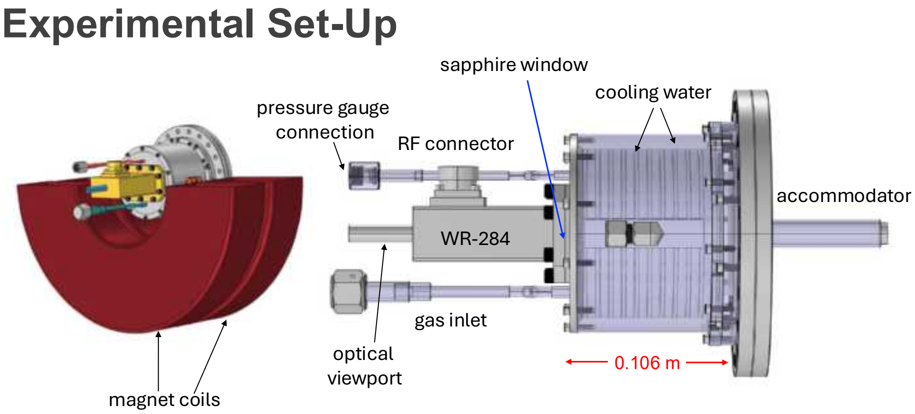

Figure 1.16-1. The Project 8 ECR source design.

Source test is expected to begin early in 2026. Among the key performance characteristics are the operating pressure, optical spectrum, dissociation fraction, beam intensity, tune shift, magnetic field profile, and persistence of operating conditions over long periods.

1.17 Low-Frequency Apparatus (LFA)

R. G. H. Robertson

The past year has seen important developments for the next phase of Project 8, the Low-Frequency Apparatus (LFA). The LFA will be a capital project that takes Project 8 to much larger physical scales and to the frontier of sensitivity of any experiment to neutrino mass.

By lowering the frequency from its current 26-GHz value in the CCA, the size of the cavity is increased, and at $\sim$ 500 MHz it reaches the cubic-meter volume that extends the sensitivity into the sub-eV range, potentially to the sensitivity of KATRIN. Our design has centered on a 5-m long, vertical TE011 cavity at 560 MHz with a 20.7-mT magnetic field. When an atomic tritium source becomes available, the apparatus becomes a highly capable neutrino mass experiment. The vertical orientation confines the atoms gravitationally if they are at a temperature of $\lesssim$ 5 mK. A 64-pole Ioffe superconducting magnet establishes a magnetic potential barrier that keeps atoms from striking the wall to the sides and bottom of the trap, while the top remains open for instrumentation, RF readout, and pumping. The estimated sensitivity of such an apparatus is about 0.5 eV in one year of running.

The atomic source design is based on production of atomic tritium in a dissociator (the ECR source) followed by magnetic capture and evaporative cooling for injection into the Ioffe trap. The success of evaporative cooling depends on the thermalization of the beam as it moves along the magnetic guide.

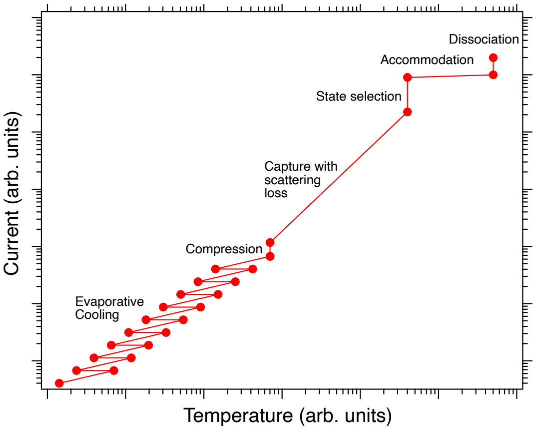

Figure 1.17-1. Conceptual cooling path for atomic tritium in Project 8.

The length of the needed beamline, in turn, depends on the scattering cross section. For atomic H the known cross section is quite small, and a very long beamline is necessary. Our Project 8 collaborator Ben Jones at the University of Texas, Arlington, discovered literature calculations that showed that the cross section for atomic T scattering was thousands of times larger owing to the chance formation of a near-threshold resonance in the more massive system. He and his student Morgan Elliott confirmed these results. This was good news unexpectedly for the design of the evaporative cooling line. Together with Jones, we wrote and published a paper on the dynamical evaporative cooling concept (Phys. Rev. A **112**, 033311 [2025]). The general scheme for evaporative cooling in Project 8 is shown in Fig. 1.17-1 in that paper.

We began to worry, however, that if the cross section was much larger, then perhaps dipolar scattering spin-flip would also be more probable. This loss mechanism is subdominant among a dozen various mechanisms that determine the overall lifetime in the Ioffe trap. Jones and Elliott undertook the full quantum-mechanical coupled-channel calculation for dipolar scattering at a non-zero temperature and a non-zero magnetic field, a calculation that had not even been done for hydrogen (Phys. Rev. A in press). Our fears were confirmed; the dipolar loss rate increased by a factor of ~ 50 and became a serious and dominant issue.

As a two-body interaction, the dipolar loss rate per atom scales with the square of the density. In the vertically oriented Ioffe trap, the gravitational confinement is synonymous with a strong vertical density gradient, and the new loss rate became unmanageably large. It looked impossible to supply enough atom current to keep up with it. An urgent action was initiated in the collaboration to consider whether a horizontal design might be possible, which would reduce the atomic demand by a factor of 3 to 4. The main difficulty is closing both ends of the Ioffe cylinder with conical Ioffe windings because one loses the open access needed for instrumentation and pumping. A joint between two independent Ioffe windings is also required, which could be a source of additional atomic losses. A possible design for a Ioffe cone with slots to allow pumping and the passage of electron beams is shown in Fig. 1.17-2.

Figure 1.17-2. Slotted Ioffe cone for a horizontal LFA apparatus.

The slots shown provide sufficient pumping speed for the interior to keep the density of ^{3}He low and to allow efficient recycling of frozen T_2. Transmitting electron beams through the slots is more difficult because the motion is non-adiabatic in the inhomogeneous Ioffe field, as described below. It may be necessary to switch off the Ioffe cone during calibration with the egun.

Significant progress has been made toward moving to a horizontal LFA, a provisional decision was reached at the October collaboration meeting, and work is progressing to address a number of open questions.

1.18 Ioffe electron-gun access simulation

Y. Yin

A controlled, well-characterized electron source is essential for the Low-Frequency Array (LFA), enabling precision frequency calibration, magnetic field mapping, and systematic studies. We performed a dedicated simulation study to evaluate electron transport through the existing Ioffe cone geometry and to determine under what injection conditions electrons can traverse the cone region with minimal change in pitch angle. Preservation of the pitch angle is critical to maintain predictable cyclotron radiation characteristics and to enable reproducible calibration measurements.

The study focused on quantifying the change in the pitch angle of the electrons, $\Delta\theta = \theta_{\text{final}} - \theta_{\text{initial}}$, after the electrons propagate through the Ioffe cone region. A total of 63,466 electrons were simulated using the `Kassiopeia` framework with full Lorentz-force tracking, *without* invoking the adiabatic approximation. Of these, 20,000 electrons were uniformly initialized in all four transverse quadrants, with an additional 43,466 electrons sampled in the first quadrant to increase spatial resolution. All electrons were monoenergetic at 18.6 keV and injected with an initial pitch angle of 180°, a.k.a. parallel to the axis. The starting longitudinal position was fixed at $z = -0.06$ m, with transverse starting positions drawn from a uniform random distribution. Electron trajectories were tracked through the Ioffe cone region spanning $z = 2.4$–3.2 m, with a maximum simulation time of 10 μs.

A series of internal consistency checks was performed to validate the simulations. Changes in kinetic energy, total momentum, cyclotron frequency, and local magnetic field magnitude were found to be negligible for the vast majority of tracks, indicating stable numerical integration and confirming that the observed effects arise primarily from geometric coupling between the particle trajectories and the magnetic field structure. The pitch angle therefore served as the primary diagnostic for the transport quality of each track.

Clear and highly structured spatial patterns emerged when $\Delta\theta$ was plotted as a function of both initial and final transverse position (Fig. 1.18-1).

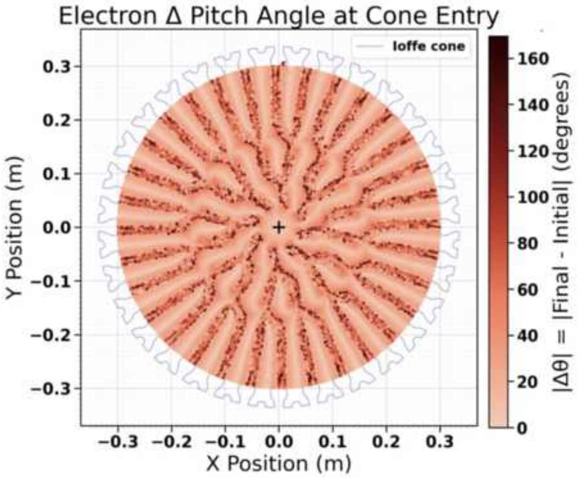

Figure 1.18-1. The change in the electron's pitch angle (z-axis) at initial and final time as a function of the x-y coordinates of its entry point on the Ioffe cone. Darker colors correspond to a larger change in pitch angle. The Ioffe geometry top-down view from the z-axis is overlaid as reference. About 2 million simulated electrons are in this plot. We can see clear spatial patterns caused by the Ioffe magnetic field. Example tracks of electrons with high or low changes in pitch angle can be seen in Fig. 1.18-2

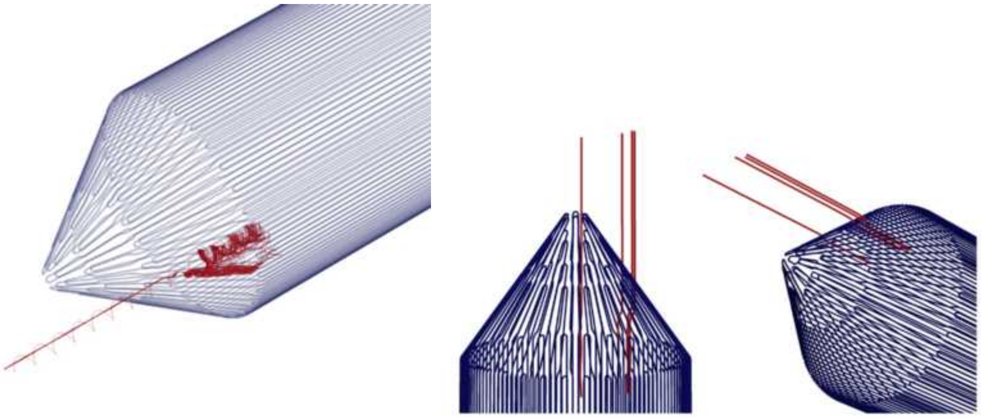

Figure 1.18-2. Example simulated electron track with high change in pitch angle ($|\Delta\theta| > 90°$). *(Left:)* The electron hits the wire of the Ioffe cone and is deflected. *(Right:)* The electrons are able to travel through the Ioffe cone region with minimal change in pitch angle.

These patterns closely follow the geometric features of the Ioffe cone and are symmetric across all four quadrants, demonstrating that the observed behavior is intrinsic to the magnetic geometry rather than a numerical artifact. The regions of favorable and unfavorable transport are sharply delineated by the cone structure.

Electrons were classified as "good" tracks if $|\Delta\theta| < 5°$. Under this criterion, 3,655 electrons (5.8% of the total) successfully traversed the cone with minimal pitch-angle distortion. A more stringent cut of $|\Delta\theta| < 1°$ identified 388 electrons (0.6%) with exceptionally straight trajectories. These electrons originate from narrow bands in the transverse phase space and propagate through the cone nearly parallel to the local magnetic field. In contrast, "bad" tracks with $|\Delta\theta| > 90°$ accounted for 5,616 electrons (8.8%), typically originating from regions strongly misaligned with the cone geometry and undergoing large pitch-angle scattering. Examples are shown in Fig. 1.18-2.

Interestingly, a subset of good tracks was found very close to the physical boundary of the Ioffe cone, indicating that favorable transport is not restricted to the cone axis and may be exploited with careful source positioning. These results demonstrate that electron access through the Ioffe cone is feasible but highly sensitive to injection position and local field geometry.

Future work will extend this study by varying the initial pitch angle and energy, incorporating realistic coil geometries and collision masking, and directly comparing full-orbit and adiabatic tracking under identical injection conditions. These efforts will further inform optimized electron-gun placement and operating conditions for LFA calibration.

Selena

1.19 4D simulation and reconstruction framework of a Selena neutrinoless double beta decay search

A.E. Chavarria and X. Ni

We developed a full four-dimensional (4D) simulation and reconstruction pipeline to model the Selena aSe/CMOS detector response to ionization tracks from neutrinoless double-beta decay. The framework combines GEANT4 charge-transport modeling, realistic detector noise, time-of-arrival (TOA) sampling, and new algorithms for noise suppression, topology reconstruction, depth estimation, and energy measurement. This allows end-to-end evaluation of detector performance at the Q-value of ^{82}Se.

To model the detector, charge ionization and recombination in aSe are modeled in GEANT4[1]. We simulate a 5 mm-thick layer of aSe with an applied field of 50 V/μm. Charge carriers drift according to standard charge transport laws, with hole lifetimes varied from 40-400 μs. We simulate pixels of 15 μm pitch, with an 8-ns TOA circuit and 10 e^- ENC. Because the pixel pitch is much smaller than the aSe thickness, the sensor operates in the small-pixel regime.

All GEANT4 energy depositions are projected onto a 15 µm × 15 µm × 15 µm voxel grid, corresponding to the pixel pitch and the 8-ns TOA sampling of drifting charge in aSe. Because only relative drift times are observable, the earliest-arriving charge defines z_0 which is subtracted from each event. The resulting event contains full 4D information (x,y,t,c), enabling 3D topology reconstruction and charge-based energy measurement.

To reduce data volume and remove background hits, we implemented a two-stage noise filter. We first remove pixels by amplitude, masking all pixels below 2σ. Then, we apply a spatial noise mask. We developed an algorithm that evaluates whether or not a pixel is a track or background by its proximity to other pixel clusters. First, we retain all pixels greater than 5σ, then remaining pixels must have a minimum number of neighbors within a radius r_{th} in order to mask isolated background pixels. The result is double beta tracks, which are significantly cleaner while preserving track topology. Tracks are shown in figure Fig. 1.19-1, which show tracks generated in GEANT4, after they have been projected onto the pixel array and after noise-masking.

We then developed a track-labeling algorithm based on connected-component growth in space. The algorithm begins by starting at a seed pixel, then expands to adjacent neighbors, and assigns sequential labels. These allow us to trace the progression of the track starting from the seed. We identify branches in the track by monitoring sudden large increases in charge-cloud variance and mark these regions. The output is a simplified “track-skeleton” that captures the connectivity and 3D topology of the ionization track.

In the simulated detector, the absolute depth of the ionization event cannot be determined from the TOA alone because the ionization start time is unknown. Thus, we developed a

---

[1] X. Li *et al*, *JINST* 16 PO6018 (2016).

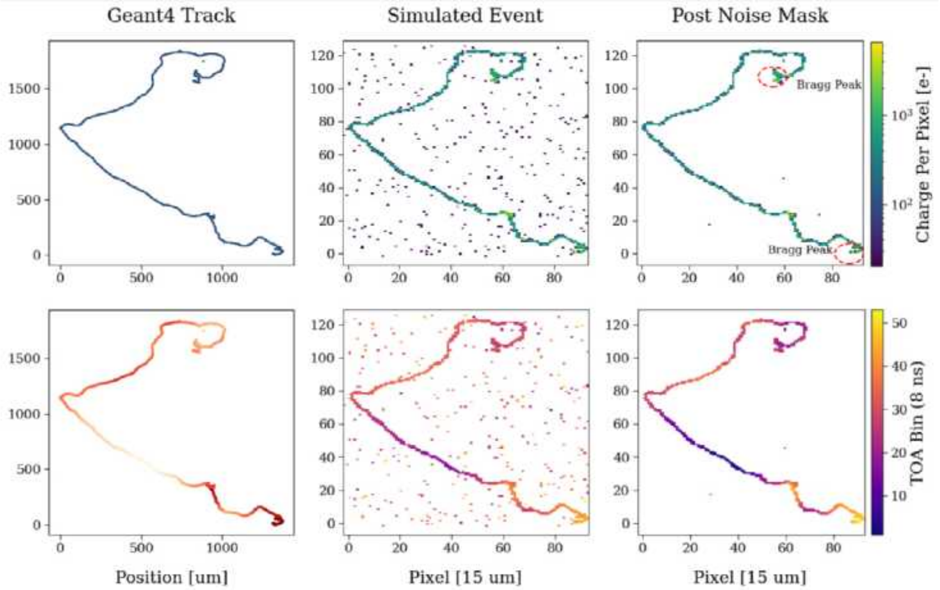

Figure 1.19-1. Simulated $0\nu\beta\beta$ track in a Selena detector. Field applied is 50 V/μm. We simulate charge diffusion and trapping with a hole lifetime of $\tau = 100$ μs. Pixel pitch is 15 μm, with 10 ENC and an 8 ns TOA circuit.

diffusion-based estimator for z_0. For a pixel group with a given track label, its charge-weighted variance satisfies $\sigma^2 \propto z_0 + z_{\rm rel}$. Using simulated tracks with known drift depths, we mapped the distribution of $\sigma^2 - z_{\rm rel}$ to the true z_0. For full $0\nu\beta\beta$ tracks, the reconstruction achieves ~1300 μm RMS. For straight-line toy tracks, we obtain ~500 μm RMS, representing the algorithm's limit.

We reconstruct the total energy of $0\nu\beta\beta$ events using a voxel-level energy function applied to all noise-filtered pixels. This function is derived from simulations of the mean energy deposited in a voxel as a function of the charge therein. Multiple detector configurations were evaluated to study the effects of noise, charge trapping, and reconstruction accuracy. A selection of results are summarized in Table 1.19-1. Approximately 7% of events reconstruct at artificially low energies due to either partial containment or highly fragmented topologies. A dominant contribution to the energy RMS is pixelation noise – discretizing a continuous charge track onto a fixed pixel grid. Varying the pixel grid alignment produces ~0.4% RMS fluctuations even under ideal conditions. Finer pixels reduce this effect but increase ENC due to a larger number of pixels contributing to the track.

Detector Case	**RMS Resolution**
0 ENC, No Trapping	1.6%
0 ENC, $\tau = 100$ μs, No Z-recon	1.8%
0 ENC, $\tau = 100$ μs, Z-recon	1.7%
10 ENC, $\tau = 100$ μs, Z-recon	1.9%

Table 1.19-1. Energy resolution of $0\nu\beta\beta$ events in a Selena detector for various different simulation cases. Pixel pitch is 15 μm.

Finally, we investigate the implications for Selena's physics reach. Double-beta reconstruction is central to the sensitivity to neutrinoless double-beta decay. Monte Carlo studies of a 100-ton-year exposure using the full reconstruction chain and background model[1] yield an expected $0\nu\beta\beta$ half-life sensitivity of $T^{0\nu}_{1/2} \approx 2.3 \times 10^{27}$ years.

1.20 Amorphous selenium coupled TopmetalSe detectors

A. E. Chavarria, M. Conde, and X. Ni

We report on the design, fabrication, and characterization of TopmetalSe sensors coupled to amorphous selenium (aSe). We successfully deposited and tested 545-µm aSe coatings on both flip-chip and wire-bonded TopmetalSe devices.

The packaging approach followed two routes, previously detailed in Ref. [2]. The flip-chip assemblies consisted of TopmetalSe bonded to the underside of a sapphire interposer with a laser-cut window directly above the pixel array. Subsequently, a 545-µm aSe layer was deposited on the entire top surface of the interposer. Wire-bonded devices consisted of TopmetalSe mounted on a PCB, which was then coated with an equivalent aSe thickness. All aSe layers were deposited by Hologic Inc., together with a proprietary hole-blocking layer and a 22-nm chromium electrode. High-voltage (HV) contact was made using a 1-mil aluminum bonding wire attached to the top electrode via carbon paint. Devices are shown in Fig. 1.20-1.

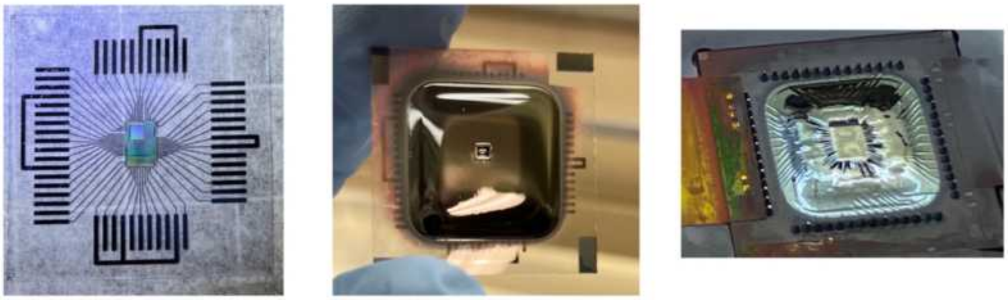

Figure 1.20-1. *Left* shows TopmetalSe flip-chip package pre amorphous Selenium coating, *Middle* shows TopmetalSe coated with 545 μm of aSe. *Right* shows wire-bonded device with aSe coating.

The testing was carried out in a vacuum chamber ($\leq 10^{-3}$ mbar) under various HV biases supplied by a CAEN DT1570ET module and using ionization sources. Both packaging schemes survived evaporation, though each presented distinct challenges. In wire-bonded devices, directional aSe deposition coated the exposed wires, creating parasitic conduction paths. Under bias, the HV supply current rose to the µA level at 1 kV – far above the expected leakage of the pristine aSe. Because this current did not appear in the pixel array, we concluded it flowed through selenium-coated wires. Prolonged operation led to HV trips and irreversible sensor damage.

---

[1] A.E. Chavarria, Proceedings of Low Radioactivity Techniques (2022).

[2] CENPA Annual Report, University of Washington (2024) p. 41.

Flip-chip devices provided more stable HV operation but revealed issues related to deposition geometry. The sharp step at the interposer window inhibited continuous Cr electrode coverage across the pixel-array pocket, leaving the top electrode over the array electrically isolated from the outer electrode. This was remedied by adding an additional HV contact directly into the pocket. The aSe layer above the array also exhibited significant sloping and thinning near the edge of the window due to the directional evaporation process. Additional failures were caused by displacement of the HV wire, which required mechanical stabilization.

After mitigation, flip-chip devices sustained HV biases between 2–5 kV (3.7–9.1 V/µm). Stable operation pulled $\leq$20 nA from the supply, with pixel-inferred leakage currents of the order of 100 pA. Leakage was concentrated near the edges of the array, consistent with aSe thinning and field distortions arising from the deposition profile.

We successfully recorded ionization signals in the flip-chip aSe devices. Figure Fig. 1.20-2 shows single beta tracks from a ^{90}Sr source and gamma interactions from a ^{57}Co source under a field of nearly 9 V/µm. Beta events exhibit long, worm-like trajectories terminating in distinct Bragg peaks, whereas gamma interactions appear as compact, high-SNR clusters. To improve track visibility, we applied a spatial noise filter that removes isolated pixel clusters that are not related to the main ionization structure.

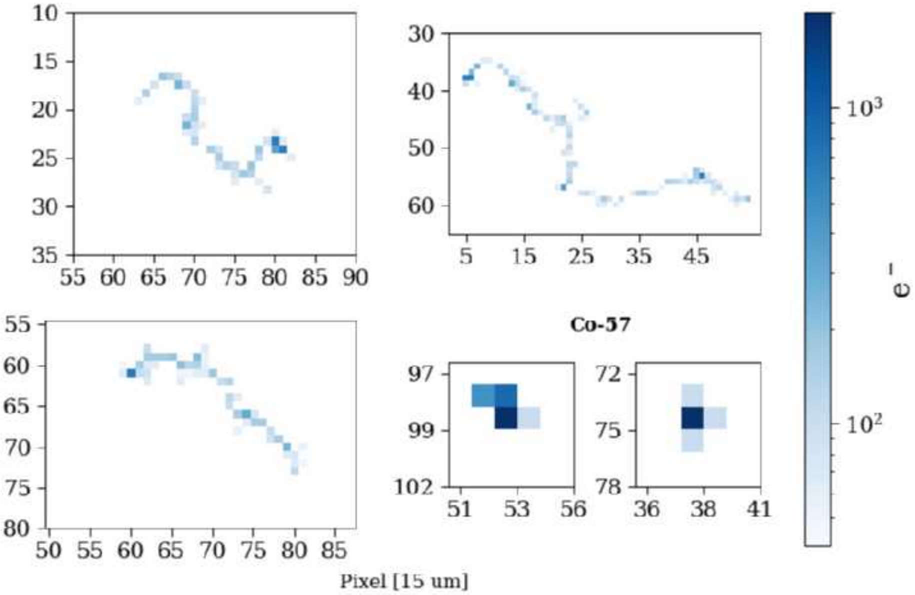

Figure 1.20-2. Ionization tracks in aSe coated TopmetalSe detectors from ^{90}Sr and ^{57}Co sources. Field applied is 9 V/μm. For single beta tracks, we apply a spatial filter to isolate the track topology.

Although ionization imaging was successful, we were unable to resolve a clear monoenergetic peak in the gamma spectrum. We attribute this to limited achievable bias fields and nonuniformities in the aSe deposition geometry. To address these issues, we are revising our packaging strategy and investigating the use of printable electronic traces[1] that would allow the pixel array to lie flush with the deposition surface, improving field uniformity and deposition quality.

1.21 Direct charge calibration of TopmetalSe

A.E. Chavarria, M. Conde, E. Solis, Y. Mei*, and X. Ni,

We provide updates to the calibration of TopmetalSe using direct charge injection.

The calibration test setup consisted of a 235 nm UV LED focused onto a 7 nm gold sputtered onto a sapphire 200 um substrate which we call the photocathode. Focusing is done with a set of silica lenses inside an adjustable lens tube. Current measurement of the photocathode and biasing between the photocathode and pixel array are done by the Source Measurement Unit (SMU). Finally, the pixel array is placed in line of sight of the LED and photocathode for charge to be injected onto it and obtain a voltage measurement from the in-pixel amplifiers.

The procedure consists in the UV light-emitting diode shining light on the photocathode to generate electrons via the photoelectric effect. This is done under atmospheric conditions, for which the free electrons emitted are short-lived as they ionize molecules they collide with[2], so that the direct charge injected would be in the form of oxygen ions.

We could verify that we are measuring ions as a result of the single pixel effect by observing a delay in the signal of irradiated pixel. With this delay, we could approximate the mobility of those ions. For our setup, the right graph in Fig. 1.21-1 shows an ion mobility of 2.60 ± 0.05 $\mathrm{cm}^2/\mathrm{Vs}$ close to the $2.3 - 2.4$ $\mathrm{cm}^2/\mathrm{Vs}$ found in the literature [3].

---

*University of Texas Arlington.

[1] A. Al Shboul, M. Ketabi, I. Ricardo, in NMDC (2021).

[2] L. M. Chanin, A. V. Phelps, and M. A. Biondi. Measurements of the attachment of low-energy electrons to oxygen molecules. Phys. Rev., 128:219–230, Oct 1962.

[3] Okuyama Y, Yasuzawa Y, Sugawara H. Effects of a small amount of H2O on negative ion mobility and ion-molecule reactions in O2 at atmospheric pressure. Electr Eng Jpn. 2023; 216:e23413. https://doi.org/10.1002/eej.23413.

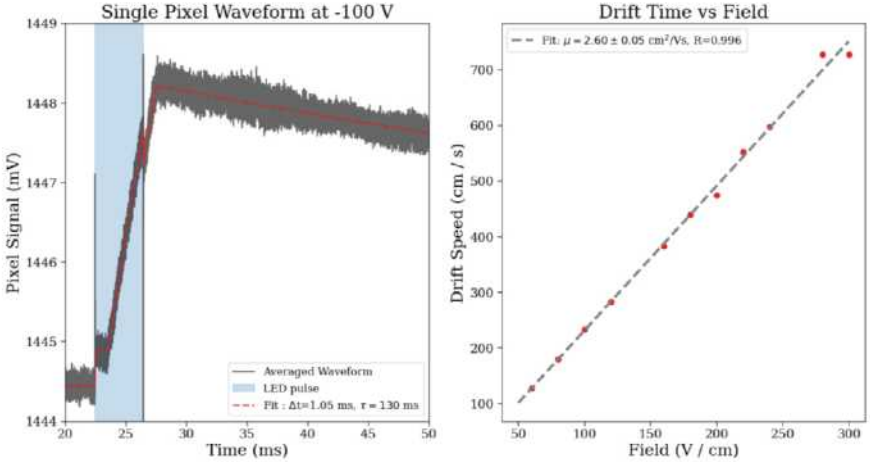

Figure 1.21-1. *Left*: Average single pixel waveform from LED pulse. *Right*: Drift speed of molecules over different electric fields

The image generated from the rolling shutter data shows the voltage measurements for each pixel at any given point in time. However, the measured signal is affected by two factors: the feedback RC circuit of in-pixel charge amplifiers causes an exponential decay of the signal over time; while readouts of each pixel correspond to signals at different points in time. To tackle these, we focus on first characterizing the signal from a pulse for just one pixel. This gave us a model for the signal we could work with for all pixels with signal as shown in the left image of Fig. 1.21-1.

We used the model to reconstruct the signal for the whole array from the rolling shutter data by pulsing the LED at different offset with respect to the beginning of pixel array. The left of Fig. 1.21-2 shows a comparison of the averaged image compared to the reconstructed image by this method.

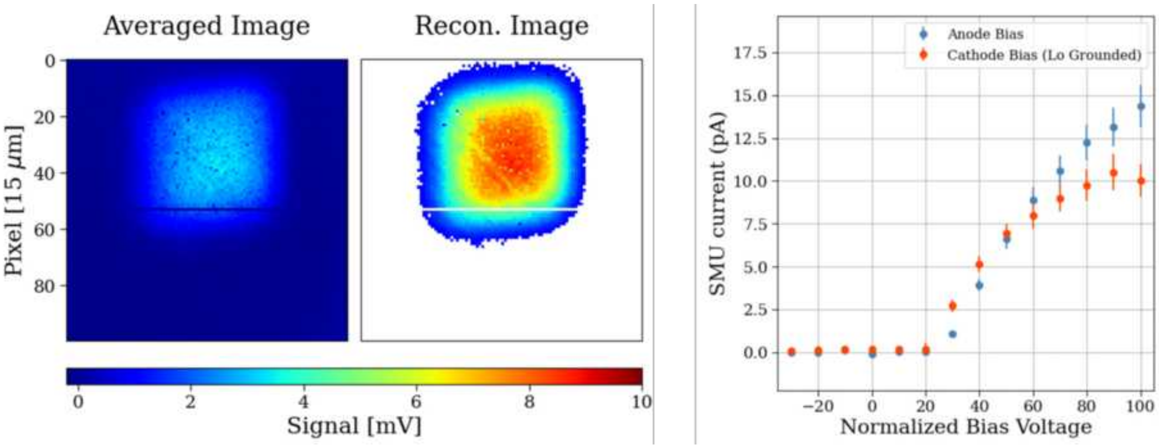

Figure 1.21-2. *Left*: Comparison of average full array image to reconstructed image from the pulse model with improved spatial uniformity and decreased signal decay. *Right*: SMU current variance due to different biasing

We estimate the amount of charge injected by using the SMU to measure the current of the photocathode and the duration of the pulse. Due to the setup of the experiment, we had some difficulty in isolating the circuit from all current sources, and we had to correct for the SMU current by considering differences in the biasing of the setup with the SMU. This introduced some uncertainty in the current value, as shown in the right graph of Fig. 1.21-2.

Taking the total signal from the reconstructed image of the pixel array and the adjusted SMU current for a given field, we are able to calculate the value of the average gain of the whole array: $G = 11 \pm 2$ uV/e^-. This is in line with the gain obtained from the Guard Ring pulse injection method of 10.9 uV/e^- for calibration. The advantage of direct charge injection over the Gring method is that the model does not depend on a simulated value of the capacitance that relies on the geometry of the pixel array; instead, we measure the gain directly and can estimate the internal capacitance from it. We are in the process of retrying the same procedure in vacuum to improve the accuracy of the method and verify that we can reproduce this with electrons directly captured by the pixel array instead of oxygen ions.

2 Precision Muon Physics Group

Muon g-2 Summary

2.1 The muon $g - 2$ experiment

C. Claessens, D. W. Hertzog, P. Kammel, J. LaBounty, B. MacCoy, and E. Swanson

Physics and motivation

The muon's anomalous magnetic moment, $a_\mu \equiv (g-2)/2$ is a special quantity because it can be both measured and predicted to sub-ppm precision, enabling a test for new physics defined by $a_\mu^{\mathrm{New}} \equiv a_\mu^{\mathrm{Exp}} - a_\mu^{\mathrm{SM}}$. As a flavor- and CP-conserving, chirality-flipping, and loop-induced quantity, a_μ is especially sensitive to new physics contributions.

The data taking for FNAL-E989 is complete and this year we reported the final results of the entire six-year campaign. Runs-4/5/6 were combined with a precision on a_μ of 139 ppb. Together with Runs 1-3, the FNAL experiment achieved 127 ppb, exceeding our 140 ppb goal. The results from the 3 published FNAL PRLs [**Muong-2:2021ojo**, **Muong-2:2023cdq**, 0] together with the final BNL-E821 result [**Muong-2:2006rrc**] are shown graphically in Fig. 1. The agreement is excellent and a testament to the very careful and precise work involved in all stages of both experiments. We remind the reader that all analyses were performed blinded.

The final world experimental average is

$$a_\mu^{EXP} = 116\,592\,071.5(14.5) \times 10^{-11}. \tag{1}$$

The test for new physics requires an equally precise Standard Model prediction. Until recently, that was based on the 2020 $g - 2$ Theory Initiative (TI) White Paper, a collective document representing the individual efforts by expert collaborations and formed into one summary recommendation. Unfortunately, an updated TI White Paper [0] in 2025 reports a SM recommendation with a much larger uncertainty and shifted central value.

$$a_\mu^{SM} = 116\,592\,033(62) \times 10^{-11}\,, \tag{2}$$

This "snapshot in time" evolved from the new lattice QCD results for the hadronic vacuum polarization (HVP) contributions and, the omission of all data-driven HVP results because of recent inconsistencies in the data sets that are not yet resolved. Accordingly, sensitivity to any new physics is diluted until the HVP situation can be resolve.

Experimental concept and status

When a muon with charge q is circulating in the horizontal plane of a magnetic storage ring, its cyclotron frequency is $\vec{\omega}_c = -q\vec{B}/m\gamma$. The muon spin precesses at frequency $\vec{\omega}_s =$

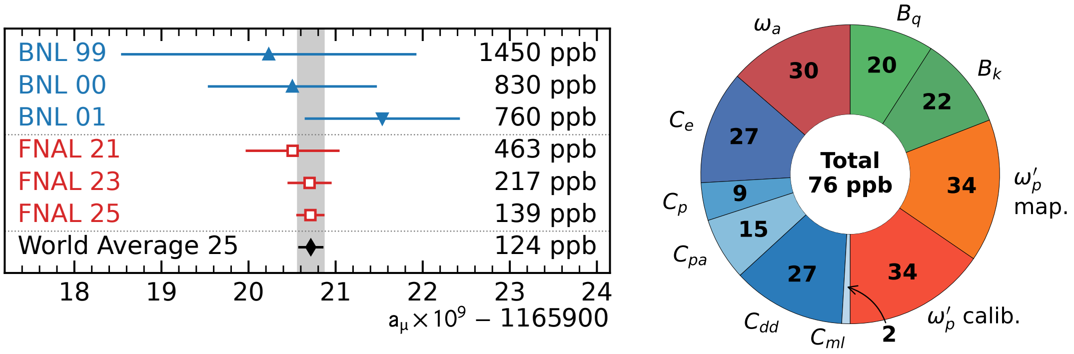

Figure 2.1-1. *Left*: Experimental values of a_μ from BNL E821 (blue triangles) and the world average (black diamond). The uncertainties combine statistics and systematics. BNL values are rounded to the nearest 10 ppb. All measurements employed positive muons except BNL-01, which ran with negative muons (inverted triangle).

$-(gq\vec{B}/2m) - [(1-\gamma)q\vec{B}/\gamma m]$. Assuming a negligible muon electric dipole moment, the anomalous precession frequency can be expressed simply

$$\vec{\omega}_a \equiv \vec{\omega}_s - \vec{\omega}_c = -\left(\frac{g-2}{2}\right)\frac{q\vec{B}}{m} = -a_\mu \frac{q\vec{B}}{m}. \tag{3}$$

Parity violation associates the decay positron energies in the laboratory frame with the average spin direction of the muon at the time of the decay, such that the highest-energy positrons are preferentially emitted when the muon spin is aligned with its momentum, and lower-energy positrons are emitted when the spin is reversed. The UW system of electromagnetic calorimeters is used to measure decay positron times and energies, and from these data determine the anomalous spin precession frequency ω_a.

The magnetic field B is measured in terms of the proton's Larmor precession frequency ω_p. The pNMR system developed at the University of Washington is critical to this effort. In addition, it is necessary to know the muon distribution in the storage ring for the muon population that contributes to the ω_a data. This distribution is folded with similarly determined azimuthally averaged magnetic field moments to give the effective magnetic field seen by the muons $\tilde{\omega}_p$ below. We obtain a_μ through the relation

$$a_\mu^{\mathrm{Exp}} = R\frac{g_e}{2}\frac{m_\mu}{m_e}\frac{\mu_p}{\mu_e}, \quad \text{where} \quad R \equiv \omega_a/\tilde{\omega}_p \tag{4}$$

is measured by our collaboration. The electron g_e factor, the muon-to-electron mass ratio, and the proton-to-electron magnetic moment ratio are known to sufficient precision.

Briefly, the experimental data taking proceeds as follows. Compact bunches of 3.1 GeV/c polarized muons are delivered into our 1.45 T superconducting storage ring at a fill rate of $\sim$ 11 Hz. Their intensity, timing, and spatial profiles are measured with UW built T0 and IBMS counter systems. Once inside the ring, muons are deflected on their first turn onto stable orbits within a defined storage volume by a magnetic outward, pulsed kick. A system of electrostatic quadrupole plates that are charged prior to injection vertically contains the particles. The positrons from muon decay curl inward, where their hit times and energies are recorded in one of the 24 UW-built electromagnetic calorimeter stations. Two calorimeter stations include in-vacuum tracking chambers just upstream of the detectors in order to provide an indirect tomography of the stored muon distribution from tracing the positron trajectories into the storage volume. The magnetic field is monitored continuously by 378 fixed NMR probes built at UW and placed above and below the vacuum chamber. The azimuthally dependent multipolarity of the field is mapped in detail in the storage volume every $\sim$ 3 days using a 17-probe NMR trolley. Representative figures for a subset of the UW-built hardware are shown in Fig. 2.1-2.

Figure 2.1-2. Representative UW-built hardware images. *Left-top:* IBMS-3 design; *bottom:* IBMS-1 detector and electronics. Three IBMS detectors measure the incoming beam intensity, trajectory, and spatial profile. Inner error bars are statistical; outer are total. *Middle:* Photo of an opened calorimeter showing crystals and SiPM readout. 24 such stations measure the decay positron times and energies. *Right-top:* An NMR probe used in the fixed and trolley systems; bottom: a typical free induction decay signal. 378 probes form the fixed-probe system and 17 are embedded in the in-vacuum trolley.

A Look to the Future

The publication of the final results of the Muon $g-2$ Experiment at Fermilab represent the culmination of more than a decade of effort by members of the Muon Physics Group. And beyond that, the completion of this analysis closes a chapter which has been nearly continuously ongoing — in some form — since the BNL experiment published its final report in 2006. The refinement of data processing techniques, analysis strategies, and the level of detail with which the data has been scrutinized over that 20 year period has been nothing short of revolutionary. Still, these techniques have not yet hit their wall.

In the next 5 years, as the theoretical situation settles, it may be that the $g-2$ measurement once again sits in some tension with theory. And in that time, it is expected that Lattice QCD calculations of the HVP contribution to $g-2$ will yield a total theoretical uncertainty on the order of 100 ppb. Should this be the case, a 50 ppb experiment may be considered using an extension of the existing techniques. Since the uncertainty budget is divided relatively evenly between beam dynamics, ω_a analysis, and magnetic field analysis systematics, each of these areas will require additional refinements to reach these goals. Beyond that, to reach the required statistical precision would require another order of magnitude more muons to be analyzed. A small group of current $g-2$ collaborators have started to gather the ideas and improvements necessary to reach this goal, and these are being compiled in an upcoming review article.

