DISSERTATION

CALIBRATION OF LIQUID ARGON TIME PROJECTION CHAMBERS USING AMBIENT RADIOACTIVITY AND ACTIVITY FROM ^{220}Rn INJECTION

Submitted by

Samuel J. Fogarty

Department of Physics

In partial fulfillment of the requirements

For the Degree of Doctor of Philosophy

Colorado State University

Fort Collins, Colorado

Summer 2026

Doctoral Committee:

Advisor: Michael Mooney

John Harton
Bob Wilson
Carmen Menoni
Yulia Maximenko

Copyright by Samuel J. Fogarty 2026

All Rights Reserved

ABSTRACT

CALIBRATION OF LIQUID ARGON TIME PROJECTION CHAMBERS USING AMBIENT RADIOACTIVITY AND ACTIVITY FROM ^{220}Rn INJECTION

Liquid argon time projection chambers (LArTPCs) are vital tools in modern particle physics, notably for neutrino and dark matter experiments. The Deep Underground Neutrino Experiment (DUNE) will leverage this technology to precisely measure neutrino oscillation parameters. It will feature an intense neutrino beam at Fermilab, a near detector to measure the neutrinos before oscillations, and a far detector (800 miles away) in Lead, South Dakota to measure the neutrinos after oscillations. Because neutrinos are neutral, their properties must be inferred by detecting the charged particles they produce upon interacting with matter. Accurate particle reconstruction requires precise detector calibration. However, because the DUNE Far Detector is located a mile underground, it is shielded from the cosmic-ray muons typically used for surface-level calibration, necessitating alternative methods.

This thesis investigates using low-energy radioactivity, both naturally occurring and artificially introduced, as an alternative calibration source for LArTPCs. While natural sources like ^{39}Ar beta decays offer potential applications for calibrations, injecting artificial sources like ^{220}Rn provides a broader range of low-energy signals. In particular, the $^{212}\text{Bi} \rightarrow {}^{212}\text{Po} \rightarrow {}^{208}\text{Pb}$ (BiPo) decay sequence produces a distinct, localized signature of both ionization charge and scintillation light from a delayed beta-alpha coincidence, making it an ideal candidate for uniform detector calibration.

In this thesis, LArTPC calibrations are demonstrated using natural low-energy radioactivity in the 2x2 Demonstrator, a DUNE Near Detector prototype at Fermilab. Then, the performance of the ^{220}Rn injection technique is assessed in a small-scale LArTPC test stand at Colorado State University. Finally, the injection technique is evaluated in the 2x2 Demonstrator, successfully performing multiple calibrations using the introduced radioactivity.

ACKNOWLEDGEMENTS

I would not have reached this point without the unwavering support from many people over the course of my PhD. Firstly, I would like to thank my advisors, John Harton and Mike Mooney, for their constant support over the years. They have been an excellent source of knowledge and advice, essential for helping me get through the many challenges I faced during this endeavor. I would also be remise not to thank the Physics Department and the High Energy Physics group at CSU for their strong support, allowing me to fully focus on performing research throughout my PhD career. Part of this support included a variety of opportunities to travel to CERN and Fermilab to participate in research activities, and to travel to various places, international and domestic, to share my research and make connections with other researchers. These experiences were critical for advancing my professional development and skills as a researcher. I am also very thankful for the opportunity to be involved in various lab activities at CSU, where I was able to taking a leading role. These experiences were incredibly important for helping me realize and develop my research interests that will surely help guide my career path in the coming years. It has been an absolute pleasure being able to work with and learn from the many students, faculty, and staff at CSU that have been involved in these laboratory activities. And of course, I would like to thank my family for their strong and constant support, which has been critical to my success over the years. They have been immensely supportive in so many ways, and I am incredibly grateful for their patience, love, and understanding while I have pursued this wild journey.

TABLE OF CONTENTS

ABSTRACT ii
ACKNOWLEDGEMENTS iii
LIST OF TABLES viii
LIST OF FIGURES ix

Chapter 1 Introduction 1
1.1 Overview of Neutrinos 1
1.1.1 Neutrinos and the Standard Model 1
1.1.2 Overview of Neutrino Oscillations 1
1.1.3 Measuring Neutrino Oscillation 2
1.1.4 Detecting Neutrinos with LArTPCs 4
1.2 DUNE Prototyping 6
1.3 Thesis Overview 7

Chapter 2 Neutrinos and Neutrino Oscillation 9
2.1 Introduction to Neutrinos 9
2.2 Neutrino Interaction Mechanisms 10
2.2.1 Overview 10
2.2.2 Charged-Current Interactions 11
2.2.3 Neutral-Current Interactions 12
2.2.4 Experimental Signatures and Implications for Detection 13
2.3 Discovery of Neutrinos 13
2.3.1 Theoretical and Conceptual Introduction of the Neutrino 13
2.3.2 Evidence of the Neutrino 15
2.4 Neutrino Flavors 17
2.4.1 More Than One Neutrino? 17
2.4.2 Discovery of the Muon Neutrino 18
2.4.3 Discovery of the Tau Particle and Tau Neutrino 20
2.5 Neutrino Oscillation 21
2.5.1 Overview 21
2.5.2 Early Solar Neutrino Experiments 22
2.5.3 Direct Evidence of Neutrino Oscillation 27
2.5.4 Two-Flavor Oscillation Theory 28
2.5.5 Three-Flavor Oscillation Theory 31
2.5.6 Measurements of Neutrino Oscillation Parameters 32

Chapter 3 The Deep Underground Neutrino Experiment (DUNE) 35
3.1 Overview of DUNE 35
3.2 DUNE's Neutrino Oscillation Goals 37
3.3 Additional Goals of DUNE 39
3.3.1 Search for Baryon Number Violation 39

3.3.2 Detection of Supernova Neutrinos . . . 41
3.4 Liquid Argon Time Projection Chambers (LArTPCs) . . . 44
3.4.1 Microphysics of Scintillation and Ionization in LAr . . . 44
3.4.2 Operating Principles and Field Dependence . . . 46
3.4.3 Electrostatic Configuration and Anode Readout Architecture . . . 48
3.4.4 Three-Dimensional Spatial Reconstruction . . . 48
3.4.5 Calorimetric Precision and Neutrino Flavor Identification . . . 49
3.5 LBNF Neutrino Beam . . . 50
3.6 The DUNE Near Detector (ND) . . . 51
3.6.1 The Purpose of the ND . . . 51
3.6.2 DUNE Liquid Argon Near Detector (ND-LAr) . . . 52
3.6.3 SAND, ND-GAr, and PRISM . . . 53
3.7 DUNE ND-LAr Technology and Prototyping . . . 55
3.7.1 Charge Readout System . . . 55
3.7.2 Light Readout System . . . 58
3.7.3 Early ND-LAr Prototyping . . . 59
3.7.4 The 2x2 Demonstrator . . . 60
3.7.5 Full-Scale Demonstrator (FSD) . . . 62
3.8 The DUNE Far Detector . . . 63
3.8.1 Overview . . . 63
3.8.2 Vertical Drift Far Detector Design . . . 64
3.8.3 Horizontal Drift Far Detector Design . . . 68
3.8.4 Prototyping at the CERN Neutrino Platform . . . 70

Chapter 4 LArTPC Calibrations and Radioactive Decays . . . 73
4.1 LArTPC Calibrations . . . 73
4.1.1 Overview . . . 73
4.1.2 Particle Energy Reconstruction . . . 73
4.1.3 Electron Lifetime Calibration . . . 74
4.1.4 Electric Field Calibration . . . 76
4.1.5 Common Calibration Techniques and Sources . . . 78
4.2 Calibrations Utilizing Sources of Radioactivity . . . 79
4.2.1 Using Ambient Radioactivity for Calibrations . . . 79
4.2.2 Injected Calibration Sources: ^{220}Rn vs ^{222}Rn . . . 80
4.2.3 Calibration Applications of the ^{220}Rn Decay Chain . . . 81

Chapter 5 Low-Energy Calibrations Using Ambient Radioactivity in the 2x2 Demonstrator . . . 84
5.1 Cryogenic System . . . 84
5.2 Low-Energy Reconstruction . . . 84
5.2.1 Charge Reconstruction Methodology . . . 84
5.2.2 Charge-Light Matching Methodology . . . 86
5.3 Charge-Light Matching Selection and Performance . . . 88
5.4 Charge Scale Calibration . . . 94
5.4.1 Reconstructed Charge Spectrum of Low-Energy Activity . . . 94

5.4.2 Fitting Procedure . . . 98
5.4.3 Fit Results . . . 100
5.5 Electron Lifetime Measurements . . . 102
5.5.1 Measurement Methodology . . . 102
5.5.2 Applying the Measurement to MC Simulation . . . 105
5.5.3 Measurements with 2x2 Demonstrator Data . . . 107
5.6 Monitoring Electric Field Distortions . . . 110
5.6.1 Electric Field Distortion Measurement Methodology . . . 110
5.7 Observations from Charge Data . . . 116

Chapter 6 Development of a LArTPC Test Stand for Evaluating the ^{220}Rn Injection Technique . . . 121
6.1 Overview . . . 121
6.2 Colorado Argon Test Stand . . . 121
6.2.1 Cryogenic System . . . 121
6.2.2 LAr Purification System . . . 125
6.2.3 High-Voltage and Power Isolation Systems . . . 127
6.2.4 Charge Readout System . . . 130
6.2.5 Light Readout System . . . 131
6.3 Charge-Light Matching . . . 134
6.3.1 Synchronization Procedure . . . 134
6.3.2 Light Readout Signal Deconvolution . . . 137
6.3.3 Charge-Light-Matching Performance . . . 140
6.3.4 Purity Monitoring System . . . 143
6.4 Dedicated Tests of ^{220}Rn Injection . . . 147
6.4.1 Injection Technique . . . 147
6.4.2 ^{220}Rn Sources . . . 148
6.4.3 Experimental Setup . . . 150
6.4.4 ^{220}Rn Injection Performance . . . 151

Chapter 7 Low-Energy Calibrations Using ^{220}Rn Injection in the 2x2 Demonstrator . . . 156
7.1 ^{220}Rn Injection System . . . 156
7.2 Activity Introduced By ^{220}Rn Injection . . . 157
7.2.1 Observations from the LRS Trigger Rates . . . 157
7.2.2 Charge-Light-Matched Activity . . . 163
7.2.3 Scintillation Light Waveform Shape . . . 163
7.2.4 Pulse Shape Discrimination (PSD) . . . 170
7.2.5 BiPo Double-Peak Identification . . . 175
7.3 Drift Speed and Drift Field Measurements . . . 181
7.3.1 ^{220}Rn Daughter Cathode Plate-Out . . . 181
7.3.2 Drift Speed Measurement Methodology . . . 182
7.3.3 Drift Speed Measurements by TPC . . . 184
7.3.4 Temperature-Dependence of Drift Speed . . . 190
7.4 Electron Lifetime Measurements . . . 195
7.5 ^{220}Rn Mixing and LAr Flow Dynamics . . . 197

7.5.1 $^{220}Rn \rightarrow {}^{216}Po \rightarrow {}^{212}Pb$ Tagging 197
7.5.2 Analysis of Charge Data 204
7.5.3 LAr Flow Dynamics During LAr Top-Off 208

Chapter 8 Summary and Conclusions 212

Bibliography 217

LIST OF TABLES

5.1 Summary of uncertainties in the 2x2 Demonstrator charge scale measurement. 101
5.2 Average spatial offsets and transverse E-field components measured in the 2x2 Demonstrator LArTPC modules using data collected at the University of Bern. The polarity of z and x indicates the corresponding side of the module where the measurement was made. "Other Locations" refers to other modules or polarities not explicitly shown in this table. 115

LIST OF FIGURES

1.1 Operating principle of a LArTPC. An incoming neutrino interacts with an argon nucleus to produce multiple charged particles, which subsequently interact to produce ionization charge and scintillation light. Image retrieved from Ref. [18]. 5

2.1 (Left) Feynman diagram for β decay, $n \to p + e^- + \bar{\nu}_e$, and (Right) Feynman diagram for inverse β decay, $\bar{\nu}_e + p \to n + e^+$. 14
2.2 Homestake neutrino detection experiment in the Homestake goldmine in Lead, South Dakota. Image retrieved from Ref. [61]. 24

3.1 Illustration of the LBNF/DUNE Project. A neutrino beam is generated at Fermilab near Chicago, Illinois, and is directed underground toward SURF in Lead, South Dakota. Image retrieved from Ref. [9]. 36
3.2 Schematic layout of the DUNE FD underground caverns and infrastructure at SURF in Phase I. 37
3.3 Sensitivity of a CP violation measurement at DUNE as a function of time (represented as calendar years of data-taking). Shown are the cases of maximal true CP violation ($\delta_{\text{CP}} = -\pi/2$, red), coverage of 50% of all true δ_{CP} values (cyan), and coverage of 75% of all true δ_{CP} values (turquoise). True normal ordering of the neutrino mass states is assumed in these projections. 38
3.4 Schematic representation of the microphysical pathways governing scintillation and ionization processes following particle interactions in LAr. The dashed vector represents omitted intermediate non-radiative relaxation steps. 44
3.5 Illustration of a single-phase horizontal-drift LArTPC. Ionization electrons produced along a charged particle track are driven through a uniform electric field toward an anode wire plane for signal readout. Image retrieved from Ref. [9]. 47
3.6 Illustration of the neutrino beamline and ND complex for the LBNF/DUNE project. Image retrieved from Ref. [9]. 51
3.7 (Left) Illustration of an ND-LAr LArTPC module. (Right) Illustration of a single row of five modules in ND-LAr. Images retrieved from Ref. [65]. 53
3.8 Illustration of the ND complex and detector components. The neutrino beam enters from the right side of the image. Image retrieved from Ref. [65]. 54
3.9 (Left) LArPix anode pixel tile for ND-LAr measuring 47.55 cm by 29.69 cm with 10,240 pixel pads at 3.72 mm pitch. (Right) LArPix pixel tile with 160 ASICs to digitize and read out signals produced on the pixels. 56
3.10 (Left) A block diagram of the LArPix charge readout system for ND-LAr. (Right) A picture of the PACMAN warm controller board. 57
3.11 (Left) A diagram showing the operating principle of a LCM photon detector module. (Right) A diagram showing the operating principle of an ACL photon detector module. 59
3.12 (Left) A diagram of a 2x2 Demonstrator LArTPC module [23]. (Right) A 3D model of the 2x2 Demonstrator. 61