2.2 Details of the Run-4/5/6 Precession Frequency Analysis

D. W. Hertzog and J. LaBounty

The Run-4/5/6 Datasets

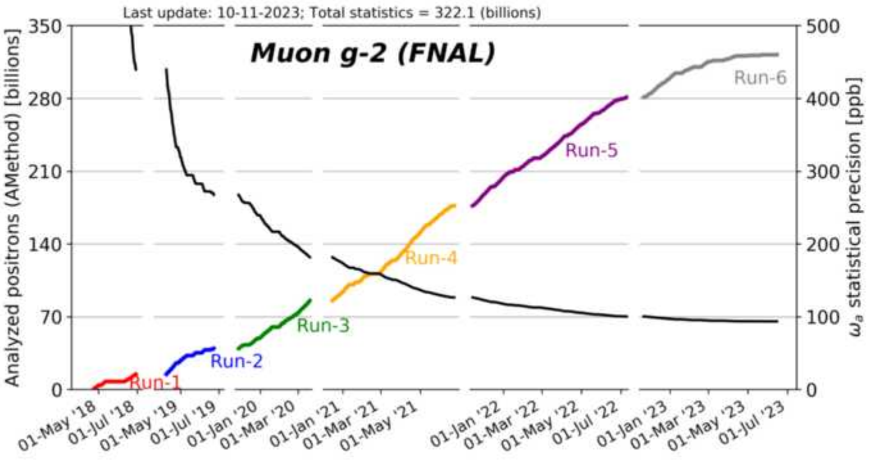

Figure 2.2-1. Integrated statistics of the Muon $g-2$ Experiment

The Run-4/5/6 dataset represented 73% of the total data taken at the E989 Muon $g-2$ Experiment. The statistical power of these datasets meant that subtle effects could be explored with unprecedented precision. Two areas where this power was crucial were the modeling of the beam dynamics terms in the calorimeter fits and the measurement of subtle gain effects. These efforts — each of which had a dedicated task force of analyzers behind them — are described below.

Beam Dynamics Modeling

The five-parameter fit to the $g-2$ precession frequency data (Figure Fig. 2.2-2,left):

$$N(t) = N_0 e^{-t/\tau_\mu}\left[1 + A_0 \cos\left(R\left(\omega_a\right)t - \phi_a\right)\right], \tag{1}$$

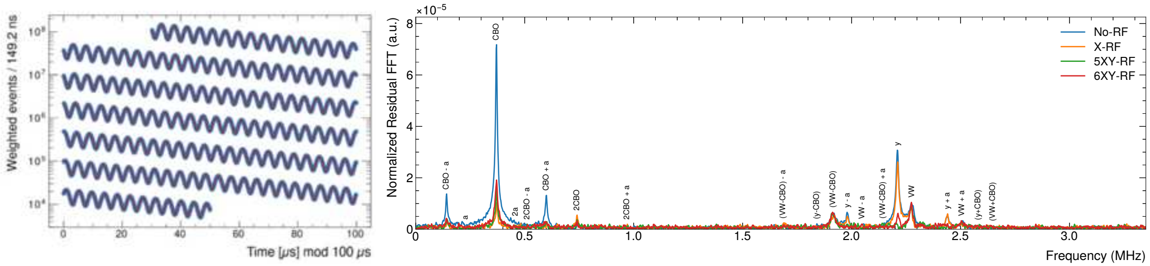

Figure 2.2-2. *Left*The five parameter fit to the Run-4/5/6 datasets (an example of which is shown on the left) yield many residual beam dynamics frequencies in the FFT (right).

describes the data well visually, but results in a χ^2 which is unacceptably large and multiple peaks from Beam Dynamics (BD) related frequencies in FFT of fit residuals (Figure Fig. 2.2-2,right). To resolve these issues, the fit function is typically extended to include these BD frequencies through the modification of the N, A, and ϕ terms:

$$\begin{aligned} N_0 \rightarrow & N_0 \times \prod_i N_i \mathcal{E}_i(t) \\ A_0 \rightarrow & A_0 + \sum_i A_i \mathcal{E}_i(t) \\ \phi_a \rightarrow & \phi_a + \sum_i \phi_i \mathcal{E}_i(t) \end{aligned}$$

where $\mathcal{E}_i$ is a function which describes the evolution of the amplitude of the frequency in the data over the course of the fill. In previous runs, this behavior was typically modeled simply by an exponential decay:

$$\mathcal{E}_i = e^{-t/\tau_i} + C_i. \tag{2}$$

However, two factors in this final analysis that made this approach intractable:

1. the increased statistics allowed analyzers to effectively see deviations from this simple model, which reduced the fit quality; and
2. the introduction of the RF system in Run-5 altered the beam dynamics to a much more complicated dynamical system.

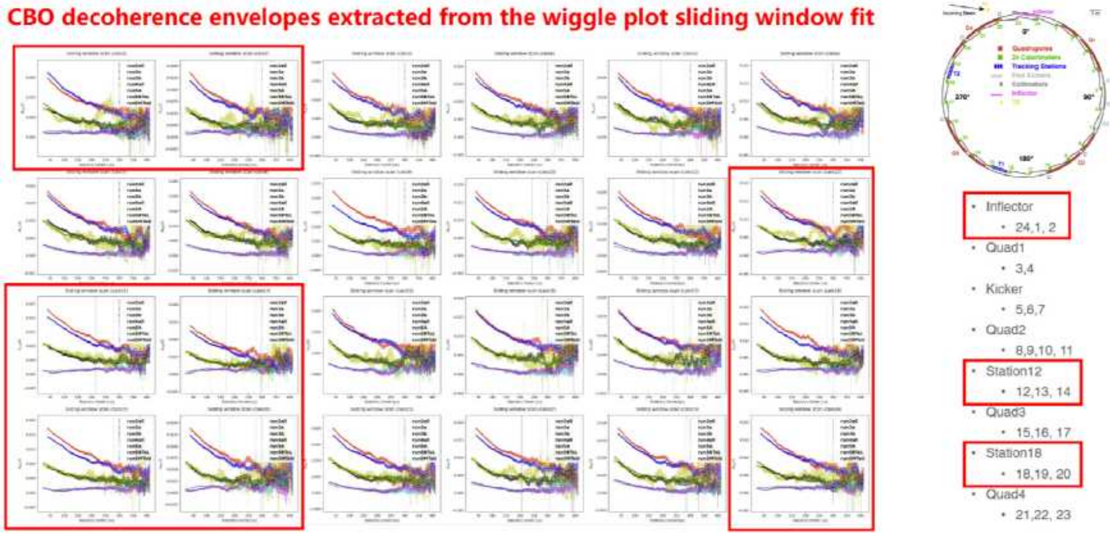

CBO decoherence envelopes extracted from the wiggle plot sliding window fit

- Inflector
  - 24,1, 2
- Quad1
  - 3,4
- Kicker
  - 5,6,7
- Quad2
  - 8,9,10, 11
- Station12
  - 12,13, 14
- Quad3
  - 15,16, 17
- Station18
  - 18,19, 20
- Quad4
  - 21,22, 23

Figure 2.2-3. Amplitude of the calorimeter by calorimeter beam dynamics envelopes for the first order CBO contribution from Run-2 to Run-6. The difference with and without the RF on can be seen in the different curves. In the RF-On case, minute differences in calorimeter acceptance yield very different behaviors which are challenging to model with a single functional form. This is magnified in the regions near the inflector and tracker stations. Reproduced from the BD Taskforce documentation.

An example of the complex dynamic evolution of these frequency envelopes can be seen in Figure Fig. 2.2-3. In response to this increase in complexity, new techniques were developed to produce more accurate models of the data. Fits with a simple model (exponential, or constant BD amplitude) were performed in a narrow 'sliding window' on a calorimeter-by-calorimeter basis. These windows could be as little as 1 ω_a period wide at the beginning of the fill, where statistics were plentiful, to as wide as 75μs late in time. From these windows, a dynamical picture of the evolution of each of the BD amplitudes was obtained and fit. Different analysis groups then set about parameterizing these in the following ways:

- Extending traditional $A(t), \phi(t)$ model with additional terms to handle subtle correlation between BD cross-terms.
- Trading the $A\cos(\omega t - \phi)$ model for $\alpha(t)\sin(\omega t) + \beta(t)\cos(\omega t)$, which eliminated some cross terms at the cost of abstracting some of the physical meaning of the fit parameters away.
- Trading analytical models of $\mathcal{E}(t)$ for spline functions or Gaussian Process Regression (GPR) models determined from the sliding windows.

In total, more than 120 pages of documentation were assembled into a task force report. Analyzers from all 7 analysis groups compared models and arrived at a consensus of recommendations which were adhered to by all analysis groups. In order to quantify the effect of modeling choices on the ω_a result, analyzers performed many bootstrap scans over the sliding window fits to construct an ensemble of models which were statistically consistent with the data. Analyzers also implemented their modeling techniques on the histograms prepared by other groups, to ensure no bias in the underlying data. The difference in the blinded ω_a from swapping between these models was used as a systematic uncertainty in the final analysis.

Residual Gain Effect

The gain stability of the Muon $g-2$ calorimeter system is monitored using a sophisticated laser system [**Anastasi˙2019**]. The system contained six laser heads which could be flashed simultaneously to illuminate the front face of each of the 1296 PbF_4 crystals. The system could also be reconfigured such that the laser light from two separate laser heads could be directed to a single calorimeter, to study gain recovery in double-pulse sequences. This system was designed to measure and correct for:

- the long term changes in gain over the weeks to months of a single dataset (OOF);
- the gain recovery from the initial ‘flash’ of particles at injection, with a recovery lifetime of $\mathcal{O}(\mu\mathrm{s})$ (IFG/LTDP);
- and the short-term SiPM pixel recovery from pulse-to-pulse, with a recovery time constant of $\mathcal{O}(\mathrm{ns})$ (STDP).

This system ran without issue throughout the six running years, and consistently was able to demonstrate gain stability at the 5×10^{-4} level over a single muon fill. However, since the first publication of the E989 $g-2$ Experiment, a persistent ‘Residual Gain-like Effect’ (RGE) was noted after the application of these standard laser-based gain corrections. Many explanations of this effect were proposed — from a true gain change to an acceptance-driven bias in the measured positron energies — but none could be conclusively proven. In the first and second round of $g-2$ publications, a conservative error estimate was included in the ω_a systematics to account for this unknown effect. In preparation for the final publication of the experiment, however, an explanation of this effect was deemed a priority and a task force was assembled to investigate.

The large statistics of the Run-4/5/6 datasets allowed for detailed investigations of the changes in the size of the RGE over many datasets. During these investigations, it was noted that:

- There was no correlation with the amplitude of the RGE with any beam dynamics parameters or effects;
- The size of the RGE vs. time in fill fell with the muon lifetime τ_μ;
- The size of RGE vs. energy bin scaled in a manner consistent with an overall energy scale change.

These results pointed to the RGE being indeed best explained by a true gain change. From there, we investigated where such an effect could manifest that would not be covered by our existing gain corrections. The following hypothesis was developed: if, in addition to the ns-scale perturbation of the STDP, the positrons also created a small μs-scale gain perturbation, then a positron-induced gain sag could build up iteratively over the course of a fill. If the impact of each individual positron was small, it could escape the analysis of the STDP and only introduce a subtle distortion to the IFG. We tested this hypothesis with a Toy Monte Carlo of the IFG recovery, and found that the distortion of the IFG recovery coming from such an effect matched what was seen in the Run-4+ IFG residuals (Figure Fig. 2.2-4).

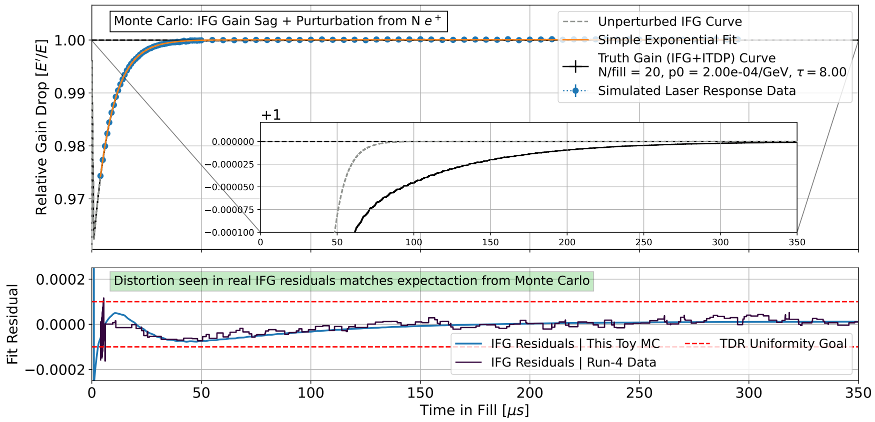

Figure 2.2-4. The addition of positron hits to the flash-recovery model creates a distortion of the pure IFG gain exponential recovery (grey, dashed) and results in the black curve. The presence of the $g-2$ oscillation can be seen in the gain curve. This also introduces a 'wave-like' shape in the MC residuals (blue) which matches what is seen in data (dark purple)

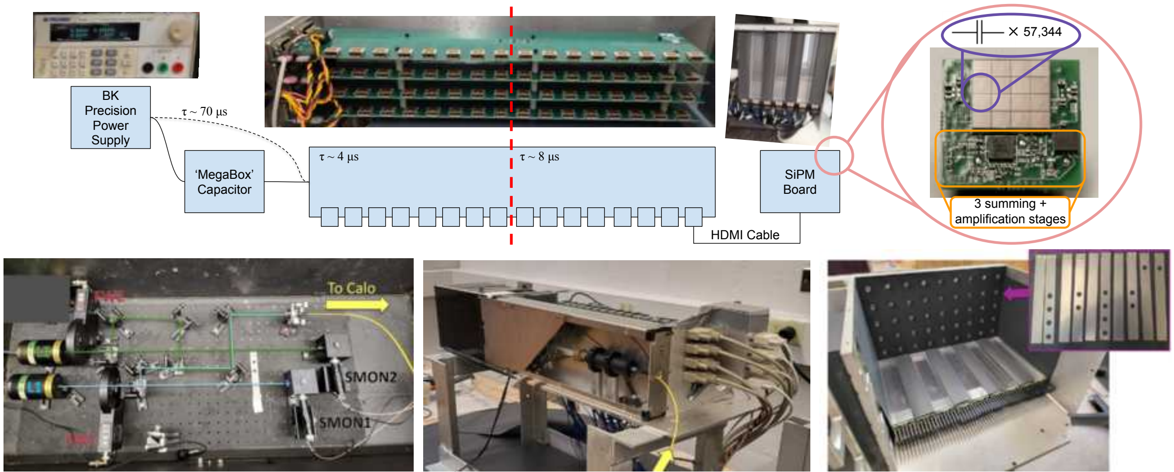

Figure 2.2-5. Electronics chain of the Muon $g-2$ calorimeters. The gain recovery time constants associated with each portion of the chain are annotated where relevant.

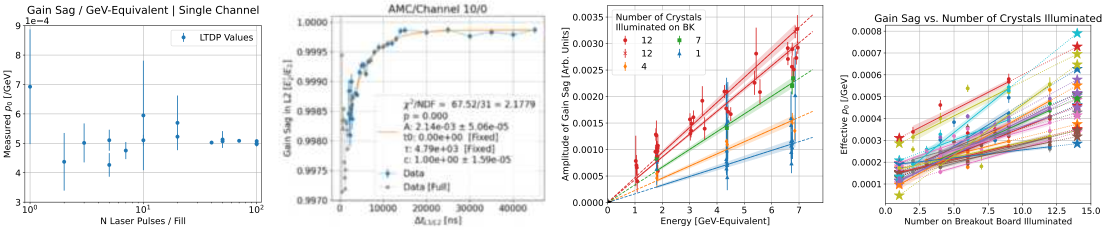

Figure 2.2-6. The progression of the ITDP measurements. First the number of pulses in the LTDP sequences is walked down from 100 to 1 (left). For the single pulse case, a high-statistics measurement is taken to establish a gain recovery curve for each crystall with all illuminated (left middle). For a single crystal, the number of crystals illuminated on a given breakout board is directly proportional to the amplitude of the gain curve (right middle). This trend is consistent across many crystals (right).

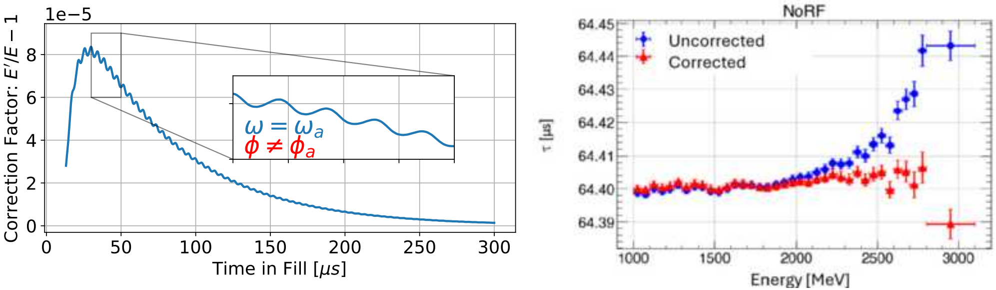

Figure 2.2-7. *Left*: The amplitude of the ITDP correction for the NoRF dataset vs. time in fill. *Right*: The muon lifetime vs. energy before and after the application of the correction. The non-physical upward trend in lifetime is removed with a corrected gain.

We attempted to measure this effect using laser-positron coincidences in existing data, but those tests proved inconclusive. Therefore, after the conclusion of the final running period of the experiment, it was decided that a single calorimeter station and portions of the laser system would be shipped to CENPA to set up a detailed gain sag test stand. This test stand can be seen in Figure Fig. 2.2-5,bottom. Two laser heads were able to illuminate the 52 crystals and 4 monitor PMTs in the setup[1]. A 3D printed laser mask system was inserted between the laser light distribution panel and the front face of the calorimeter crystals, allowing for different numbers of crystals to be illuminated per pulse depending on the configuration of metal laser masks. Laser 1 (the test laser) could be set to produce N pulses to simulate beam injection or a single positron, while Laser 2 (the probe laser) would always produce a single pulse. Filter Wheels (FW) containing nine different neutral density filters could vary the light output from L1 and L2 independently. The FW settings and relative laser timing could be altered using MIDAS ODB variables, allowing for many runs with a variety of $\Delta t_{1,2}$ in quick succession. The DAQ system used for this test was a modified version of that used for the PIONEER LYSO Testbeam, and a custom reconstruction pipeline was written in order to rule out any reconstruction bias from the existing Muon $g-2$ analysis flow.

The analysis of the laser data proceeded in 3 main steps:

---

[1] Two crystals were removed from the calorimeter enclosure, and replaced with monitor PMTs. Two monitor PMTs were added in the laser box itself.

1. Replication of the laser-based gain-sag measurements performed in-situ at Fermilab (LTDP/STDP)
2. Measurement of the existence, amplitude, and energy scaling of the ITDP effect.
3. Measurement of how the ITDP scaled with the number of crystals illuminated in any given configuration.

We found good agreement in the energy scaling and recovery lifetimes for the STDP and LTDP effects between studies at Fermilab and CENPA. From there, we took the standard LTDP mode where the number of test pulses $N_1 = 100$ and walked the number of pulses down from $100 \to 1$. Each time, we saw the size of the gain sag reduce proportionally. At each one of these steps, we varied the FW settings to verify the scaling of the measured gain sag vs. energy. We parameterized the gain sag from a single pulse as:

$$G(t) = 1 - p_0 E_0 e^{-t/\tau_i}, \tag{3}$$

where p_0 (in units of GeV^{-1}) parameterizes the amplitude of the gain sag for a given input of energy to the system. For a train of pulses, the above equation can be generalized to:

$$G(x) = 1 - (1 - e^{-p_0 E_0}) \frac{1 - e^{-(N_0 \delta)/\tau}}{1 - e^{-\delta/\tau}} e^{-x/\tau} \;\; [\text{for } t > t_N]\,, \tag{4}$$

where N_0 is the number of laser pulses in a flash, δ is the spacing between each pulse, E_0 is the energy of the first laser pulse in a train, $x \equiv t - t_N$, and t_N is the time of the N^{th} laser pulse. This series of measurements is illustrated for a single crystal in Fig. 2.2-6 left. We found that p_0 was consistent across different values of N_1, and that even with a single L1 pulse we were able to measure a gain purturbation with a lifetime consistent with the LTDP Fig. 2.2-6 center-left. We also observed that p_0 was consistent with the LTDP values obtained by Fermilab down to a single probe pulse when all crystals were illuminated.

We then began the process of masking off the various channels to determine whether this effect was entirely local to the SiPM, or whether electrically connected crystals (for instance, those sharing a breakout board as seen in Figure Fig. 2.2-5) could influence each others gain recovery. We re-ran the above tests with progressively fewer crystals at a time receiving laser light and saw that, even with the same amplitude laser pulse, a smaller gain sag was observed in a given crystal as the number of other crystals illuminated decreased (Figure Fig. 2.2-6 center-right, right). We made measurements of the reduction in the 'effective p_0' to separate out the 'self' and 'board' contributions. In general, we found that:

$$\begin{aligned} p_0^{self} &= (0.36 \pm 0.09)\; p_0^{LTDP} \\ p_0^{board} &= (0.118 \pm 0.043)\; p_0^{self} \end{aligned}$$

The measured values of p_0^{board} and p_0^{self} were highly correlated ($\approx 86\%$) and uncertainties on the applied values of the parameters were assessed as such.

Once the parameters of the gain sag were known, a crystal-level correction was applied where the energy of the positron E_i was corrected using the entire history of that fill:

$$g_I(t_2) = \sum_{t=0}^{t_2} p_0^{self} E_i \times e^{-(t_2-t_i)/\tau} + \sum_{t=0}^{t_2} p_0^{board} E_i \times e^{-(t_2-t_i)/\tau}. \tag{5}$$

This was found to be equivalent to a histogram-level correction where the contents of each histogram bin $n(E,t)$ were modified by:

$$G(t) = 1 + N_I\, e^{-t/\tau_\mu} \left[1 + A_I \cos\left(\omega_a t - \phi_I\right)\right]. \tag{6}$$

The parameters A_I and ϕ_I were determined entirely by the measured value of the gain recovery parameters (τ_{ITDP}, p_0^{self}, and p_0^{board}) while N_I was proportional to the instantaneous rate of positrons hitting the calorimeter in each fill.

With the known physical origin of the gain-effect in hand, we applied a correction factor to the previous publications: +50 ppb for Run-1 and +47 ppb in Run-2/3. This factor was determined by having the analysis groups apply the updated RGE to their existing histograms from those datasets, re-fit, and measure the shift in ω_a. The impact of the gain-correction is higher in those runs than in Run-4/5/6 because of the higher beam intensity ($N_A \propto N_{e^+}$/fill). The shift in the final value of a_μ from this correction factor can be seen in Figure Fig. 2.2-8.

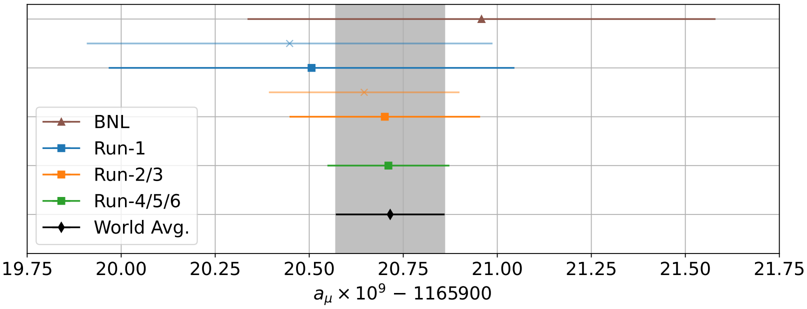

Figure 2.2-8. The final values of a_μ in the various publications of the E821 and E989 $g-2$ experiments. The blue and orange points show the shift in a_μ before (thin line) and after (thick line) the application of the updated RGE correction.

PIONEER

2.3 The PIONEER experiment

O. Beesley, Q. Buat, D. W. Hertzog, P. Kammel, J. LaBounty, B. MacCoy,
P. Schwendimann, E. Swanson, B. Taylor, and J. Yang

PIONEER is a next generation rare pion decay experiment at the Paul Scherrer Institute. The University of Washington is the leading institute of the project. We are involved in the design of the two main detectors, the calorimeter and the active target, and we are also leading the simulation effort to prototype the experiment and analyze the test beam data we collected at PSI and CENPA.

In Phase I PIONEER aims to measure the charged pion branching ratio

$$R_{e/\mu} = \frac{\Gamma(\pi \to e\nu(\gamma))}{\Gamma(\pi \to \mu\nu(\gamma))} \tag{1}$$

to a precision of 1×10^{-4}, thus providing an experimental measurement that matches the precision of the Standard Model prediction $R_{e/\mu}$ (SM) $= (1.2352 \pm 0.0001) \times 10^{-4}$ by Cirigliano and Rosell[1]. This will mark an improvement of the current PDG experimental average $R_{e/\mu}$ (Exp) $= (1.2327 \pm 0.0023) \times 10^{-4}$ precision by a factor of 20 and is an excellent test of lepton flavour universality.

In Phase II and III the goal is to measure the pion beta decay branching ratio

$$R_{\pi\beta} = \frac{\Gamma(\pi^+ \to \pi^0 e^+ \nu)}{\Gamma(\text{all})} \tag{2}$$

with an ultimate improvement in experimental precision by a factor 10. This measurement offers a theoretically pristine way to extract $|V_{ud}|$ and thus offers a handle on CKM unitarity. The current uncertainty of $|V_{ud}^{\pi}|$ is dominated by the branching ratio measurement.

An improvement of the branching ratio measurement by a factor 3 would put it competitive with the current PDG value obtained from neutron decays. With a factor of 6, the precision would match the best measurements from neutron decays that ignore tensions between different neutron measurements. The ultimate goal of a factor 10 puts the measurement on par with the value obtained from super-allowed beta decays. This is illustrated on the right panel of Fig. 2.3-1.

Both measurements are well motivated and provide sensitive probes independent of the specific details of the new physics beyond the standard model. Through quantum-loop corrections, a sensitivity to very high mass scales is reached.

In 2025, we contributed to the European Particle Physics Strategy Update (EPPSU) and submitted a short paper [0]. In this document, we discussed a comparison of PIONEER's test

---

[1] Phys. Rev.Lett. 99, 231801.

of Lepton Flavour Universality with other probes. The comparison is performed using the Standard Model Effective Field Theory (SMEFT) as recommended by the scientific committee of the EPPSU. In this frameowrk, modifications to the W boson coupling to leptons of family α are encoded by $[C_{Hl}^{(3)}]^{\alpha\alpha}$. Focusing on this class of new physics interactions, one can compare the bounds on ratios of effective couplings $g_\alpha \equiv g(1+[C_{Hl}^{(3)}]^{\alpha\alpha})$ from a variety of experiments, as shown in the left panel of Fig. 2.3-1.

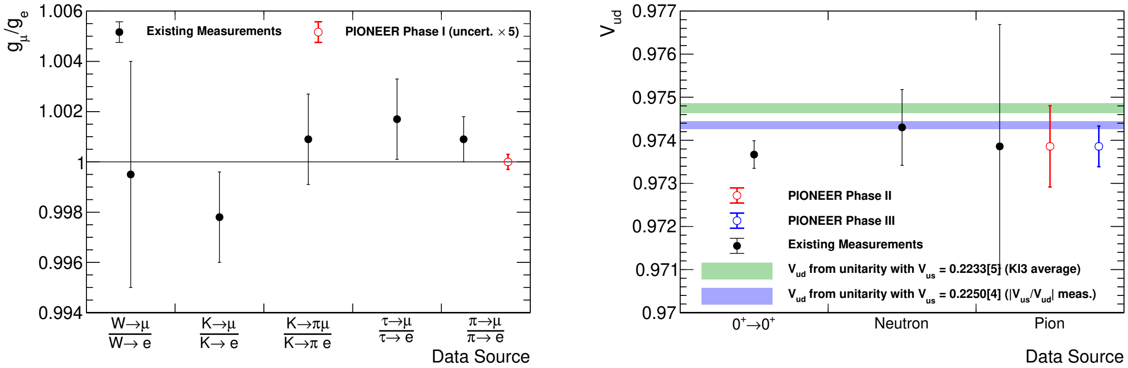

(a) Comparison of LFU tests (b) Comparison of V_{ud} probes

Figure 2.3-1. (a) Comparison of probes of LFU tests between the first and second generation. Existing measurements from Kaons, Taus and Pions are taken from Ref. [**annurev:vincenzo˙doug˙et˙al˙2022**] and from Ref. [**Aad2024ER-**] for $W \to \mu/W \to e$. (b) Comparison of the current measurements of V_{ud} with the estimate from CKM first row unitarity and the PIONEER Phase II and III projected precisions. PIONEER projections are preliminary and still under evaluation.

To drive the development of this next-generation experiment, our group leads the overall simulation and analysis efforts (Beesley, Buat, LaBounty, Schwendimann, Yang). In 2025, we have achieved the following milestones:

1. Establish the analysis strategy to define the acceptance of the PIONEER detector and work out how to correct the biases that are inevitably introduced (Buat, Beesley, Yang)
2. Develop a rule-based track reconstruction (Buat, Schwendimann, Yang) – (Sec. 2.5)
3. Establish a complete realistic event reconstruction chain (Schwendimann, Yang) – (Sec. 2.7)
4. Pioneer a new AI-driven approach to the PIONEER event reconstruction (Beesley) – (Sec. 2.6)
5. Prepare the simulation framework for large scale productions (LaBounty, Schwendimann) – (Sec. 2.8)
6. Simulate the ATAR demonstrator with surface and cloud muons (Buat, LaBounty) – (Sec. 2.11)

7. Host the first PIONEER reconstruction workshop[1] to bring together a group of ≈ 20 members of the collaboration to release the first complete realistic event reconstruction

Figure 2.3-2. First PIONEER Reconstruction workshop hosted at CENPA in July 2025. Workshop organized by Buat and Schwendimann

On the pion stopping target side, our group coordinates the overall electronic chain and is responsible for the design of the main electronic board that will host the amplifier chip (FAST-3). We also use our Van de Graaff accelerator to perform measurements over a comprehensive set of candidate sensors with varying doping profiles. In 2025, we have achieved the following milestones:

1. Publish our second test beam results (Buat, Kammel, Yang) [0]
2. Run a third test beam campaign and established a first model of the gain saturation effect (Buat, Kammel, McCoy, Taylor, Yang) – See (Sec. 2.10)
3. Successfully assembled a full electronic chain with wire bonding from a sensor to the flex and from the flex to the FAST3 amplifier.
4. Establish the feasibility of the PIONEER experiment with the SAMPIC digitizer system and perform critical tests of the ping-pong mode – See (Sec. 2.9)

Finally, our group leads the LYSO calorimeter option in the PIONEER collaboration. We perform simulation studies and coordinate the assembly of prototypes necessary to establish LYSO as a viable calorimeter. In 2025, we have achieved the following milestones:

1. We acquired tapered crystals in collaboration with the ETH-Zurich group and we evaluated the energy and time resolution of a small array of full-sized, tapered LYSO crystals (Hertzog, Beesley, LaBounty, Swanson) – See (Sec. 2.4)

---

[1] https://indico.psi.ch/event/17617/overview.

2. We published a comprehensive paper on the testing of basic LYSO rectilinear crystals, and their performance during dedicated test beam time at PSI. The array achieved an energy resolution of 1.56% at 70 MeV, 3 times better than any previous LYSO-based effort by other groups. [0]
3. The success of our LYSO developments and the progress made in reconstruction establishes that LYSO meets all the KPPs for the calorimeter and, as such, as been voted to be the baseline choice for PIONEER.

Moreover, we contribute significantly to the development of the beamline design (Kammel). In all efforts, CENPA lead engineer Roehnelt plays a crucial role in designing and fabricating small test stand equipments for prototype experiments as well as in advising on the overall detector design for the entire experiment.

2.4 Tapered LYSO Crystal Tests and Performance

O. Beesley, D. W. Hertzog, J. LaBounty, and E. Swanson

The PIONEER experiment requires a high resolution electromagnetic calorimeter capable of operating in a high rate environment. A calorimeter consisting of 311 large LYSO crystals arranged in a truncated polyhedron geometry is considered due to the attractive intrinsic and previously measured properties of LYSO crystals: high density ($\rho = 7.4$ g/cm^3), high stopping power ($X_0 = 1.14$ cm, $R_m = = 2.07$ cm), fast timing ($t_{decay} = 40$ ns), high energy resolution (1.52% at 70 MeV, measured in 2023), and high time resolution (100 ps for energies larger than 35 MeV) The stated time and energy resolutions were measured using an array of 10 LYSO crystals. In 2025, we designed and characterized an array of LYSO crystals in the same shapes and sizes that could be arranged in the final PIONEER calorimeter geometry. The array was composed of six tapered LYSO crystals (1 Pent, 5 Hex-A types) and was surrounded by NaI detectors for leakage-catching, as shown in Fig. 2.4-1. The tests were conducted at the PiM1 beam line at PSI during the last week of August 2025. The results of the test beam are as follows.

Positrons ranging in energy from 20 to 70 MeV were used to evaluate the energy and time resolution of the tapered LYSO array. Excellent energy resolutions were observed, as can be seen in Fig. 2.4-2A. Despite limited calorimeter coverage resulting in an expected large leakage contribution to the energy resolution of the array, improved energy resolutions were measured at energies below 30 MeV due to significantly improved control of PMT noise relative to tests done in 2023. An energy resolution of 2.17 % was measured at 50 MeV, and an energy resolution of 1.97% was measured at 70 MeV, exceeding the requirements for a PIONEER calorimeter. The array also saw improved time resolutions at low energies due to lower PMT noise. At 2 MeV and 0.5 MeV, improvements of two times and nearly five times over measurements in 2023 were observed, as seen in Fig. 2.4-2B. Time resolutions of approximately 100 ps for energies larger than 30 MeV were observed, matching results from 2023 and exceeding the specifications required for the PIONEER calorimeter.

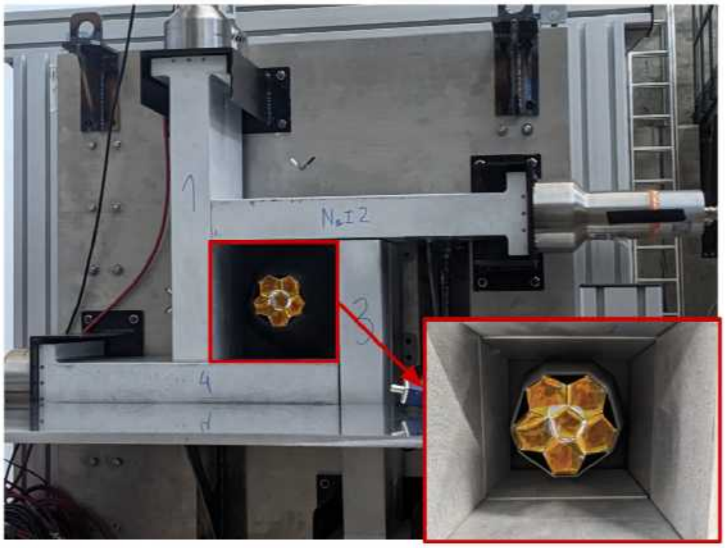

Figure 2.4-1. An array consisting of six LYSO crystals surrounded by eight NaI large detectors was used to characterize the LYSO crystal performance.

LYSO crystals were rear-mounted using cyanoacrylate bonded to a steel frame to validate a mechanical design that could be used for the final PIONEER calorimeter, as shown in Fig. 2.4-2C. The frames were then screwed into a 3D-printed holder, and a separate 3D-printed compression mechanism was used to ensure a strong coupling at the optical grease interface between the crystals and PMTs.

Energy loss due to shower back-scatter (albedo) from the front faces of the LYSO crystals will contribute to the $\pi \to e$ tail and was measured in a special configuration to validate Monte Carlo simulations of the PIONEER calorimeter. Runs testing albedo were conducted with targets of various materials (LYSO, copper, lead, air) and at different energies (30, 50, 70 MeV), and the back scatter was measured using the NaI detectors originally used for leakage detection. A very preliminary analysis shows the energy spectrum in the NaI for 70 MeV incident positron energy with and without the LYSO target, and the difference between the two spectra is shown in Fig. 2.4-2D.

LYSO crystals have recently been selected as the baseline calorimeter material for the PIONEER experiment. Our future work will focus on refining the analysis of the 2025 data, with a significant effort to analyze our albedo data. We will also continue to optimize our photosensors to ensure a stable pulse shape with minimal noise.

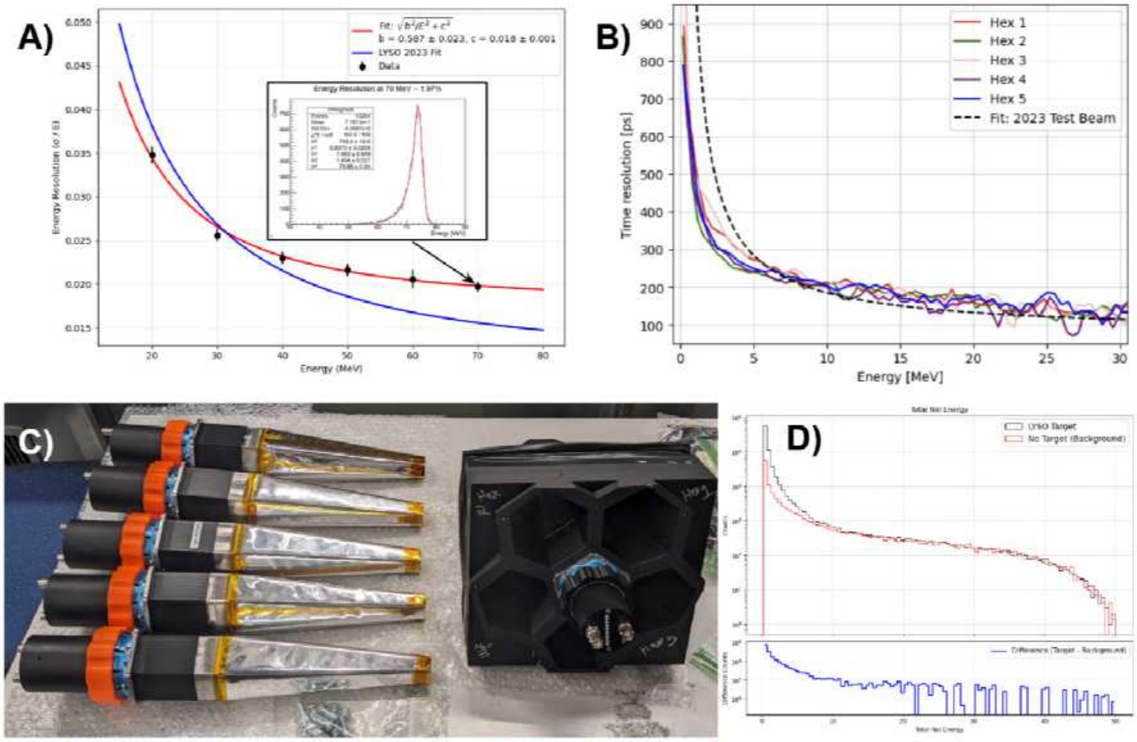

Figure 2.4-2. A) Energy resolution of the LYSO array at beam energies ranging from 20-70 MeV; the 70 MeV energy spectrum and its fit are shown. The fit of the 2025 energy resolutions is shown in red and overlaid with the fit from 2023 (in blue), demonstrating improvements in energy resolution at lower energies and worse resolutions at higher energies due to leakage from the array. B) Time resolutions for each of the Hex A crystals are shown for pulse energies in the Hex crystals ranging from 0-30 MeV. These spectra are compared with the fit from 2023, demonstrating significant improvements at low energies and comparable performance at higher energies. C) Rear-mounting structure used to hold LYSO crystals and ensure strong coupling between the crystal face and PMT surface. D) Very preliminary analysis of albedo for an incident 70 MeV positron. The energy recorded in the NaI wall is shown in black when there was a LYSO target, red when there was no target and the difference is shown in blue; the blue spectrum is therefore the albdeo from the target.

2.5 Rules-Based Reconstruction

J. Yang

PIONEER's active target (ATAR) is a highly segmented, fast-timing detector capable of nanosecond level timing resolution and sub-millimeter position resolution. To verify that the resolution is sufficient for measuring the pion branching ratio to the desired precision, a realistic reconstruction algorithm was developed and merged with the PIONEER simulation framework this year.

The ATAR reconstruction receives input in the form of ATAR hits with associated energy, time, and position information. To match the expected resolution of the ATAR, energy measurements are smeared by 10% and timing measurements are smeared by 200 ps. Position information is given only by the position of the LGAD strip in which the hit was observed. The reconstruction then proceeds as follows:

- **Time grouping:** Given a 200 ps timing resolution for each sensor and 800 ns event window, grouping together any hits within 1 ns of another successfully isolates 97% of pions from its daughter. Due to the much larger lifetime of muons relative to pions (2 µs vs. 26 ns), the probability of time-coincident muon and positron tracks is much smaller and on the order of $\mathcal{O}(10^{-4})$.
- **Spatial separation:** Once pileup is considered, a spatial separation algorithm is necessary for time-coincident tracks. This is done by first grouping hits into "strands" in each plane of the ATAR if they are within a millimeter apart. Then, strands are merged with any other strands in neighboring ATAR planes within a millimeter to form the final tracklets. The ATAR planes alternate between horizontal and vertical orientations, so spatial separation can only be determined in either the XZ or the YZ plane.
- **Endpoint finder:** For each tracklet, hits are identified as endpoints based on the average direction of its neighbors. For endpoints, the nearest neighbors are all in relatively the same direction. So, when averaging the normalized vectors from an endpoint to each of its neighbors, the resulting vector will have a magnitude close to 1. Meanwhile, neighboring vectors of a mid-point will cancel out and produce a vector with a smaller magnitude. Once endpoint candidates are identified, decay patterns can be formed between tracklets with overlapping endpoints.
- **Particle identification:** Particles are identified based on the energy deposited and track length. Pions enter with a momentum of 65 MeV and therefore deposit about 11 MeV total in the activate silicon bulk degrader (DTAR) and the ATAR. The DTAR will be designed such that the pion stops close to the center of the ATAR and therefore travels through around 25 ATAR planes. Muons born from stopped pions will always deposit around 4 MeV and travel 0.8 mm, resulting in no more than 15 hits in the ATAR. Finally, positrons will produce the most variable amounts of total energy deposit and track length. However, since they are minimum ionizing particles, they will tend to deposit 0.3 MeV or less per hit. These thresholds were selected to minimize the amount the positrons that are misidentified. While the pion identification is accurate to within a percent, muons are being misidentified as positrons around 2% of the time. Additionally, these parameters will have to be adjusted once gain suppression in the ATAR sensors is taken into account.
- **Outgoing positron direction:** The direction of the positron track is then chosen such that the track begins at the endpoint closest to the pion stop. By default, the track direction assumes positive Z motion. Thus, the reconstruction performs best for decays in the forward direction.