3.13 (Left) A picture of the inside of a 2x2 Demonstrator LArTPC module, retrieved from Ref. [23]. (Right) A picture of the 2x2 Demonstrator LArTPC modules being installed into the cryostat at Fermilab. 62
3.14 (Top) Picture of an ACL light detector module and three LCM light detector modules that comprise the light readout system for the 2x2 Demonstrator. (Bottom) Picture of both sides of a LArPix pixelated anode tile for the 2x2 Demonstrator TPC modules (left shows pixels, right shows ASICs). 63
3.15 (Left) The inside of the FSD showing the pixel planes, cathode, and light detectors. (Right) Completed FSD module staged before inserting into the cryostat. 64
3.16 Schematic of a 65.8 m (L) by 18.9 m (W) by 17.8 m (H) outer-dimension cryostat that houses a 10 kt FD module. 65
3.17 (Top) Illustration of the DUNE FD VD LArTPC concept utilizing top and bottom charge readout planes and a central cathode plane. (Bottom) A perspective view of the FD2-VD detector. Images retrieved from Ref. [22]. 66
3.18 Scintillation emission spectra of the WLS materials used in the X-ARAPUCA design, denoting the DUNE FD X-ARAPUCAs' dichroic filter cutoff wavelength. Image retrieved from Ref. [96]. 67
3.19 Picture of the inside of ProtoDUNE-VD looking down inside the active cryostat volume during construction. 67
3.20 Illustration of the DUNE FD Horizontal Drift LArTPC architecture, where "A" denotes Anode Plane Assemblies (APAs) and "C" denotes Cathode Plane Assemblies (CPAs). Image retrieved from Ref. [9]. 68
3.21 The detailed structural layout of an APA relevant for the horizontal-drift FD, highlighting the wire plane orientations and embedded photon detectors. Only a partial section of each wire plane is shown for clarity. Image retrieved from Ref. [9]. 69
3.22 Detailed schematic drawing of a DUNE FD PDS module assembly showing four constituent supercells. Image retrieved from Ref. [96]. 71

4.1 Spectra observed in the WARP LArTPC from multiple radiogenic sources, namely from β decay sources in the LAr and external γ sources. Image retrieved from Ref. [110]. Note that the detector was surrounded by 10 cm of lead shielding. 80
4.2 The α and β decays associated with the ^{228}Th decay chain. The longest half life in the chain is 10.6 h, with everything decaying eventually to the stable ^{208}Pb after several days. 81

5.1 Schematic of the 2x2 Demonstrator in the NuMI neutrino beam at Fermilab. MINERvA modules, which are downstream and upstream of the 2x2 LArTPCs and used for tracking through-going muons, is not used for the analysis presented here. Image retrieved from Ref. [115]. 85
5.2 An example visualization of the optical proximity requirement for improving the selection purity of charge-light matching. The center panel shows the position of the reconstructed light-matched charge cluster. Left/Right panels show LRS waveforms for each light detector tile in the same TPC as the cluster. Red waveforms indicate optical hits found using a simple threshold on the waveforms in a given LRS trigger. The red box indicates the physical matching window around the light detector center. . 87

5.3 Reconstructed drift time of low-energy charge-light-matched charge clusters in Module-0 (top left), Module-1 (top right), Module-2 (bottom left), and Module-3 (bottom right) using data taken at the University of Bern. 89
5.4 Selection purity of low-energy charge-light matched events in the single module tests at the University of Bern. 90
5.5 Spatial distributions of charge-light matched low-energy charge clusters in the single module test at the University of Bern with Module-0 before (left) and after (right) applying optical proximity selection. 92
5.6 Spatial distributions of charge-light matched low-energy charge clusters in the single module test at the University of Bern with Module-1 before (left) and after (right) applying optical proximity selection. 92
5.7 Spatial distributions of charge-light matched low-energy charge clusters in the single module test at the University of Bern with Module-2 before (left) and after (right) applying optical proximity selection. 93
5.8 Spatial distributions of charge-light matched low-energy charge clusters in the single module test at the University of Bern with Module-3 before (left) and after (right) applying optical proximity selection. 93
5.9 Selection purity for charge-light-matched low-energy charge clusters using data collected by the 2x2 Demonstrator at Fermilab (with optical proximity selection enforced). 94
5.10 Reconstructed low-energy charge spectrum for charge-light-matched topological blips in the 2x2 Demonstrator, compared to a simulated spectrum of ^{208}Tl γ decays distributed uniformly throughout the TPC active volume. The charge spectra associated with both data and MC simulation incorporate electron lifetime corrections. 95
5.11 Purity monitor measurements recorded during the second data-taking run of the 2x2 Demonstrator, overlaid with the interpolating cubic spline fit. Each data point denotes the mean and standard deviation of eight consecutive measurements acquired at eight-minute intervals. 99
5.12 The result of the fit to the 2x2 Demonstrator low-energy charge spectrum using a model based on a ^{208}Tl γ decay MC simulation template, smearing kernel, and background model. 100
5.13 An example of using KDEs to determine the PDF associated with a set of discrete data points. The KDE is made up of a sum of individual Gaussian kernels, each with a width determined by the bandwidth parameter h. 104
5.14 (Left) The charge distributions at different drift-time bins, (Top Right) a ISE-ISE$_{\text{min}}$ distribution for a range of τ values, and (Bottom Right) distribution of measured τ after bootstrapping (^{208}Tl simulation with $\tau = 1000$ µs). In this case, the measured electron lifetime unestimated the true value, which is accounted for in the bias correction factor. 106
5.15 Electron lifetime measurements using 208Tl MC simulation samples with a variety of simulated electron lifetime values. A bias correction is extracted from these results and applied to data measurements. 107
5.16 Empirical electron lifetime correction factor (and associated uncertainties) derived using MC simulation as a function of measured electron lifetime. This multiplicative correction factor is used to adjust electron lifetime measurements using 2x2 Demonstrator data. 108

5.17 Radiogenic low-energy electron lifetime measurements acquired using 2x2 Demonstrator data, shown before and after the empirical measurement bias correction, compared alongside independent purity monitor measurements. 109
5.18 Spatial distribution of low-energy activity in the Module-0 LArTPC using data collected at the University of Bern. Events located close to the TPC boundaries and photon detectors are shown. 112
5.19 Spatial distribution of low-energy activity in the Module-1 LArTPC using data collected at the University of Bern. Events located close to the TPC boundaries and photon detectors are shown. 112
5.20 Measurements of average spatial offset and transverse E-field component in the top left quadrant of Module-1 (left) and bottom left quadrant of Module-1 (right). 112
5.21 Spatial distribution of low-energy activity in the Module-2 LArTPC using data collected at the University of Bern. Events located close to the TPC boundaries and photon detectors are shown. 113
5.22 Measurements of average spatial offset and transverse E-field component in the top left qudrant of Module-2 (top left), top right quadrant of Module-2 (top right), and bottom right quadrant of Module-2 (bottom). 113
5.23 Spatial distribution of low-energy activity in the Module-3 LArTPC using data collected at the University of Bern. Events located close to the TPC boundaries and photon detectors are shown. 114
5.24 Spatial distribution of low-energy clusters in Module-0 of the 2x2 Demonstrator at Fermilab using only CRS data, without requiring charge clusters to be matched to a light flash in the LRS. 118
5.25 Spatial distribution of low-energy clusters in Module-1 of the 2x2 Demonstrator at Fermilab using only CRS data, without requiring charge clusters to be matched to a light flash in the LRS. 118
5.26 Spatial distribution of low-energy clusters in Module-2 of the 2x2 Demonstrator at Fermilab using only CRS data, without requiring charge clusters to be matched to a light flash in the LRS. 119
5.27 Spatial distribution of low-energy clusters in Module-3 of the 2x2 Demonstrator at Fermilab using only CRS data, without requiring charge clusters to be matched to a light flash in the LRS. 119
5.28 Comparison of the reconstructed charge spectrum at the outer edges of the TPC active volume vs. near the center of the TPC active volume. 120

6.1 (Left) The TPC mounted on the CATS cryostat lid before installing the bucket. (Right) The TPC inside the internal bucket immediately before closing and sealing the cryostat. 122
6.2 The cryogenic configuration in the CSU laboratory in preparation for a TPC test, showing the primary cryostat containing the TPC, the filter cryostat, and a LAr supply dewar. The 40 L cryostat is not shown here, as it is typically operated separately. 123
6.3 (Left) Schematic diagram of the LAr filters. (Right) Photograph of the LAr filter cryostat assembly in the CSU laboratory. 126
6.4 Recorded LAr filter temperatures during the beginning of a standard regeneration cycle. 128
6.5 Annotated photograph of the power and ground isolation system for CATS. 129

6.6 (Left) Internal components of the saturable inductor chassis. (Right) Exterior of the double-shielded isolation transformer unit. 130
6.7 (Left) The CATS TPC field cage with the photon detector installed (anode removed for visualization purposes). (Middle) The pixelated side of the anode plane. (Right) The ASIC-instrumented side of the anode plane. 131
6.8 (Left) Operational concept of the CATS photon detector. (Right) The physical photon detector assembly (bottom) and the accompanying SiPM PCB motherboard (top). . . . 132
6.9 (Top) Charge trigger timestamps and (Bottom) light trigger timestamps during a single 90 min data-taking period. The large timestamp spikes correspond to the re-initialization of the PPS, marking the time at which the timestamps are re-aligned. Outside of these spikes, the timestamps oscillate between 0 and 10^6 µs due to the PPS resets. 136
6.10 Visual alignment of charge and light triggers within a single one-second window, where each charge or light trigger signal is represented by a vertical line at its respective timestamp value. 137
6.11 Absolute timestamp difference between matched charge and light triggers recorded during a single continuous data-taking run. 138
6.12 Visual representation of a calibration waveform and the single-PE template at successive stages of the deconvolution process in the frequency domain. 139
6.13 Examples of raw and deconvolved waveforms for individual physics events. The final SiPM channel is omitted because it was physically obscured inside the detector and recorded no signal. 140
6.14 Charge-light-matched low-energy events in the CATS TPC in close proximity to the light detector. Note that the drift electric field was operated at a lower value than 500 V/cm for this test, such that the maximum drift time was around 235 µs. 141
6.15 Charge activity and light waveforms for a select number of Michel electron candidates in the CATS TPC. The light waveforms are deconvolved and summed across the five unobscured SiPM channels. The time separation between light readout peaks appear consistent with the approximately 2.2 µs muon lifetime [123]. Note also that the drift electric field was operated at a lower value than 500 V/cm for this test, such that the maximum drift time was around 235 µs. 142
6.16 (Left) Schematic diagram of the CATS purity monitor components and electrical connections. (Right) Photograph of the assembled CATS purity monitor unit. 144
6.17 (Left) Photograph of the CATS purity monitor suspended from the cryostat lid. (Right) The dedicated 40 L CATS purity monitor cryostat assembly. 145
6.18 System block diagram of the purity monitor control, trigger isolation, and data acquisition networks. 146
6.19 (Left) Averaged purity monitor waveforms under various electric field configurations. (Right) Electron lifetime measurement of ~ 22 µs. 148
6.20 ^{220}Rn sources used for tests described in this thesis. (Left) Eckert & Ziegler electroplated source. (Right) Pylon Electronics flow-through source (TH-1025). 149
6.21 Experimental configuration for preliminary testing of ^{220}Rn inline injection using an isolated light detector in LAr. 151

6.22 Recorded light detector trigger rates over time during and after inline ^{220}Rn injection using the $\sim$ 15 kBq Eckert & Ziegler source (top) and the $\sim$ 90 kBq Pylon source (bottom). 152
6.23 (Top) Zoomed view of the light trigger rates following Pylon source injection, revealing a distinct, slow increase long after ^{220}Rn stopped. (Bottom) Calculated isotope activity over time modeled using the Bateman equations. 153

7.1 Illustration of the ^{220}Rn injection manifold as installed within the gas makeup line in the 2x2 Demonstrator. 157
7.2 Top-down diagram of the 2x2 Demonstrator modules and TPCs, indicating the ^{220}Rn injection location. 158
7.3 LRS trigger rates before, during, and after the ^{220}Rn injection period in the 2x2 Demonstrator. 159
7.4 (Top) Light readout trigger rates, (Middle) cryostat relief valve open percentage, and (Bottom) internal cryostat pressure during the ^{220}Rn injection period in the 2x2 Demonstrator. The relief valve can turn on and off very quickly in response to pressure, so the graph shows an average of every two data points. 161
7.5 Reconstructed spatial distributions, in the 2x2 Demonstrator x–z plane, of charge-light-matched events before ^{220}Rn injection, over a 9.36 hour period. 164
7.6 Reconstructed spatial distributions, in the 2x2 Demonstrator x–z plane, of charge-light-matched events during ^{220}Rn injection, over a 9.36 hour period. 164
7.7 Reconstructed spatial distribution, in the 2x2 Demonstrator y–z plane, of charge-light-matched events before ^{220}Rn injection for (Left) Module-0 TPC-0 and (Right) Module-0 TPC-1. 165
7.8 Reconstructed spatial distribution, in the 2x2 Demonstrator y–z plane, of charge-light-matched events during ^{220}Rn injection for (Left) Module-0 TPC-0 and (Right) Module-0 TPC-1. 165
7.9 Reconstructed spatial distribution, in the 2x2 Demonstrator y–z plane, of charge-light-matched events before ^{220}Rn injection for (Left) Module-1 TPC-2 and (Right) Module-1 TPC-3. 166
7.10 Reconstructed spatial distribution, in the 2x2 Demonstrator y–z plane, of charge-light-matched events during ^{220}Rn injection for (Left) Module-1 TPC-2 and (Right) Module-1 TPC-3 (over a 9.36 hour period). 166
7.11 Reconstructed spatial distribution, in the 2x2 Demonstrator y–z plane, of charge-light-matched events before ^{220}Rn injection for (Left) Module-2 TPC-4 and (Right) Module-2 TPC-5 (over a 9.36 hour period). 167
7.12 Reconstructed spatial distribution, in the 2x2 Demonstrator y–z plane, of charge-light-matched events during ^{220}Rn injection for (Left) Module-2 TPC-4 and (Right) Module-2 TPC-5 (over a 9.36 hour period). 167
7.13 Reconstructed spatial distribution, in the 2x2 Demonstrator y–z plane, of charge-light-matched events before ^{220}Rn injection for (Left) Module-3 TPC-6 and (Right) Module-3 TPC-7 (over a 9.36 hour period). 168
7.14 Reconstructed spatial distribution, in the 2x2 Demonstrator y–z plane, of charge-light-matched events during ^{220}Rn injection for (Left) Module-3 TPC-6 and (Right) Module-3 TPC-7 (over a 9.36 hour period). 168