2.6 AI ATAR Reconstruction

O. Beesley

The PIONEER active target (ATAR) will drive order of magnitude sensitivity improvements through precise 5D (x,y,z,t,E) tracking. Conventional rule-based algorithms struggle with reconstruction due to low-energy positron scattering and annihilation, alternating ATAR readout geometry, and overlapping pion–muon tracks that can bias acceptance at the 10^{-4} level. Artificial intelligence (AI) is a promising alternative that enables the dynamic treatment of irregular and sparse data from the ATAR. We have started the development of a graph transfomer-based reconstruction that models ATAR hits as node labels and naturally encodes relations between hits as edge labels in a graph. The transformer's attention mechanism dynamically learns correlations across layers and times, enabling accurate vertexing and track separation in complex multi-particle topologies. This approach improves particle classification while providing robust bias control essential for achieving PIONEER's precision goal.

The current AI ATAR reconstruction architecture consists of multiple models individually trained to perform specific reconstruction tasks. The models consist of a Graph Neural Network (GNN) where hits are treated as nodes with features (2D position, energy, strip dimension) and differential relationships are edge labels. The model employs Transformer blocks with multi-headed attention to identify geometric and energy relationships (such as track kinks or directionality) and a Feed-Forward Network (FFN) to detect physics patterns like Bragg peaks or MIP-like ionization. Hits in the ATAR are grouped in time so that time groups of hits within 2 ns of each other are contained in a single object. Time groups are then processed by a chain of AI models in series as follows:

- **Particle Classifier:** This is the first model in the chain, which uses a pooling layer to aggregate the FFN output into a single vector, projected onto a probability distribution classifying the group as containing Pions, Muons, or MIPs.
- **Group Splitter:** Designed to handle complex topologies where multiple particles are contained in a single time group. This model takes multi-particle groups (e.g. a combined $\pi + \mu$ group) and splits them into distinct particle groups to ensure accurate downstream reconstruction. It also predicts the relative energies of each particle in the group.
- **Endpoint Finder:** Operates downstream of the initial classification to identify the termination points of the tracks within the ATAR. This information is critical for the subsequent matching of time groups across the detector readout.
- **Pion Stop Finder:** A specialized feature extraction model that determines the precise 3D stopping position of the pion within the active target. This is a primary event feature required to make acceptance decisions and is essential to define the decay vertex.
- **Positron Angle Finder:** Extracts the positron emission direction at the time of production and is also necessary for acceptance determination.

- **Event Builder:** This stage integrates the extracted features and classified groups to match groups that belong to a single decay chain. This step is necessary to make acceptance decisions, as it matches pion stops to positron angles.

Early benchmarks indicate performance gains over traditional reconstruction methods. The performance of the particle classifier is shown in Fig. 2.6-1A as a confusion matrix indicating the fraction of true/false positives and false negatives for each type of particles. The classifier has a better accuracy of 99.9% for pions and MIPs and approximately 99.5% for muons. Many of the failures, especially in the muon and MIP case, are due to edge cases in the definition that a time group is classified to contain a particle type when two or more hits in the group are attributed to this particle. An event display showing the capabilities of the endpoint finder is shown in Fig. 2.6-1B. In this event, we consider a time group in which a positron has undergone multiple Coulomb scattering and has produced multiple secondary electrons through Bhabha scattering. The starting point of the positron is reconstructed very precisely, whereas the end point of the group is assigned a much larger uncertainty to accommodate scattering effects. The spectra of endpoint errors for each particle type are shown in Fig. 2.6-1C.

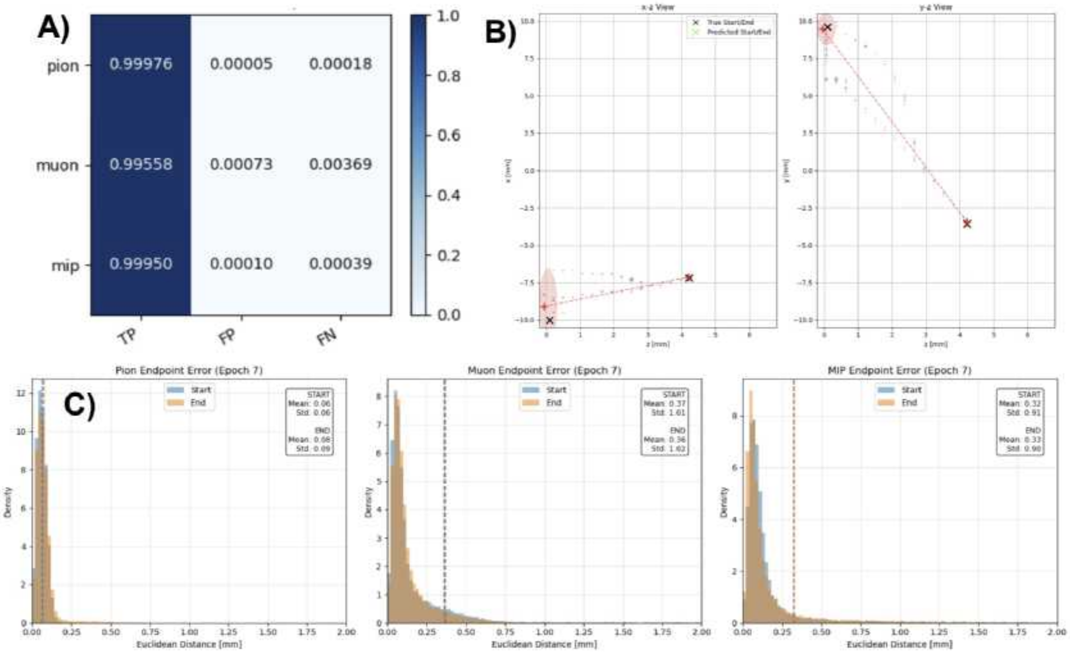

Figure 2.6-1. A) Confusion matrix demonstrating the performance (true positives, false positives, and false negatives) for different particle types of the Particle Classifier. B) Event display showing a positron/electron time group with its true endpoints in black and the endpoints and the associated uncertainties predicted by the Endpoint Finder in red. C) Spectra of distances between true and predicted endpoint positions for different particle species. Dashed lines are used to indicate the means of the distributions; muon and MIP groups have long tails in their endpoint error spectra due to groups missing hit information from one of the x/y dimensions and various scattering processes.

The AI reconstruction is pushing towards a more complete high performance pipeline that will enhance our capability to train using large datasets. Our next steps also include a campaign to integrate the AI reconstruction with the main PIONEER framework and an extension of the current AI ATAR reconstruction to a global event reconstruction that includes all detector subsystems.

2.7 The realistic reconstruction chain

P. Schwendimann

In the aftermath of the PIONEER reconstruction workshop, all components of the realistic reconstruction were assembled. On the ATAR side of the detector, this includes the track reconstruction strategies detailed in (Sec. 2.5) and pattern finding algorithms developed by collaborators at University of Kentucky and Amherst College. The LYSO calorimeter reconstruction was previously developed at UW.

After preliminary cross-checks, the most significant issue were cases of prompt muon decays where muon and positron could not be distinguished. This resulted in the kinetic energy of the muon getting added to the energy lost in ATAR by the positron, thus consequently overestimating the event energy by 4 MeV. This resulted in a significant contribution of background in the high energy bin that followed the pion lifetime, which thereby distorted the fit result - yielding 26 % more $\pi \to e\nu\,(\gamma)$ than it should have.

Above issue was easily mitigated by not considering energy losses in ATAR when computing the final energy. This translated theta-dependent material budgets in ATAR to a theta-dependence of the energy scale. This is accounted for by running an analysis that is binned in $\cos\theta$. With these strategies in place, an intermediate simulation campaign was started end of November, resulting in the largest data-set ever produced for PIONEER.

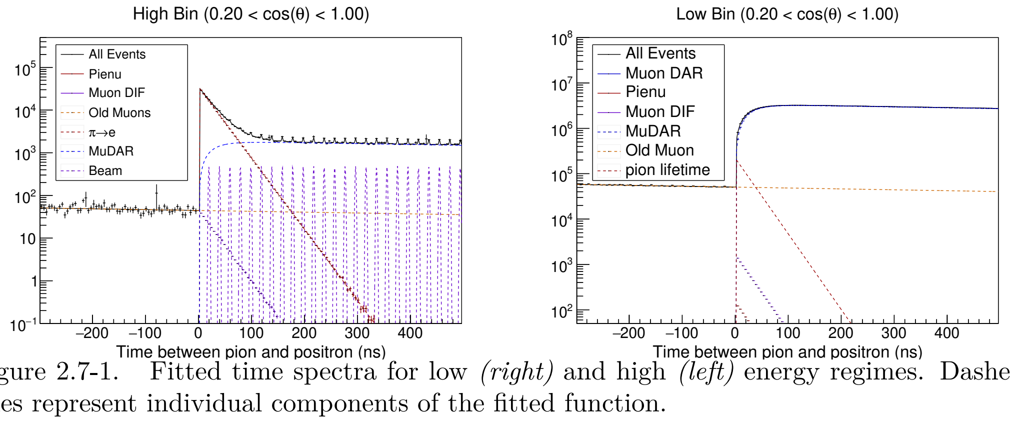

Figure 2.7-1. Fitted time spectra for low *(right)* and high *(left)* energy regimes. Dashed lines represent individual components of the fitted function.

The resulting time spectra are shown in Fig. 2.7-1. The more realistic reconstruction introduces failure modes that were expected. Most notable is the case where a positron from an old muon decay gets matched to the triggering pion. Most of these failure modes exhibit a very distinct time spectrum that can be described and fitted with a dedicated functional form as seen in Fig. 2.7-1.

Both high and low energy time spectra can be described using four distinct components:

$$
\begin{aligned}
n_1(t) &\propto \exp(-t/\tau_\pi) && t \in [0, 500]\,\mathrm{ns} && (1)\\
n_2(t) &\propto \exp(-t/\tau_\mu) - \exp(-t/\tau_\pi) && t \in [0, 500]\,\mathrm{ns} && (2)\\
n_3(t) &\propto \exp(-t/\tau_\mu) && t \in [-300, 500]\,\mathrm{ns} && (3)\\
n_4(t) &\propto \sum_{i>1} \exp\left[-0.5\left((t - i \cdot T_{\mathrm{RF}})/\sigma_{\mathrm{RF}}\right)^2\right] && t \in [0, 500]\,\mathrm{ns} && (4)
\end{aligned}
$$

The term $n_1(t)$ describes the $\pi \to e\nu(\gamma)$ signature in the high energy bin. In the low energy bin, it serves as a background fit absorbing cases where a muon under background influences is identified as positron. $n_2(t)$ describes the typical $\pi \to \mu \to e$ decay chain in the low energy region and the primary pileup contribution in the high energy regime. While the first two terms are for positive times only, the old muon term $n_3(t)$ extends to the negative time range as these events can happen prior to pion arrival. The beam modeling term $n_4(t)$ is only used in the high energy range and at positive times as it is only a preliminary approximation to the most obvious contribution.

Not all failure modes have yet been studied and accounted for. In the case of the high bin region, the current fit model yields 1.0(3) % more $\pi \to e$ events in the forward region than what a fully truth-guided reconstruction would suggest. A notable effect already identified are difficulties in the reconstruction of θ in the region around $\cos(\theta) \approx 0$, which moves events from that region more likely into the forward region than the opposite direction. This failure mode affects low and high bin region, accounting for a 0.5 % correction. More detailed investigations of other failure modes has been put on hold to fully benefit from new features to be introduced in the major simulation framework overhaul presented below in (Sec. 2.8).

2.8 Computing Upgrades

J. LaBounty and P. Schwendimann

The extended down-time of the CENPA computing cluster that triggered the migration from ROCKS to WULF was taken as good excuse to overhaul some of the computing framework of PIONEER.

Upgraded Docker Container

The Docker container underwent a major structural clean up from a linear, monolithic build that grew historically to a more modular approach. Moving to a more modern base operating system (Ubuntu 20.04 to Ubuntu 25.10) guaranteed up to date libraries and turned some of the previous build steps obsolete.

Using the modular build approach, only installed artifacts are used and intermediate build steps are not copied to the final stage. This reduced the size of the final docker image by a factor of two down to 5.6 GB. Moreover, the docker image is now distributed through the github container registry of the pioneer experiment organisation rather than the private account used previously. Additionally, dedicated builds for the ARM64 and the x86_64 architecture are provided. This increases performance for users that use e.g. a MacBook with an Apple silicon chip inside.

Notable upgrades and newly installed software include bumping ROOT from version 6.24.02 to 6.36.00 which is a key ingredient for the currently ongoing version 3 migration of the simulation framework. Geant4 only gets bumped from 11.0.1 to 11.0.4. as earlier tests performed at ETHZ indicated that an updated modelling in version 11.1.X provides questionable results for the muon stopping range. Gaudi underwent a major version jump from v36r9 to v39r4 which fixes an issue when building and linking against Gaudi libraries related to the `confdb2` files.

New additions include the GenFit package that is currently investigated by UCSC to improve the tracklet fitting performance. Also, the RESS package has been moved from a internalised build inside the PIONEER software framework to a docker-provided external dependency. As the framework currently does no longer use this package, this stage is currently skipped but can be activated easily should the need arise when revisiting the ATAR waveform reconstruction.

Moreover, the python environment was revisited, using an updated system python version 3.13.7 with some system python package installations such as numpy, scipy, pandas and matplotlib. These get extended via venv by analysis-relevant packages such as jupyter notebooks, iminuit, uproot, awkward and scikit-learn. These additional packages make analysis within the container a viable option.

Finally, the environment setup has been revisited such that the simulation-relevant environment variables get directly sourced from a dedicated `setenv.sh` script generated from within the software framework rather than hardcoded in the docker image itself. Additional scripts and functions are provided for convenience that help binding the appropriate folders to the docker container and take care of platform-dependent display forwarding as well as network forwarding to access jupyter notebooks and root-based display servers.

The new docker image - despite being available for less than a year - already features well above 100 downloads as it simplifies the onboarding of new students massively.

Job Submission Updates

As the cluster changed from using SGE to slurm, the submission scripts were largely rewritten, migrating from a python based submission script that passes a large amount of positional arguments to an extensive bash-script to be run for the job on the cluster to a more modular bash-only approach with exactly four required key-word arguments passed from submission to run script.

This reduction in arguments is mainly due to externalising all simulation-related steps to a dedicated script to be provided by the user that includes the entire simulation workflow. Thus, the new job script requires only the path to the apptainer image, the tarred input files, the output directory and the name of the user script (inside the tarred input files).

Further improvements involves a staged submission strategy that first checks for updates in the framework code and if updates have been detected, triggers a dedicated build job that provides all required build artifacts required to run all requested simulations. Additionally, after all jobs in an array have been concluded successfully, a dedicated collector job is submitted that merges all root files produced. This automation benefits mostly jobs running over night or over the weekend, as it results in data-files ready to be synced to the local machine rather than forcing this step in an interactive session that can result in an hour of idle wait time.

Diagnostics were also improved. The run script features time-stamping of all major steps along the execution such that large overheads can be identified or notoriously slow compute nodes tagged. This is also very vital input to estimate runtimes when trying to scale up the simulation. Finally, while running the user script inside the apptainer image, computing resources allocated to each process are monitored continuously. Such information is crucial to properly estimated resources to be requested for future large-scale simulations and helps detecting the presence of per-event memory leaks.

Major Simulation Framework Overhaul

At the reconstruction workshop, it became obvious that the old framework structure for detector reponse and reconstruction would soon be hard to maintain and not be future proof. Most notably, TClonesArrays were used for transient storage and TTree structures were hardcoded. This automatically limited the number algorithm instances providing a specific output to one and also allowed algorithms to update objects created by other algorithms. This was working fine in the early stages, where information was appended but can cause major problems if algorithms start overwriting information.

In the aftermath of the reconstruction workshop, a list of possible upgrades were proposed and the implementation is currently well under way. Most notable upgrades include

- migrating from root TTrees to RNTuples as TTrees are soon to be marked deprecated and will probably start phasing out before we start taking PIONEER data.
- using Gaudi::Functional algorithms. These use the Gaudi internal event storage and data objects are passed as const references to algorithms using them as inputs.
- provide python scripts for algorithm scheduling

These upgrades are currently ongoing and only parts of the framework are migrated at the current time.

2.9 ATAR digitization developments

B. Taylor

Two SAMPIC boards (Fig. 2.9-1) developed by CEA/IRFU/SEDI, Saclay, and CNRS/LAL/SERDI, Orsay, were tested at CENPA for use in the readout of the PIONEER ATAR. The ATAR digitizer must have excellent timing resolution. Additionally, since it is possible to have a pion stop in an LGAD strip, then have the muon stop in the same strip and decay, the digitizer needs to be able to record each of the pion, muon, and positron pulses. These could happen promptly or after a few microseconds (the lifetimes of the pion and muon are 26 ns and 2.2 µs, respectively). SAMPIC is known to have excellent timing resolution, but each channel has a digitization dead time of about 1 microsecond. The readout window (at our desired sampling frequency of 1.6 GS/s) is only 40 nanoseconds. Thus, many muon and positron events would be missed. The developers of SAMPIC have introduced a feature known as ping-pong mode to address this issue.

Figure 2.9-1. Two 16-Channel SAMPIC Board Digitizers

When ping-pong mode is activated, SAMPIC links two adjacent channels on the board. At first, only one is activated and waits for a trigger. After triggering and a set veto time, the second channel is activated. The second channel does not immediately trigger, but waits for the waveform to cross the threshold again. This allows for the recording of a second pulse, whether it occurs promptly, many nanoseconds later, or microseconds later. This feature has been extensively tested at CENPA.

To test SAMPIC and ping-pong mode in particular, a pulse was generated using NIM module discriminators and a small RC circuit to emulate the shape of a FAST3 pulse. (The

LGAD pulses will be amplified by FAST3, which is AC coupled and thus produces pulses with long ($> 20\,\mathrm{ns}$) overshoots. For more on FAST3, see below.) The signal was then split and one copy line was delayed and attenuated to make an a second pulse. The two pulses were combined in a fan, then the signal was split again to feed into the two ping-pong inputs. Importantly, both pulses are on the input for both channels. Ping-pong mode was then tested across many parameters, varying the delay between pulses, veto time, threshold voltage, sampling rate, pulse rate and more. The results for three pulse delays are shown in Fig. 2.9-2. Note that the two windows can be completely separated (as with the 50 ns delay), or they can overlap substantially (as with the 10 ns delay). The timing of the second window is entirely dependent on the when the second pulse occurs, not the end of the first window or veto time. If the veto time is extended so that the second pulse crosses the threshold before the veto ends, the second channel does not trigger. Similarly, if the trigger threshold is too high and the second pulse never reaches that voltage, the second channel does not trigger. Ping-pong mode functions at 100 % efficiency up to ∼ 80 kHz.

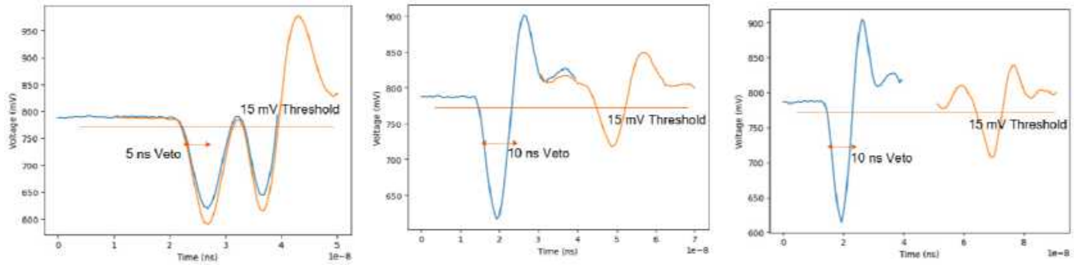

Figure 2.9-2. Ping-pong mode with 10 ns, 30 ns and 50 ns delays. The first (ping) window is shown in blue and the second (pong) window is shown in orange. The trigger threshold and veto time is also shown.

A correlation was observed between the veto time and post-trigger time (which sets the location of the trigger relative to the middle of the window) such that the longer one acted as the veto time. In order to see the entire pulse (which is necessary to detect small signals under the overshoot), a 15 ns post-trigger is necessary. However, a 15 ns veto time would cause problems for PIONEER, where the second pulse arrives too early for the second channel to activate but does not fit in the first window, as shown in Fig. 2.9-3. This issue could be fixed if the two settings were uncorrelated so that the veto time could be shorter than the post-trigger time. After discussion with the developers at CEA/IRFU/SEDI, Saclay and CNRS/LAL/SERDI, Orsay, they have implemented this change for the next version of the SAMPIC ASIC.

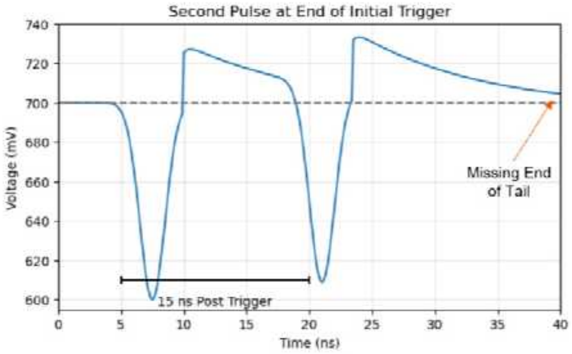

Figure 2.9-3. A simulated pion-muon-positron event recorded using ping-pong mode. Until ping-pong mode is extended to four channels, the positron could be missed.

By setting a low threshold and veto time, ping-pong mode allows SAMPIC to record two separate pulses in an LGAD, regardless of their separation. This eliminates the major issue with SAMPIC: a short (40 ns) data window and a long ($\sim 1\,\mu$s) dead time. Without this, a muon that appears more than 30 ns after a pion but less than a microsecond later would be missed. The next version of the SAMPIC ASIC will extend ping-pong mode to link four channels, allowing the recording of four successive pulses, regardless of dead time. This will also allow us to record a positron pulse, which in the current setup could be missed or truncated, as shown in Fig. 2.9-4, as well as an additional pulse (electron, background, etc.) In this new version, all four of these channels will have a single input, so an external splitter will no longer be necessary. Future tests of ping-pong mode will use real FAST3 pulses and analyze whether small positron pulses can be identified under the large overshoot.

Our SAMPIC boards were used extensively this year in various experiments. It was the primary DAQ for the LGAD sensor runs at CENPA, discussed in section 2.10, the associated LGAD betascope measurements, and for a four-plane ATAR demonstrator run at TRIUMF.

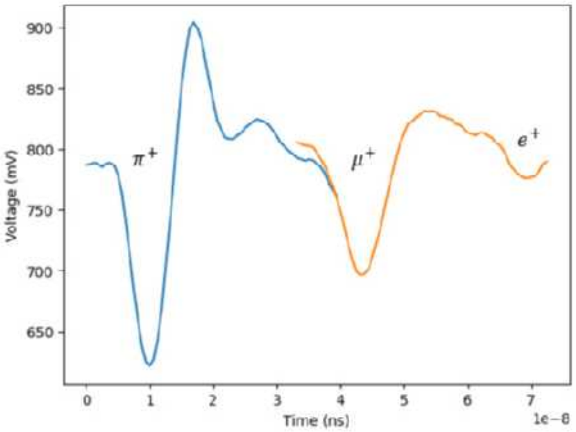

Figure 2.9-4. A simulated pion-muon-positron event recorded using ping-pong mode. Until ping-pong mode is extended to four channels, the positron could be missed.

Before digitization in SAMPIC, the LGAD signal must be amplified. The FAST3 chip has been tested at CENPA this year for use in LGAD readouts. FAST3 is a 16-channel ASIC with adjustable gain. For each channel, there are seven gain settings, plus an off setting. The gain for each channel is set using the 224-bit shift register. Each channel has a 3-bit section of this register that should assign a gain according to Fig. 2.9-1. Unfortunately, the situation is much more complicated. The setup in Fig. 2.9-5 was used to determine how to set the gain for each channel. The FAST3 ASIC is programmed using an Arduino. The resulting gain is inferred from the voltage of the pulse observed on the output. Two different transfer schemes were used, yielding different results.

Gain Setting	Bits
0	001
1	110
2	010
3	111
4	011
5	100
6	000
7	101

Table 2.9-1. Expected FAST3 Gain Settings

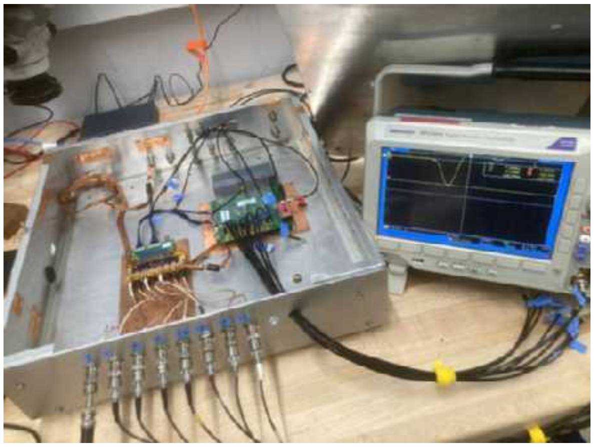

Figure 2.9-5. FAST3 Test Station

With the first transfer scheme, the results were completely stable and repeatable for all clock speeds tested (500 kHz, 250 kHz, and 125 kHz) and after power cycles. For channels 1 and 2, everything worked as expected with the gains set in the expected location according to Fig. 2.9-1. For all other tested channels, some logic is necessary to match the input bits to the expected bits necessary for a given output gain (those that are used for channels 1 and 2). Additionally, channels 4-15 have starting bits (the lowest numbered bit that affects the gain) that differ from the nominal values. This is summarized in Table 2.9-2.

Channel	True Starting Bit	Nominal Starting Bit	Logic[1]	Inaccessible Gain Options
1	209	209	2 1 0	None
2	199	199		
4	178	179	$(4 \wedge 3) \vee (3 \wedge 2)$ $(3 \wedge 2) \vee (2 \wedge 1)$ $(2 \wedge 1) \vee (1 \wedge 0)$	2 (010)
6	152	153		
8	132	133		
10	107	113	$(3 \wedge 2)$ $(2 \wedge 1)$ $(1 \wedge 0)$	7 (101)
12	86	87	$(5 \wedge 4 \wedge 3) \vee (4 \wedge 3 \wedge 2)$ $(4 \wedge 3 \wedge 2) \vee (3 \wedge 2 \wedge 1)$ $(3 \wedge 2 \wedge 1) \vee (2 \wedge 1 \wedge 0)$	2 (010) 7 (101)
14	61	67		
15	51	57		
16	41	41		

Table 2.9-2. Gain setting logic for first transfer scheme.

[1] Note that channels 4, 6, and 8 share the same behavior, as do channels 12, 14, 15, and 16. It turns out that effective bit combinations 010 and 101 are not accessible for certain channels. Since high bits without a neighbor (or two neighbors for channels 12-16) have no effect, many bit combinations can lead to the same gain. The simplest bit combinations for each gain are shown in Table 2.9-3

Gain	Channels 1 & 2	Channels 4, 6, & 8	Channel 10	Channels 12, 14, 15, & 16
0	001	00011	0011	000111
1	110	01100	1110	011100
2	010	NA	0110	NA
3	111	01110	1111	011110
4	011	00110	0111	001110
5	100	11000	1100	111000
6	000	00000	0000	000000
7	101	11011	NA	NA

Table 2.9-3. Simplest Bit Combinations to Set Gain

With the second transfer scheme we noticed that the output voltages were not stable when repeating the the same command to set the gain for a channel. Increasing the SPI clock from 250 kHz to 10 MHz greatly stabilized the output. However, gain 7 on channel 16 still had glitches at the percent level. With this transfer scheme at 10 MHz, channel 2 still works according to Table 2.9-1 and channel 4 works as it did in the first transfer scheme.

---

[1] The logic column needs some explanation. Bits are numbered with 0 as the true starting bit. Each row in each of the cells of the logic column gives a single "effective" bit, with the top line giving the most significant bit. That is, if the top line is true, the result is 1XX. The middle line gives X1X and the lowest line gives XX1. Together, these form 3-bit codes that match the nominal response and that of channels 1 and 2. For example, the lowest line of the cell for channels 4, 6, and 8 is $(2 \wedge 1) \vee (1 \wedge 0)$. This means the effective least significant bit is high if bits 2 and 1 are both high, or bits 1 and 0 are both high (the OR is inclusive). To be concrete, for channel 4, bit 0 = bit 178 and bit 1 = bit 179. If both 178 and 179 are set high, the effective least significant bit is high. Assuming these are the only high bits, the resulting effective bit code is 001, which gives gain 0. As another example, if 10110 is set for bit 182-178 (channel 4), this gives the effective bit code 011 (gain 4), since bits 1 and 2 of this code are high, activating the middle and lowest effective bits.

However, channel 12, 14, 15 and 16 with these SPI settings work with the same logic as channel 4. When the clock was set to 10 kHz, the logic observed for channel 16 with the first transfer scheme was resorted. The starting bit shift seems to be the same for both transfer schemes.

We conclude that the bit losses are not entirely related to the internal chip architecture, but also depend on the SPI properties. We tested the conjecture that SPI clock is double triggering, but did not find any change when putting R-C filters in the SPI clock and data lines. There does appear to be a pattern in the required bits, but the interpretation is unclear. Work will continue in order to define the required bitpatterns for the remaining channels and to test other FAST3 chips to see whether the pattern is chip independent.

2.10 LGAD sensor runs at CENPA

J. Yang

Three LGAD sensors fabricated by FBK were tested this year in CENPA's tandem accelerator with a proton beam ranging from 1.8 to 14 MeV. They were placed in the 24" Rutherford backscattering chamber at an angle of 110 degrees from the beamline with respect to the sub-micron thick gold foil target. Additionally, the sensors were rotated in place to scan over a range of incident angles and operated at various bias voltages. They were designed with similar breakdown voltages and therefore gain but with varying implantation depths of the gain layer. In particular, of the three sensors, two had deeper gain layers (W16 and W18), and the last one had a shallower gain layer (W9).

In the context of PIONEER, the goal of these tests was to inform LGAD design to minimize gain saturation and characterize its dependencies in the presence of highly ionizing particles. This is integral to the experiment because PIONEER's active target (ATAR), which will be composed of 48 layers of LGADs, must have sufficient energy resolution for pions, muons, and positrons in order for events to be efficiently reconstructed. Studying the behavior of LGADs in the presence of a proton beam is the first step to modeling the gain saturation expected for highly ionizing particles in the final experiment.

To properly compare the boards, the bias voltage must be selected so that each board experiences the same amount of gain and therefore collects the same amount of charge for a particular signal. So, data from W16, W9, and W18 with bias voltages of 180 V, 140 V, and 100 V, respectively, were selected for comparison such that the nominal gain of each sensor was around 14 for signals from MIPs. At the selected bias voltages, W16 (deep gain layer) exhibited greater saturation than W9 (shallow gain layer) for the same signals, while W18 (deep gain layer) exhibited less saturation. Thus, the data does not demonstrate a strong correlation between gain saturation and the depth of the gain layer (Fig. 2.10-1).

In addition to detector design, the data from this year's campaign was also used to develop a functional model for gain saturation. Physically, the saturation is a result of localized energy deposits that are amplified in the gain layer, reducing the electric field in the region and, therefore, the effective gain. Thus, one parameter of the gain saturation

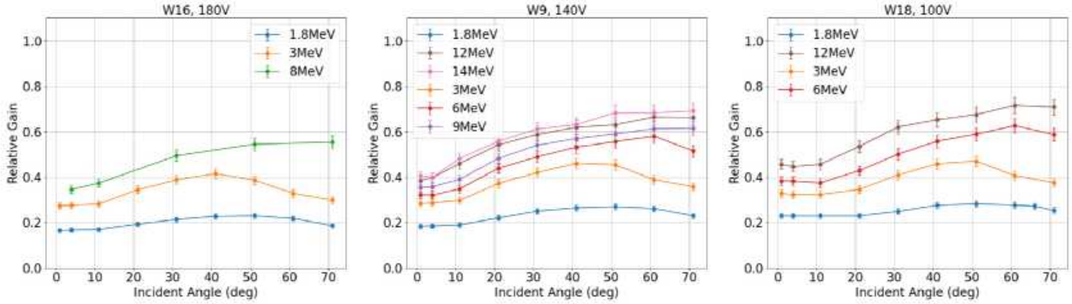

Figure 2.10-1. Relative gain is defined as the ratio between the measured gain and the nominal gain measured for MIPs. W16 and W18 each have deep gain layers while W9 has a shallow gain layer.

model is the energy deposited along the projected path of the track in the plane of the gain layer, dE/ds. Additionally, it has been observed in this and previous campaigns that the low energy protons that stop within the sensor actually experience greater saturation at the largest angles than at intermediate angles (see Fig. 2.10-1). This can be explained as a dependence on the distance z between the energy deposit and the gain layer. Since the gain layer is placed on the upstream side of the sensor, the particles that stop in the sensor with higher incident angles will deposit their Bragg peak energy closer to the gain layer. This gives the free charges less time to diffuse laterally than if the Bragg peak were deposited further from the gain layer, thereby increasing the charge density and reducing the effective gain.

With these two parameters in mind, the model for gain saturation is summarized as

$$Q_{obs} = \Sigma\kappa\left(\frac{dE}{ds}, z\right) * g * Q_{PIN,i}$$

where g represents the nominal gain measured for MIPs. The simplest model assumes that dE/ds and z are independent parameters, allowing

$$\kappa\left(\frac{dE}{ds}, z\right) = \kappa\left(\frac{dE}{ds}\right)\kappa(z)$$

where $\kappa(\frac{dE}{ds}, z)$ is the relative gain plotted in Fig. 2.10-1. A generalized formula for $\kappa(z)$ is

$$\kappa(z) = 1 - \alpha\left(1 - \frac{z}{t}\right)^{\beta}$$

where t is the thickness of the sensor. However, it was observed that this z dependence was greater for the lower energy protons with greater values of dE/ds. Thus, α is allowed to vary with respect to dE/ds

$$\alpha\left(\frac{dE}{ds}\right) = \alpha_0 * \min\left[1, \left(\frac{dE/ds}{0.03}\right)^{10}\right]$$

These formula are constructed such that the affect of $\kappa(z)$ is minimized for lower dE/ds and energy deposits further from the gain layer. The primary fit parameters are then α_0 and β.

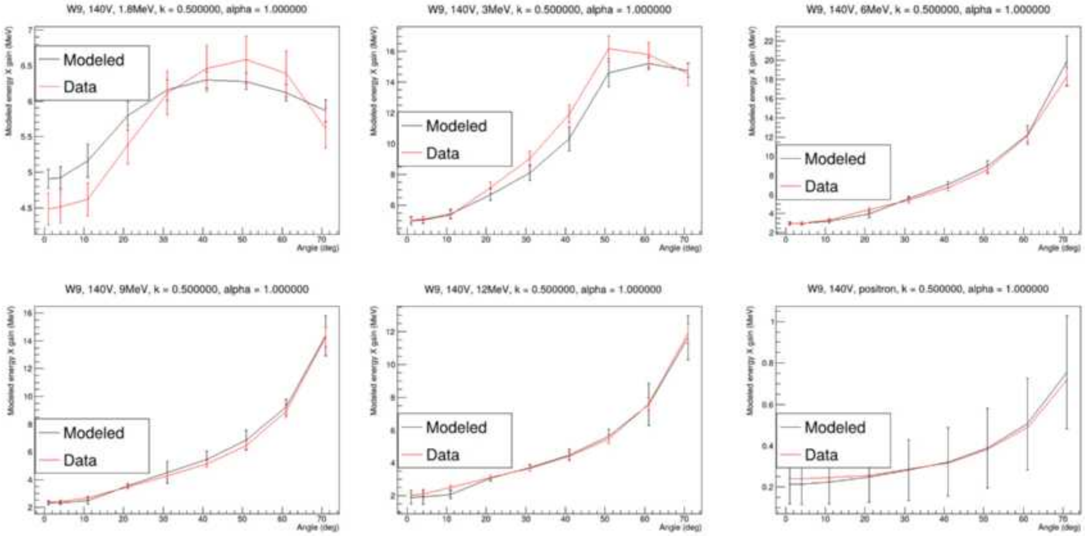

Figure 2.10-2. Comparing model to W9, 140V data with $\alpha_0 = 0.5$ and $\beta = 1$

For each choice of α_0 and α, a function for $\kappa(\frac{dE}{ds})$ can be fit to the data to complete the model. Then, the model can be applied to simulated to data and compared to measurements for validation (Fig. 2.10-2).

This model depends only on the local energy deposit and therefore is expected to apply to tracks of any particle type. Muon data from the PIONEER demonstrator in 2026 will test the validity of this model for the particles of interest to PIONEER. In the meantime, this preliminary model will be added to the simulation framework and will play a key role in the development of a realistic reconstruction algorithm and analysis strategy.

2.11 Simulations Towards an ATAR Demonstrator

J. LaBounty

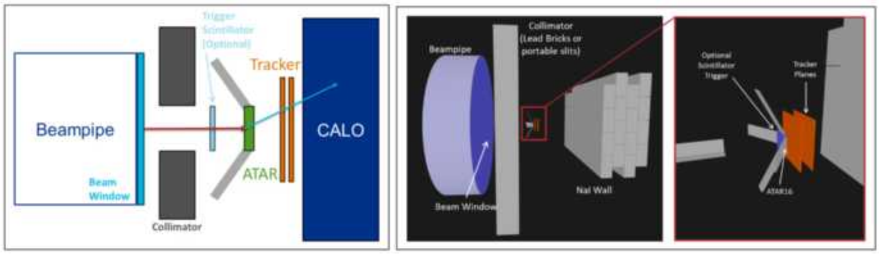

Figure 2.11-1. The layout of the PIONEER ATAR Demonstrator simulation. The left image shows the conceptual design, and the right shows to-scale views from the Geant4 simulation. The version of the Demonstrator shown consists of 16 layers.

The development of the ATAR is the linchpin of a successful PIONEER experiment. To that end, before the PSI shutdown in 2028 the PIONEER collaboration intends to conduct a series of ‘ATAR Demonstrator’ (ATAR-D) runs with increasing sophistication. The primary goals of these runs, in rough order of importance, are:

- to demonstrate the feasibility of constructing a multi-layer compact ATAR, and refining the procedures for doing so;
- to test the readout of ATAR data through the prototype readout/digitization systems;
- to demonstrate the ability of the ATAR to determine muon/pion stopping locations with adequate precision for PIONEER;
- to compare the reconstruction of the positron track direction from ATAR with that measured by an external tracking system as a function of angle and stopping location; and
- to provide realistic models of cross-talk, muon gain saturation, and other parameters for inputs to our simulation and analysis programs.

The first of these demonstrator runs is anticipated in the second half of 2026 and will be conducted using the PiM1 Muon Beamline at PSI[1]. To establish a baseline expectation for the physics goals above, a series of Geant4 simulations have been undertaken. A diagram of the setup, taken from a Geant4 simulation, is shown in Figure Fig. 2.11-1. A beam of surface or cloud muons ($p \approx 28 - 35\,\mathrm{MeV/c}$) will be brought to rest in the ATAR-D layers. Based on published PiM1 beam rates ($\mathcal{O}(10^7 - 10^8\,\mu^+/\mathrm{s})$ and spot sizes, a thick collimator upstream of the ATAR-D will be necessary to reduce background from stray muons in the calorimeter but should not impact the running of the demonstrator physics program. Two tracker planes and an NaI calorimeter will be situated downstream of the ATAR-D in order to measure the direction and energy of the outgoing positrons. The entire setup will be constructed in a ‘dual-turntable’ design, in which the ATAR-D and the tracker/calo can be rotated about the z-axis (in/out of the page in Figure Fig. 2.11-1). In this way, all combinations of incident/outgoing angles of muons/positrons will be able to be measured. The calorimeter will be built from a series of high-resolution NaI detectors used in previous PIONEER testbeams, while the tracker system will consist of two MuPIX layers developed by the Mu3e collaboration. These same tracker layers will also be used for a measurement of the PiE5 beam phase space in 2026.

A collection of simulation results can be seen in Figure Fig. 2.11-2. In these simulations, a beam of surface muons is brought to rest in the ATAR-D and allowed to decay. The momentum distribution of the muons is such that they implant evenly along the z-axis of the ATAR-D layers. From this uniform stopping distribution and the measured energy spectrum in the device, the gain suppression can be extracted in-situ. A closure test has been performed where an ad-hoc model is inserted into the simulation and then extracted using no truth information (Fig. 2.11-2, bottom). Positron decays from muons implanted in active strips will be used to develop the pion/muon stop and positron endpoint reconstruction algorithms. Information from positrons that then go on to hit the external detectors

---

[1] `https://www.psi.ch/en/sbl/pim1-beamline.`

(MuPIX and/or the calorimeter wall) will be used to validate the reconstruction. Scattering of the positrons in dead material outside the ATAR-D (such as cables) will also be studied with these events. Over the next few months, these simulations will be refined and used to construct the measurement plan for the ATAR-D testbeam.