7.15 The f_{prompt} distribution for low-energy events at the start and end of the ^{220}Rn injection period. (Left) Non-normalized to show relative rate differences, (Right) Area normalized to show shape differences. 169
7.16 Templates of β and α scintillation light waveforms used for PSD, extracted from 2x2 Demonstrator data. 171
7.17 Examples of single α, single β, and β+α (BiPo) candidates found using the PSD template fit procedure applied to 2x2 Demonstrator data. 172
7.18 Reconstructed drift time for single β, single α, and β+α (BiPo) candidates determined using the PSD selection using 2x2 Demonstrator data. 173
7.19 (Left) Charge-light-matching selection purity for single β candidates, single α candidates, and β+α (BiPo) candidates found using a template-based fitting algorithm to scintillation light waveforms. (Right) Charge spectrum of the β+α (BiPo) candidates as a function of the amplitude threshold on the β peak preceding the α peak (using 2x2 Demonstrator data). 173
7.20 (Left) The f_{prompt} variable and (Right) reconstructed charge spectra for single β candidates, single α candidates, and β+α (BiPo) candidates found using a template-based fitting algorithm to scintillation light waveforms using 2x2 Demonstrator data. 174
7.21 (Left) Fit to the β candidate charge spectrum in the 2x2 Demonstrator with $\chi^2/\text{ndf} = 50.23/47$ and (Right) fit to the BiPo candidate charge spectrum in the 2x2 Demonstrator with $\chi^2/\text{ndf} = 24.79/28$. 175
7.22 Examples of measuring the peak timing for BiPo candidates in the 2x2 Demonstrator using a sigmoid function. 177
7.23 The χ^2 distance calculated between the fitted sigmoid and the waveform peak for both the β and α peaks found in BiPo candidate waveforms. 178
7.24 (Top Left) Scan of the lower bound on the fit range, (Top Right) scan of the upper bound on the fit range, and (Bottom) scan of the minimum β peak amplitude allowed in the fit, for the measured ^{212}Po half-life from BiPo candidates in the 2x2 Demonstrator. The vertical line indicates the chosen value for the measurement and the shaded region shows the range over which this parameter was varied to evaluate the contribution to the systematic uncertainty. 179
7.25 Distribution of Δt from BiPo candidate waveforms in the 2x2 Demonstrator and corresponding fit to measure ^{212}Po half life. 180
7.26 The ^{212}Po half-life measurement using 2x2 Demonstrator data compared to measurements produced by other experiments. 181
7.27 The reconstructed drift time of near-cathode activity during ^{220}Rn injection in the 2x2 Demonstrator for Module-3 TPC-7. (Left) No additional selection requirements. (Right) Selection requirements of drift time < 2 µs and Q > 12 ke$^-$, where Q is the reconstructed charge for low-energy clusters. 183
7.28 Reconstructed ionization electron drift time for activity near the expected cathode-to-anode drift time. The fit includes a bulk activity component, a Gaussian component to capture the peak time, and a combinatorial background component. 185
7.29 (Left) Measurements of ionization electron drift speed across Module-0 TPC-0 in the 2x2 Demonstrator. (Right) Measurements of drift speed across Module-0 TPC-1. White areas are regions with too few events (< 150) to perform a measurement . . . 187

7.30 (Left) Measurements of ionization electron drift speed across Module-1 TPC-2 in the 2x2 Demonstrator. (Right) Measurements of drift speed across Module-1 TPC-3. White areas are regions with too few events ($<$ 150) to perform a measurement . . . 187
7.31 (Left) Measurements of ionization electron drift speed across Module-2 TPC-4 in the 2x2 Demonstrator. (Right) Measurements of drift speed across Module-2 TPC-5. White areas are regions with too few events ($<$ 150) to perform a measurement . . . 188
7.32 (Left) Measurements of ionization electron drift speed across Module-3 TPC-6 in the 2x2 Demonstrator. (Right) Measurements of ionization electron drift speed across Module-3 TPC-7. White areas are regions with too few events ($<$ 150) to perform a measurement. 188
7.33 Individual ionization electron drift speed measurements in each TPC of the 2x2 Demonstrator using events more than 15 pixels (17 for Module-2) from the pixel tile edges. . . 189
7.34 Ionization electron drift speed measurement in each TPC of the 2x2 Demonstrator using ^{220}Rn injection data. 190
7.35 Time dependence of the ionization electron drift speed during a brief change in cryostat pressure, shown for each TPC of the 2x2 Demonstrator. 191
7.36 NIST pressure versus temperature data for argon [133]. 193
7.37 Measurement of ionization electron drift speed dependence on temperature (%/K) in the 2x2 Demonstrator. 194
7.38 Electron lifetime measurements before, during, and after the ^{220}Rn injection period using two different sources of radioactivity. Both measurements utilize the diffusion correction determined using simulated ^{208}Tl γ decays. 196
7.39 Measured pixel-by-pixel thresholds in the 2x2 Demonstrator. The thresholds are obtained using specialized data-taking runs in which the LArPix periodic reset is disabled and the CRS is triggered by accumulated "leakage" charge. (Top left) Module-0, (Top Right) Module-1, (Bottom Left) Module-2, (Bottom Right) Module-3.) 198
7.40 (Top) Total LRS trigger rates early in the injection period and (Bottom) hit counts overtime for lower ACL detectors broken down by TPC. 199
7.41 Distribution of time differences between matched LRS triggers for RnPo candidates. The extracted half-life is consistent with the expected $\sim$ 145 ms half-life of ^{216}Po. . . . 200
7.42 Distribution of time differences between matched LRS triggers for RnPo candidates with an additional requirement of an associated charge cluster for each trigger. 201
7.43 Reconstructed charge distributions for charge–light matched RnPo candidates. Each point corresponds to an individual cluster associated with an RnPo decay. It should be noted that the measured half-life is expected to be significantly biased due to CRS inefficiencies related to variations in the CRS thresholds. 202
7.44 Spatial distribution of charge-light-matched RnPo candidates within the detector volume. 203
7.45 Background-subtracted spatial distribution of charge-light-matched low-energy charge clusters during ^{220}Rn injection in Module-0 of the 2x2 Demonstrator using a charge-only selection. (Left) TPC-0, (Right) TPC-1. 205
7.46 Background-subtracted spatial distribution of charge-light-matched low-energy charge clusters during ^{220}Rn injection in Module-1 of the 2x2 Demonstrator using a charge-only selection. (Left) TPC-2, (Right) TPC-3. 205

7.47 Background-subtracted spatial distribution of charge-light-matched low-energy charge clusters during ^{220}Rn injection in Module-2 of the 2x2 Demonstrator using a charge-only selection. (Left) TPC-4, (Right) TPC-5. 206
7.48 Background-subtracted spatial distribution of charge-light-matched low-energy charge clusters during ^{220}Rn injection in Module-3 of the 2x2 Demonstrator using a charge-only selection. (Left) TPC-6, (Right) TPC-7. 206
7.49 Comparison of charge spectra during ^{220}Rn injection versus pre-injection background for three areas inside the 2x2 Demonstrator. 207
7.50 Spatial distributions of low-energy charge activity over time during a LAr top-off for Module-1 TPC-3 in the 2x2 Demonstrator. A reference distribution is subtracted from each of the above spatial distributions to accentuate LAr flow effects. Each plot contains 15 min of data after the time shown in the title. 210
7.51 Spatial distributions of low-energy charge activity over time during a LAr top-off for Module-3 TPC-7 in the 2x2 Demonstrator. A reference distribution is subtracted from each of the above spatial distributions to accentuate LAr flow effects. Each plot contains 15 min of data after the time shown in the title. 211

Chapter 1

Introduction

1.1 Overview of Neutrinos

1.1.1 Neutrinos and the Standard Model

Neutrinos have fascinated particle physicists ever since Wolfgang Pauli first proposed them in 1930 [1,2]. Later detected experimentally in the 1950s by Reines and Cowan [3], these elusive particles were subsequently integrated into the Standard Model of particle physics [4], the framework describing the elementary particles that constitute matter and the forces governing them.

Within the Standard Model, neutrinos belong to the lepton family, alongside the electron (e), muon (μ), and tau (τ) particles. Neutrinos correspond directly to these charged leptons, appearing in three distinct "flavors": the electron neutrino (ν_e), muon neutrino (ν_μ), and tau neutrino (ν_τ). As leptons, they are fundamental particles with spin-1/2 that interact via the weak force but are completely blind to the electromagnetic and strong forces.

Traditionally, interactions involving these particles adhere to the conservation of lepton number, meaning the total number of leptons minus antileptons remains constant. Standard Model interactions also conserve individual lepton flavor numbers (L_e, L_μ, and L_τ). Neutrinos are discussed in significant detail in Chapter 2, including an overview of neutrinos (Section 2.1), the discovery of neutrinos and evidence for their existence (Section 2.3), and the discovery and theory of neutrino oscillations (Section 2.5).

1.1.2 Overview of Neutrino Oscillations

Despite the predictions of the early Standard Model, neutrinos do not strictly conserve lepton flavor number. Instead, they undergo a phenomenon known as neutrino oscillation, discovered in the late 1990s [5], which allows neutrinos to spontaneously change their flavor as they propagate through space. This behavior is described by a flavor-mixing model [6] conceptually similar to

quark mixing. In this framework, flavor states are quantum-mechanical superpositions of definite mass eigenstates. Consequently, a propagating neutrino does not possess a fixed flavor; it exists in a mixed state of electron, muon, and tau components. Its specific flavor remains fundamentally undetermined until an interaction occurs, allowing one to calculate only the quantum probability of detecting a particular flavor. Experimentally, a neutrino's flavor is identified by observing the specific charged lepton produced when the neutrino collides with matter.

Because original Standard Model formulations assumed neutrinos were entirely massless, the discovery of neutrino oscillations provided definitive evidence of physics beyond the Standard Model. This revelation has driven extensive research over the last few decades. Ultimately, physicists hope that studying these unique particles might unlock deeper cosmic mysteries, such as the profound imbalance between matter and antimatter in the universe.

1.1.3 Measuring Neutrino Oscillation

Various experiments are currently being planned or are already under construction to study neutrinos in greater detail. For instance, JUNO [7], Hyper-Kamiokande (Hyper-K) [8], and the Deep Underground Neutrino Experiment (DUNE) [9] represent next-generation projects designed to precisely measure the remaining unknown neutrino oscillation parameters, determine whether neutrinos experience CP (Charge-Parity) violation, and establish the neutrino mass ordering.

This thesis focuses on DUNE, which is engineered to accomplish these goals over its operational lifetime. The primary objectives of the DUNE experiment include conducting a comprehensive suite of precise neutrino oscillation measurements, searching for baryon number violation (such as proton decay), and detecting neutrinos from a core-collapse supernova should one occur. Additionally, DUNE will contribute to the study of solar neutrinos by providing precision measurements of neutrino-mixing parameters, the ^{8}B flux, and the day-night effect [10]. An overview of DUNE is given in Section 3.1, the specific neutrino-oscillation goals of DUNE are described in Section 3.2, and the non-oscillation goals are described in Section 3.3.

By definitively measuring the CP-violating phase, the experiment will determine whether CP

violation occurs in the lepton sector. In particular, CP violation refers to a particle whose properties change upon undergoing a charge-parity transformation (switching to its antimatter counterpart). Discovering CP violation in neutrinos would be a significant discovery on its own, and it may also offer insight into the matter-antimatter asymmetry of the universe. Determining the neutrino mass ordering is also essential for understanding the fundamental nature of neutrino mass and providing critical context for future neutrinoless double beta-decay prospects [11, 12]. Furthermore, the observation of baryon number violation would constitute a groundbreaking discovery, carrying significant implications for grand unified theories and the early evolution of the universe [13].

Measuring neutrino oscillations at DUNE will be enabled by producing an intense, well-characterized source of muon neutrinos. The experimental design incorporates a Near Detector (ND), in close proximity to the beam production point, to quantify the initial neutrino flux and flavor composition, a long baseline to allow sufficient distance for oscillations to develop, and a Far Detector (FD) to measure the mutated flux and flavor composition. Throughout the experiment, the beam polarity will alternate to produce either muon neutrinos or muon antineutrinos, allowing researchers to search for the appearance of electron neutrinos or electron antineutrinos at the FD.

The neutrino beam and ND complex will be located at Fermilab in Batavia, Illinois, whereas the FD complex will be situated 1300 km away at the Sanford Underground Research Facility (SURF) in Lead, South Dakota. Since the non-accelerator physics goals of DUNE require a low-background environment, the FD will be placed approximately one mile underground to significantly suppress the cosmic-ray background.

Due to the weakly interacting nature of neutrinos, their interaction rate within a detector is inherently low. Consequently, acquiring a dataset large enough for precise statistical measurements requires an immense detector volume. The target medium must also optimize the probability of neutrino interactions while remaining cost-effective at a multi-kiloton scale. These strict constraints dictated both the massive dimensions of the DUNE detectors and the selection of liquid argon as the interaction medium. Because argon constitutes roughly 0.9% of Earth's atmosphere, procuring the 70 kt (40 kt active mass) needed to fill the DUNE detectors is logistically feasible.

This makes it a far superior choice to alternative noble gases such as krypton or xenon, which are roughly 10,000 and 100,000 times less abundant, respectively, and would be financially prohibitive at this scale.

1.1.4 Detecting Neutrinos with LArTPCs

Because neutrinos are neutral particles, they do not interact electromagnetically with matter and instead interact solely via the weak force. Consequently, neutrino interactions with matter occur very rarely, making these particles exceptionally difficult to detect. When an interaction does take place, it produces secondary charged particles that subsequently interact electromagnetically with the detection medium. Depending on the nature of the medium, these charged particles leave behind detectable signatures, such as through ionization charge deposition, the production of scintillation light, or the production of Cherekov light. Collecting information about these secondary particles allows for the reconstruction of the properties of the original incident neutrino.

Liquid argon (LAr) was chosen as the target material for DUNE largely due to its exceptional ionization and scintillation properties. When a charged particle traverses LAr, it ionizes the argon atoms, leaving a track of free ionization electrons. In particular, producing a single electron-ion pair requires an energy deposition of approximately 23.6 eV [14]. Simultaneously, scintillation light is emitted when ionization electrons recombine with argon ions or via the de-excitation of excited argon dimers [15]. For context, a minimum-ionizing muon (such as one from cosmic-ray showers) has a stopping power of about 2.1 MeV/cm [16]. In a LArTPC operating with a drift electric field strength of 500 V/cm, depositing enough energy to produce roughly 28,000 ionization electrons and 24,000 photons per MeV [17]. The detailed microphysics involved in the scintillation and ionization processes is described in Section 3.4.1.