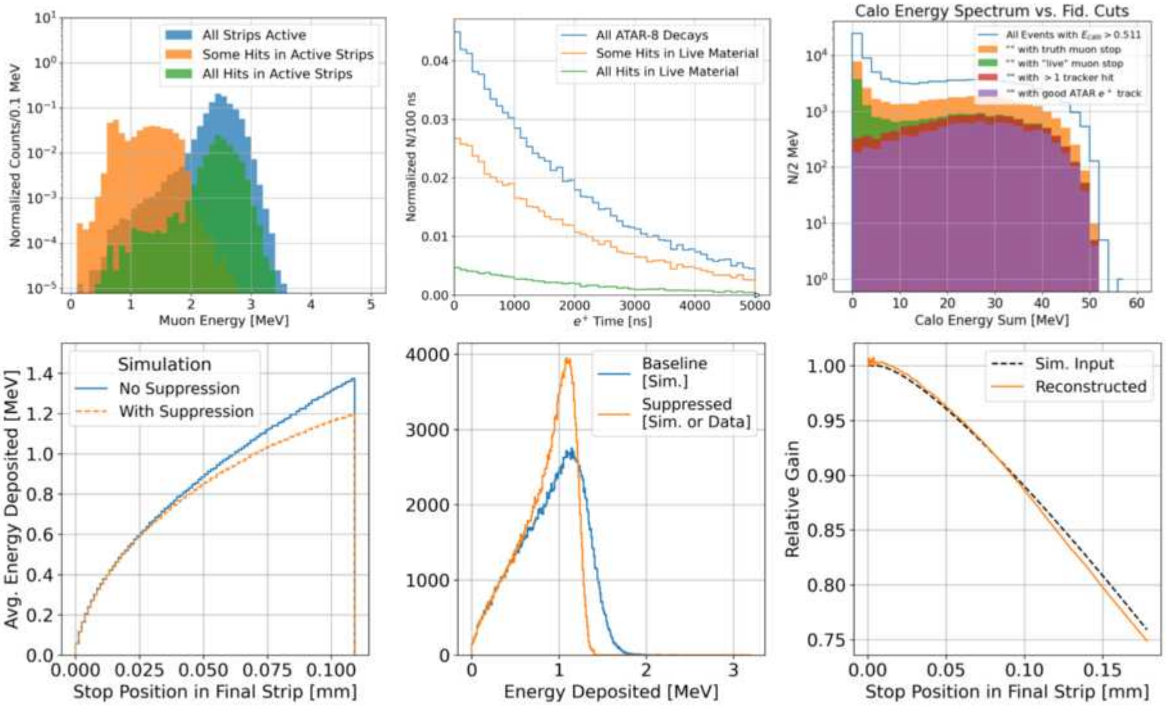

Figure 2.11-2. Preliminary simulations suggest that the PIONEER demonstrator will be able to stop muons evenly across all layers, measure the total muon energy deposition (top left), from this uniform stopping distribution extract the gain suppression of the system in-situ (bottom sequence), and measure the outgoing positron decays (top middle, top left).

3 Physics with accelerator beams

^{6}He-CRES

3.1 He6-CRES overview

L. Ahlburg, M. Blum, N. Buzinsky, B. Dodson, A. Fogel, Z. Fu, A. García, A. Goodman, S. Gopal, G. Hander, H. Harrington, L. Hayen*, P.T. Kolbeck, L. Malavasi, D. McClain†, D. Melconian†, P. Mueller‡, M. Newman, N. Oaxaca, N. S. Oblath§, G. Savard‡, D. Stancil*, D. W. Storm, H. E. Swanson, R.J. Taylor*, J. Tedeschi§, B. A. VanDevender§, C. Wiseman, F. Wietfeldt¶, B. Yommarath, and A. R. Young*

He6-CRES is progressing towards high precision measurements of the ^{6}He and ^{19}Ne beta spectra.[1] The motivation is a high-sensitivity search for chirality-flipping (scalar and tensor) interactions, which would signal physics beyond the Standard Model. Using an effective Lagrangian[2] for the charged weak current with additional scalar and tensor interactions with couplings ϵ_S and ϵ_T

$$\begin{aligned}\mathcal{L}_{\rm CC} = -\sqrt{2}G_{\rm F}^0 V_{\rm ud}\, \Big[\bar{e}^L\gamma_\mu\nu_e^L \cdot \bar{u}\gamma^\mu(1-\gamma^5)d \\ +\epsilon_S\, \bar{e}^R\nu_e^L \cdot \bar{u}d \\ +\epsilon_T\, \bar{e}^R\sigma_{\mu\nu}\nu_e^L \cdot \bar{u}\sigma^{\mu\nu}(1-\gamma^5)d\Big]\end{aligned} \tag{1}$$

the beta spectra, $\frac{dN}{dE}$, are expected to present a distortion with respect to the Standard Model prescription, $\frac{dN_{SM}}{dE}$, proportional to m/E (the ratio of the mass to total relativistic energy of the beta):

$$\frac{dN}{dE} = \frac{dN_{SM}}{dE}\,\left(1 + b\,\frac{m}{E}\right). \tag{2}$$

The so-called Fierz interference coefficient, b, is directly proportional to the couplings ϵ_S and ϵ_T.

Fig. 3.1-1(left) exemplifies the type of expected distortion, and Fig. 3.1-1(right) gives the sensitivity for $b < 10^{-3}$ measurements in comparison to other determinations.[3] Significantly, we are aiming for searches that are beyond the level of sensitivity of any previous experiment.

---

*Department of Physics, North Carolina State University, Raleigh, NC.

†Texas A&M, College Station, TX.

‡Argonne National Laboratory, Lemont, IL.

§Pacific Northwest National Laboratory, WA.

¶Department of Physics, Tulane University, New Orleans, LA.

[1] Byron et al., Phys. Rev. Lett. **131**, 082502 (2023).

[2] Bhattacharya et al., Phys. Rev. D **85**, 054512 (2012).

[3] M. Gonzlez-Alonso, O. Naviliat-Cuncic, and N. Severijns, Prog. Part. Nucl. Phys. **104**, 165 (2019).

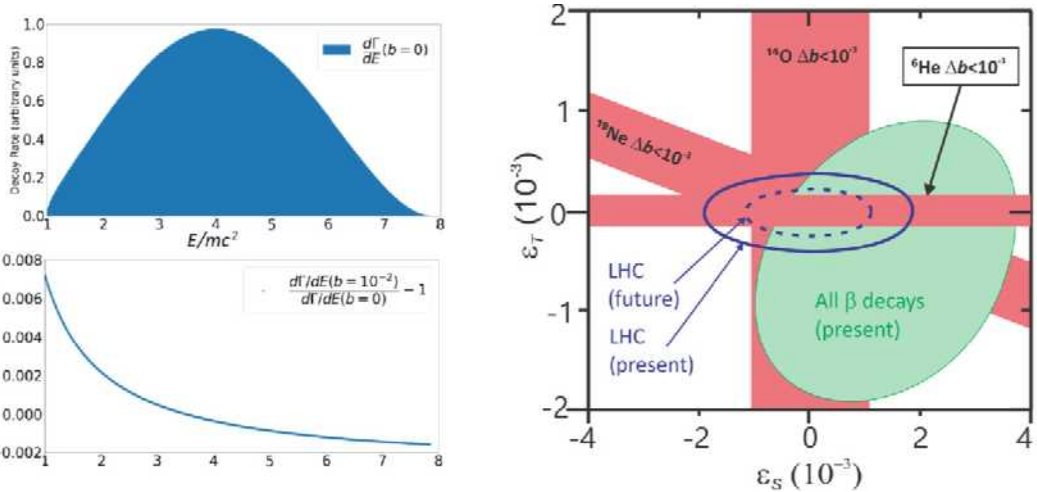

Figure 3.1-1. *Left:* example of beta-spectrum distortion expected from non-zero Fierz coefficient b. *Right:* sensitivities to scalar and tensor couplings from upper limits on b expected from the He6-CRES nuclei in red. In green, we show the result of a combination of all beta decay measurements to date, including neutron decays. The blue lines indicate limits from the LHC.

We use a technique called Cyclotron Radiation Emission Spectroscopy (CRES), developed by the Project 8 collaboration to determine the neutrino mass using the tritium endpoint. The basic idea is to arrange for the beta decays to occur within a magnetic field. The betas undergo cyclotron motion with frequency

$$f = \frac{qBc^2}{2\pi\ E}. \tag{3}$$

For $E \approx 1$ MeV in a ≈ 1 T field, the frequency is ≈ 20 GHz. Letting radioactive decays occur within an RF waveguide generates RF excitations at the cyclotron frequency. Typically, we digitize RF waves in slices of $\approx 3.4\ \mu$s, calculate FFTs, and produce spectrograms, called ‘waterfall’ plots, which allow extraction of E at the point of decay, before any significant loss of energy. Precise energy resolutions, combined with low backgrounds and a first-principles connection between the measured frequency and the energy, are very attractive for a method aimed at probing new physics.

Our present RF-bandwidth covers 18-20.3 GHz, and the coverage of the beta spectra is achieved by scanning the magnetic field, as shown in Fig. 3.1-2.

Fig. 3.1-3 shows the main RF components. Careful attention was paid to the distribution of heat in the design to allow optimal temperatures at the Low Noise Amplifiers (LNAs), while keeping the decay cell warm enough to avoid freezing ^{83}Kr, which was used as a calibration source.

He6-CRES is organized into phases with milestones that are presented in Table 3.1-1. We report here on progress since Jan 2025. We have finished Phase I and are starting Phase II. An overhaul of the track and event reconstruction software plus several sequential hardware improvements have led to a much improved system for identification of initial frequency for

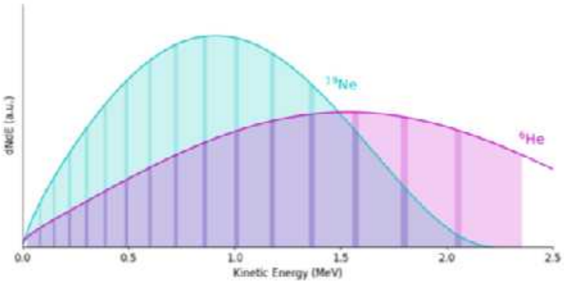

Figure 3.1-2. Sketch of ^{6}He and ^{19}Ne beta spectra with bands indicating the regions measured simultaneously by a $18 - 20.3$ GHz RF-bandwidth system. (Credit: N. Buzinsky.)

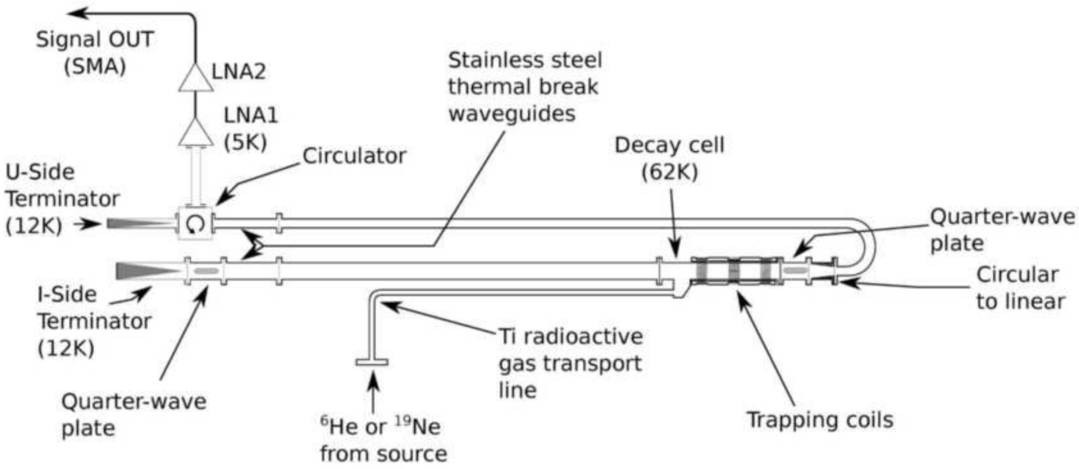

Figure 3.1-3. Sketch of RF components for He6-CRES. The radioactivity is compressed by a turbo pump into a decay cell (at ≈ 70 K), which is part of an RF guide. The guide was connected on one end to an RF Low Noise Amplifier (at ≈ 6 K) and at the other end to a terminator (at ≈ 12 K.) An ambient receiver box, not shown, mixes down the LNA signals to be digitized at several GHz. (Credit: H. Harrington.)

Phase I: *proof of principle*;
Observe ^{83}Kr lines;
Understand RF issues and spectra;
Study RF power distribution;
Develop ^{19}Ne source.
Show detection of cycl. radiation from ^{6}He;
Phase II: *first measurement ($b < 10^{-2}$)*;
^{6}He and ^{19}Ne measurements;
Develop design of ion-trap systems.
Phase III: *ultimate measurement ($b < 10^{-3}$)*;
Develop and build ^{14}O source;
^{14}O measurements;
Use ion-trap for no limitation from geometric effect.

Table 3.1-1. He6-CRES phases with proposed milestones.

events. Details are described in (Sec. 3.2). Due to the limited RF bandwidth (≈ 2 GHz) we scan the magnetic field to achieve a complete beta decay spectrum. The separate portions are normalized and stitched together using a 'beta monitor' system, which determines the radioactive atom density by measuring beta intensity simultaneous with the CRES events detection. The challenge for this beta monitor is to be stable at the level of a few hundred parts per million. Details of recent improvements are given in (Sec. 3.3). Another important aspect of the experiment is the ability to empty the CRES volume of betas to minimize pileup, which can lead to misidentifications. A sequence of tests of our trap emptying "$E \times B$ system" using ^{19}Ne data and contrasting the trap emptying efficiency versus beta energy led us to a better system. Details are given in (Sec. 3.4). In (Sec. 3.5) we summarize how we will quantify uncertainties using both Frequentist and Bayesian perspectives on fluctuations. An important consequence of trapping betas is that they undergo axial oscillations, with amplitudes of up to several RF wavelengths. As a result, the amplifier receives signals from different locations with different phases and the net effect is the appearance of *sidebands*, separated from the main frequency by the axial frequency. The power in these sidebands depends on the modulation index, which is a function of the beta's pitch angle at the time of creation. A summary of analysis of the consequences are related in (Sec. 3.6).

We are presently working on a search for b_{Fierz} with sensitivity < 0.01, which requires approximately three days of running time with each isotope. In July of 2025, we took ^{19}Ne data covering the full spectrum, which we summarize in (Sec. 3.7). Attempts to get ^{6}He data in August and December of 2025 were unsuccessful due to problems with the FN tandem.

Anticipating that the next level of precision will require addressing possible contaminants, we are developing methods to place upper limits on them: (Sec. 3.8) and (Sec. 3.9) summarize the work.

(Sec. 3.10) summarizes work on production of ^{79}Kr, which yields conversion electron lines of up to 130 keV, useful for calibrations.

(Sec. 3.11) summarizes considerations on using different transforms to the data that may be better suited for our events, which present a constant chirp with slopes limited in range for a given magnetic field.

As we move to higher precision, we want to take data faster and (Sec. 3.12) summarizes our work on doubling the frequency bandwidth.

Studying the effects of the $E \times B$ trap emptying system, we realized we could use a Penning trap, instead of the magnetic one, with a potential advantage of having a more harmonic trap, which would have a smaller distribution of axial frequencies. The work is summarized in (Sec. 3.13).

(Sec. 3.14) summarizes work done in preparation for shimming our magnet, which we believe will be important as we move toward a measurement with sensitivity $b_{\mathrm{Fierz}} < 0.001$.

3.2 Reconstruction upgrades

N. Buzinsky, A. Gorman, and H. Harrington

He6-CRES uses Katydid, a C++-based analysis framework developed by Project 8. The track and event reconstruction (T&ER) aims to produce a list of event start frequencies, corresponding to beta start energies, for CRES events in our noisy data.

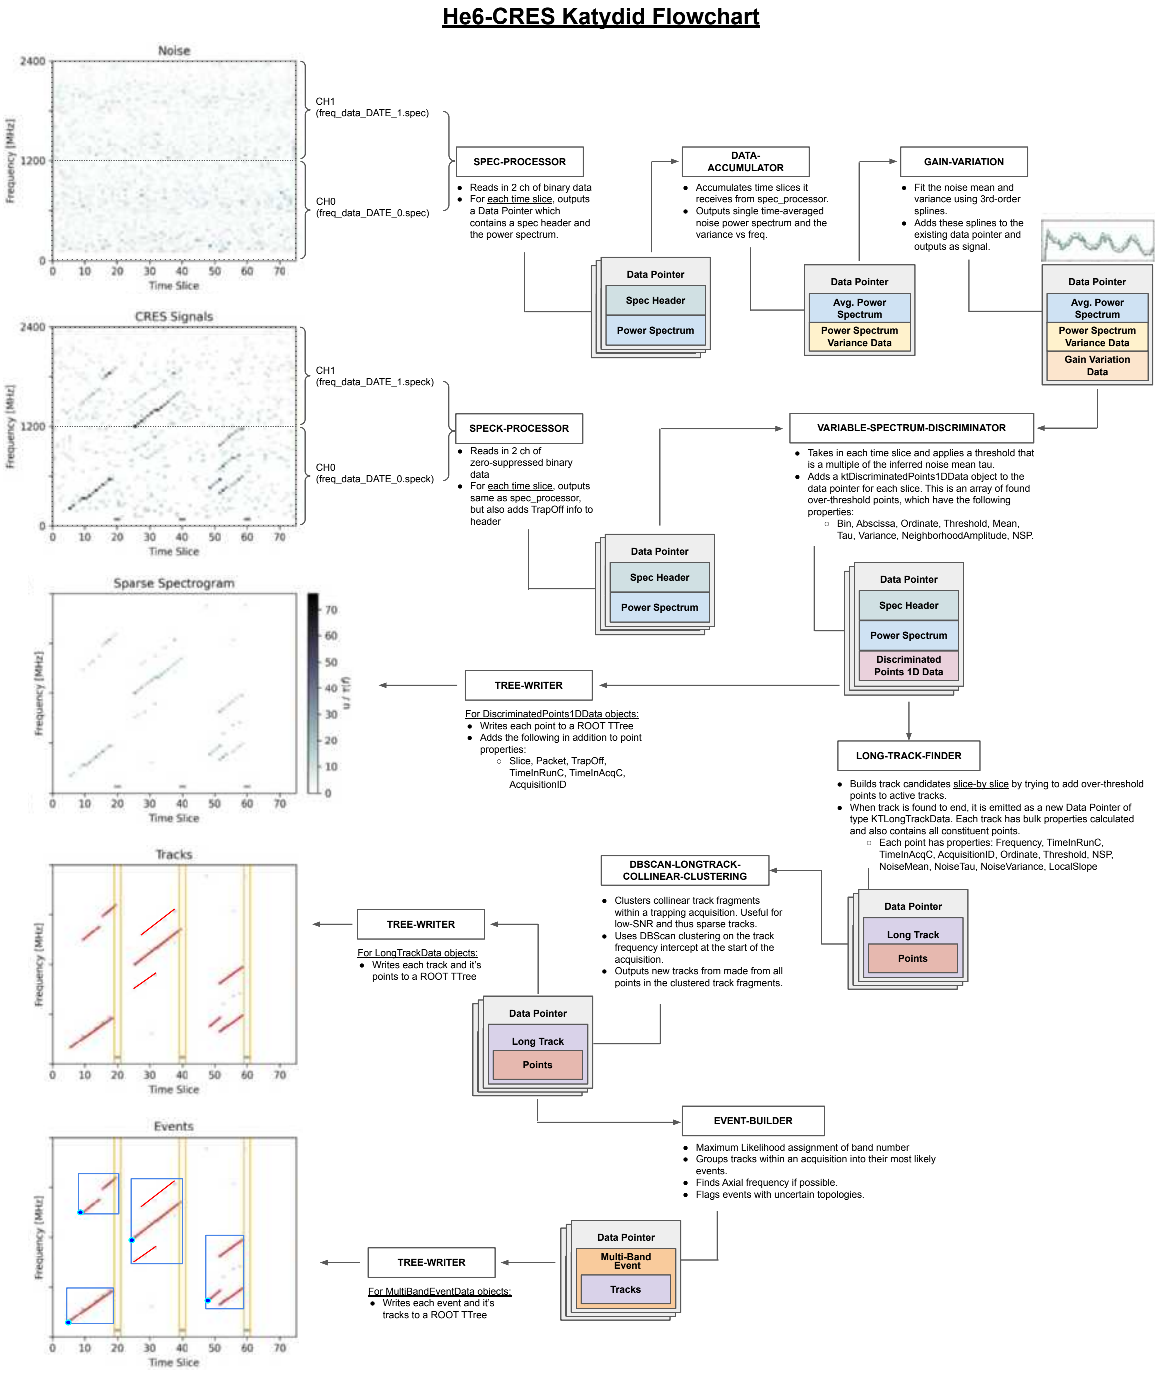

SPEC-PROCESSOR

- Reads in 2 ch of binary data
- For each time slice, outputs a Data Pointer which contains a spec header and the power spectrum.

DATA-ACCUMULATOR

- Accumulates time slices it receives from spec_processor.
- Outputs single time-averaged noise power spectrum and the variance vs freq.

GAIN-VARIATION

- Fit the noise mean and variance using 3rd-order splines.
- Adds these splines to the existing data pointer and outputs as signal.

SPECK-PROCESSOR

- Reads in 2 ch of zero-suppressed binary data
- For each time slice, outputs same as spec_processor, but also adds TrapOff info to header

VARIABLE-SPECTRUM-DISCRIMINATOR

- Takes in each time slice and applies a threshold that is a multiple of the inferred noise mean tau.
- Adds a ktDiscriminatedPoints1DData object to the data pointer for each slice. This is an array of found over-threshold points, which have the following properties:
  - Bin, Abscissa, Ordinate, Threshold, Mean, Tau, Variance, NeighborhoodAmplitude, NSP.

TREE-WRITER

For DiscriminatedPoints1DData objects:

- Writes each point to a ROOT TTree
- Adds the following in addition to point properties:
  - Slice, Packet, TrapOff, TimeInRunC, TimeInAcqC, AcquisitionID

LONG-TRACK-FINDER

- Builds track candidates slice-by slice by trying to add over-threshold points to active tracks.
- When track is found to end, it is emitted as a new Data Pointer of type KTLongTrackData. Each track has bulk properties calculated and also contains all constituent points.
  - Each point has properties: Frequency, TimeInRunC, TimeInAcqC, AcquisitionID, Ordinate, Threshold, NSP, NoiseMean, NoiseTau, NoiseVariance, LocalSlope

DBSCAN-LONGTRACK-COLLINEAR-CLUSTERING

- Clusters collinear track fragments within a trapping acquisition. Useful for low-SNR and thus sparse tracks.
- Uses DBScan clustering on the track frequency intercept at the start of the acquisition.
- Outputs new tracks from made from all points in the clustered track fragments.

TREE-WRITER

For LongTrackData objects:

- Writes each track and it's points to a ROOT TTree

EVENT-BUILDER

- Maximum Likelihood assignment of band number
- Groups tracks within an acquisition into their most likely events.
- Finds Axial frequency if possible.
- Flags events with uncertain topologies.

TREE-WRITER

For MultiBandEventData objects:

- Writes each event and it's tracks to a ROOT TTree

Figure 3.2-1. High-level flow chart of the track and event reconstruction done using the Katydid analysis framework, highlighting key processors, data objects, inputs, and outputs.

Multiple complicating factors contribute to ambiguities in determining the set of CRES events that caused our observed data, the totality of which the reconstruction algorithm must handle with minimal errors. For instance, high-power thermal noise fluctuations can

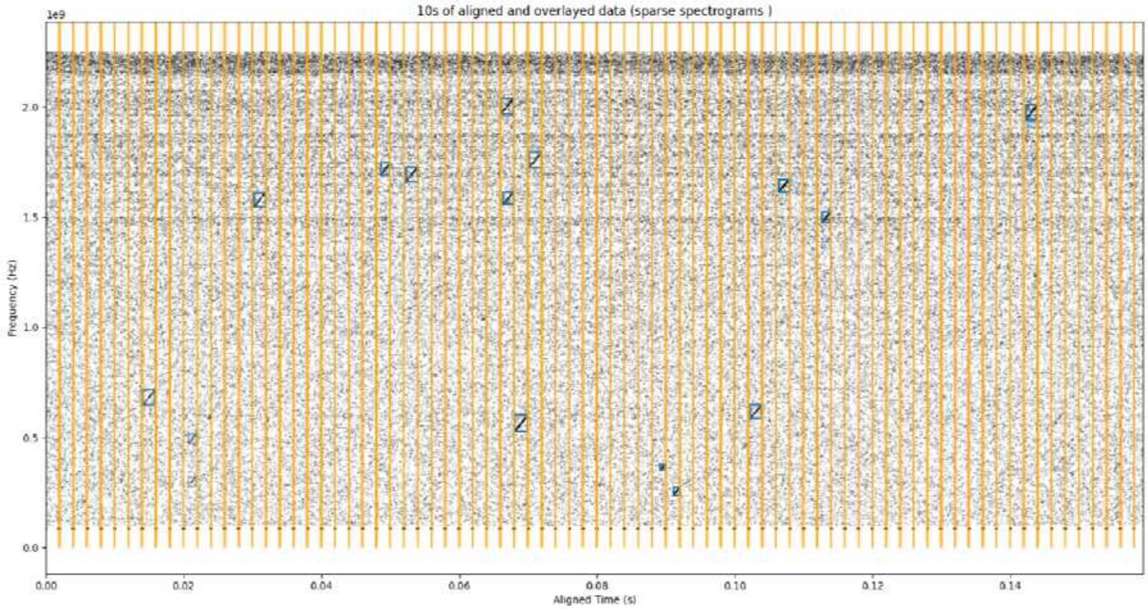

Figure 3.2-2. 10 overlayed sparse spectrograms with orange lines indicating the reconstructed trap emptying from the injected Vaunix tone, and blue boxes around the reconstructed tracks found by Katydid.

both obscure and distort true CRES signals or mimic false CRES events entirely. Successive scattering of a CRES beta off residual gas within the decay cell can result in a single beta producing multiple short-duration CRES tracks, which can be readily misinterpreted as being from multiple events. A single beta also experiences Doppler shifting as it travels axially within the magnetic trap, resulting in multiple simultaneous tracks; we must find the single correct clustering of the reconstructed bands into an event (Sec. 3.6), otherwise we get false events and/or incorrect start frequencies. The time-dependence of the cyclotron frequency is also non-linear, and in practice unpredictable, depending on resonances at mode crossings of higher harmonics. Finally, reflections in the apparatus result in frequency-dependent SNR, resulting in some tracks being detected later than their true start times, or not at all. Low SNR is seen as the ultimate source of our compounding challenge, described above, as we attempt to tally the start frequencies of all mainbands, prior to any scatters.

Hardware upgrades over the last several years have iteratively diminished the challenges faced by the T&ER. Upgrades to the RF system between 2022 – 2024 have reduced reflections within our system which had contributed to frequency-dependence in the SNR by approximately an order of magnitude[1]. The ExB trap clearing (Sec. 3.4), developed and characterized over a similar time-span, has opened up a new regime for the experiment, in which we can rapidly empty the trap, with little (∼ 0.2 ms) associated dead-time. Frequent trap clearing reduces reconstruction challenges associated with scattering (depending on how the timescales compare), frequency-dependent SNR, pileup from multiple events, and non-linearities in the tracks (Fig. 3.2-2).

---

[1] CENPA Annual Report 2024, Section 5.5.

Iterative upgrades to the hardware of the apparatus to improve the signal have also been paralleled by changes to the T&ER. A big part of overhauling the analysis pipeline was the new track finder developed by A. Gorman and H. Harrington. KTLongTrackFinder is designed to identify long-duration, variable-slope tracks in the sparse spectrogram data. It operates on sequential KTDiscriminatedPoint1DData objects emitted by the discriminator. These are single time slices containing over-threshold points. This algorithm is diagrammed in Fig. 3.2-3.

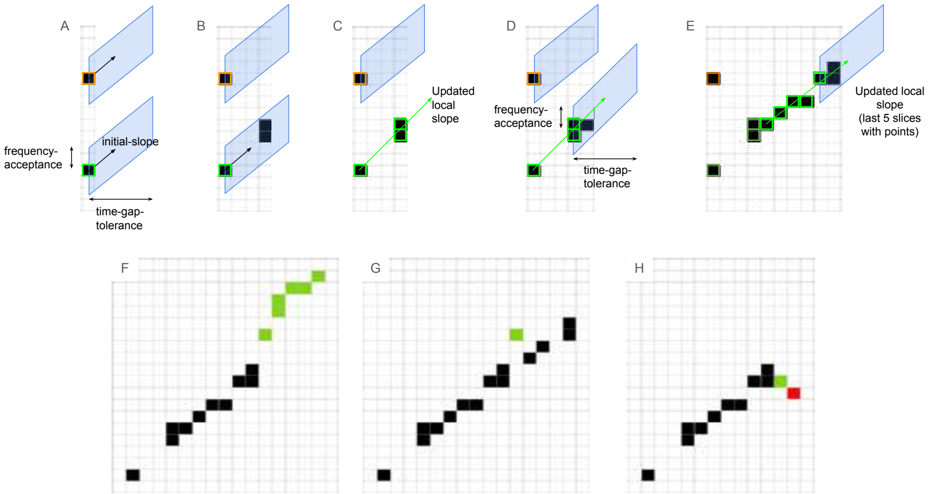

Figure 3.2-3. Cartoon of key features of the Long Track Finder algorithm. A) the initialization of two new track candidates from over-threshold points with assigned initial slope. B) Both tracks remain active for time-gap-tolerance looking for points that match the projected frequency range. In slice 3, two points are found that match the lower track. C) Both these points are added to the lower track. The local slope of this most recent slice is set using the most recent n-slope-slices worth of points, which in this case is all three points in the track. D) The median frequency of the added points and the new local slope are used to search for matching points. E) After time-gap-tolerance time has passed, no new points have been added to the upper track. As it is has fewer than min-points, it is not emitted. The lower track continues to add more points. F) Illustration of how a high point could be the start of a resonance. G) Illustration of how a high point could just be a close noise point. H) Exclusion of a point that would have made the local slope lower than a set minimum.

In the new analysis pipeline, we also improved our SNR estimation. One way to estimate the SNR is from the observed power in over-threshold points in a track. We can then use a maximum likelihood fit the the PDF. The second way is from the density of over-threshold points in the track, which is proportional to the CDF. Precise estimates of the SNR are important for improving the signal-noise discrimination, increasing the quantity of detected events for a given false alarm rate. RF improvements minimizing reflections has also spurred refinements of our SNR definition as we struggle to distinguish our data from a reflection-free model.

With the fully updated analysis chain, were were able to very successfully identify tracks in our data up to visual inspection. We are in the process now of producing fake spectrograms from Monte Carlo to run the analysis on to quantitatively assess the efficiency and lineshape response of the reconstruction.

3.3 Beta monitor

A. García, H. Harrington, B. Maxwell, D. McClain*, D. Melconian*, and C. Wiseman

When running the experiment, the radioisotope source intensities fluctuate in time. Differences in the total activity within the decay cell as we sequentially take data at different magnetic fields (according to Fig. 3.1-2) can obscure any potential BSM signature present in our data, if unaccounted for. To normalize out fluctuations in the total activity reaching the decay cell, we use the total rate observed by a silicon photomultiplier (SiPM)-based beta monitor immediately before the transport of the radioactive gas into the CRES apparatus. In 2025, we rigorously characterized this beta monitor, as any drifts or instabilities in the beta monitor efficiency lead to corresponding systematic uncertainties in the measurement of b. For a proposed 10^{-2}-level measurement campaign on b, we want systematic uncertainties associated with the beta monitor to be at or below $\sim \mathcal{O}(10^{-3})$.

The monitor consists of two thin EJ-212 plastic scintillators each with two Hamamatsu SiPMs optically coupled to them. SiPMs are advantageous over PMT's for this application as the beta monitor sits immediately beside the bore of the CRES main magnet and PMTs are very sensitive to magnetic fields. Having a pair of scintillators allows us to make a trigger in coincidence to discriminate against photon backgrounds.

In SiPMs, the gain, noise, and photon detection efficiency all depend strongly on temperature and voltage. For example, a temperature drift changes the breakdown voltage which likewise changes the overvoltage, if the operating voltage is not adjusted to compensate. Temperature dependence is currently the most dominant source of instability in the beta monitor and it varies between individual SiPMs.

To precisely measure the breakdown voltage for each SiPM at a reference temperature (T_{ref}), we extract the gain from the charge spectrum at various overvoltages and extrapolate to zero gain with a linear model. This analysis was done with the help of Sam Borden with the SiPMStudio2024 software.[1] From the extrapolated breakdown voltages, the Hamamatsu C11204-01 power supplies were able to be programmed for automatic temperature compensation of individual SiPMs.

We performed stability tests with a 0.1 mCi Sr-90 source. This resulted in detector coincidence rates of around 30 kHz, comparable to data-taking with the Ne-19 or He-6 gas sources. For these data, we compare the hardware coincidence with an analog logic of ((`A1` OR `A2`) AND (`B1` OR `B2`)) to a coincidence with a logic of ((`A1` AND `A2`) AND (`B1` AND `B2`)) from an offline (software) analysis. Note that the software coincidences are a subset of

---

[*] Texas A&M, College Station, TX.

[1] Sam Borden and Nick Ruof. Sipm studio, 2024.

the hardware coincidences. The results of the hardware and software coincidences are shown in Fig. 3.3-1.

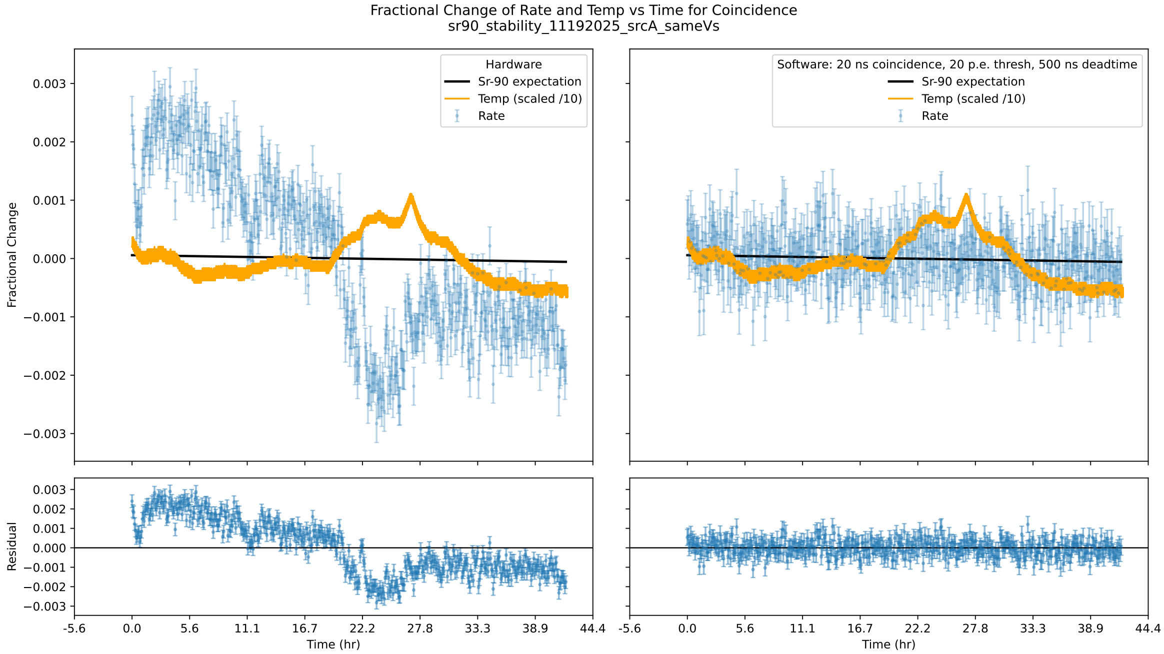

Figure 3.3-1. Fractional change in rate and temperature vs time for hardware (left) and software (right) coincident event triggers. 5 minute binned rates shown with Poisson uncertainties. The fractional change in temperature is scaled by 1/10. Black line is the expected drop in the fractional rate of the decay of the ^{90}Sr source ($\tau_{1/2} = 28.91$ yr), corresponding to 6.6×10^{-5}/day. Bottom plots show the 5 minute fractional rates with the decay expectation subtracted.

For the software coincidence, the charge spectra for each SiPM were fit to a model to get the gain for each hour, they were then normalized by the gain to get the charge spectra in units of number of photoelectrons. A threshold was applied and then coincidences found with an expanding window approach on time intervals between sequential events across all 4 channels to find and cluster those that occurred within a coincidence time window (20 ns). We then keep only events with at least one SiPM signal from scintillator A and scintillator B and assign a multiplicity number (MLT) between 2 and 4 to indicate how many SiPM pulses were in that event. Finally, the analysis enforces a non-paralyzable dead-time (500 ns) on coincident events.

The offline coincidence clearly outperforms the hardware coincidence in terms of stability. The main difference between the two comes in with the coincidence calculation of the re-thresholded counts, and the MLT=4 requirement as the gain was very stable. This improvement in the offline software coincidence stability is also indicated in the overlapping Allan Deviation (OADEV), which is defined via:

$$\sigma^2_{\mathrm{OADEV}}(\tau) = \frac{1}{2(N_y - (2m-1))} \sum_{n=m}^{N_y - m} (\bar{y}_{n+m}(m) - \bar{y}_n(m))^2. \tag{1}$$

where τ is m times the time step of y.

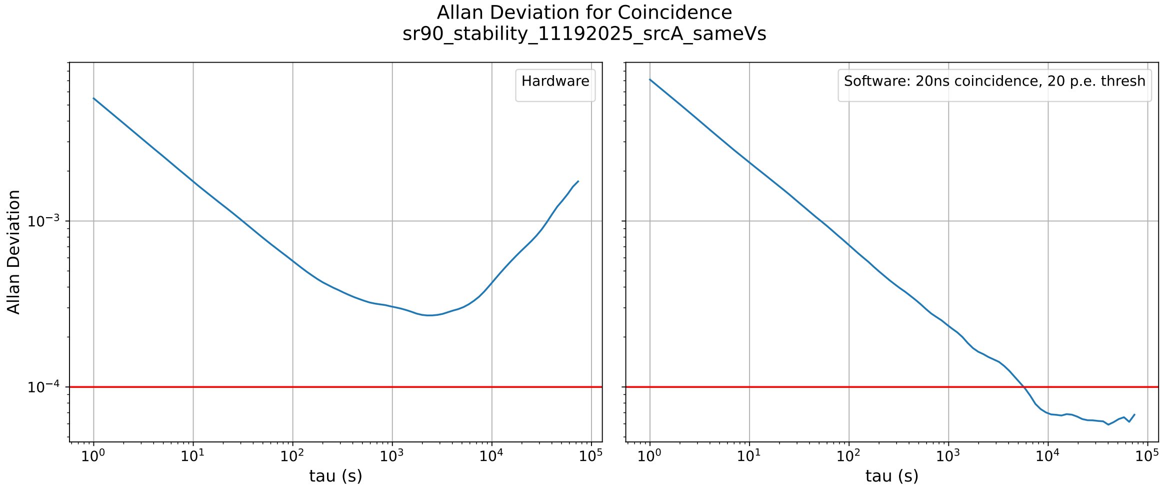

Figure 3.3-2. Overlapping Allan Deviation (OADEV) vs averaging time τ for hardware (left) and softare (right) coincidences.

The OADEV is expected to decrease linearly with the averaging time τ, asymptoting (or increasing) at time-scales over which instabilities are present. The hardware OADEV does not get below 3×10^{-4} but with the software coincidence and the MLT=4 requirement, we easily exceed the target 10^{-4} OADEV target for 1 hour up to 24 hour averaging times.

We can compare the measured fractional standard deviation in the rates vs binning time, σ_{obs}, to that expected statistically from a Poisson process, σ_{stat}. The difference in quadrature can be attributed to the fractional systematic uncertainty on the rate (σ_{sys}). From this analysis, we find that the software coincidence requiring MLT=4 gives an upper limit on the fractional systematic uncertainty on the beta monitor rate of $\sigma_{sys} = 6 \times 10^{-4}$.

3.4 $E\times B$ system for trap emptying

W. A. Byron, N. Buzinsky, B. Dodson, A. García, A. Goodman, A. Gorman, H. Harrington, L. Hayen*, P.T. Kolbeck, B. Maxwell, L. Malavasi, D. McClain†, D. Melconian†, P. Mueller‡, N. Oaxaca, N. S. Oblath§, G. Savard‡, D. Stancil*, D. W. Storm, H. E. Swanson, R.J. Taylor*, J. Tedeschi§, B. A. VanDevender§, C. Wiseman, F. Wietfeldt¶, and A. R. Young*

An $E\times B$ trap emptying system was developed and installed in 2024.[1] The system works

---

*Department of Physics, North Carolina State University, Raleigh, NC.
†Texas A&M, College Station, TX.
‡Argonne National Laboratory, Lemont, IL.
§Pacific Northwest National Laboratory, WA.
¶Department of Physics, Tulane University, New Orleans, LA.