The LArTPC technology utilized by DUNE leverages these properties of LAr to measure particle properties with high fidelity. In a typical configuration, a LArTPC establishes a uniform electric field across a volume of LAr. This electric field drifts the ionization electrons in a straight line away from their point of creation toward an anode plane. An array of wires or a grid of pix-

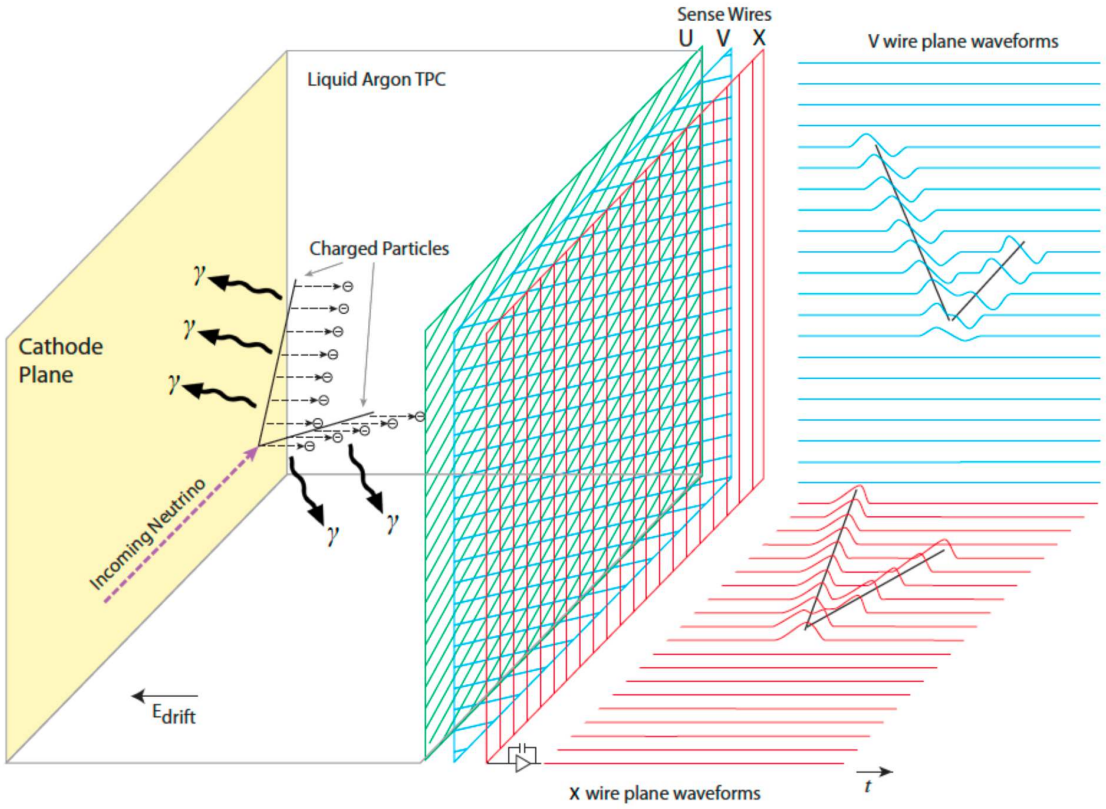

Figure 1.1: Operating principle of a LArTPC. An incoming neutrino interacts with an argon nucleus to produce multiple charged particles, which subsequently interact to produce ionization charge and scintillation light. Image retrieved from Ref. [18].

els at the anode then records the arriving charge. The signals collected at the anode enable the reconstruction of the positions and trajectories of the interacting charged particles with millimeter-scale resolution. The total measured ionization charge, along with the spatial data, can be used to reconstruct the energies and momenta of incident particles.

Determining the drift coordinate, the distance the charges traveled from their point of origin, requires establishing the precise time of the initial interaction. This is achieved by detecting the prompt scintillation light emitted during the interaction. Because ionization electrons drift to the anode on the millisecond timescale, while the scintillation light is detected on a nanosecond timescale, the light signal serves as an effective time zero (t_0) marker. Figure 1.1 illustrates the operating principles of a LArTPC.

1.2 DUNE Prototyping

To test the DUNE technology before the construction of the final detectors, a comprehensive prototyping campaign has been conducted over the last decade. ProtoDUNE-SP served as a large-scale prototype of the single-phase, horizontal-drift-style DUNE FD module and operated from 2018 to 2019 at the CERN Neutrino Platform [19, 20]. This prototype underwent a second operational run in 2024 under the name ProtoDUNE-HD (Horizontal-Drift), which incorporated various structural changes and improvements, including an upgraded photon detection system.

While the HD-style FD module design was the first to undergo large-scale prototyping, it will actually be implemented as the second module built for the final experiment. The first module will utilize a vertical-drift (VD) design, in which ionization electrons drift vertically toward a horizontally oriented charge readout plane. The VD technology evolved from the dual-phase (DP) concept originally planned for use in DUNE, which was tested in ProtoDUNE-DP [21]. The technology for the VD module was prototyped and tested at the CERN Neutrino Platform from 2025 to 2026 under the name ProtoDUNE-VD [22]. The VD-style and HD-style FD designs are described in Sections 3.8.2 and 3.8.3, respectively.

The DUNE ND prototyping program has similarly involved the operation of multiple detectors. The 2x2 Demonstrator is an ND liquid argon (ND-LAr) prototype operated at Fermilab. It consists of four separate modules arranged in a 2-by-2 array, each housing an optically isolated LArTPC. Each module was initially constructed and tested at the University of Bern in Switzerland before being shipped to Fermilab [23]. The purpose of the 2x2 Demonstrator is to perform large-scale integration testing of the ND-LAr technology and to conduct a technical evaluation under realistic beam neutrino interactions. The first operational run of the 2x2 Demonstrator in 2024 utilized a neutrino beam from the NuMI beamline, making the 2x2 Demonstrator the first DUNE prototype to successfully detect neutrinos. This was followed by a second run in 2025 to further validate the cryogenic system and evaluate various calibration sources. A subsequent run within the NuMI beamline is planned for 2026/2027.

Because the 2x2 Demonstrator design does not represent the final technical layout intended for

ND-LAr, a prototype of the full-scale detector design, named the Full Scale Demonstrator (FSD), was tested at the University of Bern in 2024. The purpose of the FSD is to systematically test and validate the final detector technology ahead of final detector manufacturing, installation, and commissioning for the final 5×7 grid of LArTPC modules in ND-LAr. The overall prototyping program also serves to provide the essential data streams required to develop the reconstruction and analysis software tools necessary for delivering the ultimate DUNE physics measurements. An overview of the ND-LAr detector design is provided in Section 3.6.2, while the remaining non-LAr detectors that make up the ND complex are described in Section 3.6.3. The specific details of the ND-LAr technology, including the charge and light readout system designs, are described in Sections 3.7.1 and 3.7.2, respectively. Finally, the ND-LAr prototyping program is described, including early prototyping, the 2x2, and the FSD in Sections 3.7.3, 3.7.4, and 3.7.5, respectively.

1.3 Thesis Overview

This thesis presents the results of analyses performed using data from the ND-LAr detector prototypes, notably the 2x2 Demonstrator, alongside a small-scale LArTPC operated at Colorado State University (CSU). One of the primary focal points is the calibration of ND-LAr using low-energy radioactivity. Specialized reconstruction and analysis software was developed to process the ND-LAr prototype data and isolate low-energy signals. Detailed analyses utilizing both charge and scintillation light signals are presented, demonstrating the monitoring of electric-field distortions, measurements of electron lifetime, and calibrations of the absolute charge scale. Subsequently, various tests of the ^{220}Rn injection technique are described, including in a local, small-scale LArTPC test stand at CSU and in the 2x2 Demonstrator. The technique was first tested in the Colorado Argon Test Stand to provide initial validation using two different ^{220}Rn sources in LAr before eventual deployment in a larger detector. The test in the 2x2 Demonstrator provided a first evaluation of the technique in a ton-scale LArTPC detector, which provided critical data to demonstrate calibrations using the ^{220}Rn decay chain. In particular, this thesis describes measurements and monitoring of various aspects of a LArTPC, including the electron drift speed and E-field spatial

variations in the detector, temperature-dependence of electron drift speed, electron lifetime, and LAr flow patterns.

The structure of this thesis is organized as follows: Chapter 2 provides an introductory description of neutrino physics and LArTPC technology. Chapter 3 describes the DUNE experiment and its specific LArTPC configurations. Chapter 4 includes a conceptual discussion of useful LArTPC calibrations, a comprehensive overview of radioactive decays in LArTPCs, the logistical details of ^{220}Rn injection. Chapter 5 details the analysis and calibrations executed with data collected by the 2x2 Demonstrator utilizing ambient, low-energy radioactivity. Chapter 6 outlines a small-scale LArTPC test stand and details preliminary verification tests of the ^{220}Rn injection technique. Chapter 7 presents the performance evaluation of ^{220}Rn injection into the 2x2 Demonstrator, showcasing multiple calibrations that utilize the activity from the ^{220}Rn decay chain. Finally, Chapter 8 summarizes the conclusions and broader implications of these studies.

Chapter 2

Neutrinos and Neutrino Oscillation

2.1 Introduction to Neutrinos

Neutrinos are neutral, low-mass elementary particles that belong to the lepton family within the Standard Model of particle physics [1, 24]. These particles interact with matter almost exclusively via the weak nuclear force, rendering them incredibly difficult to detect. Due to this elusive nature, early theoretical skepticism existed regarding the feasibility of ever detecting neutrinos experimentally. However, the neutrino was eventually discovered in 1956 [25], and experimental investigations into its properties have since become a cornerstone of modern physics.

Subsequent discoveries over the decades have revealed that neutrinos exist in three distinct types or "flavors," electron, muon, and tau (ν_e, ν_μ, ν_τ), corresponding to their heavier charged-lepton counterparts. Furthermore, the discovery of neutrino oscillation demonstrated that neutrinos spontaneously change flavors as they propagate through space. Because neutrino oscillation was not predicted by the original framework of the Standard Model, probing the neutrino sector has become a primary avenue for searching for evidence of physics beyond the Standard Model (BSM).

A number of profound, unanswered questions remain in particle physics, such as the true identity of dark matter, which constitutes the majority of the universe's mass [16, 26], and the mechanism driving the matter-antimatter asymmetry in the cosmos [27, 28]. Many physicists anticipate that neutrinos will provide answers to these puzzles or at least offer guidance toward new theoretical frameworks.

Over the years, a diverse array of experiments has successfully demonstrated evidence of neutrino oscillation and constrained the associated oscillation parameters using various neutrino sources:

- **Solar Neutrinos:** Projects such as Homestake [29], SAGE [30], GALLEX [31], and the Sudbury Neutrino Observatory (SNO) [32] investigated the historic solar neutrino problem

and measured solar neutrino fluxes.

- **Reactor Neutrinos:** Experiments including Daya Bay [33], KamLAND [34], RENO [35], and Double Chooz [36] utilized electron antineutrinos from nuclear reactors to measure mixing parameters.

- **Atmospheric Neutrinos:** Observatories like IceCube [37] and Super-Kamiokande [38] studied neutrinos produced in atmospheric cosmic-ray showers to investigate oscillation behavior over long baselines.

- **Accelerator Neutrinos:** Long-baseline experiments using directed beams from accelerator-based sources, including K2K/Super-Kamiokande [39], MINOS [40], and NOvA [41], have provided highly precise measurements of neutrino oscillation properties.

A number of planned next-generation experiments aim to make even more precise and conclusive measurements, including the DUNE [9], Hyper-Kamiokande [8], and the Jiangmen Underground Neutrino Observatory (JUNO) [7]. Collectively, these historical and upcoming projects form a global effort to definitively resolve the remaining unknowns within the neutrino sector.

2.2 Neutrino Interaction Mechanisms

2.2.1 Overview

To contextualize the operational principles and detection strategies of past, present, and next-generation neutrino experiments, it is necessary to examine the fundamental mechanisms by which neutrinos interact with matter. As previously noted, neutrinos do not carry an electric charge, meaning they interact exclusively via the weak force. The weak force is mediated by the massive gauge bosons: the charged $W^{\pm}$ bosons and the neutral Z^0 boson. These interactions are broadly categorized into two distinct interaction channels: Charged-Current and Neutral-Current. The kinematic thresholds, final-state particles, and experimental signatures of these two modes dictate how detectors must be engineered to observe them.

2.2.2 Charged-Current Interactions

Charged-current (CC) interactions are mediated by the exchange of a charged W^+ or W^- boson between the incoming neutrino and a target lepton or quark. The defining characteristic of a CC interaction is the conservation of lepton flavor at the interaction vertex. When a neutrino (ν_ℓ) or antineutrino ($\bar{\nu}_\ell$) of a specific flavor ($\ell \in \{e, \mu, \tau\}$) undergoes a CC interaction, it transforms into its corresponding charged lepton companion ($e^\pm, \mu^\pm,$ or $\tau^\pm$).

For a generic interaction with a nucleon (N), the process can be written as:

$$\nu_\ell + N \xrightarrow{W^+} \ell^- + N' \tag{2.1}$$

$$\bar{\nu}_\ell + N \xrightarrow{W^-} \ell^+ + N'', \tag{2.2}$$

where N' and N'' represent the hadronic final states resulting from the interaction. From an experimental standpoint, CC interactions are highly advantageous because they provide an unambiguous method for identifying the flavor of the incoming neutrino. By measuring the properties of the resulting charged lepton, such as the highly penetrating track of a muon (μ^-) or the electromagnetic shower generated by an electron (e^-), detectors can reconstruct both the flavor and the kinematics of the incident neutrino.

However, CC interactions are fundamentally constrained by kinematic thresholds due to the rest mass of the produced lepton. By conserving energy and momentum, the minimum neutrino energy E_ν^{thr} required in the laboratory frame to initiate a CC interaction on a stationary target nucleon of mass M is given to a first approximation by:

$$E_\nu^{\text{thr}} \approx \frac{m_\ell^2 + 2m_\ell M}{2M}, \tag{2.3}$$

where m_ℓ is the mass of the final-state charged lepton and:

$$m_e = 0.511 \text{ MeV}/c^2$$

$$m_\mu = 105.66 \text{ MeV}/c^2$$

$$m_\tau = 1776.86 \text{ MeV}/c^2.$$

While the production threshold is negligible for electron neutrinos and relatively low for muon neutrinos, it presents a significant barrier for tau neutrinos, which requires a very large incident neutrino energy to produce a τ^- lepton.

2.2.3 Neutral-Current Interactions

In contrast, neutral-current (NC) interactions are mediated by the exchange of the uncharged Z^0 boson. Discovered experimentally by the Gargamelle bubble chamber collaboration in 1973 [42], NC interactions do not alter the flavor or charge of the interacting particles. Instead, the incident neutrino scatters off a target particle (a quark within a nucleon, or an atomic electron), imparting a fraction of its kinetic energy to the target and emerging from the interaction vertex as a neutrino of the exact same flavor:

$$\nu_\ell + T \xrightarrow{Z^0} \nu_\ell + T^* \tag{2.4}$$

where T represents the target and T^* represents the recoiling target particle or broken hadronic fragments.

Because the incoming neutrino does not change into a heavy lepton, NC interactions are completely flavor-blind. This means that the cross-section for a neutral-current interaction is identical for electron, muon, and tau neutrinos. Furthermore, because the final state always includes an invisible, escaping neutrino, NC events cannot be used on an event-by-event basis to identify the neutrino flavor.

Instead, the experimental signature of an NC interaction consists solely of the recoiling target: a low-energy recoiling electron, a knocked-out proton, or a diffuse cluster of hadronic activity (pion production, nuclear de-excitation γ rays, or spallation products). Because there is no heavy lepton to create, NC interactions lack the strict mass-dependent energy thresholds governing CC

interactions, allowing neutrinos of any energy to participate.

2.2.4 Experimental Signatures and Implications for Detection

The stark contrast between CC and NC interactions governs how different detector technologies interpret neutrino data:

- **Water-Cherenkov Detectors (e.g. Super-Kamiokande):** These rely primarily on CC interactions to distinguish flavors. An electron neutrino creates an electron that generates a fuzzy, shower-like Cherenkov ring via bremsstrahlung. A muon neutrino creates a muon that travels cleanly, yielding a sharp, well-defined Cherenkov ring. NC interactions, meanwhile, generally produce isotropic gamma rays or low-energy hadronic debris that act as a difficult-to-classify background, or are used collectively for total flux measurements.

- **Liquid Argon Time Projection Chambers (LArTPCs, e.g. MicroBooNE, DUNE):** LArTPCs provide exquisite, millimeter-scale imaging of interaction vertices. CC interactions are identified by matching a clear lepton track or shower to the vertex. NC interactions are readily isolated because they lack an associated electron shower or muon track originating from the primary vertex, appearing instead as an isolated proton stub or pion cascade.

- **Scintillation Detectors (e.g. SNO, KamLAND):** Scintillators are exceptionally proficient at capitalizing on both modes. For example, the Sudbury Neutrino Observatory (SNO) famously leveraged NC interactions on deuterium to measure the total active solar neutrino flux ($\nu_e + \nu_\mu + \nu_\tau$), comparing it directly against the ν_e-only CC flux to decisively confirm solar neutrino oscillation [43].

2.3 Discovery of Neutrinos

2.3.1 Theoretical and Conceptual Introduction of the Neutrino

The history of the neutrino began with its conceptualization by Wolfgang Pauli in 1930 [1, 2], who initially termed the proposed particle a "neutron." This hypothesis emerged following James

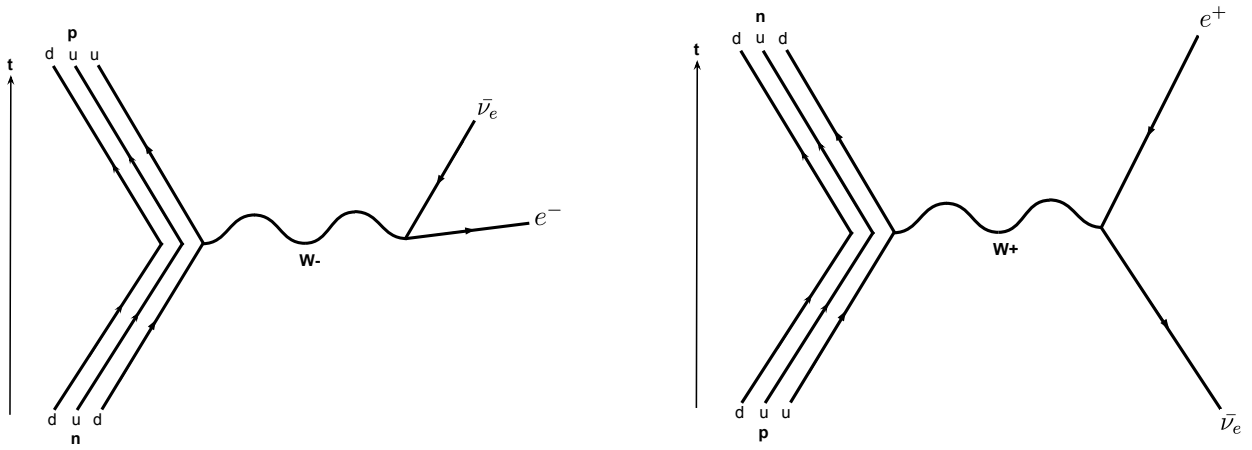

Figure 2.1: (Left) Feynman diagram for β decay, $n \rightarrow p + e^- + \bar{\nu}_e$, and (Right) Feynman diagram for inverse β decay, $\bar{\nu}_e + p \rightarrow n + e^+$.

Chadwick's observation in 1914 that β decay, i.e. the decay of a neutron into a proton plus an electron (as illustrated in Figure 2.1), resulted in a continuous electron energy spectrum rather than the discrete peaks observed in α and γ decays [44]. This observation introduced a significant paradox in contemporary nuclear physics: because only the electron from the β decay process was detectable, it was expected that such an apparent two-body decay would yield a discrete energy spectrum consistent with the conservation of energy and momentum.

Assuming that energy conservation held true in these interactions, Pauli proposed that an additional, undetected particle was involved in the process: a light-weight, electrically neutral particle that carried away the missing energy. In 1934, Enrico Fermi published a comprehensive theory of β decay that formally incorporated this particle, which he renamed the neutrino [45]. In this case, the β decay process also included an electron antineutrino in the final state (Figure 2.1, left). Fermi's theory not only mathematically reinforced Pauli's proposal but also provided quantitative, testable expressions for calculating the β decay lifetime and predicting the exact shape of the electron energy spectrum.

Shortly after Fermi published his framework, Hans Bethe and Rudolf Peierls used it to demonstrate that neutrinos must interact exceptionally weakly with matter [2]. This calculation led to a widespread consensus, shared by Bethe and Peierls themselves, that experimental detection of the

neutrino would be nearly impossible due to its negligible interaction cross-section.

2.3.2 Evidence of the Neutrino

Inverse Beta Decay

In 1946, Bruno Pontecorvo suggested that a viable method for verifying the existence of the neutrino was through its absorption by an atomic nucleus via a process known as inverse β decay [46] (Figure 2.1, right). In this type of interaction, the target nucleus transforms into a radioactive isotope. Consequently, evidence for the inverse β decay process could be established by detecting the delayed radioactivity emitted by the resulting unstable nucleus, thereby confirming the existence of the neutrino.

Pontecorvo outlined several logistical requirements for materials suitable for tracking this radioactivity [46]. First, the selected target element had to be inexpensive, because an immense volume of material was required to achieve a statistically significant number of interaction events. Second, the produced radioactive isotope needed to possess a sufficiently long half-life (at least one day) to allow researchers enough time to chemically extract the radioactive atoms from the bulk medium. Because neutrino detection technology was in its infancy, real-time detection within the bulk material was impossible; instead, the activity could only be measured after isolating the radioactive atoms and placing them inside a specialized detector. Additionally, the chemical separation process had to be straightforward, and the rate of background events mimicking the neutrino signal had to be minimal. Pontecorvo specifically proposed ^{37}Cl, ^{79}Br, and ^{81}Br as suitable target candidates.

Detecting Neutrinos from a Nuclear Reactor

The first experimental detection of neutrinos was achieved by Frederick Reines and Clyde Cowan in the 1950s [3, 25]. Expanding upon Pontecorvo's suggestions, they designed an experiment to observe the delayed radioactivity produced by inverse β decay. Their setup utilized a large tank filled with an organic liquid scintillator containing a dissolved cadmium compound. Cadmium was introduced because of its high cross-section for capturing neutrons produced during inverse β

decay, a process that subsequently triggers γ ray emission. Within the liquid, these neutrons would thermalize for approximately 5 µs before being captured by a cadmium nucleus. Meanwhile, the positron produced alongside the neutron would rapidly annihilate with a surrounding electron in the scintillator, generating prompt γ rays.

An array of photomultiplier tubes (PMTs) surrounded the tank to detect the scintillation light emitted. Reines and Cowan sought to identify a distinct delayed coincidence signature: prompt γ rays from positron annihilation, followed microseconds later by γ rays from neutron capture. During their initial trial, the observed rate of delayed coincidences significantly exceeded expectations. This surplus was later attributed to cosmic rays interacting within the apparatus's shielding, generating spurious neutrons and γ rays. To mitigate this background, Geiger counters were positioned around the apparatus to act as a cosmic-ray veto system. This adjustment yielded a delayed coincidence rate of 0.41 ± 0.20 counts per minute above the background measured while the reactor was offline. Though highly suggestive, the result retained a substantial background component and carried a 5% statistical probability of being consistent with a null signal.

Following this inconclusive test, Reines, Cowan, and their collaborators refined the apparatus to significantly suppress background rates for a subsequent experiment [25]. This definitive study demonstrated that the prompt signal was indeed produced by positrons and that the delayed pulse originated from neutron captures. The investigators also detailed how alternative background sources mimicking this signature were systematically eliminated. For example, installing extra shielding around the detector, which would block backgrounds but leave the neutrino flux unaffected, produced no statistically significant change in the signal rate (1.69 ± 0.17 hr^{-1} with shielding versus 1.74 ± 0.12 hr^{-1} without shielding). Ultimately, these findings provided the first definitive proof of the neutrino's existence, though many of its anomalous properties remained to be discovered.

2.4 Neutrino Flavors

2.4.1 More Than One Neutrino?

Following the experimental confirmation of the neutrino in the Cowan–Reines experiment, several fundamental questions remained unanswered. One prominent question concerned whether multiple distinct neutrino types existed. As Reines later remarked, physicists were uncertain whether the neutrino emitted during β decay was identical to the neutral particle produced alongside the muon as a result of pion decay. A compelling solution was proposed: two distinct neutrino types existed, each associated with a specific charged lepton, which would be referred to as electron neutrinos and muon neutrinos.

Confirming this two-neutrino hypothesis required the use of very high-energy (> 100 MeV) interactions. Such energies vastly exceeded those produced by conventional nuclear reactors, introducing the major technical challenge of engineering an intense source of high-energy neutrinos. Bruno Pontecorvo suggested testing the hypothesis by searching for the following distinct inverse processes [47]:

a. $\bar{\nu}_\mu + p \rightarrow \mu^+ + n$

b. $\bar{\nu}_e + p \rightarrow e^+ + n.$

Pontecorvo noted that reaction (b) could be observed at energies below 100 MeV using the liquid scintillation techniques developed by Reines and Cowan. Conversely, reaction (a) was entirely unobservable at reactor energies or with existing detection setups. Because the muon is roughly 200 times heavier than the electron, possessing a mass of about 105.7 MeV$/c^2$, low-energy reactor neutrinos lacked the necessary kinetic energy to produce a muon under the laws of energy conservation.

To overcome this limitation, Pontecorvo [47] and Melvin Schwartz [48] independently suggested that sufficiently energetic neutrinos could be produced by directing a relativistic proton beam from a high-energy accelerator into a stationary target. These collisions would generate

charged pions, which would subsequently decay into muons and neutrinos:

$$\pi^{\pm} \rightarrow \mu^{\pm} + (\nu/\bar{\nu}). \tag{2.5}$$

The pions could be focused into a directed beam using magnetic horns prior to their decay. A thick absorber shield could then block all secondary background particles. Given the extremely low probability of neutrino interactions with matter, only the neutrinos would pass through the shielding unimpeded. By placing a target detector downstream of this pure neutrino beam, researchers could observe which type of charged lepton was preferentially produced: muons or electrons. Observing an exclusive or significant excess of muons would prove that muon neutrinos, distinct from electron neutrinos, were being produced in the decay of the pion. This design for a "neutrino beam" would become the basis for beams used in long-baseline neutrino oscillation experiments conducted multiple decades later, such as the DUNE experiment discussed in this thesis.

2.4.2 Discovery of the Muon Neutrino

By 1960, the Alternating Gradient Synchrotron (AGS) at Brookhaven National Laboratory (BNL) became operational. Utilized by Schwartz, Leon Lederman, Jack Steinberger, and their collaborators, the AGS was powerful enough to reach the energies required to test the two-neutrino hypothesis [49]. Protons accelerated by the AGS would strike a beryllium target to produce pions, and the resulting secondary particles would be directed into a thick steel barrier engineered to absorb everything except the neutrinos.

Directly behind this barrier sat a spark chamber that served as both a particle detector and a tracking medium [50]. Generally, spark chambers consisted of a stack of parallel metal plates sealed within an enclosure filled with an inert gas such as helium or neon. As an incoming charged particle traversed the gas, it left a trail of ionization electrons. Applying a high-voltage pulse across the stacked plates triggered visible electrical sparks along the ionized track. With each pair of plates acting as an independent readout channel, the device could effectively image the trajectories of interacting charged particles. A fast electronic trigger system was typically paired

with external counters to selectively pulse the high voltage at specific times, thereby minimizing cosmic-ray or random backgrounds.

Following data collection with the active beam, 113 candidate interaction events were identified. Thirty-four of these images contained either a clean single track, a multi-track vertex event, or a localized "shower." A shower was considered an irregular, cascading track characteristic of an electron or photon interaction. The cosmic-ray background during the operational run was estimated to be 5 ± 1 events by collecting out-of-beam data while the accelerator was offline.

Neutron backgrounds were systematically ruled out because the existing steel shielding was designed to reduce their flux by a factor of 10,000. To experimentally verify this, a four-foot section of the iron shielding was temporarily removed; the event rate did not exhibit the hundred-fold increase expected if the signals were caused by neutrons. Furthermore, the absence of neutral pions, together with the single-track events, provided additional evidence against neutron contamination.

The possibility that these single-track events were actually pions was also eliminated. Because pions have greater mass than muons, they exhibit greater stopping power and would penetrate a much shorter distance into the detector; the observed tracks were too long to match this profile. Consequently, the remaining 29 background-subtracted tracks were definitively identified as muons.

While this confirmed that the incoming neutrinos were producing muons, it did not automatically exclude the possibility that electron and muon neutrinos were identical ($\nu_\mu = \nu_e$). If the two neutrino types were the same particle, the beam would have produced an equal abundance of electrons. However, separate observations of electrons with comparable kinetic energies demonstrated that true electron tracks did not resemble the six shower-like images found in the dataset. The total number of observed shower-like interactions was far below what would be expected if the neutrinos were flavor-blind. It was therefore concluded that neutrinos exclusively produced muons in pion-decay, establishing the existence of at least two distinct neutrino generations ($\nu_\mu \neq \nu_e$). This groundbreaking discovery at BNL ultimately earned Schwartz, Lederman, and Steinberger the 1988 Nobel Prize in Physics.