[1] CENPA Annual Review 2024, Section 5.7.

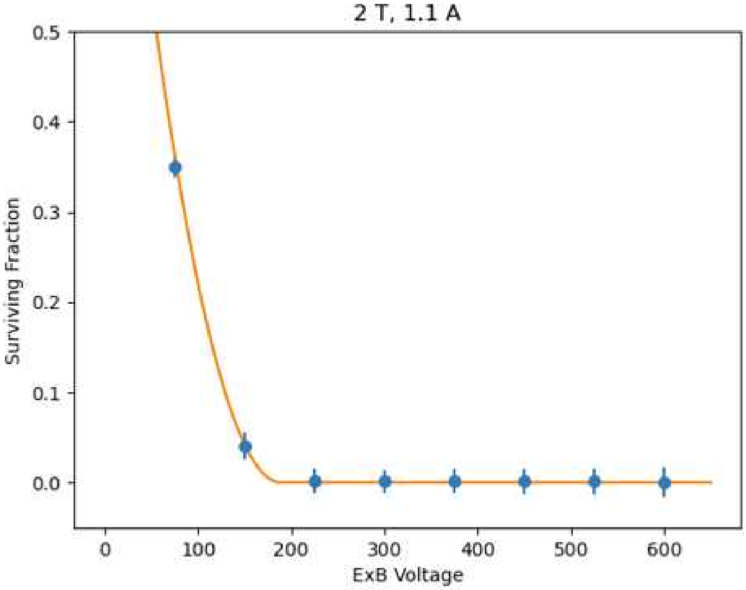

Figure 3.4-1. Surviving fraction of trapped positrons versus $E \times B$ electrode voltage. Main field was set to 2 T and the trap field coils had a current of 1.1 A. Above 200 V the surviving fraction is zero, indicating complete trap emptying.

by applying a 600 V bias to a half-cylinder electrode in the middle of the CRES decay cell which creates an electric field, which in combination with the large axial magnetic field, drifts charged particles to the decay cell walls. Two electrodes sitting opposite each other are charged one after the other to prevent stable trapping regions from occurring during the emptying period. This system was determined to effectively (>∼99.9%) empty positrons in the 2–3.6 T range as shown in Fig. 3.4-1 where the surviving fraction of trapped Ne19 positrons is plotted as a function of $E \times B$ voltage. Grad-B drift counteracts $E \times B$ drift in a circular region of decay volume, which at low voltages increases the surviving fraction of events. The surviving fraction can be predicted as the relative area of the emptying and trapped circular regions:

$$\frac{N_\mathrm{f}}{N_\mathrm{i}} = \left(1 - C\frac{V}{\nabla B}\right)^2 , \tag{1}$$

where N_i and N_f are the initial and final population sizes respectively, and C is a constant describing the relative drifts due to the electric field strength (E) and the magnetic field gradient (∇B). This equation was fit to the trap emptying data and shows the $E \times B$ trap emptying versus voltage is consistent with this expectation.

This year during a physics data run it was discovered that emptying was <≈ 90% efficient at fields less than 2 T. The currently favored explanation is that at the lower field strengths the electric field and magnetic field form two smaller axial traps at the edges of the normal trapping volume. As shown in Fig. 3.4-2 the side traps only form when the electrode is repulsive with respect to the betas: positively charged for positrons and negatively charged for electrons. Since ExB is most efficient in the middle of the trap, these side traps create high survival rates during trap emptying.

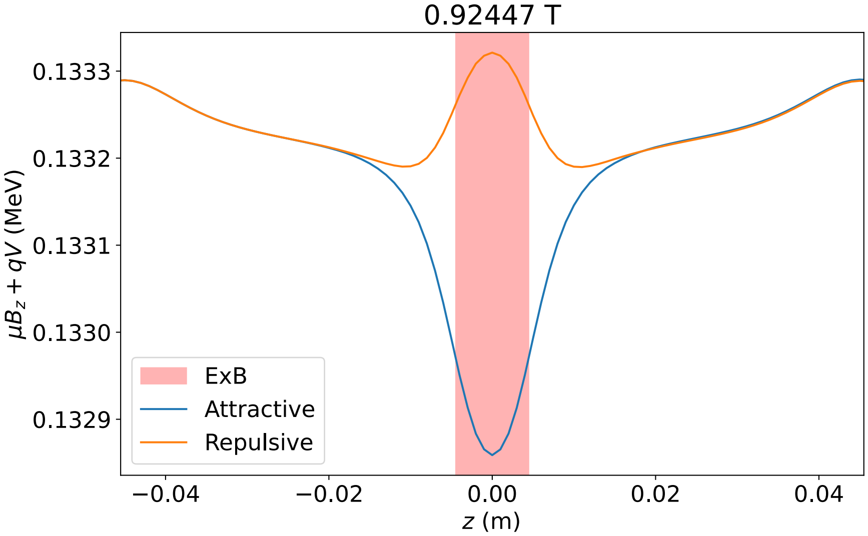

Figure 3.4-2. Magnetic and electric potential in the z-direction for an attractive and and repulsive electrode in the He6-CRES decay cell for standard trapping magnetic (1.5 ppt) and emptying electric fields (600 V) at 0.92447 T. The region of efficient $E \times B$ emptying is shaded in the middle. Electrodes with a repulsive polarity with respect to the betas result in side traps, displaced axially from the $E \times B$ electrodes.

The side traps can be easily prevented by always applying an attractive potential to the electrodes. The original ExB pulse driver was designed to only apply a positive potential to the electrodes, so it cannot effectively empty the magnetic trap for positrons. To address this issue we have acquired a new ExB pulse driver capable of also applying negative voltages to the electrodes (built by the CENPA electronics shop). Further tests of the ExB electrodes are planned during the next He6-CRES data run (early 2026) to confirm the existence of these side traps. We expect to achieve 100% trap emptying at all fields with the additional capabilities of the new ExB pulse driver.

3.5 Frequentist and Bayesian inference of b

N. Buzinsky, A. Fogel, A. García, H. Harrington, L. Malavasi, and R.J. Taylor[*]

He6-CRES is aimed towards measuring b_{fierz} from observations of the beta spectra shapes of ^{6}He and ^{19}Ne with CRES. CRES is advantageous to the extent that it naturally exhibits low backgrounds, good energy resolutions, and symmetric detector responses between electrons and positrons. However, the detection efficiency is a complicated and poorly known function which can mask the desired BSM signature. We describe advances from both the frequentest and the Bayesian perspectives. These efforts place our inference of b given observations of ^{6}He and ^{19}Ne, with shared but unknown detection efficiencies, on a more rigorous footing.

In the frequentist view, the total log-likelihood for He6-CRES is a sum of the CRES and beta monitor measurements, each of which are Poisson distributed. In our model, the

---

[*] Department of Physics, North Carolina State University, Raleigh, NC.

expected CRES counts for isotope i, frequency bin j, and magnetic field k are given by

$$\lambda_{ijk} = A_i\, A^{\beta}_{ik}\, \varepsilon_{jk}\, f_{ijk}\,(b_i)\, \Delta E_{jk}, \tag{1}$$

where b_i is the Fierz interference parameter for isotope i. Our nuisance parameters include A_i, defining the overall gas-density-to-CRES rate scaling for each isotope, A^{β}_{ik}, characterizing the mean (full-spectrum) activity within the decay cell (measured by the beta monitor), and ε_{jk} defining the shared, but unknown CRES detection efficiency.

Despite potentially thousands of nuisance parameters (with the unknown efficiencies ε_{jk}), we greatly simplify the problem by computing the profile likelihood, in which the nuisance parameters are analytically maximized for each b_i. This yields:

$$\log \mathcal{L} = y_{0jk} \log(r_{jk}) - (y_{0jk} + y_{1jk}) \log(1 + r_{jk}), \tag{2}$$

where y_{ijk} is the number of observed CRES counts in isotope i, frequency bin j, magnetic field k. The resulting parametric model of the ratio $r_{jk} = \lambda_{0jk}/\lambda_{1jk}$ is independent of ε_{jk}. Monte Carlo tests with $\sim 10^4$ pseudo-experiments indicate that the frequentist approach is unbiased at the $\sim 10^{-3}$ level.

In the Bayesian view, inference on b is based on its marginal posterior distribution, obtained from a weakly informative (broad) Gaussian prior on b together with the likelihood derived from the observed CRES count data. The likelihood from the Bayesian perspective is the same as in the frequentist analysis. We take a joint Poisson likelihood of the observed counts (y_{0jk}, y_{1jk}) conditional on the ratio of expected counts r_{jk} and a nuisance parameter representing the shared efficiencies ε_{jk}. After marginalizing over ε_{jk}, this method gives the same likelihood in Equation 2, up to a normalization constant.

The posterior distribution is evaluated using Markov chain Monte Carlo (MCMC) sampling techniques, and the marginal posterior distribution of b is obtained by numerically integrating over nuisance parameters.

3.6 Sideband clustering

N. Buzinsky and H. Harrington

While the ExB system (Sec. 3.4) clears the CRES decay cell of betas, limiting the maximal time duration of events, multiple events may be present simultaneously, to some probability.

In Fig. 3.6-1, we show the Monte Carlo distribution of band powers as a function of band order and the axial frequency of the underlying beta. Given a minimum detection threshold (for false alarm rejection) at approximately 0.75 fW, we get different “classes” of detectable event topologies at different axial frequencies.

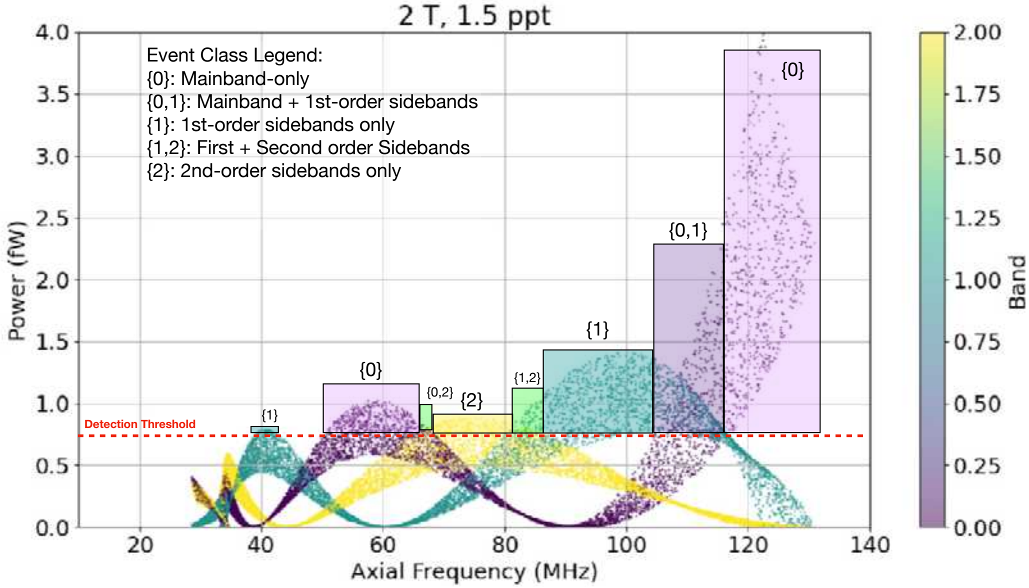

Figure 3.6-1. Monte Carlo simulation of sideband power vs. axial frequency of the underlying beta. Given some (SNR-limited) detection threshold at approximately 0.75 fW, we expect a correlation between the observed band structure of events and the axial frequency of the beta.

In conjunction with the relatively wide ($\sigma \sim 50$ MHz) distribution of axial frequencies, and the variety of sideband topologies that may be detected, an ambiguity arises in which multiple reconstructed bands may have multiple allowed reconstruction interpretations, with more or less hypothesized CRES events (with less or more sidebands per event, respectively). In the simplest case, two reconstructed tracks within the same trap acquisition may be due to either two independent CRES events, each with a single visible band, or one CRES event, with (some multiple) of the axial frequency given by the frequency spacing.

To decide, we choose the more probable hypothesis, found by assessing the relative probabilities of getting the number of hypothesized CRES events and the probability of the events having the hypothesized band topologies.

The total number of possible ways to group N reconstructed tracks in a single trap acquisition into events grows like $\sim N^N$. For $N < 10$, we are able to use brute-force, with some automatic pruning of events that put far-separated tracks in the same event.

The number of events per trap acquisition are Poisson distributed, with the rate parameter λ estimated via the fraction of trap acquisitions with any (≥ 1) events, yielding $\lambda \approx 0.025$.

Despite broad agreement, there exist multiple $\mathcal{O}(1\%)$-level discrepancies between Monte Carlo and data in the distribution of start frequency differences between pairs of reconstructed bands within a trap acquisition. For instance, predictions for the relative probabilities and the axial frequencies for $\{0, 2\}$, $\{1, 2\}$, and $\{1\}$ events are in tension with data. These data provide a unique probe of our magnetic field model, which we aim to refine further in 2026. Our magnetic trap model would be further constrained via a lineshape measurement as in

(Sec. 3.10).

3.7 2025 data-taking results

N. Buzinsky, B. Dodson, A. García, H. Harrington, L. Malavasi, M. Newman,
N. Oaxaca, R.J. Taylor*, and C. Wiseman

2025 saw the He6-CRES collaboration's first attempt at a 10^{-2} data campaign. The goals were as follows:
1) Measuring b_{Fierz} with CRES to $\sim \mathcal{O}(10^{-2})$ sensitivity
2) Collecting a contiguous spectrum from 0.02 – 2.35 MeV, with overlap of spectrum segments to demonstrate stitching/monitor normalization
3) Collecting CRES data beyond the endpoint of ^{19}Ne, constraining contaminant backgrounds
4) Refining the b_{fierz} analysis to use our high-precision energy measurements, demonstrating the cancellation of energy-dependent efficiencies between isotopes
5) Construction of a preliminary systematic budget
6) Scouting out a feasible path for a subsequent 10^{-3} measurement

The run plan for this data campaign consisted of 7,000 s of data per field over 16 magnetic fields between 0.75 T – 3.6 T, for each isotope. This corresponds to 112,000 s of total data per isotope (> 1 day) and was projected to give a $\approx 6.7 \times 10^{-3}$ statistical sensitivity to b_{Fierz}.

---

*Department of Physics, North Carolina State University, Raleigh, NC.

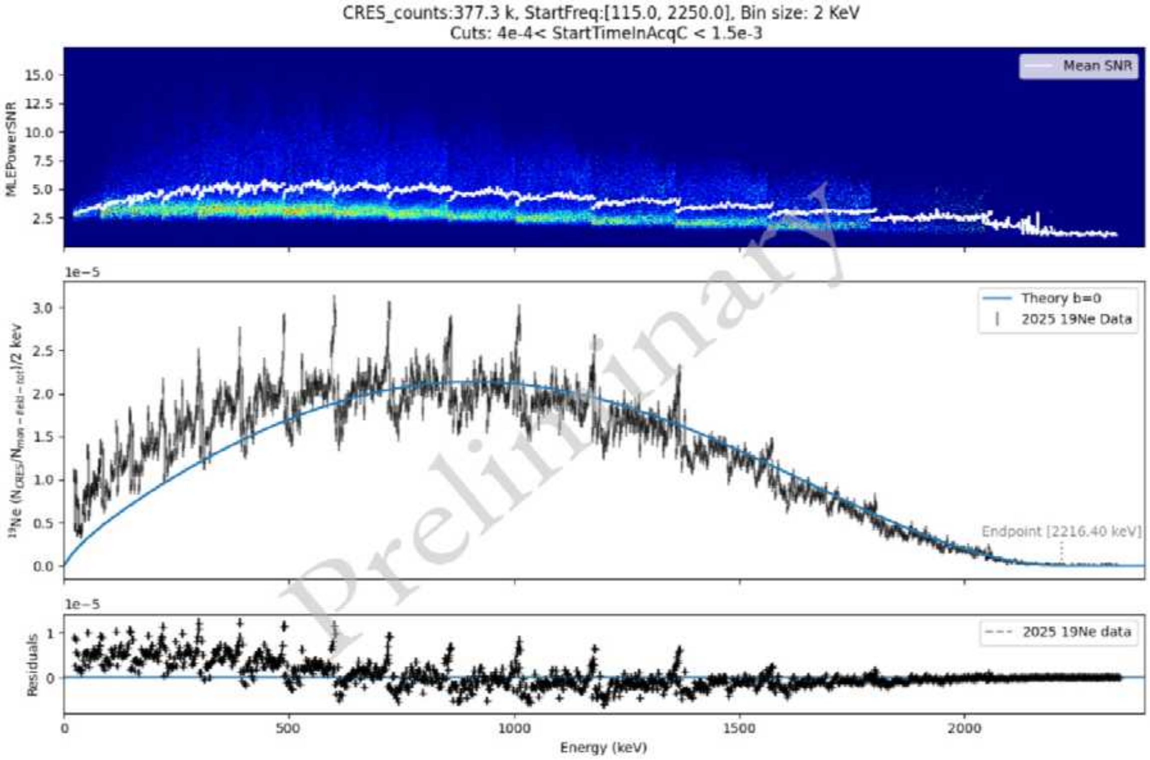

Figure 3.7-1. Top: distribution and mean of measured SNR in each kinetic energy bin. Middle: Preliminary ^{19}Ne spectrum with start frequency, start time, and reconstructed band number cuts, overlaid with an unfitted Standard Model ^{19}Ne beta spectrum. Bottom: residuals between data and the SM spectrum.

In July 2025, we successfully completed the 16-field science run, described above, as well as several systematic studies, with CRES events from decays of ^{19}Ne. However, issues with the CENPA Tandem Van de Graaff accelerator prevented us from collecting the corresponding ^{6}He dataset, necessary for a b_{Fierz} measurement.

The incomplete science run did reveal useful insights, as there were unexpected failures in the trap emptying at low fields, prompting a rapid response diagnosing and potentially resolving the issue (see (Sec. 3.4)). In combination, these factors motivate re-taking the ^{19}Ne dataset in early 2026, along with the ^{6}He data.

Our systematic studies investigated the normalized beta monitor rate with respect to the magnetic trap depth, the linearity of the CRES rate to the beta monitor rate, and the reproducibility of our system, measuring at a given magnetic field. The latter is shown in Fig. 3.7-2 and was performed by collecting CRES data at a single field on separate days, re-ramping the magnet in between.

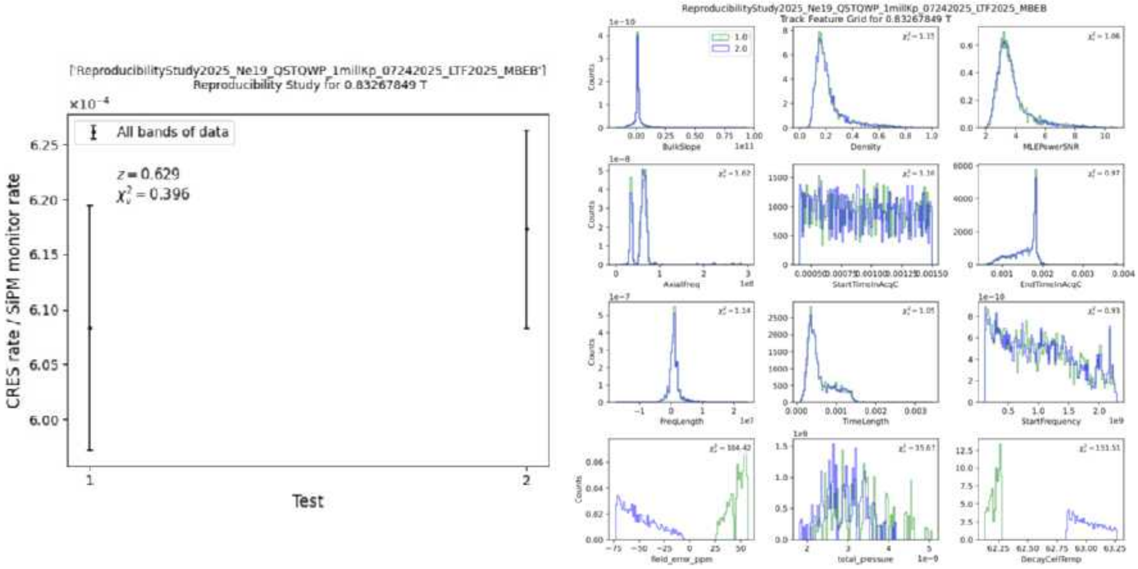

Figure 3.7-2. Summary of reproducibility test repeating a spectrum segment measurement at a single field (0.8326 T). Left: total CRES rate / beta monitor rate for each test. Right: comparison of reconstructed track features between the two measurements.

These datasets have provided useful testing grounds for our reconstruction, motivating improvements to band identification/clustering algorithm and to our cuts on ill-behaved tracks, robust to noise fluctuations. The data, though incomplete and imperfect, has also spurred work in improving the Monte Carlo modeling of our system, which will be need to be well-understood for the eventual fits of b_{fierz}.

3.8 Bounding beta-decaying contaminants in ^{19}Ne Data

M. Blum, Z. Fu, A. García, G. Hander, and H. Harrington

To estimate the extent to which the measurement of b_{fierz} will be biased from beta decays of other radioactive contaminants created in concurrence with ^{19}Ne/^{6}He production, it is essential to place an upper limit on these contaminants within the decay cell. To some extent, CRES measurements have some capability of detecting beta-producing radioactive contaminants via distortions in the observed energy spectra (particularly beyond the endpoint of ^{19}Ne). However, given the unknown energy-dependent detection efficiency, even the potential presence of beta-decaying contaminants can greatly reduce our sensitivity to fits of b_{Fierz}.

To avoid this, we consider a complementary study in which the target is sequentially hit, and then not hit, by the beam with durations T_{loading} and T_{decay}, respectively. One therefore gets a handle on the composition of the produced isotopes (intended + contaminant) by observing the time-dependence of decays during T_{decay}, owing to the different half-lives among the isotopes. Having previously completed the analogous study for ^{6}He, we aim to do so for ^{19}Ne. For these purposes, we require a fitting procedure capable of estimating the contaminant-to-^{19}Ne rate ratio $R_{\mathrm{cont}/^{19}\mathrm{Ne}}$ immediately after the ^{19}Ne loading process.

The reaction used to produce ^{19}Ne is $^{19}\mathrm{F}(p,n)^{19}\mathrm{Ne}$, where a 12 MeV proton beam is fired into gaseous SF_6. Based on the Brookhaven National Laboratory's Q-Value Calculator, six radiative isotopes are identified as the predominant sources of contamination of the ^{19}Ne β energy spectrum. The time-binned total population of all beta emission sources is modeled as a linear combination of exponential decay functions corresponding to ^{19}Ne and these six main contaminant isotopes, together with a constant background contribution. With the decay constants fixed, calculation of the isotope populations at the time of interest reduces to a linear fitting problem. To ensure physical interpretability, non-negative linear least squares (NNLS) is used to impose that the fit only consider non-negative populations of each isotope. With the model specification and the non-negativity constraint in place, the fit calculates the isotope populations, required for estimating $R_{\mathrm{cont}/^{19}\mathrm{Ne}}$ at the time of interest. An example decomposition of a time-binned simulated dataset into ^{19}Ne, contaminants, and background components is shown in Fig. 3.8-1.

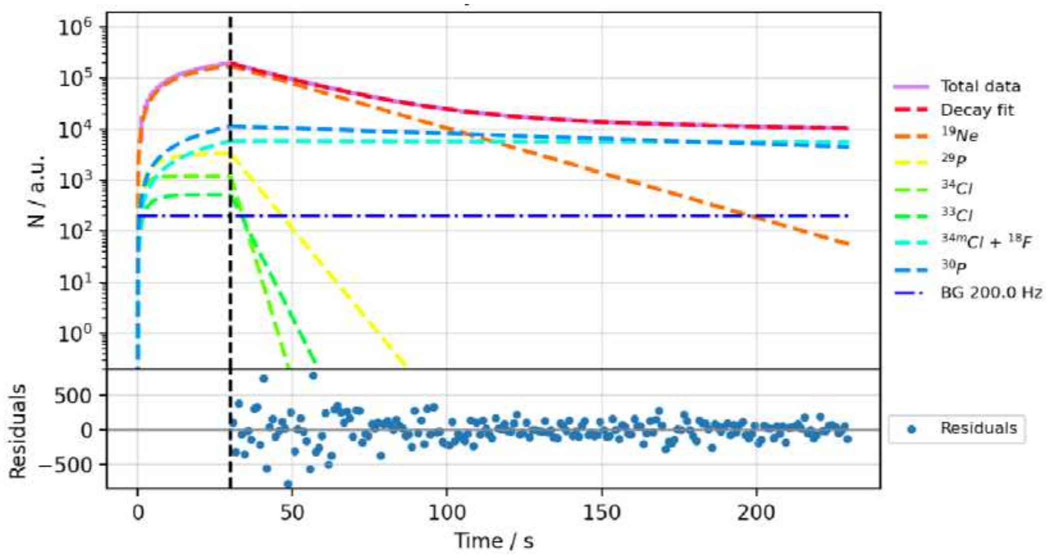

Figure 3.8-1. An example decomposition of population curve for each isotope, from simulated data. We aim to extract upper bounds on these radioisotopes, with respect to ^{19}Ne, at the end of the loading phase.

With the non-negativity constraint on fitted parameters, the cost function can deviate from a local quadratic form and potentially cause asymmetric confidence intervals. Therefore, uncertainties for the calculated populations of isotopes are obtained using profile-$\Delta\chi^2$ (MINOS-style) intervals rather than symmetric covariance estimates. These calculated errors for each isotope are then propagated to calculate the uncertainty for the $R_{\mathrm{cont}/^{19}\mathrm{Ne}}$ (rate ratio error), which is used to optimize the run configuration of the experiment.

A forward simulation framework was implemented to generate synthetic time-binned datasets that emulate the experimental acquisition structure. These datasets are generated with known inputs (loading rate, decay constants, and background assumptions) and can be processed through the same fitting and uncertainty workflow used for experimental data. A Monte Carlo type sanity check algorithm is applied based on this framework to test uncertainty coverage and calculation stability.

In addition, an experimental configuration optimization function is also created based on the same simulation framework mentioned above. For a specified total run-time limit and a defined set of contaminant isotopes, the optimizer minimizes the expected uncertainties in $R_{\mathrm{cont}/^{19}\mathrm{Ne}}$ over timing choices $(T_{\mathrm{loading}}, T_{\mathrm{decay}})$.

Following the planned collection of this dataset in early 2026, the next step will be to apply the validated analysis pipeline to the real time-binned decay data, extracting upper bounds on contaminants with uncertainty estimates. $(T_{\mathrm{loading}}, T_{\mathrm{decay}})$ can be further optimized for subsequent runs given information gained in initial runs (as the isotope composition becomes better known).

3.9 Estimating systematics errors from contamination

N. Buzinsky and A. Fogel

In contrast to statistical errors, which arise from statistical fluctuations in (finite) datasets, systematic errors arise from uncertainties or errors in one's fitting model. To investigate this, we consider fits of Monte Carlo datasets in which the simulated dataset includes some additional feature that is unaccounted for in the fitting model. We consider how unmodeled contamination, firstly by a second beta-decaying contaminant gas, and secondly by incorrect sideband reconstruction, can result in biases in fits to b_{Fierz}. These studies inform the degree to which we need to limit beta-decaying contaminants within our decay cell (Sec. 3.8), or the degree to which sideband tagging must be accurate (Sec. 3.6).

For these studies, as a first step, we considered a single-isotope measurement only (with known detection efficiency), simulating events drawn from a Standard Model spectrum of ^{19}Ne ($b_{\mathrm{fierz}} = 0$). Fitting the resultant data for b_{fierz} then informs us about the expected bias that will arise, depending on the type and the prevalence of the contamination. The number of simulated events were orders of magnitude larger than expected for the 1% science campaign, so that the systematic uncertainty was the dominant source of discrepancy from the Monte Carlo true value of b_{fierz}. Systematic errors should be compared with respect to the targeted $\sigma_b \approx 0.01$ for the upcoming campaign.

First, we examined the impact of contamination of the magnetic trap by radioactive isotopes other than the isotope of interest (see Fig. 3.9-1).

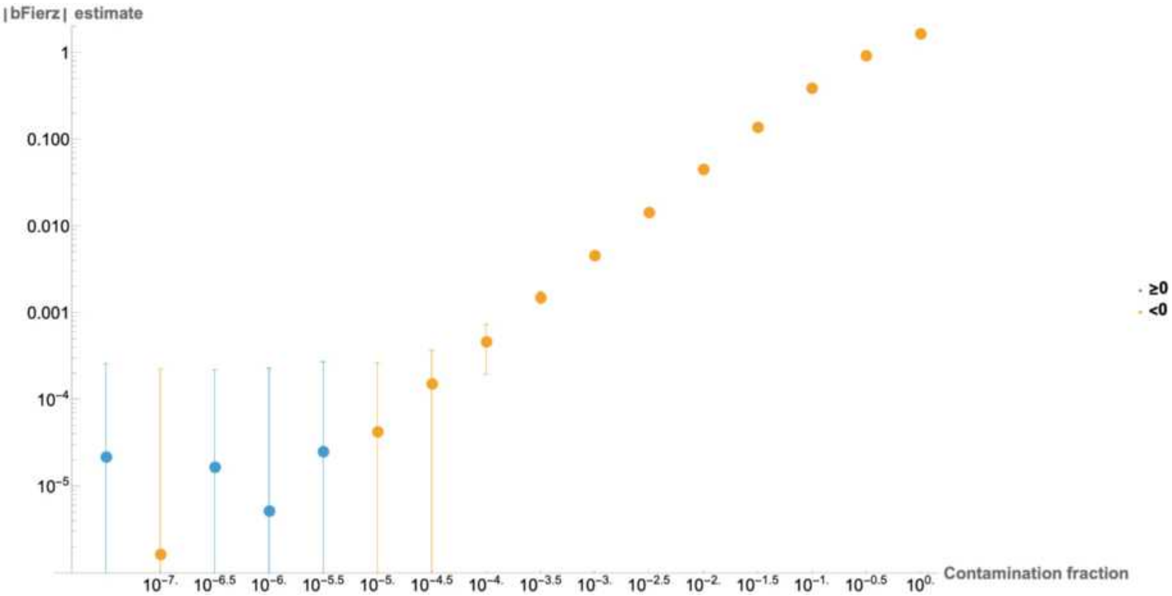

Figure 3.9-1. Calculation of b_{Fierz} in the presence of contamination of ^{19}Ne by ^{34}Cl for various contaminant fractions, assuming no contamination for the fit.

The preliminary result is that for a finding of $|b_{\mathrm{Fierz}}| < 10^{-3}$, the proportion of total events resulting from beta decays of ^{34}Cl must be $<\sim 10^{-4.5}$.

The second source of systematic error we examined was the contamination of sidebands (Sec. 3.6) within our list of reconstructed event start frequencies. For the simulations implementing this preliminary analysis, we made the following assumptions and algorithmic choices:

1. Every decay event generated 3 bands at f_c and $f_c \pm f_a$ (mainband, upper sideband and lower sideband, respectively)
2. The reconstruction of mainbands, upper sidebands, and lower sidebands as being event start frequencies was parameterized by an efficiency parameter for each band type. Ideally, $m = 1$, $u = \ell = 0$.
3. The efficiencies chosen for each simulation were constrained for fixed $m + u + \ell$, so that we only compared studies with equal total efficiency.
4. The mainband frequency for each generated decay event was chosen at random using a probability distribution over energy consistent with $b_{\mathrm{Fierz}} = 0$.

The generated statistics from each simulation was used to calculate a range for b_{Fierz} as in the isotope contaminant simulations. The results are shown in Fig. 3.9-2.

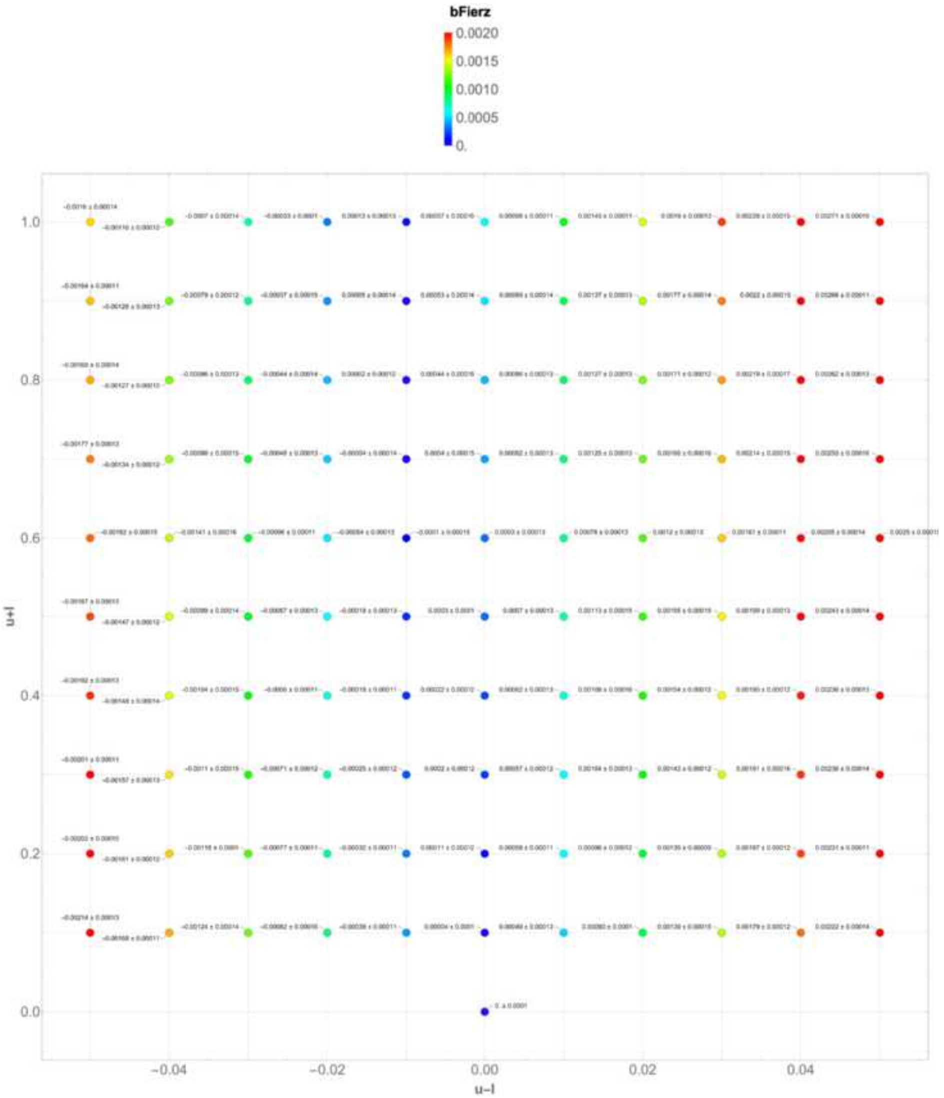

Figure 3.9-2. Calculation of b_{Fierz} in the presence of sideband spectral contamination for various detection efficiencies, assuming no contamination for the calculation. The horizontal axis represents the difference between the upper and the lower 1st-order sideband detection efficiencies. For each data point, b_{Fierz} values were calculated and averaged over 32 simulation iterations, with each iteration consisting of 10^{10} generated decay events.

The key outcome of this study is that biases in b_{Fierz} arising from sideband misidentification are primarily a function of $u - \ell$, that is misidentification probability *differences* between upper and lower sidebands. While misidentifying both upper and lower sidebands as mainbands add variance to our lineshape, asymmetries bias b_{Fierz} estimates. From this study, we find that we can measure b_{Fierz} to 10^{-3}, without a correction for sideband contamination, provided that upper vs. lower sideband efficiencies differ by $< 5\%$.

3.10 Towards a ^{79}Kr source

A. García, N. Oaxaca, D. W. Storm, H. E. Swanson, and B. Yommarath

To characterize the CRES detection system over a wide range of energies, we are working on the production, purity, and transport of a ^{79}Kr calibration source. ^{79}Kr produces conversion electrons of $\approx$ 11 keV, 31 keV, 130 keV, 604 keV, as well as a β^+ spectrum with a $\approx$ 604-keV endpoint. We are using the $^{79}\mathrm{Br}(p,n)^{79}\mathrm{Kr}$ reaction by firing a proton beam from the tandem accelerator at 13 MeV at a potassium bromide (KBr) target through a 2 mil Al window. Previous production tests had resulted in high levels of vacuum system contamination at the level of $\approx$ 0.3 Torr. Less than $< 10^{-9}$ Torr of the total pressure was due to ^{79}Kr which we calculated using gamma-ray measurements from a germanium detector. The measurements revealed that there were no observable radioactive contaminants that resulted from ^{79}Kr production.

To decrease contamination during ^{79}Kr production, we added an LN2 cold trap in the transport line between the production cell and our vacuum system in Cave 2. After we installed the cold trap, we tested again to confirm that we were producing ^{79}Kr and to verify that the cold trap improved the pressure in the vacuum system. We measured the same gamma rates with our germanium detector and found that the pressure in the vacuum system was $\approx 10^{-6}$ Torr after 5 minutes of production. For our next test, we transported the radiation to our low background vacuum chamber in Cave 1 while we continuously monitored gamma rays from the source. We observed gammas from ^{79}Kr, ^{79m}Kr, and ^{81m}Kr with our germanium detector while the proton beam continuously hit the target in Cave 2 as shown in Fig. 3.10-1. ^{81m}Kr has a half-life of 13.10 s and emits electrons with 190.44 keV which is useful for calibration.

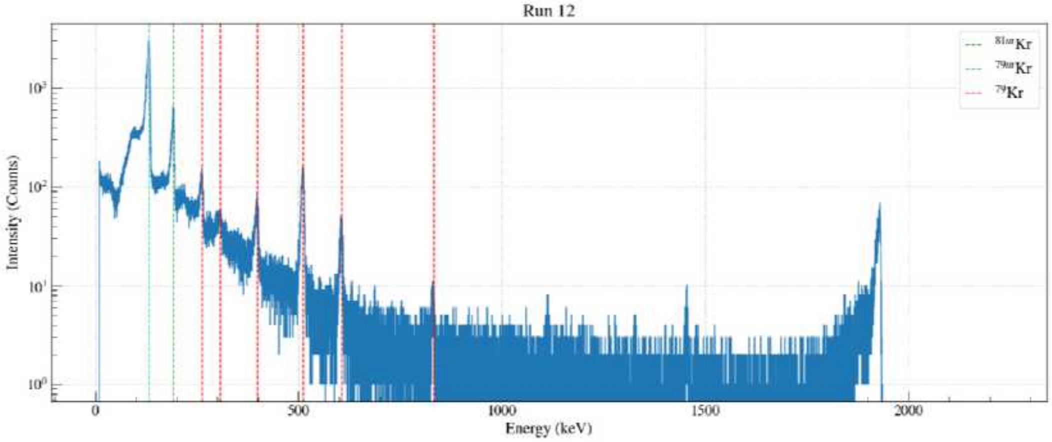

Figure 3.10-1. Gamma spectrum showing branches from ^{79}Kr, ^{79m}Kr, and ^{81m}Kr.

During irradiation of the KBr target, the plastic mount melted while we bombarded the target with a proton beam at 12 MeV and 6 μA for a duration of around 20 minutes. The mount has to be electrically insulating and able to withstand 72 W, so we remade the mount out of Macor. For our next steps, we plan to send the source to the Cave 1 RGA to test its purity as it enters the low background vacuum system and take CRES data of the monoenergetic electrons emitted by ^{79m}Kr and ^{81m}Kr while heating the CRES cell above ≈ 75 K so that the krypton does not freeze out on the waveguide walls.

3.11 Alternative time-frequency approaches

L. Malavasi

To study signals with time-varying spectral content, such as CRES signals, it is standard to use a time-frequency representation (TFR) of the signal, a transformation that maps a signal in the time-domain $s(t)$ to a 2D space that encodes both time and frequency-dependent characteristics. CRES experiments currently apply the windowed (or short-time) Fourier transform to received cyclotron radiation, the simplest and most commonly used TFR.

The construction and interpreting of TFRs is known in signal processing as time-frequency analysis. The Gabor-Heisenberg uncertainty limit $\sigma_t \sigma_f \geq \frac{1}{2}$ imposes the inability of a linear and unbiased transform to represent an arbitrary signal with perfect simultaneous time and frequency localization; selecting appropriate techniques to represent and analyze a signal despite this uncertainty is a fundamental challenge of time-frequency analysis. A heuristic interpretation of the uncertainty principle is that there can be no universally valid or preferred time-frequency representation of a signal. As Auger et al. (2013) emphasize, the result of the transform depends jointly on the mathematical properties of the transform itself and the signal being analyzed. Thus, time-frequency representations should not be understood as objective “truths” about a signal, but rather as measurement tools selected according to the analysis being performed. Many TFRs have been developed for a variety of applications, including the familiar windowed (or short-time) Fourier transform, wavelet transforms, and more exotic representations.

There is some preexisting work, especially from Project 8 analysis efforts, exploring the possible use of other TFRs to CRES signals. Representations explored include wavelet transforms, dechirped transforms such as the fractional STFT, chirplet/warblet transforms, Wigner-Ville distributions, and related transforms (broadly forming a family of TFRs known as Cohen’s class). Other signal processing techniques outside of time-frequency analysis have also been considered, such as matched filtering, polyphase filter banks, and machine learning-based signal reconstruction. Thus far, the advantage of alternate TFRs for CRES over the STFT has not been conclusively established.

Modern advances in time-frequency analysis have established a technique known as synchrosqueezing, a special case of a “reassignment method”. Synchrosqueezing constructs a new time-frequency representation from the STFT or wavelet transform representations in which spectral energy (initially smeared according to the Gabor-Heisenberg uncertainty) is reassigned along the frequency axis to the estimated true instantaneous frequency of the signal, inferred from the complex phase information. In principle, the signal “collapses” to the true instantaneous frequency, with theoretically perfect time-frequency localization for an AM-FM modulated signal, provided the signal is sufficiently well-behaved. An additional benefit of

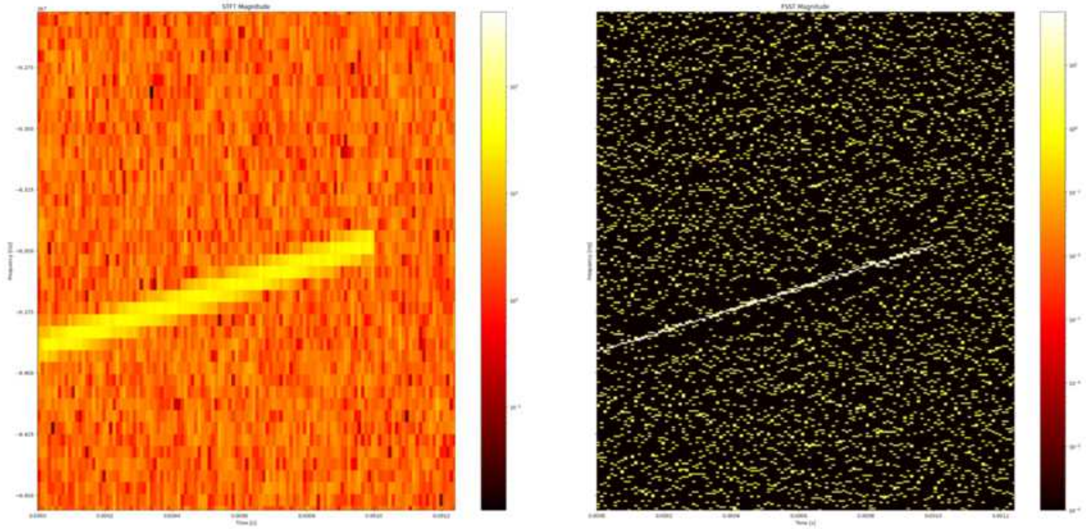

Figure 3.11-1. *Left*: A short-time Fourier transform representation of a Project 8 CRES event from Phase II data; *Right*: A synchrosqueezed Fourier transform representation of the same event.

the synchrosqueezed transform is that the signal combines coherently at the instantaneous frequency, while noise adds incoherently elsewhere, creating a denoising effect.

We are exploring the potential advantages of using synchrosqueezed Fourier transforms for CRES event reconstruction for He6-CRES. In particular, our aim is to address the following questions:

- Can synchrosqueezed CRES signals be reconstructed with a more precise energy resolution? This could reduce uncertainty when constructing the spectrum and inferring b_{Fierz}.
- Does the synchrosqueezing of CRES signals improve their detection due to denoising and enhanced SNR? This could improve the effective event rate by increasing the sensitivity to low-power events.

An example of a traditional short-time Fourier transform compared to a synchrosqueezed Fourier transform of real CRES event time-series data from Project 8 is shown in Fig. 3.11-1. Synchrosqueezed wavelet transforms for CRES are also being explored in parallel by a team in the Project 8 analysis group.

3.12 Activating additional ROACH channels

A. Goodson and L. Malavasi

Several upgrades to data acquisition (DAQ) hardware are underway. The motivation for these upgrades is primarily to double the DAQ bandwidth from 2.3 GHz (between 18.0-20.3 GHz) to 4.3 GHz by adding another readout channel in the 20-22.3 GHz range. In the He6-CRES apparatus, RF signals from the cryogenic receiver stage are downmixed with a 17.9 GHz local oscillator (LO) to the 0-2.4 GHz baseband and filtered by an ambient receiver stage, then transmitted to the CASPER ROACH-2 FPGA platform (ROACH). The ROACH digitizes the baseband signal and performs the windowed Fourier transform before streaming the digital data on two parallel 10 GbE lines to a network interface card for online compression and storage. Historically, He6-CRES received RF signals from both sides of the waveguide system on two identical DAQ channels. We now read one side and place a waveguide termination on the other to reduce reflections in the waveguide leading to interference. As such, one of the two ambient receiver channels and Xilinx Virtex-6 FPGAs in the ROACH is unused. With minimal hardware adjustments, this unused second channel can be repurposed to process signal content in the 20.2-22.5 GHz band.

The necessary changes are splitting the single-channel input signal into the current channel and the new channel, replacing the input 18-20 GHz bandpass filter with a 20-22 GHz filter, and adding an independent LO at 19.9 GHz to downmix the higher frequency band to the same 0-2.4 GHz baseband. After downmixing, the ambient receiver and the ROACH hardware are identical between channels. A schematic summary of changes is shown below.

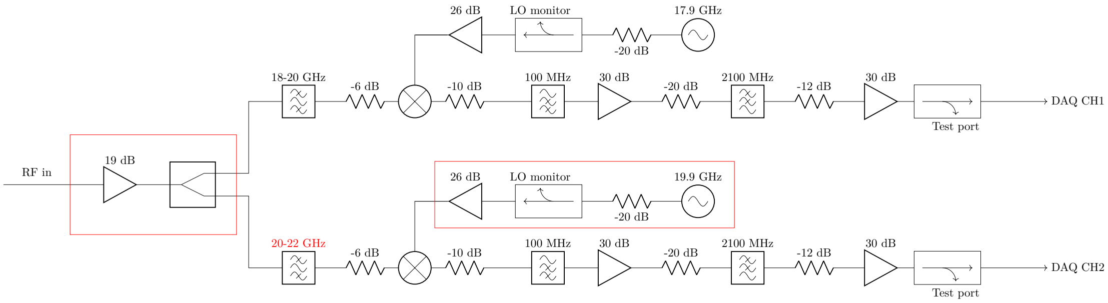

Figure 3.12-1. A schematic of the ambient receiver DAQ upgrades to double the receiver bandwidth, with modified components highlighted.

Furthermore, we are developing a revision of the ROACH firmware to record the real and imaginary components of the Fourier transformed signal, rather than just the complex magnitude-squared. The phase information of the CRES signals could enable more precise reconstruction algorithms. To overcome network bottlenecks anticipated by both doubling the bandwidth and writing the phase information, we are developing an online zero suppression routine to be performed on the ROACH.

3.13 Penning Trap

N. Buzinsky, A. García, H. Harrington, L. Malavasi, and R.J. Taylor*

Ongoing efforts to increase the detector bandwidth (Sec. 3.12) and improve the gas transport efficiency portend a potentially massive increase in the CRES rate in the near-future. If the conjectured 100× increase in the event rate from the Cryo-II upgrade comes to fruition, sideband clustering (Sec. 3.6), becomes significantly more difficult. The event rate would grow from $\sim$ 30 Hz to $\sim$ 3 kHz, and we could expect $\sim$ 6 events per trap acquisition. For a reasonable length b_{fierz} campaign, we can expect at least one trap acquisition with $\sim$ 20 simultaneous events. The difficulty of clustering reconstructed bands into events will get more pronounced as further detector optimizations continue to increase the event rate.

Ambiguities in assigning reconstructed tracks to events could be greatly mitigated if all betas had a fixed axial frequency, as opposed to the high-variance ($\sigma_{f_a} \sim 50$ MHz) axial frequency distribution that we have currently. To achieve this, we consider *electrostatic*, as opposed to *magnetostatic* traps for CRES.

For the harmonic, electrostatic trapping potential

$$V(z) = V_0 \left(\frac{z}{L}\right)^2, \tag{1}$$

we find that the motion is approximately given by a harmonic oscillator, with axial frequency and turning point position

$$\omega_a = \sqrt{\frac{2qV_0}{\gamma_t m L^2}} \tag{2}$$

$$z_t = \pm\sqrt{z_b^2 + \frac{L^2 \gamma_t m v_b^2 \cos^2\theta_b}{2qV_0}} \tag{3}$$

where z_b, v_b, θ_b are the position, velocity, and pitch angle of the beta at birth, and $\gamma_t = \left(1 - v_\perp^2/c^2\right)^{-1/2}$. The result is that axial frequencies are projected to vary by $\sim$ 0.01 MHz for a electrostatic trap, for a given energy, at similar estimated trapping efficiencies (Fig. 3.13-1).

The proposed Penning trap would use the existing ExB electrodes charged with the opposite sign with respect to the betas in the cell. As inferred in (Sec. 3.4), we believe that this would trap the betas in the z-direction. To clear the system one (or both) electrodes may be discharged, eliminating the Penning trap.

The Penning trap is not without new unknowns; the approach imposes a floor on the energy resolution attainable (which depends on the position of the beta decay, which cannot be inferred from CRES signals). In addition, the approach faces the challenge of requiring large voltages for historically comparable trapping efficiencies for MeV-scale betas.

Preliminary feasibility tests of the proposed Penning trap CRES system are scheduled for early 2026 following the completion of the 1% b_{fierz} measurement campaign.

---

*Department of Physics, North Carolina State University, Raleigh, NC.

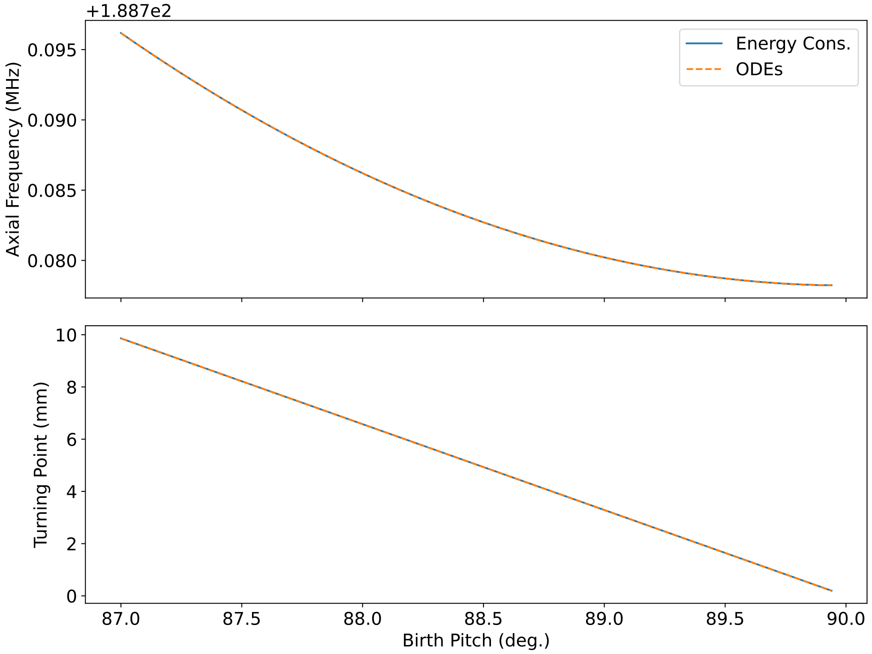

Figure 3.13-1. Axial frequency (upper) and turning point (lower) vs. birth pitch angle for betas with $z_b = 0$, $\gamma_b = 1.5$. All pitch angles are within ~ 0.01 MHz, so it is well-described by a tautochrone, for fixed energy. Even at $\gamma_b = 5.5$, the Penning trap is tautochronic to within ~ 0.01 MHz.

3.14 Shim coil-control for improvement in magnetic field uniformity

A. García, S. Gopal, and H. Harrington

CRES relies on the precise measurement of the cyclotron frequency from betas confined in our magnetic trap. Since the cyclotron frequency is directly proportional to the magnetic field, unaccounted spatial or temporal variations in the field will result in biases in the estimated beta energies, with adverse implications for the measurement of b_{fierz}. For this reason, magnetic field uniformity and stability are fundamental to the performance of CRES.

Currently, the un-shimmed field uniformity has been mapped to be 15 ppm/cm. Our AMI superconducting magnet is equipped with eight superconducting shim coils that generate magnetic field corrections to B_z with different multipoles proportional to (X, Y, Z, Z^2, YZ, XZ, XY, and $X^2 - Y^2$). Magnetic field shimming allows us to fine-tune our background magnetic field as desired, overcoming intrinsic imperfections in the magnet or geometric constraints that affect the uniformity. A well shimmed field produces a uniform background field, with well-understood beta motion, a prerequisite for Monte Carlo modeling of the system. A system to map the field with and without current on each shim, and use these maps to calculate optimal shimming currents has already been developed and implemented. We would now like to automate the application of the calculated shim currents to reach <1 ppm/cm uniformity which until now has been done manually (and therefore, infrequently).

To modify the magnetic field associated with a particular multipole moment, the shim heater linked to that specific shim coil will be turned on, putting the specific shim coil out of its superconducting state so that a current can be applied. Once the target current is reached, the heater is turned off, allowing the shim coil to return to its superconducting state, with the optimized current persisting.

The first phase of the autoshimming development involved replacing the previous Arduino controller that managed all of the shim heaters collectively, as it offered no flexibility for independent control of individual heaters and was not suited to be remotely controlled. The previous micro-controller was then replaced with a Raspberry Pi (model 3B+) which was set up to perform the exact functions of the Arduino while allowing each shim heater to be controlled remotely and independently.

In the next phase of this project, a programmable variable power supply will be integrated with the Raspberry Pi-controlled heater system to deliver current to the shim coils. Communication between both power supply and shim heaters will be incorporated into a python script. The Raspberry Pi coordinates heater activation which allows us to apply current to the shim coil when it is out of the superconducting state.

During this entire process, the currents are ramped up in small, controlled increments to avoid excessive induced voltages and reduce the risk of quenches. When the magnetic field correction is no longer needed, another script responsible for ramping down signals the Raspberry Pi to activate the heater of the relevant shim coil and the power supply safely ramps the current down to zero.

In the future, we will perform an in situ validation of the effectiveness of this system on the magnetic field uniformity via observations of the lineshape of ^{79}Kr (Sec. 3.10).

4 Dark matter searches

ADMX (Axion Dark Matter eXperiment)[1]

4.1 Searching for QCD axions between 800 MHz and 2 GHz

M. Carhart, K. Enzian, M. Guzzetti, C. Goodman, C. Hanretty, O. Kim, L. J. Rosenberg, G. Rybka, J. Sinnis, and D. Zhang

The Axion Dark Matter eXperiment (ADMX) is an axion haloscope[2] housed at the University of Washington's Center for Experimental Nuclear Physics and Astrophysics. It consists of a 136-liter cylindrical, copper-plated, resonant microwave cavity with a copper tuning rod that is immersed in a 7.6 Tesla superconducting solenoid. The magnetic field stimulates the decay of axions into microwave photons, producing a signal which, when on resonance with the cavity frequency, becomes amplified, making something once thought to be undetectable detectable.

Axions were first proposed as a solution to a long-standing mystery related to the quantum chromodynamics (QCD) Lagrangian, known as the strong CP problem, by Peccei & Quinn in 1977[3]. Not long after the axion was introduced, physicists realized that its properties such as electric neutrality, feeble interactions with standard model particles, and abundant production in the early universe make it a compelling dark matter candidate as well.

There are two benchmark models that describe the QCD axion: the Kim-Shifman-Vainshtein-Zakharov (KSVZ)[4,5] model and the Dine-Fischler-Srednicki-Zhitnitsky (DFSZ)[6,7] model. Many find the more weakly coupled of the two, DFSZ, to be more compelling due to its compatibility with Grand Unified Theories. However, reaching this level of sensitivity is nontrivial. ADMX was the first group to reach DFSZ[8], and to date remains one of only two groups[9] that have reached it at all.

So far, ADMX has excluded both KSVZ and DFSZ axions between 645-800 MHz (2.66-3.3 μeV) over the course of two data taking runs[10,11] Additionally, ADMX excluded KSVZ and DFSZ axions between 800-1015 MHz (3.3-4.2 μeV) and 943-991 MHz (3.9-4.1 μeV) respectively during its third data taking run[12]. A fourth data taking run excluded KSVZ axions

---

[1] ADMX is supported by the DOE Office of High-Energy Physics and makes use of CENPA resources by recharge to the cost center.

[2] P. Sikivie, Phys. Rev. Lett. **51**, 1415 (1983).

[3] R.D. Peccei and H. Quinn, Phys. Rev. Lett. **38**, 1440 (1977).

[4] J.E. Kim, Phys. Rev. Lett. **43**, 103 (1979).

[5] M.A. Shifman, A.I. Vainstein, and V.I. Zakharov, Nucl. Phys. **B166**, 493 (1980).

[6] M. Dine, W. Fischler, and M. Srednicki, Phys. Lett. **B104**, 199 (1981).

[7] A.R. Zhitnitsky, Sov. J. Nucl. Phys. **31**, 260 (1980).

[8] N. Du et al. (ADMX Collaboration), Phys. Rev. Lett. **120**, 151301 (2018).

[9] Andrew K. Yi et al. (ADMX Collaboration), Phys. Rev. Lett. **130**, 071002 (2023).

[10] N. Du et al., Phys. Rev. Lett. **120** 151301 (2018).

[11] T. Braine et al., Phys. Rev. Lett. **124** 101303 (2020).

[12] C. Bartram et al. (ADMX Collaboration), Phys. Rev. Lett. **127**, 261803 (2021).

between 1.1 and 1.3 GHz ending in December 2024[1]. Over the last year the ADMX group has made improvements to the experiment including incorporating a new flux bias JPA allowing us to reach DFSZ axions in the 1 to 1.4 GHz range. This fall we prepared for our next data taking run which we will begin in late December.

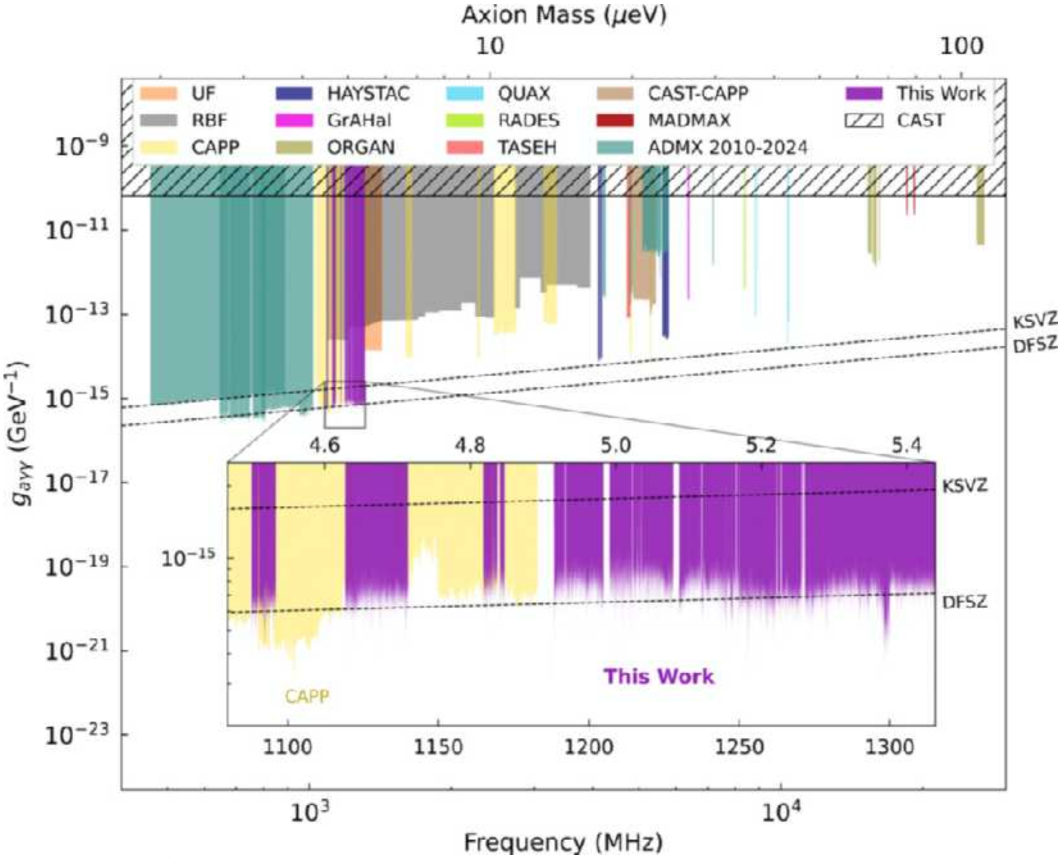

Figure 4.1-1. Axion-photon coupling constant sensitivity plot. The vertical axis is the axion-photon coupling constant, $g_{a\gamma\gamma}$, with smaller values indicating higher sensitivity. The two horizontal axes are axion mass and frequency respectively. Exclusion limits previously set by ADMX are shown in teal. Sensitivity for ADMX's most recent data taking run is shown in purple.

4.2 Making, Training, and Repairing a superconducting dipole magnet for ADMX Orpheus

G. Rybka and J. Sinnis

Orpheus is a pathfinder axion haloscope project that uses a novel cavity design to search for axions in mass ranges generally above the values that traditional tunable haloscopes are sensitive to. Orpheus searches for axions with the 18th order TEM mode of a Fabry-Perot interferometer. The frequency of this mode is tuned from about 15.8-16.8 GHz, which

---

[1] N. Du et al. (ADMX Collaboration), Phys. Rev. Lett. **135** 191001 (2025).

corresponds to axion masses of roughly 65-70 μeV. Generally, the use of the higher order modes of a resonator in an axion search is not feasible because the axion signal power scales with the effective volume V_{eff}, which follows this proportionality relation:

$$V_{eff} \propto \int \vec{E} \cdot \vec{B} d^3x \tag{1}$$

where $\vec{E}$ is the electric field of the mode used to search for the axion and $\vec{B}$ is an externally applied magnetic field. The direction of $\vec{B}$ is fixed, so using a higher order mode for $\vec{E}$ introduces oscillations into the integral which cancel out. To avoid this cancelation of the effective volume, dielectric plates are placed along the length of the resonator with even spacing, so that select oscillations of $\vec{E}$ have diminished amplitudes. This reduces the cancelation caused by the oscillations in the integral.

In the last year, we completed the helium recovery system for Orpheus, which will allow us to take data for longer periods of time within our limited budget, and we completed the two-coil superconducting dipole magnet. The original design of the magnet was intended to operate with 100 A current, but this is a rough and optimistic estimate. The maximum current allowed for a magnet is limited by defects in the winding, potting, and wire metallurgy. The magnet was wound and potted at CENPA by hand, so it is difficult to quantify how any unnoticed errors would determine the maximum current.

The magnet was completed last year and was trained for the first time this spring. It is a NbTi magnet that requires liquid helium to become superconducting. The process of training a superconducting magnet is inherently risky and involves quenching multiple times to reach the maximum current. We experienced two quenches during our training, one at 40 A and a second at 50 A. The second quench damaged one of the magnet coils. This forced us to stop training to diagnose the damage.

We found that the damage to the magnet was severe but fixable. The wires on one side of the magnet melted and were shorted to its stainless steel frame. The damage all occurred on the edge of the magnet, so the melted wires are all accessible. We have joined the broken wires with a low-resistance (but not superconducting) joint by shorting them all together with a large piece of OFHC copper which has its other end immersed in the helium bath. This copper piece is clamped against the broken exposed wires. Indium is used to improve surface contact between the broken wires and the copper joint. This design is intended to reduce electrical resistivity and function as a heat sink to prevent the Ohmic heating in the low-resistance joint from causing another destructive quench.

The ultimate sensitivity of Orpheus will be determined by the current that we are able to achieve in the repaired magnet. A rough estimate of the mean exclusion value as a function of the magnetic field strength B is

$$g_{a\gamma\gamma} = \left(\frac{1.0\text{T}}{B}\right) \times 4.0 \times 10^{-12} GeV^{-1}. \tag{2}$$

The magnetic field strength in Tesla is roughly equal to one-tenth of the current in Amperes. Therefore, in the most optimistic scenario, in which the repaired magnet achieves the same max current as before the destructive quench, we would expect a sensitivity of about $8.0 \times 10^{-12}\mathrm{GeV}^{-1}$. A more realistic expectation is that the magnet will reach a current of at least 10A. In that case, we would reach a sensitivity of roughly

$$g_{a\gamma\gamma} \sim 4.0 \times 10^{-11}\mathrm{GeV}^{-1}. \tag{3}$$

In this case, Orpheus would still set a new limit on $g_{a\gamma\gamma}$ in its frequency range. We are currently preparing for our first axion data taking run, and we hope to take this data in the spring of 2026.

Figure 4.2-1. The Orpheus insert and magnet being inserted into its helium dewar in preparation for magnet training.

DAMIC

4.3 World-leading sub-GeV dark matter search results from DAMIC-M prototype modules

K. Aggarwal, A.E. Chavarria, H. Lin, M. Traina, and outside collaborators*

As previously reported[1], the Low Background Chamber (LBC)—a test stand operated at the Modane Underground Laboratory (LSM) for the development of DAMIC-M—was upgraded with two prototype CCD modules fabricated at UW. The detector achieved the lowest background from ionizing radiation in a CCD detector to date, and an improvement by a factor of 50 in the single-electron rate, to a value of $<0.5\,e^{-}\,\mathrm{mm}^{-2}\,\mathrm{day}^{-1}$. This enabled an improved search for interactions of sub-GeV dark matter particles with electrons, which was performed in a 1.3 kg-day exposure acquired between October 2024 and January 2025. The analysis was based on the search for clusters of contiguous pixels in the CCD images with a total ionization signal between 2 and $4\,e^{-}$, inclusive. The observed number of clusters is consistent with background, which results in the exclusion limits shown in Fig. 4.3-1. These results[2] demonstrate sensitivity to dark matter-electron interaction cross sections almost two orders of magnitude lower than before. The results exclude, for the first time, predictions by hidden-sector dark matter models over a wide range of sub-GeV dark matter masses.

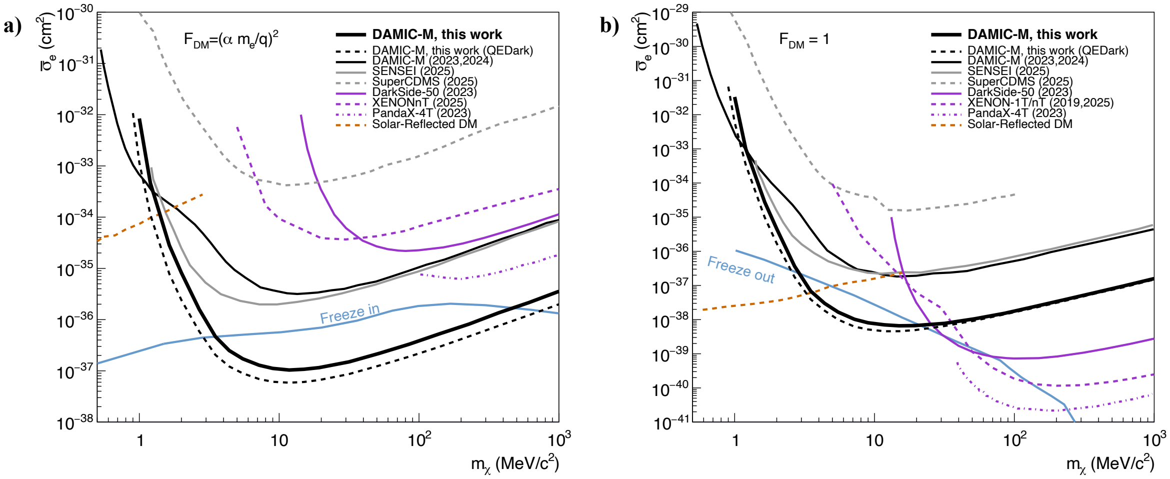

Figure 4.3-1. Exclusion limits on dark matter-electron interaction cross section $\bar{\sigma}_e$ versus dark matter mass m_χ. Parameter space above the lines is excluded by the different experiments in the legend. The latest results from DAMIC-M are given by the thick black lines. Left (right) panel presents results for the case of an ultra-light (heavy) hidden-photon mediator. The blue lines labeled Freeze In/Out are baseline predictions of models where a single hidden-sector particle constitutes all the dark matter in the universe.

---

*DAMIC-M Collaboration.

[1] CENPA Annual Report, University of Washington (2024) p. 125.

[2] DAMIC-M Collaboration, Phys. Rev. Lett. **135**, 071002 (2025).

4.4 DAMIC-M detector construction

K. Aggarwal, A.E. Chavarria, M. Conde, M. Huehn, H. Lin, R. Roehnelt, M. Traina, and outside collaborators*

The fabrication and testing of 28 CCD modules for the DAMIC-M detector was completed at UW in late 2024[1]. The modules were shipped by truck and boat to LSM in a 16-ton steel shield to prevent cosmogenic activation and were delivered in May 2025. The modules were promptly brought to the underground lab, where a new clean room was setup with radon-free storage, humidification, a CCD handling station and a CCD test system for quality control. Final, exhaustive testing and characterization confirmed the excellent performance of 26 modules, which now await installation in the detector.

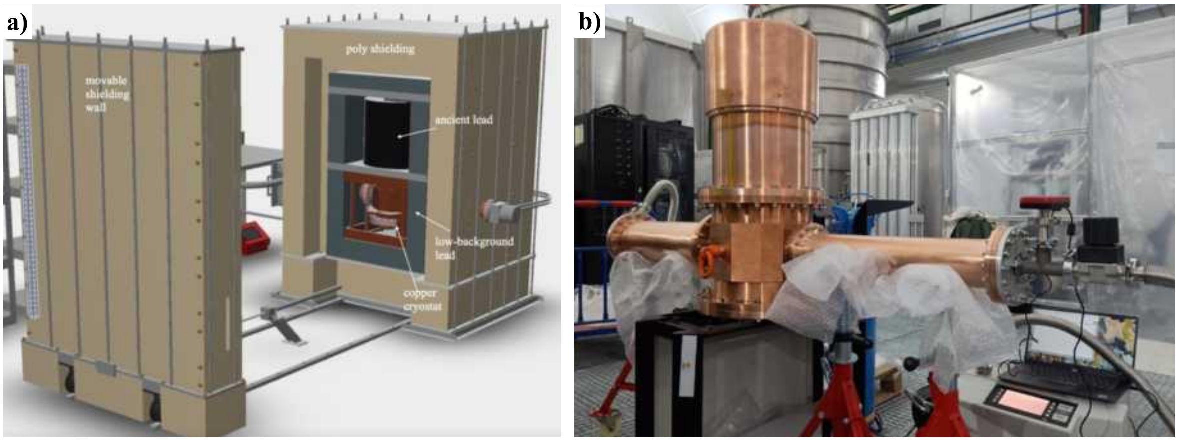

Figure 4.4-1. **a)** CAD view of the DAMIC-M detector inside its movable shield in the clean room at LSM. **b)** Vacuum test of the DAMIC-M cryostat at LSC. The top can is made from the can of a MAJORANA Demonstrator detector module electron-beam welded to a high-purity copper flange procured from Luvata.

The overall design of the DAMIC-M detector remains as previously reported[2], with the CCD array housed inside a radiopure copper cryostat and protected from environmental radiation by multiple layers of shielding. Over the past year, the detailed design of the external shield, its support structure and layout inside the clean room were completed, as shown in Fig. 4.4-1a. All components of the shield were ordered and most were already delivered, including the inner 5 cm layer of ancient lead cast by MTH Metall-Technik in Germany, confirmed to have an activity of 30 mBq/kg of ^{210}Pb. The outer shield layers, including lead bricks from the MAJORANA Demonstrator, are stored in LPSC Grenoble, to be cleaned before installation in LSM. The copper components for the array support structure, infrared shield and cryostat, including critical pieces machined at Sanford Underground Research Facility (SURF) from ultra-high radiopurity copper grown underground, are currently stored at Canfranc Underground Laboratory (LSC) in Spain. Fig. 4.4-1b shows the cryostat assembled

---

*DAMIC-M Collaboration.

[1] H. Lin, M. Traina *et al.*, arXiv:2509.06943 (2025).
[2] CENPA Annual Report, University of Washington (2024) p. 118.

for the vacuum test at LSC. A final cleaning campaign will take place in March 2026 before the copper is shipped to LSM. The 16-ton steel shield is also used for the transportation of the copper. Other critical detector tasks have been completed, including the production of the readout electronics in France, the vacuum feedthrough by Johns Hopkins University, and the commissioning of the vacuum pump, cryocooler and slow control at The University of Chicago.

The CCD array assembly procedure was developed and validated at CENPA with 13 mock CCD modules. The procedure for the installation of the CCD array in the cryostat, including radon mitigation precautions and the "hot wire"[1] for calibration are under development. We foresee all detector components to be on-site by mid-2026, and for installation of the CCD array this summer, with detector commissioning to follow. Once ready, DAMIC-M will acquire kg-year exposures with a radiogenic background $<1\,(\mathrm{keV\ kg\ day})^{-1}$ and an energy threshold of a few electron volts. This will enable unparalleled sensitivity to the interactions of sub-GeV dark-matter particles with electrons, improving on our latest results (Sec. 4.3) by three orders of magnitude.

4.5 Modeling the DAMIC-M CCD detector response

K. Aggarwal and A.E. Chavarria

4.5.1 Modeling Nuclear Recoil Quenching in DAMIC-M Simulations

Accurate modeling of the detector response to nuclear recoils is critical for interpreting low-energy signals in DAMIC-M. In DAMIC-M CCDs, nuclear recoils arising from elastic scattering of silicon nuclei produce reduced ionization compared to electron recoils from electromagnetic interactions, as a significant fraction of the deposited energy is lost to lattice excitations rather than electron–hole pair production. This suppression of ionization, commonly referred to as nuclear recoil quenching, is particularly important at nuclear recoil energies below a few keV. During the reporting period, a comprehensive nuclear recoil quenching framework was implemented within the WADERS detector-response simulation to provide a physically consistent description of nuclear recoil ionization, including suppressed charge yield and energy-dependent statistical fluctuations.

The quenching model is based on a composite parameterization of the nuclear recoil ionization yield in silicon, constructed to ensure continuity across the full energy range relevant to DAMIC-M. At low recoil energies, the yield follows extrapolated recent collaboration experimental measurements[2], transitions smoothly through the keV-scale data regime[3], and asymptotically approaches a Lindhard-based description at higher energies. This approach

---

[1] CENPA Annual Report, University of Washington (2024) p. 119.

[2] Cuevas-Zepeda, J. J., "Skipping to the Finish Line: Sub-Electron Resolution Measurements with DAMIC-M CCDs for Dark Matter Detection," Ph.D. thesis, University of Chicago (2024), doi:10.6082/uchicago.13992.

[3] Sarkis, Y., Aguilar-Arevalo, A., and D'Olivo, J. C., "Ionization efficiency for nuclear recoils in silicon from about 50 eV to 3 MeV," Phys. Rev. A **107**, 062811 (2023).

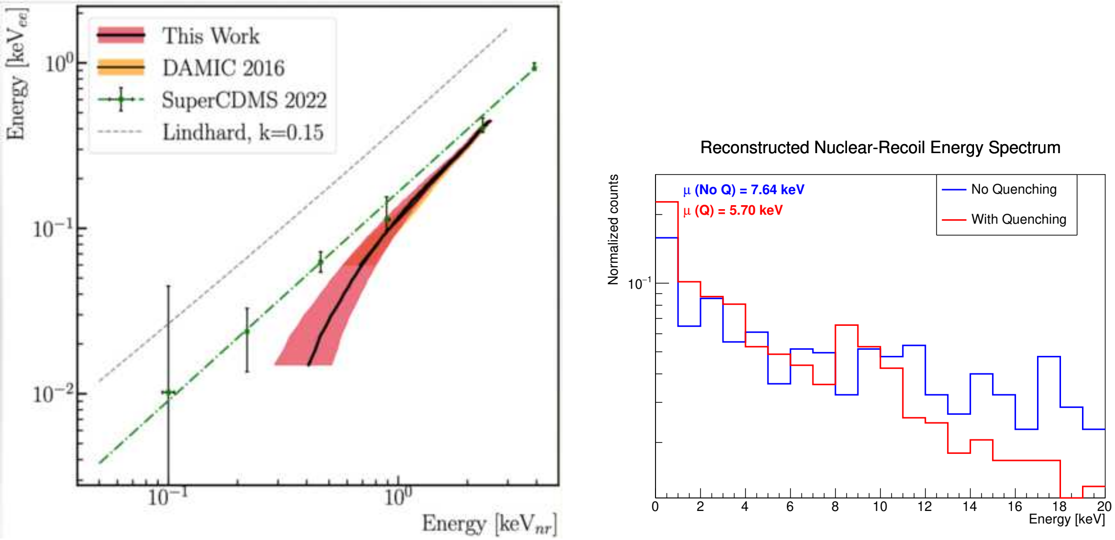

Figure 4.5.1-1. Nuclear recoil quenching in DAMIC-M simulations. *Left:* Ionization yield for nuclear recoils in silicon, illustrating the suppression of charge production at low recoil energies, particularly below the keV scale where quenching effects are most significant. *Right:* Reconstructed nuclear-recoil energy spectra from simulated muon-induced events with and without nuclear recoil quenching. When quenching is enabled, the nuclear-recoil spectrum shifts toward lower reconstructed energies, consistent with the reduced ionization yield expected for nuclear recoils.

avoids artificial discontinuities in reconstructed energy spectra while preserving consistency with existing measurements.

In addition to the mean ionization yield, energy-dependent fluctuations in the number of electron–hole pairs produced by nuclear recoils were incorporated through a nuclear-recoil specific Fano factor. Two statistically equivalent implementations were developed: a Poisson-scaled model that preserves the discrete nature of charge production at low energies, and a Gaussian approximation applicable at higher energies.

The quenching framework was integrated consistently with existing detector-response components, including charge diffusion, pixelization, and readout noise. Quenching is applied exclusively to nuclear recoil energy depositions, ensuring that mixed electron- and nuclear-recoil events are treated without double counting.Configuration parameters controlling the yield model, fluctuation treatment, and activation of quenching are managed through the configuration file.

Validation studies were performed using simulated muon-induced backgrounds and dedicated nuclear recoil samples. As expected, the application of quenching shifts reconstructed nuclear recoil energy spectra toward lower energies while leaving electron recoil spectra unaffected. An illustrative comparison demonstrating the impact of nuclear recoil quenching in DAMIC-M simulations is shown in Fig. 4.5.1-1.

4.5.2 Charge Diffusion and Coulomb Repulsion Modeling in DAMIC-M

A correct understanding of lateral charge transport in DAMIC-M CCDs is essential for reconstructing low-energy events and for correctly describing the depth-dependent cluster morphology used in background rejection and signal characterization. During the reporting period, a physically motivated charge diffusion model, including Coulomb repulsion, was implemented and validated within the WADERS detector-response framework.

The model is based on the charge transport formalism developed by Sofo Haro, Fernández Moroni, and Tiffenberg,[1], which describes the combined effects of thermal diffusion, drift in the bulk electric field, and mutual Coulomb repulsion among charge carriers in fully depleted, high-resistivity silicon CCDs. In this framework, the lateral extent of a charge cloud at collection is governed not only by diffusion during the drift time but also by electrostatic repulsion, which becomes increasingly important for dense charge packets and for interactions occurring deep in the detector bulk. The present implementation was validated by reproducing the depth and charge dependence reported in Ref. [1] and was subsequently extended to the thicker DAMIC-M CCD geometry and operating conditions.

The charge transport is parameterized using the carrier mobility μ and the effective dopant charge density ρ, which together determine the drift time and diffusion length as a function of interaction depth. For DAMIC-M operating conditions, a hole mobility normalized to $\mu_0 = 0.135\ \mathrm{m^2\,V^{-1}\,s^{-1}}$ and an effective space-charge density of $\rho_0 = 4 \times 10^{17} e\,\mathrm{m^{-3}}$ were adopted as reference values, with dimensionless scaling factors applied to account for device-to-device variations. The substrate bias voltage, detector thickness, and operating temperature are treated as configurable parameters, allowing the model to be tuned to specific DAMIC-M running conditions.

To efficiently capture Coulomb repulsion effects while maintaining computational tractability, the charge cloud is represented using a set of concentric Gaussian shells that evolve during transport. This effective-shell approach reproduces the mean and variance of the full many-body solution while reducing the numerical scaling from quadratic to linear in the number of charge elements per time step, enabling efficient large-scale simulation without loss of physical fidelity.

The diffusion-plus-repulsion model was validated against multiple calibration and science datasets. Simulated lateral charge spreads reproduce the depth and energy dependence observed in ^{55}Fe calibration data taken at the University of Washington, as well as measurements during LBC science runs at LSM. In particular, the model accurately reproduces the observed increase in cluster width with deposited energy at fixed depth, a feature that is not captured by diffusion-only[2] descriptions. A representative comparison between the previous diffusion-only model and the updated diffusion plus repulsion model is shown in Fig. 4.5.2-2.

---

[1] Sofo Haro, A., Fernández Moroni, G., and Tiffenberg, J., "Studies on small charge packet transport in high-resistivity fully depleted CCDs," IEEE Trans. Nucl. Sci. **67**, 1940 (2020).