2.4.3 Discovery of the Tau Particle and Tau Neutrino

By the 1960s, the electron, muon, and their corresponding neutrinos had been experimentally confirmed. In the late 1960s and early 1970s, the theoretical motivation for a third lepton generation was developed, featuring key contributions from Antonino Zichichi and from Yung-su Tsai [51]. Theoretical discussions focused on the possibility of creating new heavy leptons via electron-positron (e^+e^-) annihilation. In this process, the colliding particles produce a virtual photon (γ), which subsequently materializes into a lepton-antilepton pair (l^+l^-):

$$e^+ + e^- \rightarrow \gamma \rightarrow l^+ + l^-. \tag{2.6}$$

According to the principle of mass-energy equivalence, this transformation requires the center-of-mass energy of the collision to exceed the mass threshold of the new lepton pair.

Zichichi conducted an initial search for a third-generation lepton in 1969 at the ADONE facility in Italy, which housed an early e^+e^- collider. However, that accelerator lacked the necessary energy to produce these heavy states. A higher-energy accelerator, the Stanford Positron Electron Asymmetric Ring (SPEAR), became operational at the Stanford Linear Accelerator Center (SLAC) in 1972. Utilizing SPEAR, an experimental campaign led by Martin Lewis Perl and collaborators, working alongside Tsai's colleagues, searched for evidence of a third lepton generation [52].

The investigators observed 64 anomalous events that were best explained by the production of a third lepton flavor, designated as τ. These events matched the signature:

$$e^+ + e^- \rightarrow e^\pm + \mu^\mp + \text{(at least two undetected particles)}. \tag{2.7}$$

The team proposed that this topology resulted from the creation and subsequent leptonic decay of a tau pair:

$$e^+ + e^- \rightarrow \tau^+ + \tau^- \rightarrow e^\pm + \mu^\mp + 4\nu. \tag{2.8}$$

Follow-up experiments at SLAC and DESY-Hamburg precisely measured the mass and spin of

the tau particle. Martin Perl ultimately shared the 1995 Nobel Prize in Physics with Frederick Reines, in recognition of Perl's mid-1970s discovery of the tau particle and Reines's discovery of the electron neutrino in 1956.

Following the discovery of the tau lepton, questions naturally arose regarding whether it possessed a unique, corresponding neutrino flavor. Because neutrino detection technology required significant refinement to isolate such a heavy-flavor signature, a considerable amount of time passed before the tau neutrino could be directly observed. In the 1990s, the Direct Observation of the Nu Tau (DONuT) experiment was established at Fermilab. DONuT utilized high-energy protons accelerated by the Tevatron to produce a tau neutrino flux via the decay of charmed mesons [53].

The experiment collected data in 1997, and the definitive discovery of the tau neutrino was announced in July 2000. The initial analysis identified four signal events, a result significantly above the expected background of 0.34 ± 0.05 events, corresponding to a statistical significance of approximately 3.5σ. The discovery achieved widespread consensus, and a final report in 2008 further solidified the observation by identifying nine tau neutrino candidate events out of 578 total interactions [53]. This discovery marked the detection of the final lepton predicted by the Standard Model. The entire framework was ultimately completed in 2012 with the discovery of the Higgs boson by the ATLAS and CMS collaborations at CERN [54, 55].

2.5 Neutrino Oscillation

2.5.1 Overview

Following the experimental discovery of all three neutrino generations, the Standard Model approached structural completion during the 1990s and early 2000s. The Standard Model originally described neutrinos as strictly massless particles. This structural assumption was implemented because initial experimental measurements were entirely consistent with a zero-mass profile, and the mathematical framework could not accommodate non-zero neutrino masses without substantial extensions.

However, theorists noted that if neutrinos were capable of undergoing quantum-mechanical oscillations, spontaneously changing flavor as they propagate, it would fundamentally imply that neutrinos possess a non-zero mass. Consequently, confirming the existence of neutrino oscillation would prove that neutrinos are not massless, demonstrating that the original Standard Model description of the lepton sector was incomplete.

Mixing between quantum states and flavor oscillations had previously been observed in hadronic systems, such as kaon and B-meson decays. Furthermore, quark flavor mixing stood as a foundational feature of the Standard Model, mathematically formalized through the Cabibbo-Kobayashi-Maskawa (CKM) matrix [16, 56]. By the late 20th century, multiple experimental anomalies had emerged in the neutrino sector, providing the initial indications of neutrino oscillations. These anomalies manifested as the "solar neutrino problem" [57] and the atmospheric neutrino anomaly [58], both of which involved observing significantly fewer neutrinos of specific flavors than theoretically predicted.

Definitive proof of neutrino oscillation was ultimately delivered by Super-Kamiokande in 1998 [38], and the Sudbury Neutrino Observatory (SNO) in 2001–2002 [59], with subsequent long-baseline confirmation from accelerator experiments such as MINOS [40]. This discovery necessitated extensive extensions to the Standard Model. The following section outlines the comprehensive history of the discovery of neutrino oscillation.

2.5.2 Early Solar Neutrino Experiments

Homestake

In the 1950s, theoretical models detailing the interior of the Sun predicted that neutrinos would be generated in vast quantities as a byproduct of nuclear fusion pathways [1], such as the primary proton-proton (pp) chain:

$$4\,^{1}\mathrm{H} \rightarrow {}^{4}\mathrm{He} + 2e^{+} + 2\nu_{e}. \tag{2.9}$$

Validating these solar models required direct experimental verification. Photons generated within the solar core undergo millions of years of scattering before escaping, rendering them ineffective

for probing real-time core dynamics. Conversely, weakly interacting neutrinos could propagate from the solar core to Earth essentially unimpeded. Comparing measurements of the solar neutrino flux and flavor composition with theoretical expectations offered a direct method for evaluating the accuracy of solar interior models.

Pontecorvo's initial proposal suggested using ^{37}Cl in the form of carbon tetrachloride as a target medium, an approach implemented by Raymond Davis Jr. during the 1950s and 1960s to measure the solar neutrino flux [29,60]. The ^{37}Cl isotope provided several distinct advantages for solar neutrino detection. First, the reaction threshold was sufficiently low (0.814 MeV) to permit sensitivity to high-energy neutrinos originating from ^{8}B decay. Second, the inverse β decay product, ^{37}Ar, possessed a stable half-life of 35 days, granting researchers sufficient time to chemically extract it and count individual decay events. Finally, carbon tetrachloride was readily available in multi-ton quantities, rendering the construction of a massive detector logistically feasible.

Davis initially constructed a pilot detector using 3,800 liters of carbon tetrachloride buried 19 feet underground to provide shielding against cosmic rays. While this configuration established an upper limit on the solar neutrino flux of 40,000 SNU (1 SNU $= 10^{-36}$ captures per target atom per second), the theoretically predicted value was considerably lower. This disparity indicated that a significantly larger and more sensitive apparatus was required.

This realization led to the construction of a scaled-up detection system featuring a 100-fold increase in volume (390,000 liters) filled with liquid tetrachloroethylene. To suppress cosmic-ray backgrounds, the tank was situated 4,850 feet underground within the Homestake gold mine in Lead, South Dakota, as shown in Figure 2.2. This deep underground location was critical because cosmic-ray muons interacting within the detector produced secondary protons that could generate spurious signatures that mimic ^{37}Ar, making cosmic-ray suppression paramount. Background rates from all other environmental processes were verified to be minimal and well below the expected solar neutrino signal.

Davis initiated data collection in 1965, operating the experiment for several decades and eventually reporting an average neutrino capture rate of 2.56 ± 0.25 SNU (approximately 25 μHz) [57].

Figure 2.2: Homestake neutrino detection experiment in the Homestake goldmine in Lead, South Dakota. Image retrieved from Ref. [61].

Contemporary theoretical calculations predicted a rate of 8.1 $\pm$ 1.2 SNU (approximately 80 μHz), highlighting a severe discrepancy with Davis's empirical measurements. Assuming the experimental apparatus was free of systematic tracking errors, the data indicated a significant deficit in the observed neutrino flux relative to expectations from standard solar models. Although contemporary critics voiced skepticism due to the detector's lack of directional or real-time energy resolution, the persistent deficit formalized what became known as the "solar neutrino problem." The Sun appeared to emit far fewer electron neutrinos than predicted, suggesting either an incomplete understanding of stellar astrophysics or an unmapped property within neutrino physics itself.

Gallium Experiments

The deficit observed by the Homestake experiment was subsequently investigated by alternative experimental methods. Around the time of Davis's research, discussions intensified regarding improved target materials for capturing solar electron neutrinos. In 1966, Vadim Kuzmin proposed using gallium (Ga) as a detection medium [62], specifically leveraging the inverse β decay reaction:

$$^{71}\mathrm{Ga} + \nu_e \rightarrow {}^{71}\mathrm{Ge} + e^{-}. \tag{2.10}$$

This reaction features a threshold energy of 0.237 MeV, which is significantly lower than the 0.814 MeV threshold of the chlorine-based Homestake reaction. Consequently, ^{71}Ga detectors were expected to be sensitive to the low-energy, primary *pp* solar neutrinos that eluded the chlorine experiment, yielding a higher overall capture rate and a more robust dataset.

Several other physics properties made gallium an appealing target. For one, the cross-section for electron neutrino capture on ^{71}Ga is exceptionally high. Additionally, the half-life of the resulting ^{71}Ge isotope is 11.4 days; while shorter than the 35-day half-life of ^{37}Ar, this duration still granted researchers sufficient time to chemically extract and isolate the germanium atoms. Furthermore, the electron capture decay of ^{71}Ge releases an energy of 10.4 keV (with an Auger electron peak near 1.2 keV), providing a distinct energy signature compared to the lower 2.8 keV release from ^{37}Ar. This higher energy deposition enabled pulse-shape discrimination to systematically filter out low-energy background electron signals. Despite these distinct experimental advantages, a concrete detector design was not formally proposed until 1978 by John Bahcall and collaborators, and the first empirical results were not published until the 1990s [31, 63].

During the early 1990s, two independent gallium-based experiments set out to measure the solar ν_e flux. The first was the Soviet-American Gallium Experiment (SAGE), which operated deep within a 4-km-long tunnel beneath Mount Andyrchi in the Caucasus Mountains to suppress the cosmic-ray background [30, 63]. Following a protocol similar to the Homestake experiment, the detector ran continuously for three to four weeks to allow the neutrino flux to transmute a measurable population of gallium atoms into germanium. The germanium was then chemically extracted from the liquid metallic gallium target, and its subsequent electron-capture decays were recorded inside a proportional counter over several months.

SAGE employed pulse-shape discrimination based on signal rise times to cleanly isolate the germanium decay events from environmental backgrounds. After running a series of data-taking cycles over several years, they reported an avg. capture rate of $73^{+18}_{-16}\,(\text{stat})\,^{+5}_{-7}\,(\text{syst})\,\text{SNU}$. This result fell significantly short of the contemporary theoretical expectation of about 132 SNU [63].

A parallel experiment, GALLEX (Gallium Experiment), operated during the same era within

the underground Gran Sasso National Laboratory in Italy [64]. Unlike SAGE's use of metallic gallium, GALLEX utilized 30 tons of gallium in an acidic aqueous solution. The GALLEX collaboration measured a solar neutrino capture rate of $79 \pm 10\,(\text{stat}) \pm 6\,(\text{syst})$ SNU [31], remarkably consistent with the SAGE measurement. Because both independent gallium experiments confirmed a substantial neutrino deficit in a lower energy regime than Homestake, these results provided undeniable confirmation of the solar neutrino problem.

Kamioka Nucleon Decay Experiment

The Kamioka Nucleon Decay Experiment (Kamiokande) was a water-Cherenkov detector operated deep within the Kamioka mining complex in Japan. Although originally designed to search for proton decay to investigate the validity of Grand Unified Theories, the detector was subsequently upgraded to achieve sensitivity to solar neutrinos to probe the solar neutrino problem.

Kamiokande was a significant improvement over the Homestake and gallium-based radiochemical experiments because it enabled real-time detection of events. It accomplished this by using a large volume of purified water as its active medium. When an incoming neutrino interacts with an electron or atomic nucleus in the water, it accelerates a secondary charged lepton to relativistic speeds. If this lepton moves faster than the speed of light in water, it emits a directional cone of Cherenkov radiation.

The detector's inner volume was lined with approximately 1,000 photomultiplier tubes (PMTs) to detect the Cherenkov photons. Reconstructing the hit patterns and relative arrival times across the PMT array allowed researchers to calculate the arrival time, directional vector, and kinetic energy of individual neutrinos.

Kamiokande's solar neutrino data confirmed a severe deficit relative to standard solar models, confirming the validity of the puzzle. Because the detector provided directional tracking, it confirmed that the detected neutrino flux originated directly from the Sun's core. Additionally, the experiment achieved a landmark triumph by recording the real-time neutrino burst from the core-collapse supernova SN 1987A. Ultimately, the combined data from Homestake, SAGE, GALLEX, and Kamiokande established that the solar neutrino deficit was an indisputable physical reality

requiring new theoretical insights.

2.5.3 Direct Evidence of Neutrino Oscillation

Anomalous data across multiple experiments indicated that the traditional understanding of neutrino physics was incomplete. The leading hypothesis, which was eventually proven correct, proposed that neutrinos undergo quantum-mechanical flavor oscillations as they propagate across space.

The Super-Kamiokande (Super-K) experiment provided the first definitive, statistically overwhelming evidence of this phenomenon in 1998. Super-K was designed as a massive, international water-Cherenkov observatory tasked with exploring the solar neutrino problem, searching for proton decay, and analyzing atmospheric cosmic-ray neutrino fluxes. Featuring roughly 15 times the water volume and ten times the number of PMTs as the original Kamiokande detector, Super-K accumulated statistics at a vastly accelerated rate.

Super-K began taking data in 1996. In 1998, the collaboration announced groundbreaking evidence for neutrino oscillation, primarily through the analysis of atmospheric neutrinos rather than solar neutrinos [58]. The data revealed a stark up-down asymmetry in the atmospheric muon neutrino (ν_μ) flux: the number of upward-going ν_μ was significantly smaller than the number of downward-going ν_μ.

These atmospheric neutrinos are generated when primary cosmic rays collide with the upper atmosphere, creating hadronic cascades that yield secondary pions and muons, which subsequently decay into neutrinos. Downward-going neutrinos are produced in the atmosphere directly above the detector and travel only a few tens of kilometers before reaching the tank. Conversely, upward-going neutrinos are generated in the atmosphere on the opposite side of the Earth and travel thousands of kilometers through the planet's interior. This pronounced deficit of upward-going ν_μ provided an unambiguous signature of flavor change over long baselines. Because the oscillation length of these multi-GeV neutrinos matched the diameter of the Earth, the long propagation distance gave the muon neutrinos sufficient time to oscillate into tau neutrinos (ν_τ). However, these

tau neutrinos typically would not have sufficient kinetic energy to produce a tau lepton in water, meaning they could not be directly detected by water-Cherenkov detectors.