[2] Aguilar-Arevalo, A., Amidei, D., Bertou, X., Butner, M., Cancelo, G., Castañeda Vázquez, A., *et al.* (DAMIC Collaboration), "Search for low-mass WIMPs in a 0.6 kg day exposure of the DAMIC experiment at SNOLAB," *Phys. Rev. D* **94**, 082006 (2016).

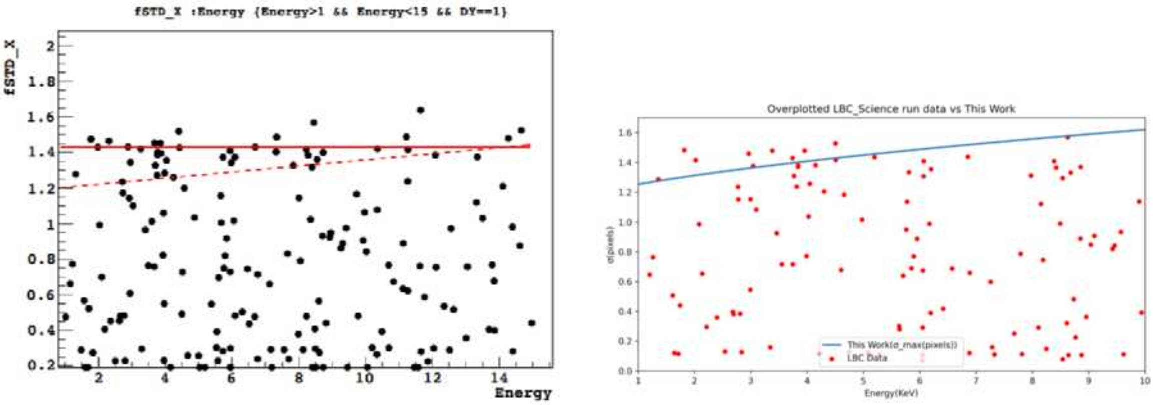

Figure 4.5.2-2. Comparison of lateral charge transport models used in DAMIC-M simulations. models. *Left: Previous diffusion-only implementation, which predicts an approximately linear dependence of the effective lateral charge spread on deposited energy. Right: Updated diffusion plus Coulomb repulsion model implemented in this work, showing a non-linear energy dependence of the maximum lateral charge spread at the collection plane. The inclusion of Coulomb repulsion improves agreement with calibration data and LBC science run measurements.*

5 Gravity and fundamental symmetries

5.1 Tests of the Equivalence Principle

S. K. Apple, J. H. Gundlach, I. A. Paulson, and M. P. Ross

We completed a science campaign with our rotating torsion balance to search for novel long range interactions between dark matter and normal matter. The possibility of sub-gravitation strength dark matter interactions is an under explored aspect of the nature of dark matter.

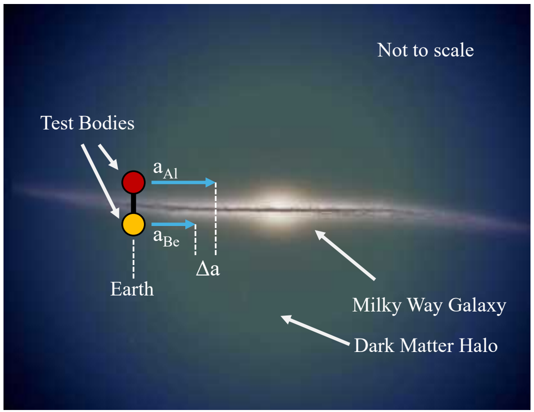

Figure 5.1-1. A schematic of the experimental setup in the galactic frame. The experiment searches for differential accelerations of two different materials towards the dark matter in the Milky Way galaxy. Adapted from ESA/Gaia/DPAC, Stefan Payne-Wardenaar.

Our torsion balance consists of a pendulum suspended from a 22-µm thick fused quartz fiber of 1-m long. The pendulum has eight test bodies, half made of beryllium and the others made of aluminum, in a dipole orientation. The entire apparatus is suspended from a rotating turntable that modulates the signal of interest at the specified revolution speed. We searched for a signal towards dark matter over density at the center of the Milky Way, shown in Figure Fig. 5.1-1, which is additionally modulated by the rotation of the Earth.

We recorded data from July 7, 2024 to July 7, 2025 with significant downtime due to the demolition of the dorms above the lab. We set strict upper limits on the strength of a novel long-range dark matter interaction relative to gravity of:

$$\eta_{DM,Be-Al} \leq 2.6 \times 10^{-5} \quad \text{(95\%-confidence).} \tag{1}$$

This is a factor of four improvement over the previous limits and shows that dark matter appears to interact only gravitationally. A publication detailing this has been submitted to Science (preprint: M. P. Ross et al., arXiv 2509.10701 (2025).).

Additionally, we search the same data set for Equivalence Principle violations towards the Sun to set the limits (publication in progress):

$$\eta_{Be-Al,\odot} \leq 2.1 \times 10^{-13} \text{ (95\%-confidence)} \tag{2}$$

Further, with this data, we set limits on the coupling strength to B-L for ultra-light vector dark matter candidates with masses between 1.3 $\times 10^{-22}$ and 1.9 $\times 10^{-18}$ eV, reaching a peak sensitivity of $g_{B-L} \leq 9 \times 10^{-26}$. (preprint: M. P. Ross et al., arXiv 2510.21764 (2025).).

5.2 Inverse Square Law Test

S. K. Apple, J. H. Gundlach, I. A. Paulson, and M. P. Ross

We have begun the development of a test of gravity in short ranges that attempts to measure the force of gravity at separations of 20 μm, while the previous experiment achieved separations of 52 μm. The apparatus is still in its prototyping stages as we are testing a new pendulum-attractor mass geometry that will allow us to achieve shorter distance measurements. The primary concern that our prototyping aims to fix is the noise on the pendulum due to patch contact potentials on the surfaces of the pendulum and the attractor mass. We mitigate these by using a control loop with electrostatic plates to keep the pendulum as still as possible. We are currently characterizing the patch potential noise and hope to start on a full-fledged experiment soon.

Figure 5.2-1. Photograph of the pendulum sitting in front of the attractor mass with two golden electrostatic plates to control the pendulum twist and a third silver electrode injecting a signal to help measure SNR.

By measuring gravity at small separations, we hope to explore an unprobed parameter space that could reveal the presence of a Yukawa potential only visible at small distances. This Yukawa potential is commonly present in many string theories aiming to unify General Relativity and Quantum Mechanics.

5.3 Instrumentation for LIGO

S. K. Apple, J. H. Gundlach, I. A. Paulson, and M. P. Ross

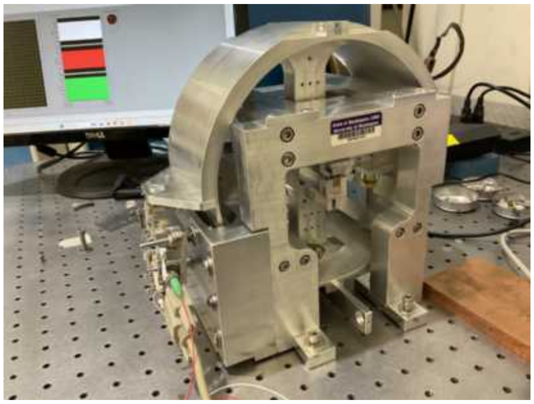

Figure 5.3-1. The Newtonian noise rotation sensor.

We have begun work on development of a rotation sensor to help mitigate Newtonian noise in LIGO and other future gravitational wave detectors. Newtonian noise is the noise caused by small fluctuations in the gravitational fields near the LIGO mirrors. While there are many sources of Newtonian noise, it is expected that the main source in surface based detectors like LIGO will be passing Rayleigh waves.

By measuring the ground tilt from a Rayleigh wave, given by:

$$\tau_x = \partial\xi_z = ik\cos(\phi)\xi_z e^{i(\vec{k}\cdot\vec{\varrho}_0-\omega t)} \tag{1}$$

we can estimate the Newtonian noise created by that wave using

$$\delta a_x(\vec{\varrho}_0, t) = 2\pi\gamma G\rho_0\xi_z e^{-hk} e^{i(\vec{k}\cdot\vec{\varrho}_0-\omega t)}\cos(\phi) \tag{2}$$

where γ, G, and ρ_0 are all known constants.

Using techniques developed to construct other rotation sensors for LIGO [1], we have constructed a prototype Newtonian noise rotation sensor (NRS) pictured in figure Fig. 5.3-1. The NRS has two main components, a cylindrical aluminum proof mass, and a support system which the proof mass is suspended from using 20-30 μm flexures. On either side of the NRS there is an interferometer which measures the distance from the support, which rotates with the ground, and the proof mass, which remains stationary.

Unlike previous rotation sensors developed by this group, the NRS has a higher resonant frequency (∼0.1 Hz) as Rayleigh waves tend to have a frequency of $\sim 10-30$ Hz. Additionally, the NRS is small enough to fit under the LIGO chambers which hold the mirrors, and have condensed optics.

Using this sensor, we can measure ground tilts on a scale of 10 picoradians, which should allow the detection of Rayleigh waves that cause Newtonian noise. There are currently plans to install the NRS at the Hanford LIGO site from May to July for testing.

---

[1] M. P. Ross et al., Review of Scientific Instruments 94, 094503 (2023).

6 Accelerator and Ion Sources

6.1 Van de Graaf Accelerator Overview

B. Dodson, M. Newman, and C. Wiseman

The Tandem Van de Graaf accelerator (VdG) remains a cornerstone of CENPA, continuing support of CRES beta decay research (Sec. 3.1), PIONEER detector characterization (Sec. 2.10), and source production for LEGEND (Sec. 1.6). The VdG Pelletron charging chains ran a total of 667.6 hrs, and the Direct Extract Ion Source (DEIS) ran 753.5 hrs. The accelerator provided 32 days of science operations, usage by project described below:

- Accelerator, Ion Sources total operations time:
  - Chains: 667.6 hrs
  - DEIS: 753.5 hrs
  - Sputter Ion Source (SpIS) : 0 hrs
- Users : 1 day = 8 hours of experiment preparation and experiment use, 32 days total
  - CRES beta decay science runs: 16 days total
    - ^{19}Ne CRES: 14 days
      - Run Dates
        1. 10/3/24-10/4/24 (day runs)
        2. 4/13/25
        3. 5/18/25
        4. 7/24/25
        5. 7/28/25 - 7/30/25 (24 hour run)
    - ^{6}He CRES: 2 days
      - Run Dates
        1. 8/12/25
        2. 8/15/25
    - ^{79}Kr Production: 3 days
      - Run Dates
        1. 11/20/24
        2. 3/14/25
        3. 7/23/25
  - PIONEER RBS LYSO detector characterization: 11 days
    - Run Dates
      1. 8/20/25 - 8/22/25 (day runs)
      2. 8/25/25 - 8/29/25 (day runs)

        3. 9/3/25 - 9/5/2025 (day runs)

  - LEGEND: 2 days
    - ^{56}Co Production: 7/29/25 (performed during the ^{19}Ne run, same energy & beamline, no additional time used)
    - Gas cell window stress test: 12/3/2025
    - ^{41}Ar Production: 12/4/25

It is worth noting tandem stability at high voltage (above 8.0 MV) was not achieved for continuous running needed for He6-CRES experiment operations. 8.8 MV was achieved for short periods of time that allowed some preliminary ^{6}He production. These days allowed experimenters to gauge fluctuations in the target vacuum with an active beam, observe CRES tracks at low production rates, and adjust Li target stir rod configurations. The accelerator operations team was able to begin tuning optimization on the R45 target while monitoring vacuum and resulting gate valve status. The machine experienced excess sparking without obvious cause during all He6-CRES run attempts. The tank gas was consistently at 225 psi, 80/20 CO_2/N_2 gas mixture, and low dew point. Historical charge conditioning techniques (50 - 100 kV increases every 10 - 15 mins above 6.0 MV, ensuring 6 xrays/min or less on the Low Energy (LE) cup before charge increase) were applied without success.

This difficulty in charge stability motivated two tank openings this year. The first tank entry occurred after the accelerator failed to recover full voltage after a series of sparks. It was discovered upon opening, a resistor jumper had blown off among other damage that could've been the cause for stability issues at +8 MV. After repairs of obvious damage and procedural tank cleaning, another attempt to charge above 8 MV was made with similar outcome. Another opening focusing on resistor and column damage was initiated. During this opening, a great effort was made towards cleaning column glass, spark gaps between planes, and accessible metal surfaces (Figure 6.1-1). Thorough investigation found some column cracking/surface tracking between planes on the LE columns. It is unclear when this damage occurred but further degradation will be mitigated with half value resistance. More detail about both openings are presented in the maintenance section (Sec. 6.2). The tank will be closed in February of 2026, and another attempt to condition at high voltage will be made.

During these operational downtimes, great progress was also made toward other critical maintenance and upgrades needs. Successful replacement of intercooler coils on the Ingersoll-Rand tank gas compressor and improvements on various water cooling systems, most critically the analyzing magnet, was accomplished. New sensors digitized to SlowDash readout have improved diagnostics capabilities spanning multiple systems including the analyzing magnet water systems, tank gas parameters, and tank conditioning monitoring. Control systems upgrades have made significant progress, including the replacement of the control room LE parameter knob box and development of the vacuum test stand to explore PLC systems for eventual VAX and GORDO systems replacement. More details are presented in the controls upgrade section (Sec. 6.3). Diligence building the VdG Archive Google Drive with both historical documentation and current project outlines paired with regular ELOG usage has reinforced centralized structure for keeping institutional knowledge available to all staff.

Accelerator engineering staff attended SNEAP 2025 at Brookhaven National Laboratory (BNL) this year. This 5 day convention hosts a longstanding group of electrostatic accelerator personnel and is an unmatched environment for exchanging challenges and advice within the community. VdG charging challenges and other maintenance projects were presented to the group which yielded a list of verification checks and recommendations involving gas mixture and cleanliness, as well as the opportunity to provide excellent guidance to TUNL on their gas compressor intercooler replacement. Engineers received insight on vacuum repair conventions at BNL, especially regarding cryopump servicing which informed decisions about our equipment maintenance. Additionally, a special session on foil floating methods was offered by Arizona Carbon foils representative and BNL engineer, further developing floating techniques.

Figure 6.1-1. Polished terminal shell hardware (left), dusting between spark gaps (right)

6.2 Accelerator Maintenance

B. Dodson, M. Newman, and C. Wiseman

Substantial work was carried out this year on the accelerator supporting hardware water systems. There was particular focus on magnet cooling reliability due to the discovery of a partially shorted coil in the analyzing magnet. Efforts focused on mitigating the risk of increased shorting from heat due to insufficient water flow.

Initial maintenance included performing blow-downs on the power supply and magnet cooling lines. The power supply blow-down process was subsequently reworked to avoid bypassing portions of the cooling loop, improving overall effectiveness. To further address contamination, iron separators were added to each power supply return water line before the flow switch. These have proven to increase the lifetime of the flow switch and prevent large contaminants from entering other equipment. Magnets inside of 3D printed enclosures were used inside the iron separators to prevent clogging. The flow rates have held steady in the magnet power supplies and both flow switches have been integrated back into the interlock system.

Significant redesign work was also performed on the AN magnet flow switch manifold. Pipe diameters were increased, the number of tight turns was reduced, and the flow switches were removed and cleaned. These combined changes successfully restored water flow to previous levels on both magnet coils. The bottom coil of the analyzing magnet saw the highest increase from 0.4 GPM to 0.7 GPM.

To improve diagnostics and long-term monitoring, a dedicated AN magnet data acquisition system was developed. Temperature sensors and pressure sensors were installed on both the water supply and return lines. A microcontroller with Ethernet capability was used to transmit UDP data packets to a control room PC that records the data to a database file visible by SlowDash. This enables continuous remote monitoring of magnet conditions as well as an indication of the total water loop health. The next iteration of this device will also continuously read and send magnet coil voltage to calculate resistance using the current setpoint.

Additional infrastructure improvements were completed in collaboration with facilities plumbing. A new magnet manifold was created and set to be installed to replace existing steel piping, incorporating a large filter located closer to the magnet coils. The system was transitioned to PVC piping, with the long-term goal of limiting magnet cooling loops to plastic and copper materials to reduce corrosion and contamination.

Water system work was also carried out in Cave 1, including coordination with the He6-CRES team on cryocooler operations and transitions between city water and condenser water. A new filter was installed on the city water connection to prevent sediment from migrating to the cryo-compressors when switching from condenser water. Also, old cooling lines to and from the L30 turbo pump were removed and replaced, and new connections were added to support the LEGEND target cooling loop. The new L30 lines are a larger consistent diameter of new tubing and are able to deliver multiple times the flow rate needed to cool a turbo and a target. Critical steps were made in completing the intercooler coil replacement on the Ingersoll-Rand tank gas compressor. The original intercooler coils, removed in 2024 due to rust damage, were replaced with newly designed coils developed in-house. These coils were bent, x-ray inspected, and galvanized off-campus in 2024, with an in-house rig constructed to support the galvanizing process. In 2025, the coils were threaded, and prepared internally

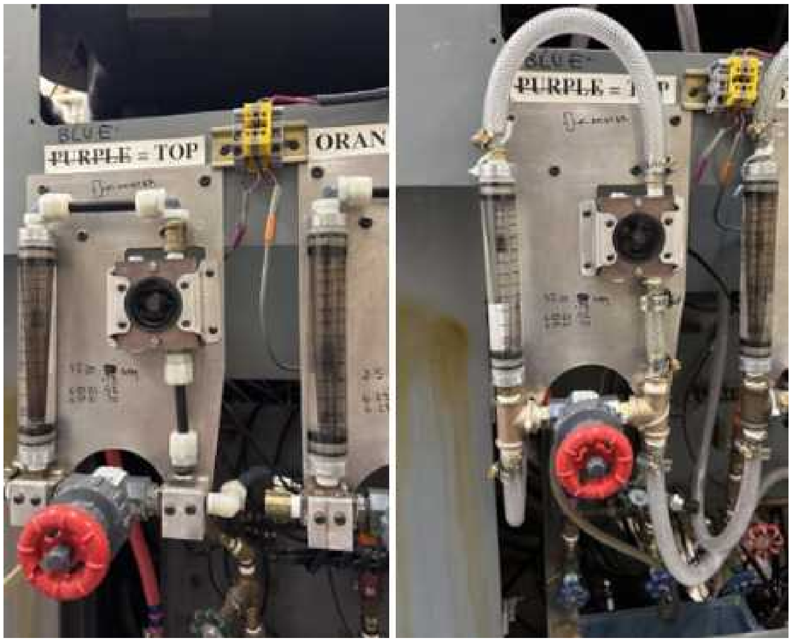

Figure 6.2-1. The old manifold section (left) versus the new layout (right).

by being acid-flushed and oil-filled to clean and protect the internal surfaces. The intercooler tub itself was recoated with marine epoxy in 2024 to extend its life and improve durability. Installation of the new coils was completed in 2025. Since then the compressor has been used many times to perform gas transfers and tank pump-outs. The design files, fabrication complines, and pre-installation procedures were shared with TUNL accelerator engineers.

This year the accelerator tank has been opened twice to perform scheduled maintenance, execute repairs, and document conditions prohibiting stable +8 MV runs. The first opening included repairing the damaged resistor limiting charge capacity after sparking, installing new stripper foils, replacing a fiber optic cable, replacing a shredded LE chain idler pulley, securing a loose terminal hardware screw, reparing a grounding wire on the corona point, testing high-voltage steerers, and performing the standard tank-closure checklist. We purchased a hand-held CO_2 sniffer unit designed for refrigeration systems that is able to detect small tank gas leaks when the tank is pressurized with a nitrogen and CO_2 gas mixture. The sniffer was able to detect the port-holes contributing to the gas loss in the tank. The sniffer is now incorporated with the tank closure checklist to determine if the port-hole gaskets have been properly sealed. The second opening occurred after numerous long attempts to condition to 8.8 MV for He6-CRES. Flashing visible in the LE column port and instability in the terminal voltage was a clear indication that the tank needed to be entered. The team pumped out and started assessing the hardware's condition. Many loose resistors were found and tightened or replaced. The flashing was a resistor that had loosened to the point where it was sparking from the resistor to the tube. Each ring was cleaned and moved so the column and all the resistors could be inspected, documented, and cleaned. Especially notable damage in the form of column cracks across planes were addressed with half value resistance to decrease potential difference across cracks to mitigate voltage breakdown and further damage. Due

Figure 6.2-2. Newly fabricated Ingersoll-Rand coil acid flush before installation.

to the spiral beam inclination, resistors located a half-cycle away from the damaged planes are also reduced. This adjustment was made in the following locations: 55-56, 61-62, 72-73, 81-82 on beamtube 1, 83-84, 90-91 on beamtube 2 and 24-25, 33-34 on beamtube 4.

6.3 Accelerator Controls Upgrade

B. Dodson, M. Newman, and C. Wiseman

As described in last year's Annual Report, the VDG team has renewed efforts to update the accelerator's control system, with the goal of improving beam uptime, analytics, and reliability. The work this year has been accomplished during a busy schedule of operations and maintenance. New Ethernet switches and shielded cable runs connecting the control room and the tunnel have formed a backbone to which many devices have been added. This includes webcams, environmental monitoring sensors (water temperature, pressure, etc), and passive readout of vacuum gauges, the GVM voltage reading, and more. Figure Fig. 6.3-1.

The overall goal of the project is to replace the "ancient" VAX, PDP-11, and IBM PC computers which represent single points of failure for the system, with additional passive diagnostics and readbacks where none previously existed. The replacement of the Vacuum Satellite LSI-11+ machine remains the primary near term goal. Towards that end, a vacuum system test stand continues development in the High Bay. Two turbopumps, mechanical pumps, and gate valves, will be operated analogously to the HE Object and HE Image pumping stations on the accelerator, simulating power failures and various operating conditions

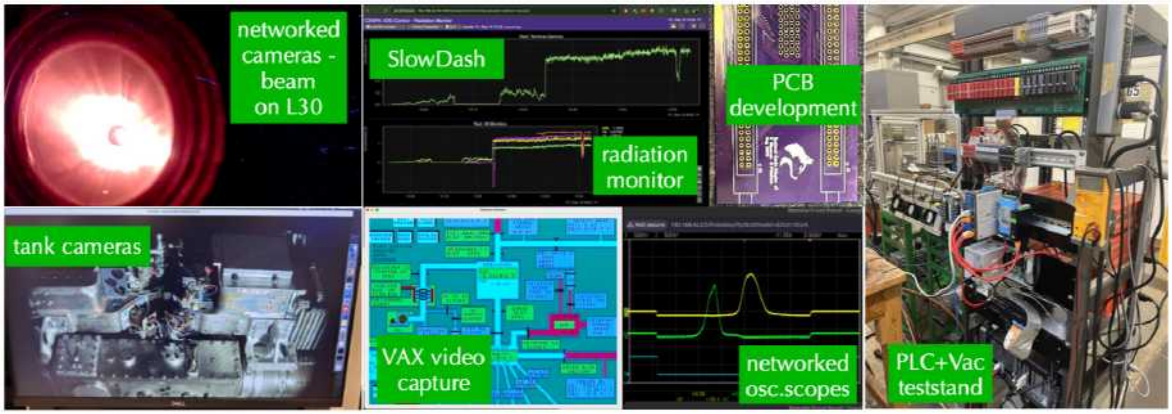

Figure 6.3-1. Overview of controls projects. Left: Networked cameras provided new diagnostic imaging options.

before changing over to the new system. Two PLC systems are being evaluated, as are the options to interface to the "Gordo"-like opto-isolated relay racks.

We are exploring the use of SlowDash, a slow controls data acquisition framework under development by S. Enomoto. A flexible web framework and Python backend allows integration of many devices and sensors. This work has already yielded renewed insight into the performance of the VDG machine. For instance, a set of paper "chart recorders" was historically used to track terminal voltage (GVM) and HE/LE ion gauge activity during conditioning. After installing a set of analog sensors, this essential functionality was restored. During conditioning of the machine to high voltage, the readout of the LE ion gauge in particular proved to be very sensitive to conditioning on the accelerator beam tubes, allowing the team to exercise additional care when raising the terminal voltage. A similar system attached to the existing radiation monitor readout proved useful. A new set of network connected webcams was created and installed on the tank viewports to watch for sparking events, and was immediately used to isolate a sparking column resistor on the LE side during 8+ MV beam operations.

7 Education

7.1 Use of CENPA facilities in education and coursework at UW

<u>B. W. Dodson</u>, A. Garcia, and C. Wiseman

CENPA continues to present immersive opportunities for students to participate in a wide range of applied methods in modern physics experimentation. Graduate and undergraduate student contributions are a critical part of daily lab operations, offering interdisciplinary skill building across hardware and software applications. All students receive specialized safety training mandated through UW's Environmental Health & Safety (EH&S) training programs[1]. Access to the electronics shop, student machine shop, and a variety of other laboratory spaces give students advantageous skills relevant in engineering and technical industries. CENPA also continues to offer a graduate-level laboratory course on nuclear physics with special interest in accelerators and radiation detection, PHYS 575/576(Sec. 7.2).

CENPA offers educational tours to external organizations and schools at a range of levels to inspire interest in experimental physics. CENPA organized group tours for the UW Physics Program for Advanced Training in Hands-On Science (PATHS))[2], a program for community college students and prospective UW undergraduates to gain hands-on experience in physics labs at UW. Other visitors included two high-school groups and students from Seattle University.

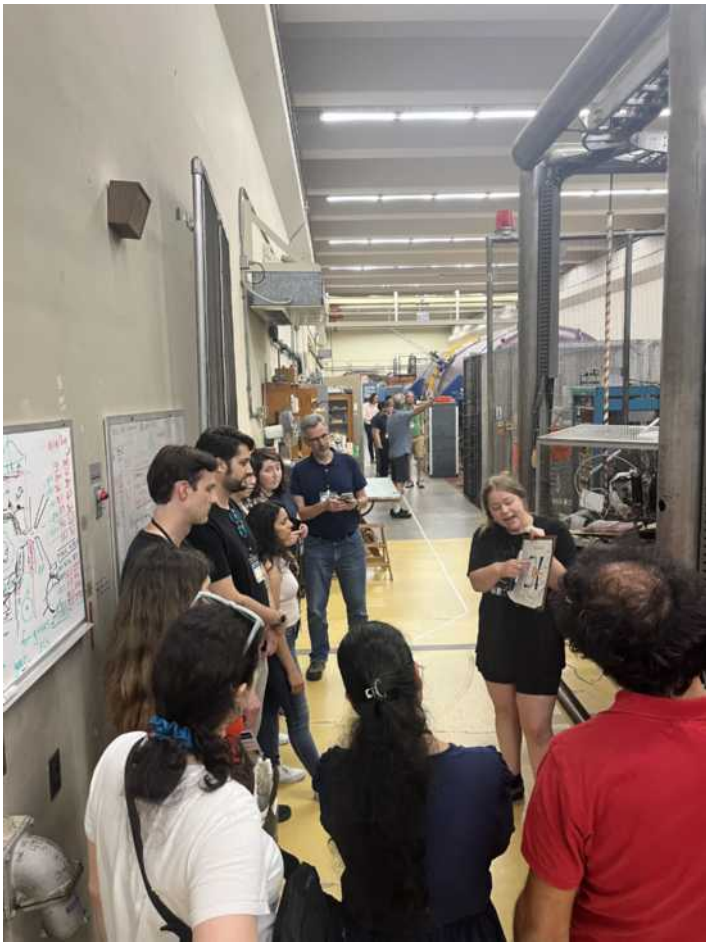

Figure 7.1-1. Accelerator Engineer Brittney Dodson shows the Direct Extraction Ion Source

---

[1] Environmental Health & Safety, https://ehs.washington.edu/training

[2] , https://sites.uw.edu/uwpaths/

7.2 Accelerator-based lab class in nuclear physics

B. W. Dodson, A. García, and C. Wiseman

The Accelerator team continues to support UW Physics evening master's program PHYS 575: Selected Topics In Applications Of Physics - Nuclear Physics, radiation, and detection. The class offers accessible and uniquely immersive education in nuclear theory, research, and application in the local workforce. Emphasis on relevant applied methods of nuclear physics experimentation supports professional development for students already working or interested in growing aerospace and technical industries.

The class is typically offered biennially, this year's class occurred Autumn quarter 2025. Laboratory sessions offer hands-on experience for each student with experimental equipment including the Van De Graaff accelerator, the associated Direct Extraction Ion Source (DEIS), an array of particle detectors, nuclear instrumentation electronics, and data acquisition systems. Lectures are 90 minutes once a week followed by 90-minute small group lab sessions. Lectures and lab instruction videos were offered online for students additionally. Supported by professional engineers and research scientists, the following list of subjects are covered:

Supported by professional engineers, the following list of subjects are covered:

1. The atomic nucleus, orders of magnitude, basic nuclear physics, nuclear energy and decays
2. Emission of photons and interactions, attenuation of gamma radiation
3. Laboratory practice 1: Deciphering a mystery γ spectrum measured using a Ge solid-state detector
4. Ranges of ions and electrons. The weak interaction. Radioactivity, radiation damage, and health risks (α, β, γ, and neutron activity)
5. Counting statistics
6. Lab Practice: 3 duoplasmatron ion source function, mass analyzing with magnetic fields
7. Ion beam accelerators
8. Lab Practice 4: Low Energy ion beam optics
9. Rutherford back-scattering (RBS) spectra
10. Fission and fusion. Nuclear reactors, nuclear astrophysics (nucleosynthesis in stars)
11. Lab Practice 5: RBS
12. Basics of the Shell Model. Approximate predictions using simple models, using quantum numbers, predicting decay rates.
13. Double beta decay in search for Lepton number violation

Lab 3 allowed students to use the DEIS and mass analyzing magnet to find different isotopes within deuterium gas by measuring peak current in the deck image faraday cup. Students then derived an equation that relates the needed magnetic field for selecting an ion of particular atomic mass at constant potential. In lab 4, students were able to tune a similar deuterium beam by adjusting Low Energy (LE) steering and optics parameters while observing the effects on beam transmission with faraday cups and oscilloscopes. This process involved allowing each student to find an optimized beam from the DEIS to the LE cup.

8 Facilities

CENPA Engineering

8.1 CENPA Engineering

B. Dodson, M. Huehn, M. Kallander, N. Miedema, M. Newman, and R. Roehnelt

Organization Purpose

CENPA Engineering is a support organization for every experiment. This organization is able to design, analyze, procure, and fabricate experimental hardware. The goal is to make anything hardware related easier such that principal investigators and their students can focus on the physics.

Engineering Software

For design work the primary CAD package used is Autodesk Inventor, in part as it has an excellent tie-in with Autodesk Vault for revision control and a central PDM file repository on the CENPA computer cluster. Specifically for Legend-1000, Solidworks CAD is used. There is a Legend specific Solidworks PDM, also on the CENPA computer cluster. For non-PDM tied tasks, Autodesk Fusion is sometimes used. Unlike other CAD used at CENPA, Fusion has the capability to output g-code to directly run CNC machines. This is an important capability to machine precise parts repeatably.

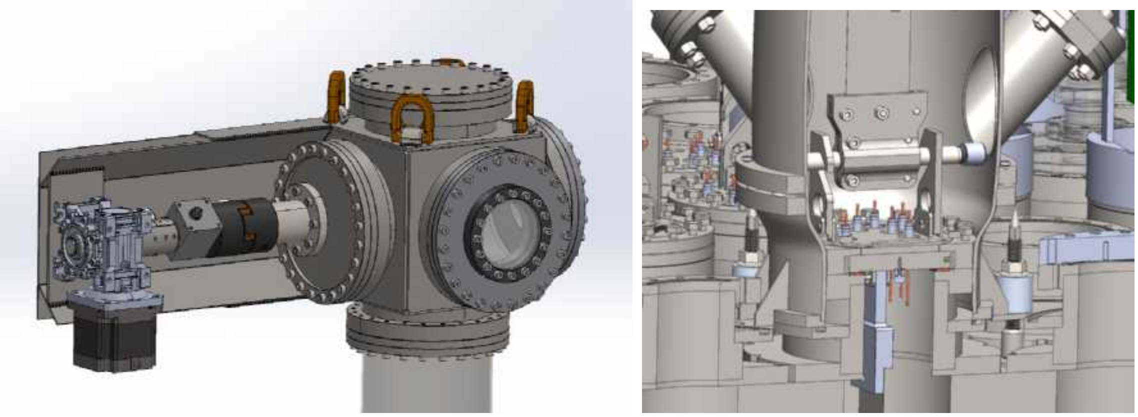

Figure 8.1-1. Legend-1000 Lock Design in Solidworks. Left: Winch drive train. Right: Cross section of magnetic coupling concept.

In some special cases other CAD packages are used for component or assembly design and then brought into the parent PDM system. One instance of this is the ITEM frame assembly used in the LYSO cluster cart designed for use at Paul Scherrer Institute (PSI). This cart needed to be able to be craned into place at PSI, be wheeled into position, have adjustable height for optical components, and be able to lock into a stable position for use. The cart was designed in the ITEM online CAD system, exported in two configurations as a step file, and brought into the Inventor PDM system for revision control and file retention.

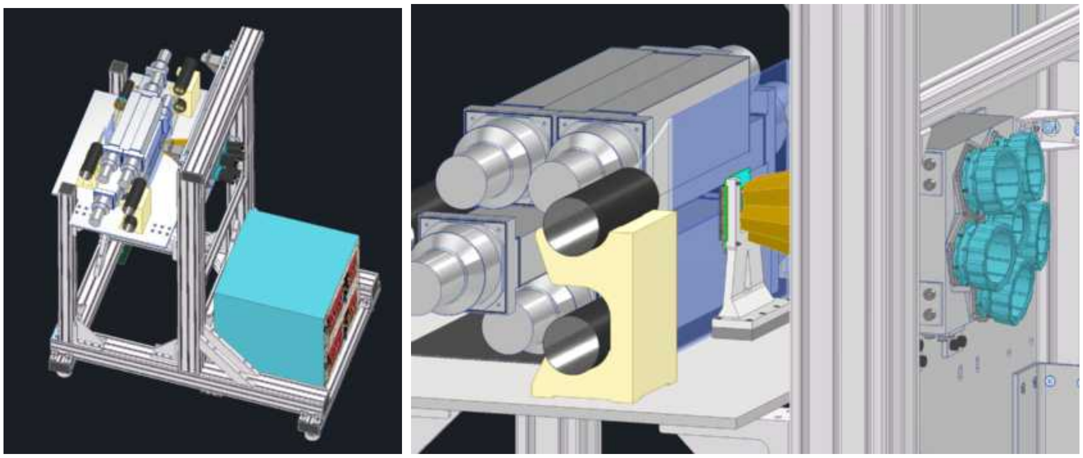

Figure 8.1-2. LYSO resolution test cluster design. Designed in collaboration with Paul Scherrer Institute (PSI), for use at PSI.

For analysis the engineering group uses COMSOL, MolFlow, the finite element (FEA) tools available in Inventor and Solidworks, and analytical methods.

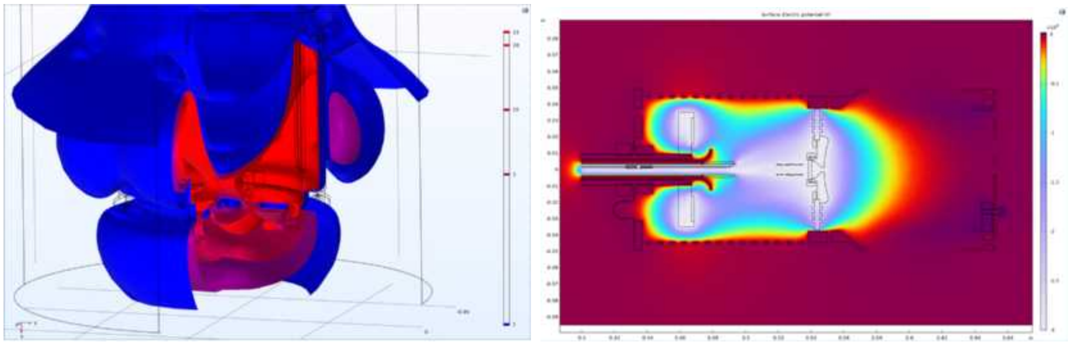

Figure 8.1-3. COMSOL Analysis Screen Shots

Fabrication

An important part of facilitating hardware for CENPA physics experiments are all of procurement, fabrication, and assembly. Procurement follows University of Washington procedures using either Workday or a university credit card. Fabrication is done using varying methods. If a simple item needs to be machined then engineering staff is capable and can use the machine tools in the student shop for that process. More complex or difficult to machine items are tasked to the professional instrument maker. When plastic parts are needed they are often printed, rather than machined, using either the Prusa Mk3S+ or Prusa Mk4 3D printers. Parts are often prototyped using 3D printing, design is verified, and then the final part can be machined.

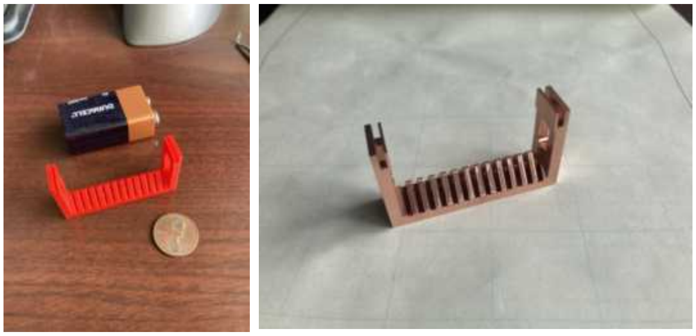

Figure 8.1-4. DAMIC tip spacer - Prototype and Final

Some of the assembly operations are more complex than others. On the more complex side is the Project-8 electron gun. Over the past several years the design has been iterated to better perform to specifications. This assembly will need to achieve 30kV, live in vacuum, and be non-magnetic as it will be inside the magnetic field of the MRI housing the resonant cavity of the Project-8 experiment.

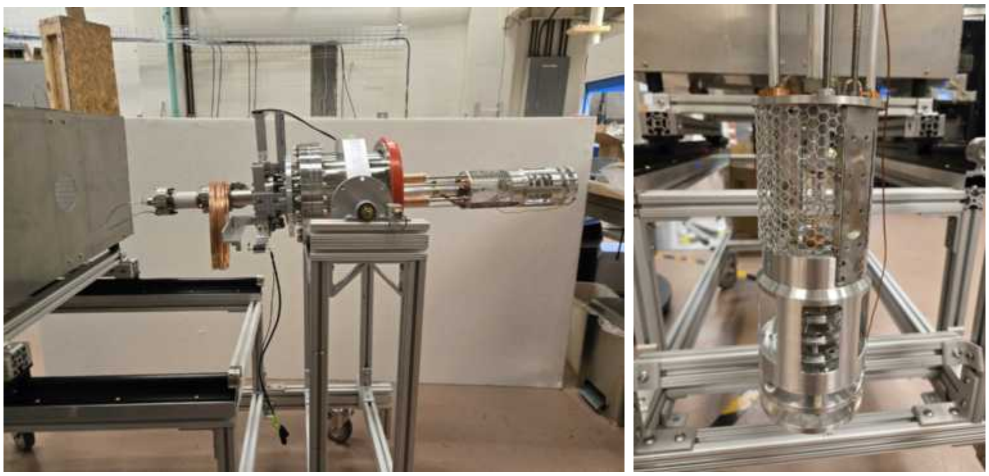

Figure 8.1-5. Project-8 e-gun. Left: Ready for insertion into the MRI. Right: Head detail.

Figure 8.1-6. M. Kallander assembling the Project-8 electron gun.

Another complex assembly was the LYSO cluster stand for PSI. This stand was designed by CENPA with input from PSI. The stand itself was designed in ITEM online software, pieces cut to size by ITEM, and shipped from Germany to PSI directly. Parts used to hold the LYSO and NaI were machined, welded, and/or 3D printed at CENPA. Other machined parts were machined at ETH Zurich. Assembly of the parts was performed at PSI by personnel from all participating institutions.

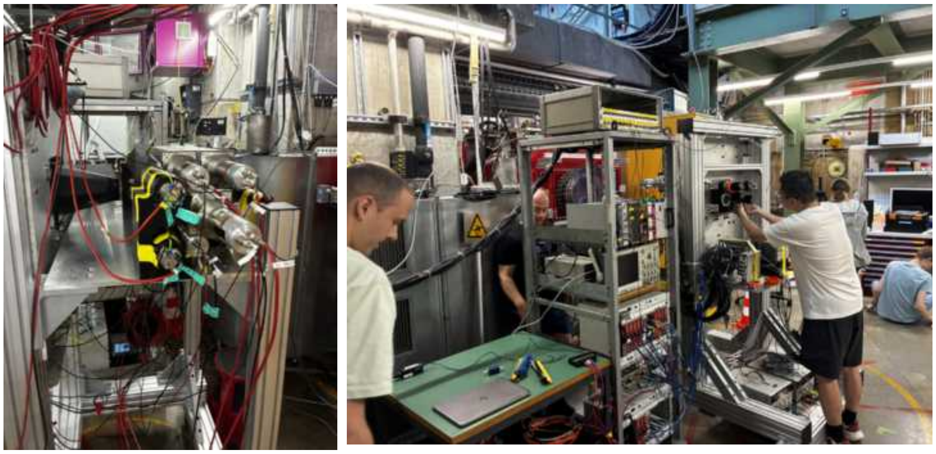

Figure 8.1-7. LYSO resolution test cluster as fabricated and assembled at PSI.

Often there is collaboration between the electronics shop and hardware engineering. This can range from the power supply of the electron gun to smaller items, such as jigs for holding small PCBs. These jigs can assist with wire bonding, soldering, test, or even shipping. The ability to output PCB designs in a CAD friendly format and the use of 3D printing at CENPA has greatly facilitated this process.

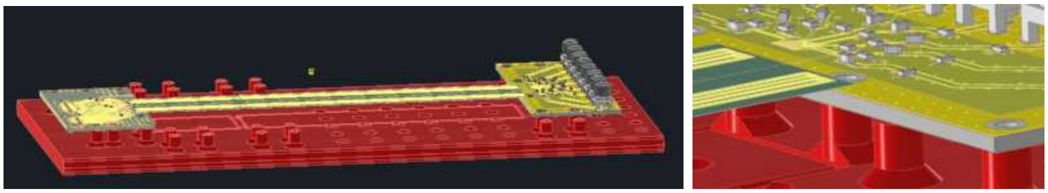

Figure 8.1-8. Pioneer PCB assembly mounted on test and shipping jig

8.2 Instrument Shop

N. Miedema and R. Roehnelt

The CENPA Instrument shop is manned by a highly skilled instrument maker. With the help of our engineering team, the shop fabricates ultrahigh vacuum systems, cryogenic components, equipment operating in high magnetic fields, and mechanical devices for high and low voltage applications, hydraulics and gas handling systems. The shop assists faculty, staff, and students. It provides advice and suggestions to ensure what they want is feasible and cost effective. There are also training courses that students can attend to learn machining, CNC machining, welding, soldering/brazing, and sheet metal work. Over the past year we have acquired a new laser welder to assist with welding very thin material. We are also in the process of clearing out redundant and outdated equipment with the hopes of updating the shop's capabilities. Together we solve mechanical problems and give CENPA the ability to be a leading force in Physics. In addition to supporting CENPA we also machine parts for other groups at the University and occasionally other organizations.

Figure 8.2-1. Pioneer mock-up crystals

Figure 8.2-2. Non-magnetic cooling jacket.

Figure 8.2-3. Project-8 electron gun ion curtain

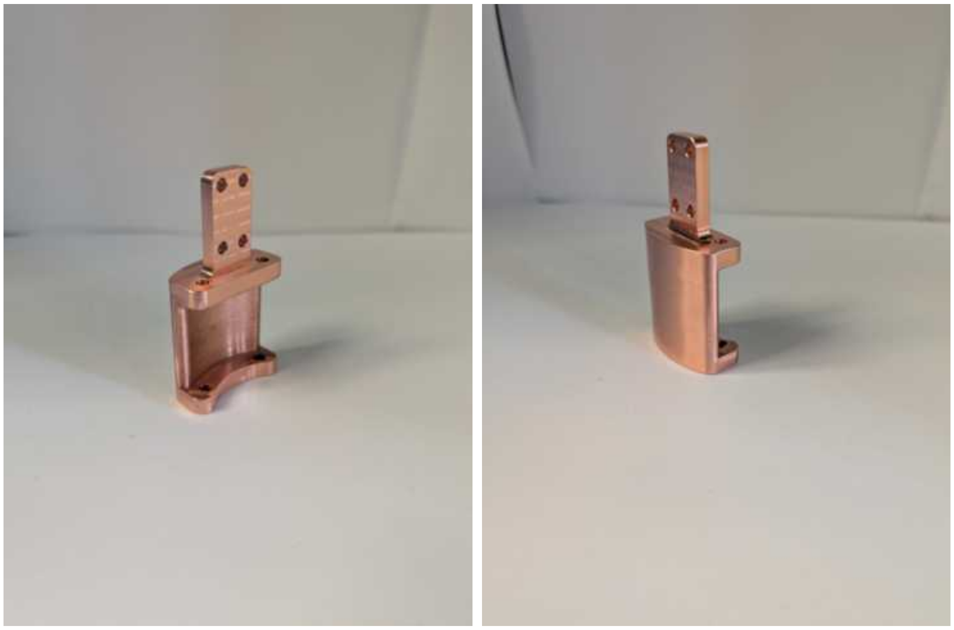

Figure 8.2-4. Cernox Coldhead Non-magnetic Mount

8.3 CENPA electronic shop

D. A. Peterson and T. D. Van Wechel

The electronics shop is responsible for the design and construction of new laboratory electronic equipment as well as the maintenance and repair of existing CENPA electronics. Projects undertaken by the electronics shop in the past year include the following:

1. For the Pioneer ATAR project significant progress has been made in the development and testing of the Front End Boards(FEB) with bare die FAST3 chips and of the flex2025 FLEX cables that connect the FEB to the LGAD detector board. We expect to receive packaged BGA FAST3 chips soon and have begun the design of a new FEB that combines the FEB and FLEX cable into a single rigid flex board.

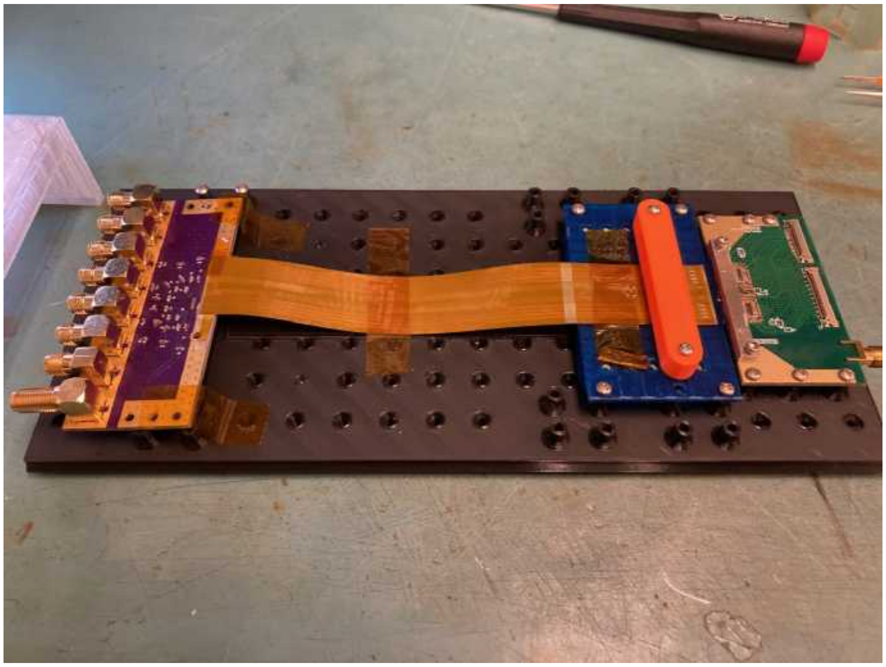

Figure 8.3-1. FAST3 FEB and flex2025 FLEX Cable Test Setup.

2. For the Pioneer ATAR project significant progress has been made in the development and testing of the Front End Boards(FEB) with bare die FAST3 chips and of the flex2025 FLEX cables that connect the FEB to the LGAD detector board. We expect to receive packaged BGA FAST3 chips soon and have begun the design of a new FEB that combines the FEB and FLEX cable into a single rigid flex board.

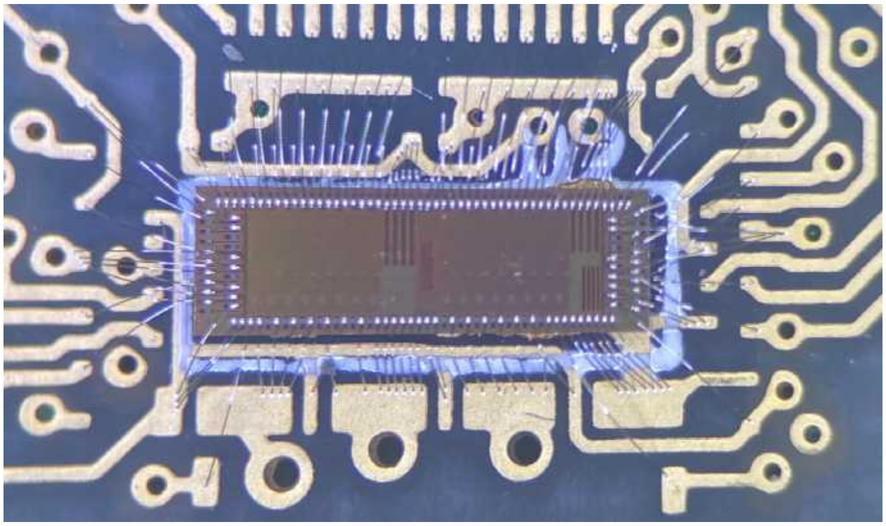

Figure 8.3-2. Closeup of bare die FAST3 chip and wirebonds on the FAST3 FEB.

3. For Project 8, a Low Noise, High Gain Wire Detector Preamplifier and a Current Shutdown Latch for the Spellman egun high voltage power supply were designed and

constructed.

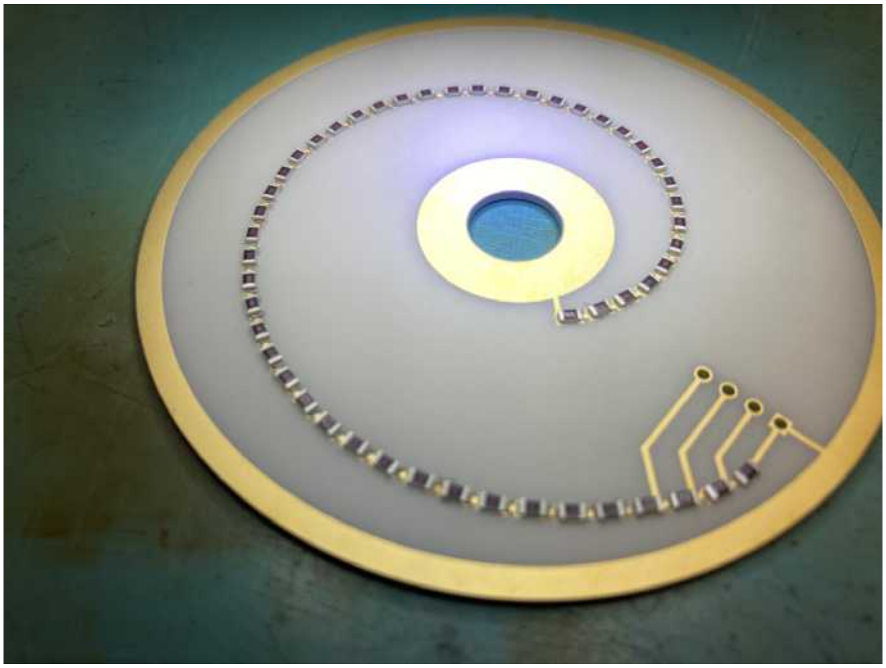

Figure 8.3-3. Project 8 spiral voltage divider board.

4. The Bipolar ExB High Voltage Plate Driver was designed and constructed for He6 CRES.
5. Construction of a flip chip carrier board for SELENA.
6. Construction of a Pirani Readout board for LEGEND.

8.4 CENPA IT & compute cluster

G. Holman and C. Wiseman

CENPA is a mixed bag of Windows 7 (some legacy), 11, Mac OS X, and various Linux distributions. Windows 11 is installed on new workstations, FreeDOS 1.2 on accelerator controllers, and an embedded Win 98 for a mechanical shop mill. As in every year, IT focus was on server consolidation, network security, process documentation, and removal of redundant processes. We continue to utilize virtualization for Autodesk Vault versions, ELOGs, wikis, collaboration calendars, and document servers. We now use Proxmox for virtualization. The CENPA website and the research group web pages run on a Drupal 7 web framework. The NPL mail server still provides NPL presence, but all email is relayed to UW email hardware. The workstations are connected to the UW Delegated Organization Unit (OU), which mainly eliminates the need to run a dedicated domain or LDAP server.

Two Dell 510 20-TB servers (Lisa and Marie) continue to offer user storage and as a Proxmox server backup set. Linux, Windows, and Mac workstations are backed up to the 20-TB Marie raid, which is backed up off-site using the UW Trivoli backup service. Lisa runs

Proxmox Backup Server (PBS), which supports all operating systems and provides differential and encrypted backups. We now use the updated UW Crash Plan Pro suite license. Whereas workstations rely on Crash Plan Pro for backups, all servers utilize rsnapshot in addition to the virtualization snapshots to PBS. Marie provides 20 TB for research, user, and shared group data.

The CENPA network has undergone a major restructuring. The new appliance is a Dell PE620 with eight 10G network interfaces running PFsense gateway/firewall on FreeBSD. Many DLink switches are still in place; however, some have been replaced with updated Aruba 48P switches. We have isolated workstations and labs through VLANs. Although we are still on a 1G-ethernet backplane, our uplink is now two 10G-connections with traffic shaping.

Previously we ran a very dated Rocks high-performance computing (HPC) cluster on CentOS 7 (http://www.rocksclusters.org/) with PBS (previously Torque/Maui). This year, the HPC has now been fully upgraded.

The HPC is now managed by Warewulf (https://warewulf.org/) as the resource manager and Slurm (https://slurm.schedmd.com/) as the workload manager. The overall goal was to place our software in the UW Hyak-compatible operational space. This allows CENPA jobs to run on our HPC and also on UW Hyak (and other national HPCs similarly). We monitor the system and its usage through the ganglia (http://warewulf.npl.washington.edu/ganglia). This Warewulf/Slurmd HPC uses stateless-provisioning for deployment and jupyter-, virtual-container functionality. We are exploring GPU hardware upgrades for the processing and parallel processing significant speedups compared to CPUs. This will be a hybrid HPC setup with sequential and parallel CPU/GPU operations.

There are currently 112 HPC nodes: 1 login node; 1 control node; 3 raid nodes; 107 compute nodes. We are working to maximize the nodes to the rack capacity of 160. The compute nodes are mostly Dell PowerEdge R620 with Theoretical Calculations (Floating-point Operations per Clock: Each core can perform 16 floating-point) For 107 compute nodes (current): 47,936 peak GFLOPS (48 TFLOPS) single precision. For 107 compute nodes (current): 23,968 peak GFLOPS (24 TFLOPS) double precision. In general, compute nodes each have 256GB each and a backplane network: up to 1 Gbps (Gigabit per second) per port, copper. The raids house a total of 450TB capacity. The operating system for all nodes is AlmaLinux 9.5 (Teal Serval.

We have been able to operate the cluster inexpensively by relying on the infrastructure (and some of the compute nodes) left from the previous Athena cluster, which was installed in 2009. This infrastructure includes power and power distribution, cooling and racks, and physical setup. Two years ago, we reported an increase in our compute nodes with NERSC surplus nodes[1]. These nodes and our original Athena nodes are all more than ten years old, and over the past year, a number of them have gone offline. We acquired 150 Dell R620 servers from the generous donation by Bryce Livingston at The Institute for Health Metrics and Evaluation (IHME). These were acquired through an inventory transfer UW form 1024.

---

[1] CENPA Annual Report, University of Washington (2019) p. 161.

Lastly, some of the Power Connect switches show signs of age. We plan to replace these with refurbished units for less than \$200 each or upgrade the backplane to 10GBE.

At over 1000 cores and significant data storage, cenpa-rocks offers a powerful computing platform at very low cost. As much of the infrastructure is quite old, there are a number of possible failure points. Among these are the switches, power strips, and cooling fans. These can be replaced at low cost. We expect to be able to maintain and, with new nodes obtained from surplus, enhance the cenpa-rocks cluster as a computing resource for CENPA over the next year and, hopefully, beyond.

9 CENPA Personnel

9.1 Faculty

Eric Adelberger[1]	Professor Emeritus
Quentin Buat[1]	Assistant Professor
Alvaro Chavarria	Assistant Professor
Jason Detwiler	Professor
Sanshiro Enomoto	Research Professor
Alejandro Garcia	Professor
Jens Gundlach[1]	Professor
David Hertzog	Professor; Director
Peter Kammel[2]	Professor Emeritus
Elise Novitski	Assistant Professor
Hamish Robertson	Professor Emeritus
Leslie Rosenberg[1]	Professor Emeritus
Gray Rybka	Associate Professor
Claudio Savarese	Research Assistant Professor
Derek Storm[1]	Professor Emeritus

9.2 Postdoctoral Research Associates

Nicholas Buzinsky	He-6 CRES
Conner Gettings[1,3]	Gravity
Stephen Kuenstner[4]	Project 8
Joshua Labounty	Muon G-2/Pioneer
Brynn MacCoy[5]	Muon G-2/Pioneer
George Marshall[6]	Legend
Patrick Schwendimann	Muon G-2/Pioneer
Yu-Hao Sun[7]	Project 8
Michaelangelo Traina[8]	DAMIC/Selena
Louis Varriano	Legend
Dan Zhang[1]	ADMX

---

[1] Not supported by DOE CENPA grant.

[2] Retired June, 2025.

[3] Departed March, 2025.

[4] Arrived January, 2025.

[5] Departed September, 2025.

[6] Arrived March, 2025.

[7] Departed October, 2025.

[8] Departed January, 2025.

9.3 Predoctoral Research Associates

Kush Aggarwal	DAMIC/Selena
Shoshana Apple[1]	Gravity
Omar Beesley	Pioneer
Samuel Bellman[2]	Legend
Samuel Borden[3]	Legend
Madeleine Carhart[1,4]	ADMX
Kieran Enzian[1]	ADMX
Michaela Guzzetti[1,5]	ADMX
Heather Harrington	He-6 CRES
Paul Kolbeck	Project 8/He6-CRES
Ali Kurmus[6]	Project 8
Yuan-Ru Lin	Legend
Luciano Malavasi	He-6 CRES
Emily Miller[7]	Legend
Christian Nave	Legend
Harry (Xiaochen) Ni[1,3]	Selena
Ian Paulson	Gravity
James Sinnis[1]	ADMX
Edwin Solis[1,2]	Selena
Grace Song	Legend
Bradley Taylor[8]	Pioneer
Meg Wynne	Project 8
Jessie Yang	Pioneer

---

[1] Not supported by DOE CENPA grant.

[2] Arrived January, 2025.

[3] Departed December, 2025.

[4] Arrived April, 2025.

[5] Departed September, 2025.

[6] Arrived March, 2025.

[7] Departed June, 2025.

[8] Promoted from Undergrad RA to PDRA March, 2025.

9.4 Undergraduates

Leo Ahlburg	He6-CRES - Garcia
Moritz Blum	He6-CRES - Garcia
Amelia Bossom	ADMX - Rybka
Ari Callahan	Pioneer - Hertzog
Charli Chambers	Gravity - Gundlach
Quanrui Chen	ADMX - Rybka
Jason (Junwen) Diao	Legend - Detwiler
Lara Dodd	Pioneer - Hertzog
Ryan Farber	Pioneer - Hertzog
Matt Foresi	Legend - Detwiler
Gabriel Hander	He6-CRES - Garcia
Senna Keesing	ADMX - Rybka
Skylar Kemper	ADMX - Rybka
Isaac Kunen	Legend - Detwiler
Eliz Locke	Gravity - Gundlach
Gabrielle Loop	ADMX - Rybka
Marcos Marroquin	ADMX - Rybka
Sophia Michael	Pioneer - Hertzog
Anika Paudel	Gravity - Gundlach
Shriya Prasanna	Pioneer - Hertzog
Ian Wang	Legend - Detwiler
Yiwen Zhang	Legend - Detwiler

9.5 Professional and Technical staff

The professional staff are listed with a description of their recent major efforts.

Brittney Dodson	RS/E3	Accelerator Engineer
Cyrus Goodman	RS/E2	ADMX Liquification Operator
Charles Hanretty	RS/E3	ADMX Engineer Project Manager
Gary Holman	Associate Director	Computer Systems
Michael Huehn	RS/E2	Jr. Engineer, Project 8, DAMIC
Matthew Kallander	RS/E2	Jr. Engineer, Project 8
On Kim	RS/E3	ADMX
Nate Miedema[1]	RS/E2	Machine Shop
Mason Newman	RS/E2	Jr. Accelerator Engineer
David Peterson	Eng Tech 3	Electronics Technician
Ryan Roehnelt	RS/E4	Sr. Mechanical Engineer
Michael Ross	RS/E3	Gravity
Erik Swanson[2]	RS/E4 Sr	Precision experimental equipment
Timothy Van Wechel	RS/E4	Analog and Digital Electronics Design
Clint Wiseman	RS/E3	Accelerator, He6-CRES, Computer Systems

9.6 Administrative staff

Angela Lawson	Program Operations Specialist
Kyle Fitzsimmons	Fiscal Specialist 2

9.7 Hourly staff

Marcel Conde[3]	Temporary RS/E2
Landon Davidson	Undergrad RA
Gale Guisinger	Student Assistant
Armina Moghadasi[4]	Student Assistant
Alex Morgan[4]	Undergrad RA
Nicolas Oaxaca	Research Aide 2
Peter Park[5]	Research Aide 2
Brooke Yommarath	Student Assistant
Adam Yusupov	Student Assistant

---

[1] Promoted from Instrument Maker (Technical Staff) to RS/E2 March, 2025.

[2] Departed June, 2025.

[3] Departed FTE RS/E2 Role; Started Temporary RS/E2 Role July, 2025.

[4] Departed June, 2025.

[5] Departed February, 2025.

Bibliography

Publications 2024-2025

Published papers

1. Arnquist, I. *et al.* The DAMIC-M Low Background Chamber. *J. Instrum.* **19,** T11010. arXiv: `2407.17872 [physics.ins-det]` (2024). DOE Supported.
2. Aggarwal, K. *et al.* Probing Benchmark Models of Hidden-Sector Dark Matter with DAMIC-M. *Phys. Rev. Lett.* **135,** 071002. arXiv: `2503.14617 [hep-ex]` (2025). DOE Supported.
3. Saldanha, R. *et al.* Removal of spallation-induced tritium from silicon through diffusion. *Phys. Rev. D* **112,** 052011. arXiv: `2506.06568 [physics.ins-det]` (2025). DOE Supported.
4. Acharya, B. *et al.* Solar fusion III: New data and theory for hydrogen-burning stars. *Rev. Mod. Phys.* **97,** 035002. `https://link.aps.org/doi/10.1103/8lm7-gs18` (3 2025). DOE Supported.
5. Swanson, E. *et al.* The fixed probe storage ring magnetometer for the Muon g-2 experiment at Fermi National Accelerator Laboratory. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **1075,** 170338. ISSN: 0168-9002. `https://www.sciencedirect.com/science/article/pii/S0168900225001391` (2025). DOE Supported.
6. Buzinsky, N *et al.* Larmor power limit for cyclotron radiation of relativistic particles in a waveguide. *New Journal of Physics* **26,** 083021. `https://doi.org/10.1088/1367-2630/ad6d85` (2024). DOE Supported.
7. Aker, M. *et al.* Direct neutrino-mass measurement based on 259 days of KATRIN data. *Science* **388,** 180–185 (Apr. 2025). DOE Supported.
8. Aker, M. *et al.* First Constraints on General Neutrino Interactions Based on KATRIN Data. *Phys. Rev. Lett.* **134,** 251801 (June 2025). DOE Supported.
9. Acharya, H. *et al.* Measurement of the inhomogeneity of the KATRIN tritium source electric potential by high-resolution spectroscopy of conversion electrons from 83m-Kr. *Eur. Phy. J. C* **85,** 727 (July 2025). DOE Supported.
10. Acharya, H. *et al.* Sterile-neutrino search based on 259 days of KATRIN data. *Nature* **648,** 70–75 (Dec. 2025). DOE Supported.
11. Aguillard, D. P. *et al.* Measurement of the Positive Muon Anomalous Magnetic Moment to 127 ppb. *Phys. Rev. Lett.* **135,** 101802. arXiv: `2506.03069 [hep-ex]` (2025). DOE Supported.
12. Aliberti, R. *et al.* The anomalous magnetic moment of the muon in the Standard Model: an update. *Phys. Rept.* **1143,** 1–158. arXiv: `2505.21476 [hep-ph]` (2025). DOE Supported.
13. Adelmann, A. *et al.* European Strategy for Particle Physics Update – PIONEER: a next generation rare pion decay experiment. arXiv: `2504.06375 [hep-ex]` (Apr. 2025). DOE Supported.
14. Beesley, O. *et al.* Measurements of a LYSO crystal array from threshold to 100MeV. *Nucl. Instrum. Meth. A* **1075,** 170320. arXiv: `2409.14691 [physics.ins-det]` (2025). DOE Supported.

15. Yang, J. *et al.* Measured gain suppression in FBK LGADs with different active thicknesses. *JINST* **20,** C08028. arXiv: `2502.02244 [physics.ins-det]` (2025). DOE Supported.
16. Hertzog, D. W. & Hoferichter, M. The anomalous magnetic moment of the muon: status and perspectives. arXiv: `2512.16980 [hep-ph]` (Dec. 2025). DOE Supported.
17. Ashtari Esfahani, A. *et al.* Dynamics of magnetic evaporative beamline cooling for the preparation of cold atomic beams. *Phys. Rev. A* **112,** 033311. arXiv: `2502.00188 [physics.ins-det]` (2025). DOE Supported.
18. Ashtari Esfahani, A. *et al.* The real-time data processing and acquisition system for Project 8 Phase II. *Nucl. Instrum. Meth. A* **1081,** 170843. arXiv: `2506.22392 [physics.ins-det]` (2026). DOE Supported.
19. Astaschov, M. *et al.* Calorimetric wire detector for measurement of atomic hydrogen beams. *Eur. Phys. J. D* **79,** 60. arXiv: `2501.01268 [physics.ins-det]` (2025). DOE Supported.
20. Ashtari Esfahani, A. *et al.* Antenna arrays for neutrino mass measurements with cyclotron radiation emission spectroscopy. *Phys. Rev. C* **112,** 045506 (2025). DOE Supported.
21. Abe, S. *et al.* Combined Pre-supernova Alert System with KamLAND and Super-Kamiokande. *Astrophys. J.* **973,** 140. arXiv: `2404.09920 [hep-ex]` (2024).
22. Arnquist, I. J. *et al.* Assay-based background projection for the Majorana Demonstrator using Monte Carlo uncertainty propagation. *Phys. Rev. C* **110,** 055804. arXiv: `2408.06786 [nucl-ex]` (2024). DOE Supported.
23. Abe, S. *et al.* Search for charged excited states of dark matter with KamLAND-Zen. *Phys. Lett. B* **855,** 138846. arXiv: `2311.09676 [hep-ex]` (2024).
24. Borden, S., Detwiler, J. A., Pettus, W. & Ruof, N. W. Characterization of Silicon Photomultiplier Photon Detection Efficiency at Liquid Nitrogen Temperature. *JINST* **19,** P12014. arXiv: `2405.01529 [physics.ins-det]` (2024). DOE Supported.
25. Arnquist, I. J. *et al.* Search for charge non-conservation and Pauli exclusion principle violation with the Majorana Demonstrator. *Nature Phys.* **20,** 1078–1083. arXiv: `2203.02033 [nucl-ex]` (2024). DOE Supported.
26. Arnquist, I. J. *et al.* Exotic Dark Matter Search with the Majorana Demonstrator. *Phys. Rev. Lett.* **132,** 041001. arXiv: `2206.10638 [hep-ex]` (2024). DOE Supported.
27. Barbeau, P. S. *et al.* Accessing new physics with an undoped, cryogenic CsI CEvNS detector for COHERENT at the SNS. *Phys. Rev. D* **109,** 092005. arXiv: `2311.13032 [hep-ex]` (2024). DOE Supported.
28. An, P. *et al.* Measurement of Electron-Neutrino Charged-Current Cross Sections on I127 with the COHERENT NaIνE Detector. *Phys. Rev. Lett.* **131,** 221801. arXiv: `2305.19594 [nucl-ex]` (2023). DOE Supported.
29. Arnquist, I. J. *et al.* Constraints on the Decay of Ta180m. *Phys. Rev. Lett.* **131,** 152501. arXiv: `2306.01965 [nucl-ex]` (2023). DOE Supported.
30. An, P. *et al.* Measurement of Pbnat(νe,Xn) production with a stopped-pion neutrino source. *Phys. Rev. D* **108,** 072001. arXiv: `2212.11295 [hep-ex]` (2023). DOE Supported.

31. Arnquist, I. J. *et al.* Energy calibration of germanium detectors for the Majorana Demonstrator. *JINST* **18,** P09023. arXiv: `2306.08661 [physics.ins-det]` (2023). DOE Supported.
32. Adamski, S. *et al.* Evidence of Coherent Elastic Neutrino-Nucleus Scattering with COHERENT's Germanium Array. *Phys. Rev. Lett.* **134,** 231801 (2025). DOE Supported.
33. Acharya, H. *et al.* First Results on the Search for Lepton Number Violating Neutrinoless Double-β Decay with the LEGEND-200 Experiment. *Phys. Rev. Lett.* **136,** 022701. arXiv: `2505.10440 [hep-ex]` (2026). DOE Supported.
34. Arnquist, I. J. *et al.* Rare multinucleon decays with the full data sets of the Majorana Demonstrator. *Phys. Rev. C* **112,** L022501. arXiv: `2412.16047 [nucl-ex]` (2025). DOE Supported.
35. Arnquist, I. J. *et al.* Final Results of the Majorana Demonstrator's Search for Double-Beta Decay of Ge76 to Excited States of Se76. *Phys. Rev. Lett.* **134,** 242501. arXiv: `2410.03995 [nucl-ex]` (2025). DOE Supported.
36. Abe, S. *et al.* Limits on the Low-energy Electron Antineutrino Flux from the Brightest Gamma-Ray Burst of All Time. *Astrophys. J.* **981,** 192. arXiv: `2410.01996 [astro-ph.HE]` (2025).
37. Abe, S. *et al.* Search for Majorana Neutrinos with the Complete KamLAND-Zen Dataset. *Phys. Rev. Lett.* **135,** 262501. arXiv: `2406.11438 [hep-ex]` (2025).
38. Acerbi, F. *et al.* Flow and thermal modelling of the argon volume in the DarkSide-20k TPC. *Journal of Instrumentation* **20,** P06046. arXiv: `2503.08468` (2025).
39. Agnes, P. *et al.* Fermionic sub-GeV dark matter from evaporating primordial black holes at DarkSide-50. *Phys. Rev. D* **112,** 123058. arXiv: `2505.13093` (2025).
40. Agnes, P. *et al.* Impact of extreme ultraviolet radiation on the scintillation of pure and xenon-doped liquid argon. *Phys. Rev. D* **111,** 102001. arXiv: `2410.22863` (2025).
41. Acerbi, F. *et al.* Quality assurance and quality control of the 26m^2 SiPM production for the DarkSide-20k dark matter experiment. *The European Physical Journal C* **85.** arXiv: `2412.18867`. `https://doi.org/10.1088/1748-0221/20/02/P02016` (2025).
42. Acerbi, F. *et al.* Benchmarking the design of the cryogenics system for the underground argon in DarkSide-20k. *Journal of Instrumentation* **20,** P02016. arXiv: `2408.14071 [hep-ex]`. `https://doi.org/10.1088/1748-0221/20/02/P02016` (2025).

Papers submitted or to be published

43. Lin, H., Traina, M., *et al.* First Production of Skipper-CCD Modules for the DAMIC-M Experiment. arXiv: `2509.06943 [physics.ins-det]` (2025).
44. Aggarwal, K. *et al.* Daily Modulation Constraints on Light Dark Matter with DAMIC-M. arXiv: `2511.13962 [hep-ex]` (2025).
45. Zhang, X., Detwiler, J. & Wiseman, C. The lowest-radiation environments in the Solar System: new opportunities for underground rare-event searches. arXiv: `2411.09634 [hep-ex]` (Nov. 2024).
46. Abgrall, N. *et al.* The MAJORANA DEMONSTRATOR experiment's construction, commissioning, and performance. arXiv: `2501.02060 [physics.ins-det]` (Jan. 2025). DOE Supported.

Whitepapers and Proceedings

47. Arnquist, I. J. *et al.* Majorana Demonstrator Data Release for AI/ML Applications. arXiv: `2308.10856 [cs.LG]` (Aug. 2023). DOE Supported.
48. Akimov, D. *et al.* COHERENT Collaboration data release from the measurements of CsI[Na] response to nuclear recoils. arXiv: `2307.10208 [physics.ins-det]` (July 2023). DOE Supported.
49. Jacobs, P. M. *et al.* White Paper on Software Infrastructure for Advanced Nuclear Physics Computing. arXiv: `2501.00905 [nucl-th]` (Jan. 2025).

Invited talks

50. Chavarria, A. E. *The Search for Dark Matter with DAMIC-M* Invited talk, Laboratory of Nuclear and High-Energy Physics (LPNHE) Seminar, Sorbonne University, Paris, France. 2025.
51. Chavarria, A. E. *The Search for Dark Matter with CCDs* Invited talk, The International Workshop on Light Dark Matter at Accelerators (LDMA), Genoa, Italy. 2025.
52. Garcia, A. *Cyclotron radiation for nuclear spectroscopy* Radioactive Molecules Workshop, Oak Ridge, TN, 2024. DOE Supported.
53. Garcia, A. *Novel techniques for searching for new physics in nuclear beta decays* Plenary talk, Department of Atomic Energy Symposium on Nuclear Physics, Roorke, India. 2024. DOE Supported.
54. Garcia, A. *Searching for new physics using an FN tandem accelerator* Seminar, Inter Universities Accelerator Center (IUAC), Dehli, India. 2024. DOE Supported.
55. Garcia, A. *Probing physics beyond the TeV scale via nuclear beta decays* Plenary Talk, International Nuclear Physics Conference, Daejon, South Korea. 2025. DOE Supported.
56. Buzinsky, N. *He6-CRES: Precision Beta Spectroscopy with Microwaves* Nuclear Physics Seminar, University of Notre Dame, South Bend, IN. 2025. DOE Supported.
57. Harrington, H. *Nuts Bolts of Beta Spectroscopy with Microwaves* FRIB Theory Alliance topical program: New Directions in Nuclear Beta Decay. 2025. DOE Supported.
58. Enomoto, S. *Direct Measurement of Neutrino Mass with the KATRIN Experiment* Kobayashi-Maskawa Institute Colloquium, University of Nagoya, Nagoya, Aichi, Japan. Dec. 2025. DOE Supported.
59. Novitski, E. *Direct neutrino mass measurements* Two invited lectures, International Neutrino Summer School, Fermilab, Batavia, IL. 2025. DOE Supported.
60. Novitski, E. *Project 8: precisely probing neutrino mass with Cyclotron Radiation Emission Spectroscopy* Invited talk, Precision Measurements and BSM Physics Workshop, APS DNP meeting, Boston, MA. 2024. DOE Supported.
61. Novitski, E. *Cryogenic* ^{3}*He NMR for the Project 8 Neutrino Mass Experiment* Invited talk, ^{3}He Magnetometry Workshop, Ann Arbor, MI. 2024. DOE Supported.
62. Kuenstner, S. *Searching for dark matter, the neutrino mass, and more - adventures at the quantum limit (and beyond)* Department Seminar, Santa Clara University Department of Physics and Engineering Physics, Santa Clara, CA. 2025. DOE Supported.

63. Novitski, E. *A new approach to measuring neutrino mass* Physics Seminar, University of Washington-Bothell, Bothell, WA. 2025. DOE Supported.
64. Novitski, E. *A new approach to measuring neutrino mass* Physics Department Colloquium, University of Kentucky, Lexington, KY. 2025. DOE Supported.
65. Novitski, E. *A new approach to measuring neutrino mass* Nuclear Science Seminar, Facility for Rare Isotope Beams, East Lansing, MI. 2025. DOE Supported.
66. Novitski, E. *A new approach to measuring neutrino mass* KICP Seminar, Kavli Institute for Cosmological Physics, Chicago, IL. 2025. DOE Supported.
67. Chavarria, A. E. *Development of Selena* Invited talk, 15th Conference on the Intersections of Particle and Nuclear Physics (CIPANP 2025), Madison, WI. 2025. DOE Supported.
68. Detwiler, J. A. *Searching for the Creation of Matter without Antimatter* INTURN Summer Lecture, Institute of Nuclear Theory, Seattle, WA, USA. Aug. 2024. DOE Supported.
69. Detwiler, J. A. *Searching for the Creation of Matter (Without Antimatter) in* Majorana *and LEGEND, Physics Seminars* T. D. Lee Institute and Shanghai Jiao-Tong University, Shanghai, China. Sept. 2023. DOE Supported.
70. Detwiler, J. A. *Searching for Neutrinoless Double-Beta Decay with KamLAND-Zen and LEGEND, Invited Talk* MAYORANA Workshop, Modica, Italy. July 2023. DOE Supported.
71. Detwiler, J. A. *0 Experiments: Global Overview* PandaX-xT Open Meeting, T. D. Lee Institute and Shanghai Jiao-Tong University, Shanghai, China. Apr. 2025.
72. Detwiler, J. A. *Reactor Neutrino Oscillations: Discovery with KamLAND, Status, and Future Outlook* Workshop on Ghost Particle Hunting, University of Hawaii at Manoa, Honolulu, HI. May 2025.
73. Detwiler, J. A. *LEGEND: New Results and Future Prospects in the Search for Neutrinoless Double-Beta Decay in 76Ge* Experimental Particle Physics Seminar, University of Pennsylvania, Philadelphia, PA. May 2025.
74. Detwiler, J. A. *Recent Results and Future Prospects in the Quest to Observe the Creation of Matter without Antimatter* Physics Colloquium, Drexel University, Philadelphia, PA. May 2025.
75. Marshall, G. G. *Neutrinoless Double Beta Decay: Status and Future* Physics in Collision Conference, Zewail City of Science, Cairo, Egypt. Oct. 2025.
76. Savarese, C. *New Tech for New Physics - An "how-to" guide on tomorrow detectors* Enrico Fermi Institute Colloquium, University of Chicago, Chicago, IL. Jan. 2025.
77. Savarese, C. *Two lectures on Experimental Dark Matter Searches* Tri-Institute Summer School on Elementary Particles, TRIUMF, Vancouver, Canada. June 2025.
78. Savarese, C. *The exceptional impact of liquid argon on rare-event experimental physics* Physics Astronomy Colloquium, University of California Riverside, Riverside, CA. Oct. 2025.
79. Varriano, L. J. *The search for Majorana neutrinos: an experimental overview* APS Global Summit 2025. Anaheim, CA. 2025.
80. Varriano, L. J. *First results from the LEGEND-200 search for* $0\nu\beta\beta$ *decay* Pacific Northwest National Laboratory seminar. Richland, WA. 2025.

Abstracts and contributed talks

81. Harrington, H. *Recent Results from He6-CRES Towards a Percent Level Measurement of bfierz* Contributed talk, APS-DNP meeting, Chicago. 2025. DOE Supported.
82. Taylor, R. *He6-CRES ExB Sweeper: a new decay cell for emptying trapped CRES events* Contributed talk, APS-DNP meeting, Boston. 2024. DOE Supported, NSF Supported.
83. Novitski, E. *Development of a resonant-cavity-based detector and magnetic field calibration for Cyclotron Radiation Emission Spectroscopy* Contributed talk, APS DNP meeting, Chicago, IL. 2025. DOE Supported.
84. Robertson, R. G. H. *Integrated Magnetic Field Design for the Project 8 Experiment* Contributed talk, APS DNP meeting, Chicago, IL. 2025. DOE Supported.
85. Sun, Y.-H. *High-Resolution Measurement of Internal Conversion Lines of 83m-Kr* Contributed talk, APS DNP meeting, Chicago, IL. 2025. DOE Supported.
86. Sun, Y.-H. *Analytic solution for the cyclotron radiation of an electron in a simple harmonic magnetic trap* Contributed talk, APS Global Summit, Anaheim, CA. 2025. DOE Supported.
87. Novitski, E. *Project 8: precisely probing neutrino mass with Cyclotron Radiation Emission Spectroscopy* Contributed talk, APS DNP meeting, Boston, MA. 2025. DOE Supported.
88. Novitski E., S. J. V. D. P. W. & Peña, J. f. t. P. c. *Project 8: results and prospects* Contributed poster, The XXXI Intl. Conf. on Neutrino Physics and Astrophysics, Milan, Italy. 2024. DOE Supported.
89. Claessens, C. & Weiss, T. E. *Sensitivity analysis for the neutrino mass experiment Project 8* Contributed poster, The XXXI Intl. Conf. on Neutrino Physics and Astrophysics, Milan, Italy. 2024. DOE Supported.
90. Ni, X. *TopmetalSe, an open-source low-noise CMOS charge imager for Selena* Contributed talk, APS Global Physics Summit 2025, Anaheim, CA. 2025. DOE Supported.

Ph.D. degrees granted

91. Ni, X. *TopmetalSe: Development of a CMOS Imager for Selena* PhD thesis (University of Washington, 2025). DOE Supported.
92. Hostiuc, A. *Search for Neutrinoless Double Beta Decay Using the Full MAJORANA DEMONSTRATOR Dataset* PhD thesis (Washington U., Seattle, Aug. 2023). DOE Supported.
93. Borden, S. *Quasi-Background-Free Neutrinoless Double-Beta Decay Searches with LEGEND: Statistical Methods and Cryogenic SiPM Characterization* PhD thesis (Washington U., Seattle, Dec. 2025). DOE Supported.
94. Song, G. *The CAGE Scanner: Development of a Novel Surface Event Rejection Technique in High Purity Germanium Detectors* PhD thesis (Washington U., Seattle, Dec. 2025). DOE Supported.
95. Guzzetti, M. *Designing, Operating and Analyzing: The Quest for Axion Dark Matter with ADMX* PhD thesis (Washington U., Seattle, Oct. 2025).