Another milestone experiment, the Sudbury Neutrino Observatory (SNO), provided complementary and definitive confirmation in the solar sector. Located 2,100 meters underground in Sudbury, Ontario, Canada, the SNO detector utilized 1,000 tons of high-purity heavy water (D_2O) as its target medium [32]. Heavy water incorporates deuterium, a stable isotope of hydrogen containing one proton and one neutron. Because deuterium provides a loosely bound neutron target, it opens distinct interaction channels, specifically NC and CC interactions, giving the experiment a significantly higher sensitivity to non-electron neutrino flavors compared to light water. An array of 9,600 PMTs recorded the resulting Cherenkov light.

SNO provided definitive proof of solar neutrino flavor transformation in 2001 and 2002 [59]. By measuring both the electron-neutrino flux via CC interactions and the total active neutrino flux ($\nu_e + \nu_\mu + \nu_\tau$) via NC interactions, SNO demonstrated that while the electron neutrino flux was heavily suppressed, the total active neutrino flux matched the predictions of the standard solar model perfectly. Roughly two-thirds of the ^{8}B solar electron neutrinos had transformed into muon or tau flavors during transit.

The combined achievements of SNO and Super-K conclusively established that neutrino oscillation solves the historic solar neutrino problem. Recognizing these milestones, the 2002 Nobel Prize in Physics was awarded to Masatoshi Koshiba (shared with Raymond Davis Jr. and Riccardo Giacconi) for pioneering work in astrophysics and neutrino detection. Subsequently, the 2015 Nobel Prize in Physics honored Takaaki Kajita of Super-Kamiokande and Arthur B. McDonald of SNO for the historic discovery of neutrino oscillations, establishing that neutrinos possess non-zero mass.

2.5.4 Two-Flavor Oscillation Theory

The solar neutrino problem provided the first empirical hints of neutrino flavor transformation. While the original Standard Model did not predict neutrino oscillations, the confirmation of this

phenomenon, under the structural assumption that neutrinos were strictly massless, proved that neutrinos possess nonzero masses.

This quantum-mechanical mechanism can be modeled straightforwardly by considering the simplified case of two-flavor oscillation. Although physical neutrinos exist in three distinct flavor states, the two-flavor approximation remains a highly instructive framework for illustrating the core physics of neutrino mixing and its dependence on experimental parameters. Furthermore, this approximation yields highly accurate predictions in specific experimental scenarios where the oscillation dynamics are dominated by a single mass-squared splitting.

The phenomenon of neutrino flavor oscillation arises from a fundamental mismatch between the flavor eigenstates ($|\nu_e\rangle$ and $|\nu_\mu\rangle$), which participate in weak interactions, and the mass eigenstates ($|\nu_1\rangle$ and $|\nu_2\rangle$), which govern free-particle propagation. In the two-flavor approximation, these distinct bases are related by a unitary rotation matrix parameterized by a single mixing angle θ:

$$|\nu_e\rangle = \cos\theta|\nu_1\rangle + \sin\theta|\nu_2\rangle \tag{2.11}$$

$$|\nu_\mu\rangle = -\sin\theta|\nu_1\rangle + \cos\theta|\nu_2\rangle. \tag{2.12}$$

The time evolution of the mass eigenstates is governed by the time-dependent Schrödinger equation, yielding the state vectors $|\nu_i(t)\rangle = e^{-iE_it/\hbar}|\nu_i\rangle$ for $i \in \{1, 2\}$. Substituting this temporal dependence into the flavor basis expressions determines the time evolution for states that are initially produced in a pure flavor state:

$$|\nu_e(t)\rangle = \cos(\theta)e^{-\frac{iE_1t}{\hbar}}|\nu_1\rangle + \sin(\theta)e^{-\frac{iE_2t}{\hbar}}|\nu_2\rangle \tag{2.13}$$

$$|\nu_\mu(t)\rangle = -\sin(\theta)e^{-\frac{iE_1t}{\hbar}}|\nu_1\rangle + \cos(\theta)e^{-\frac{iE_2t}{\hbar}}|\nu_2\rangle. \tag{2.14}$$

The transition probability from an initial flavor state $|\nu_\alpha\rangle$ to a final flavor state $|\nu_\beta\rangle$ is defined by the amplitude $P(\nu_\alpha \to \nu_\beta) = |\langle\nu_\beta|\nu_\alpha(t)\rangle|^2$. Assuming ultrarelativistic neutrinos, the propagation time t can be approximated by the baseline distance L divided by the speed of light ($t \approx L/c$).

Using a Taylor expansion, the energy difference between the mass states can be approximated as:

$$E_2 - E_1 \approx \frac{\Delta m^2 c^3}{2E}, \tag{2.15}$$

where $\Delta m^2 = m_2^2 - m_1^2$ represents the mass-squared splitting and E is the average neutrino energy. Evaluating the resulting inner products yields the flavor conversion, or appearance, probability:

$$P(\nu_e \to \nu_\mu) = P(\nu_\mu \to \nu_e) = \sin^2(2\theta) \sin^2\left(\frac{\Delta m^2 c^3 L}{4\hbar E}\right). \tag{2.16}$$

By invoking the conservation of total probability, the corresponding survival, or disappearance, probabilities are expressed as:

$$P(\nu_e \to \nu_e) = P(\nu_\mu \to \nu_\mu) = 1 - \sin^2(2\theta) \sin^2\left(\frac{\Delta m^2 c^3 L}{4\hbar E}\right). \tag{2.17}$$

To evaluate these physical quantities using practical experimental parameters where Δm^2 is given in eV^2, the baseline distance L in km, and the neutrino energy E in GeV, the natural units must be explicitly converted to SI dimensions. This yields the standard operational formulas for neutrino flavor appearance and disappearance:

$$P(\nu_\mu \to \nu_e) = \sin^2(2\theta) \sin^2\left(1.27 \cdot \frac{\Delta m^2 \cdot L}{E}\right) \tag{2.18}$$

$$P(\nu_\mu \to \nu_\mu) = 1 - \sin^2(2\theta) \sin^2\left(1.27 \cdot \frac{\Delta m^2 \cdot L}{E}\right), \tag{2.19}$$

where the coefficient 1.27 absorbs the numerical conversion factor $\frac{c^3}{4\hbar} \approx 1.26693 \times 10^9$ $\text{eV}^{-2}\text{km}^{-1}\text{GeV}$.

Physically, $P(\nu_\mu \to \nu_e)$ represents the fraction of an initial ν_μ population that oscillates into the ν_e flavor over a propagation distance L. Crucially, the mathematical formulation dictates that if the mass eigenstates are degenerate or zero ($\Delta m^2 = 0$), the oscillation probability vanishes completely. The discovery of neutrino oscillation therefore provided definitive proof that at least two neutrino mass eigenstates are non-zero, requiring a structural extension of the original Standard

Model to accommodate neutrino mass terms and mixing matrices.

As shown in Equation 2.20, the maximum oscillation amplitude is achieved when the kinematic argument of the spatial oscillatory term equals $\pi/2$:

$$1.27\Delta m^2 \frac{L\,[\text{km}]}{E\,[\text{GeV}]} = \frac{\pi}{2}. \tag{2.20}$$

Because Δm^2 and θ are fixed constants of nature, the baseline distance L and the beam energy E are the primary parameters that researchers can adjust to tune the oscillation probability. For instance, an experiment generating a pure ν_μ beam can look for the appearance of ν_e or the disappearance of un-oscillated ν_μ at a fixed distance L from the source. If an experiment is designed with sub-optimal parameters, the resulting oscillation probability may be too small to detect over environmental and systematic backgrounds. Such low statistics would impose prohibitively strict requirements on detector volume and spatial resolution. Consequently, experimental designs strategically optimize the ratio L/E to maximize sensitivity to the underlying oscillation parameters.

2.5.5 Three-Flavor Oscillation Theory

As established previously, physical neutrinos exist in three distinct flavor generations. Consequently, the simplified two-flavor approximation cannot provide a complete picture of the neutrino oscillation phenomenon, nor can it accurately model flavor transitions across all experimental conditions. The three neutrino flavor eigenstates ($|\nu_e\rangle$, $|\nu_\mu\rangle$, $|\nu_\tau\rangle$) are mathematically described as linear combinations of the three discrete mass eigenstates ($|\nu_1\rangle$, $|\nu_2\rangle$, $|\nu_3\rangle$):

$$\begin{pmatrix} |\nu_e\rangle \\ |\nu_\mu\rangle \\ |\nu_\tau\rangle \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{pmatrix} \begin{pmatrix} |\nu_1\rangle \\ |\nu_2\rangle \\ |\nu_3\rangle \end{pmatrix}. \tag{2.21}$$

This 3×3 unitary mixing matrix is known as the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix. It is conventionally parameterized by three distinct mixing angles (θ_{12}, θ_{13}, and θ_{23}) alongside a complex phase factor (δ_{CP}):

$$
\begin{aligned}
U &= \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} \\
&= \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_{\text{CP}}} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta_{\text{CP}}} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta_{\text{CP}}} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta_{\text{CP}}} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta_{\text{CP}}} & c_{23}c_{13} \end{pmatrix},
\end{aligned}
\tag{2.22}
$$

where $s_{ij} = \sin(\theta_{ij})$ and $c_{ij} = \cos(\theta_{ij})$, and δ_{CP} represents the Dirac Charge-Parity (CP) violating phase.

By applying this full three-generation framework, the exact quantum-mechanical transition probability from an initial flavor state $|\nu_\alpha\rangle$ to a final flavor state $|\nu_\beta\rangle$ can be written as:

$$
\begin{aligned}
P(\nu_\alpha \to \nu_\beta) = \delta_{\alpha\beta} &- 4\sum_{i>j} \text{Re}\big(U_{\alpha i}U^*_{\beta i}U^*_{\alpha j}U_{\beta j}\big)\sin^2\left(\frac{\Delta m^2_{ij}\,L}{4E}\right) \\
&+ 2\sum_{i>j} \text{Im}\big(U_{\alpha i}U^*_{\beta i}U^*_{\alpha j}U_{\beta j}\big)\sin\left(\frac{\Delta m^2_{ij}\,L}{2E}\right),
\end{aligned}
\tag{2.23}
$$

where $\Delta m^2_{ij} = m^2_i - m^2_j$ represents the respective mass-squared differences. This complete framework demonstrates that the exact probability of flavor conversion depends fundamentally on a combination of fundamental constants, specifically the three mixing angles, the mass-squared differences, and the CP-violating phase, interfering as a function of the experimental baseline distance L and the incident neutrino energy E.

2.5.6 Measurements of Neutrino Oscillation Parameters

Because absolute neutrino masses are remarkably small, their direct laboratory measurement presents an extraordinary experimental challenge. However, calculating quantum-mechanical neutrino oscillation probabilities requires only knowledge of the relative mass-squared differences (Δm^2_{ij}) and the mixing angles (θ_{ij}). Over the past several decades, a combination of atmospheric,

solar, reactor, and accelerator-based neutrino experiments has successfully constrained these parameters.

The Kamioka Liquid Scintillator Antineutrino Detector (KamLAND) was an underground observatory situated at the Kamioka Observatory in Japan designed to detect electron antineutrinos ($\bar{\nu}_e$) emitted by surrounding commercial nuclear reactors. KamLAND investigated the disappearance of $\bar{\nu}_e$ over an average flux-weighted baseline of approximately 180 km. Utilizing data collected over a 515-day initial operational period, the experiment recorded a significant deficit in the expected reactor antineutrino flux, a result determined to be highly consistent with two-flavor neutrino oscillation [34]. In 2008, the collaboration released a highly precise, combined global measurement of the solar mass-squared splitting:

$$\Delta m^2_{21} = (7.59 \pm 0.21) \times 10^{-5}\ \mathrm{eV}^2. \tag{2.24}$$

Because this value is defined as positive, it dictates that the mass eigenstate m_2 is heavier than m_1.

Complementary measurements from the Super-Kamiokande collaboration established the absolute scale of the atmospheric mass-squared splitting:

$$|\Delta m^2_{31}| = (2.44 \pm 0.06) \times 10^{-3}\ \mathrm{eV}^2. \tag{2.25}$$

Notably, the mass-squared difference between the m_1 and m_3 states is roughly 30 times larger than the splitting separating the m_1 and m_2 states. However, determining the explicit sign of Δm^2_{31} remains an active and vital frontier of particle physics research.

The leptonic mixing angles governing the PMNS matrix have also been tightly constrained by global data [16]. The reactor mixing angle, θ_{13}, is measured to be:

$$\sin^2(2\theta_{13}) = 0.093 \pm 0.008, \tag{2.26}$$

while the solar mixing angle, θ_{12}, is constrained to:

$$\sin^2(2\theta_{12}) = 0.846 \pm 0.021. \tag{2.27}$$

The value of $\sin^2(2\theta_{13})$ was definitively established by short-baseline reactor experiments, including the Daya Bay, RENO, and Double Chooz collaborations. Conversely, $\sin^2(2\theta_{12})$ was resolved by combining KamLAND reactor data with solar and long-baseline accelerator neutrino datasets.

The remaining mixing angle, θ_{23}, is known to be relatively close to maximal mixing ($\approx 45^\circ$). However, current data has not yet definitively resolved its octant, specifically, whether θ_{23} is slightly less than (lower octant) or slightly greater than (upper octant) 45°. Determining whether θ_{23} is exactly maximal or behaves asymmetric between flavor couplings has profound theoretical implications, as it directly influences Grand Unified Theories (GUTs) and underlying models of neutrino mass generation.

The value of the Dirac CP-violating phase, δ_{CP}, also remains an unresolved parameter. Pinning down this phase factor will determine whether CP symmetry is violated within the lepton sector. Measuring a non-zero leptonic CP violation is highly sought after, as it may provide the necessary theoretical insights to explain the primordial matter-antimatter asymmetry of the universe.

Furthermore, resolving the sign of Δm^2_{31} is required to settle the neutrino mass ordering. Two potential scenarios exist: the normal ordering ($m_1 < m_2 < m_3$) and the inverted ordering ($m_3 < m_1 < m_2$). The former scenario is designated as "normal" because it mirrors the hierarchical mass patterns observed within the charged lepton and quark sectors of the Standard Model. Although global neutrino fits place statistical bounds on these mass ordering options, a definitive, high-significance discovery has not yet been achieved. Next-generation long-baseline and large-scale experiments are engineered to definitively measure these remaining unknowns, including DUNE in the United States [9], Hyper-Kamiokande in Japan [8], and JUNO in China [7].

Chapter 3

The Deep Underground Neutrino Experiment (DUNE)

3.1 Overview of DUNE

DUNE is a flagship project with the primary goal of measuring key neutrino oscillation parameters [9]. The Long Baseline Neutrino Facility (LBNF) will generate an intense beam of ν_μ and $\bar{\nu}_\mu$ (at separate times, depending on the beam polarity) across a long baseline to allow neutrinos to oscillate from one flavor to another. Figure 3.1 illustrates the key elements of the LBNF/DUNE project.

The neutrino beam will be produced at Fermilab near Chicago, Illinois. To measure the initial fraction of ν_μ ($\bar{\nu}_\mu$) relative to ν_e ($\bar{\nu}_e$) prior to oscillations, a Near Detector (ND) will be employed. The ND complex will house three distinct yet complementary detectors, which are collectively tasked with measuring the neutrino beam flux, energy, and flavor composition at a distance of 574 m from the beam source [65]. Consequently, the ND plays a crucial role in constraining systematic uncertainties in oscillation measurements.

To enable neutrino oscillation measurements, the information collected by the ND must be compared to data collected by the Far Detector (FD). The FD will be located approximately 1.6 km (1 mile) underground at the Sanford Underground Research Facility (SURF) in Lead, South Dakota, minimizing cosmic-ray backgrounds. Situated about 1300 km (800 miles) from the neutrino beam source to allow ample distance for neutrino oscillation, the detector will comprise the largest LArTPCs ever built. It will contain a total of 40 kt of fiducial LAr (70 kt total mass) spread across four independent cryostats, each with internal dimensions of 18.9 m (W) $\times$ 17.8 m (H) $\times$ 65.8 m (L). Figure 3.2 illustrates the layout of the underground detector caverns at SURF which shows the first two FD cryostats in place.

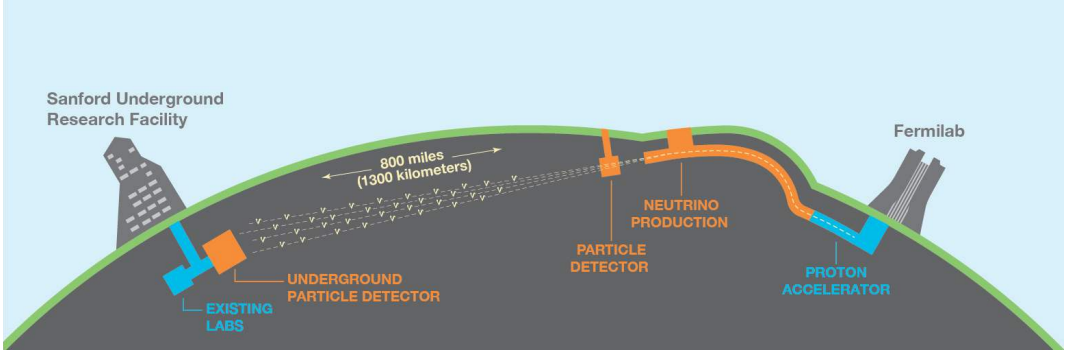

Figure 3.1: Illustration of the LBNF/DUNE Project. A neutrino beam is generated at Fermilab near Chicago, Illinois, and is directed underground toward SURF in Lead, South Dakota. Image retrieved from Ref. [9].

The execution of the DUNE physics program is structured as a dual-phase deployment strategy to systematically scale its experimental sensitivities while maximizing early physics returns. In Phase I, the experiment establishes its long-baseline neutrino framework, utilizing the LBNF wide-band neutrino beam operating at a proton beam power of 1.2 MW. This phase deploys the first two 10 kT fiducial-mass LArTPC FD modules underground at the SURF, specifically the vertical-drift (FD1) and horizontal-drift (FD2) modules, complemented by a suite of ND subsystems at Fermilab. To reach the full statistical and systematic precision required to definitively resolve leptonic CP violation and execute high-sensitivity searches for baryon number violation (BNV), DUNE will transition to Phase II. This definitive phase elevates the LBNF beam power to a high-intensity 2.4 MW configuration via the Fermilab Proton Improvement Plan II (PIP-II) upgrades. Concurrently, Phase II doubles the FD fiducial mass by constructing two additional FD modules (FD3 and FD4), while fully upgrading the ND complex including the deployment of the specialized ND-GAr gaseous argon TPC.

Both the FD and ND will utilize LArTPC technology as their primary method for measuring neutrino interactions. Utilizing the same technology across both detectors, and specifically the same neutrino interaction medium of liquid argon, significantly reduces systematic uncertainties in the final measurements. Furthermore, LArTPCs provide fine-grained imaging of charged particles resulting from neutrino interactions within the detector, alongside excellent calorimetry performance necessary for robust particle energy measurements and particle identification. The efficacy

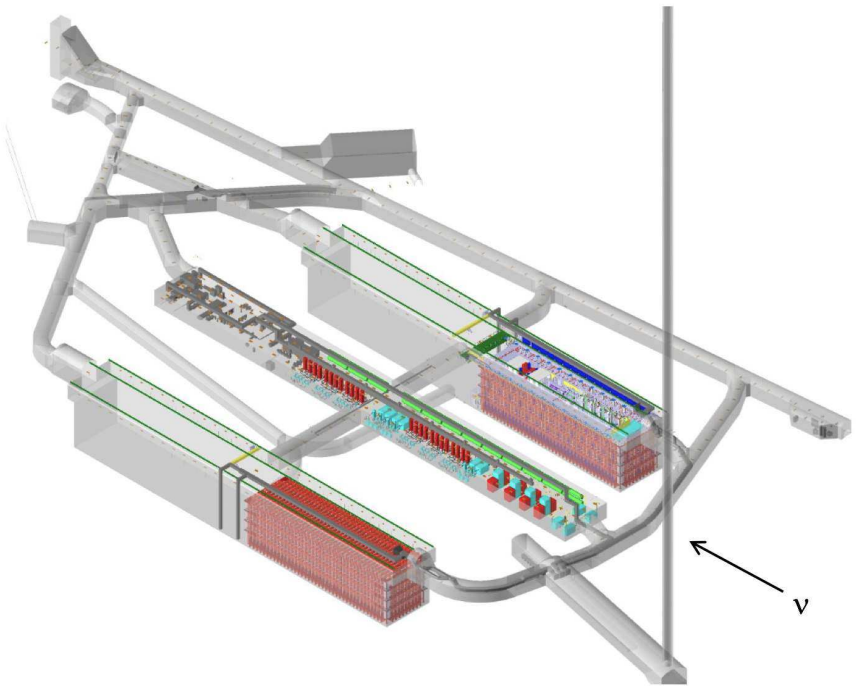

Figure 3.2: Schematic layout of the DUNE FD underground caverns and infrastructure at SURF in Phase I.

of this technology has been successfully demonstrated by experiments such as ICARUS [66] and MicroBooNE [67], as well as DUNE itself with the successful operation of FD LArTPC prototypes [19, 20].

3.2 DUNE's Neutrino Oscillation Goals

One of the primary scientific objectives of DUNE is to search for CP violation in the neutrino sector and to precisely measure the CP-violating phase, δ_{CP}. Specifically, DUNE is designed to achieve a 3σ statistical significance for discovering CP violation, thereby excluding the null hypothesis of no CP violation, over approximately 75% of all possible δ_{CP} values after an initial running period. Ultimately, the experiment will accumulate enough exposure to provide a definitive $> 5\sigma$ discovery over a broad range of values [9, 68]. Figure 3.3 illustrates this CP violation sensitivity as a function of calendar years and exposure milestones. If experiments confirm that $\sin(\delta_{\text{CP}}) \neq 0$, it will establish the existence of CP violation in the lepton sector. Figure 3.3 shows the DUNE sensitivity to a CP violation measurement as a function of calendar years of data-taking.

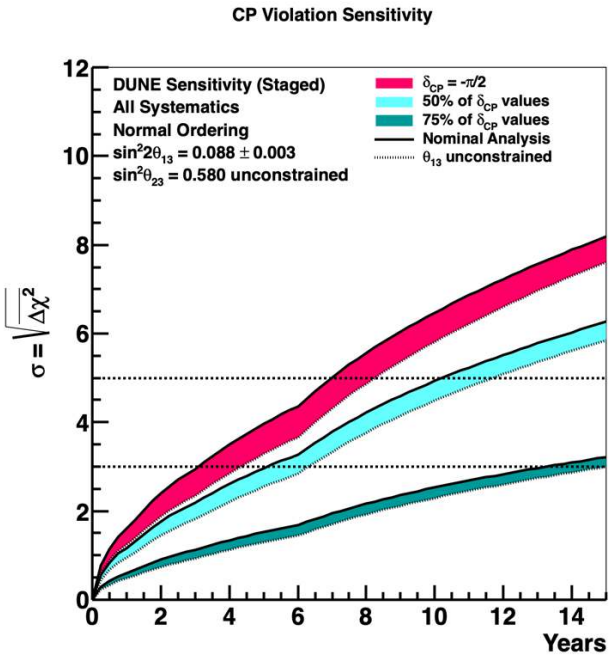

Figure 3.3: Sensitivity of a CP violation measurement at DUNE as a function of time (represented as calendar years of data-taking). Shown are the cases of maximal true CP violation ($\delta_{\mathrm{CP}} = -\pi/2$, red), coverage of 50% of all true δ_{CP} values (cyan), and coverage of 75% of all true δ_{CP} values (turquoise). True normal ordering of the neutrino mass states is assumed in these projections.

There is far more matter than antimatter in the universe, implying that a mechanism must have tipped the scales toward matter in the early universe. A prominent theoretical framework called leptogenesis [69] proposes a process that would have generated an asymmetry in the leptons in the early universe, which could subsequently produce an asymmetry in the number of baryons (baryogenesis). While CP violation has already been discovered in quarks, it has been determined to be far too small to explain the generation of the observed excess of baryons over antibaryons [56]. Consequently, leptogenesis remains the favored theory for explaining this matter-antimatter asymmetry.

To allow for leptogenesis, a set of requirements known as the Sakharov conditions must be satisfied, as identified by Andrei Sakharov in 1967 [70]. Because one of these conditions mandates the presence of CP violation in the lepton sector, discovering a non-zero value of δ_{CP} in the neutrino sector would satisfy a fundamental requirement for early-universe leptogenesis [69]. The remaining Sakharov conditions include BNV, C-symmetry violation, and a departure from

thermal equilibrium to "freeze in" the matter-antimatter asymmetry before it can be undone. The discovery of BNV would prove that matter is ultimately unstable and would validate GUTs that seek to unify the strong, weak, and electromagnetic fundamental forces. Accordingly, another key goal of DUNE is to search for BNV, such as proton decay and neutron-antineutron oscillation, by taking advantage of the immense size of its detectors, a low-background environment, and the high-resolution LArTPC technology [9]. BNV will be discussed further in the context of DUNE in Section 3.3.1.

DUNE also aims to resolve the neutrino mass ordering by definitively establishing whether the neutrino masses follow a normal ($m_1 < m_2 < m_3$) or inverted ($m_3 < m_1 < m_2$) hierarchy. It will simultaneously determine the octant of the atmospheric mixing angle θ_{23} by measuring whether the mixing is maximal ($\theta_{23} = 45°$), or if it lies in the first octant ($\theta_{23} < 45°$) or second octant ($\theta_{23} > 45°$). To realize these goals, reducing statistical uncertainties by collecting more data is simply not enough; it is equally critical to significantly minimize overall systematic uncertainties. This requirement highlights one of the key roles of the ND complex within the DUNE neutrino oscillation measurement program.

3.3 Additional Goals of DUNE

3.3.1 Search for Baryon Number Violation

Baryon number B is a quantity that has historically been understood to be conserved in nature. However, the extreme excess of matter over antimatter in the universe implies that it must be violated through early-universe processes. BNV is a hypothetical process where the total number of baryons is not conserved, and multiple potential decay or oscillation modes have been theorized [71, 72]. The observation of BNV would be a monumental discovery, providing critical support for GUTs and satisfying one of the fundamental Sakharov conditions required for leptogenesis to have occurred in the early universe.

While BNV has not yet been observed experimentally, current limits on specific BNV modes, such as the proton decay lifetime, have eliminated or highly constrained several theoretical models

predicting the phenomenon [9, 73, 74]. DUNE will be sensitive to multiple modes of BNV, including the phenomenon of neutron-antineutron (n–$\bar{n}$) mixing alongside the specific decay modes of protons and neutrons [9]. To increase the likelihood of observing a BNV process, DUNE features a very large sensitive mass (40 kt of active LAr) and is projected for a long exposure timeline spanning over a decade. Furthermore, because the DUNE FD will be located deep underground, it benefits from a low-background environment uniquely suited for rare-event BNV searches.

DUNE is uniquely poised to search for BNV due to its exceptional sensitivity to proton decay modes yielding kaons or electrons in the final state. In water Cherenkov detectors, final-state kaons typically travel below the Cherenkov threshold, making them difficult to detect. Conversely, they can be readily identified in a LArTPC detector owing to its ability to accurately track particles and identify their species via their unique energy-deposition signature per unit length (dE/dx). Consequently, DUNE is highly sensitive to kaon-producing channels, such as the decay $p \to K^+\bar{\nu}$.

Simulation studies demonstrate that the overall tracking efficiency, defined as how often a simulated particle is successfully associated with a reconstructed track in the detector, is 58.0% for positive kaons (K^+) with kinetic energies up to 200 MeV [9]. For kaons above approximately 100 MeV, this efficiency exceeds 80%, while the tracking efficiency for the daughter muons resulting from K^+ decay reaches 90%. These benchmarks indicate that DUNE can efficiently reconstruct K^+ particles and their subsequent decay products, yielding excellent sensitivity to this channel. Cosmogenic backgrounds can be further suppressed by enforcing a fiducial volume cut that requires no activity near the boundaries of the LArTPC, leaving atmospheric neutrinos as the dominant remaining background source.

In contrast, stringent lower limits on the decay lifetime of protons into electrons via the channel $p \to e^+\pi^0$ have already been established by Super-Kamiokande [73]. Given these tight pre-existing constraints, sensitivity studies for DUNE have primarily focused on the kaon channels where the experiment holds a distinct technological advantage. Nevertheless, DUNE remains sensitive to electron-channel decays due to the high-fidelity particle reconstruction capabilities provided by LArTPCs, enabling it to independently constrain the proton lifetime.

Sensitivity studies indicate that a DUNE exposure of 400 kt·years, corresponding to a 40 kt active LAr mass over ten years of data-taking, would yield a lifetime sensitivity of approximately 10^{34} years for this channel. Doubling this exposure could yield a sensitivity comparable to the current limits set by Super-Kamiokande (1.6×10^{34} years) [9]. Additional studies have evaluated DUNE's sensitivity to the decay of a bound neutron into a charged lepton and a meson, specifically $n \to e^+ K^+$. The projected sensitivity for a 400 kt $\cdot$ years exposure is 1.1×10^{34} years, representing an improvement of two orders of magnitude over the current best global estimate. Finally, detailed studies have been performed regarding DUNE's sensitivity to neutron-antineutron mixing [9]. This process is identified by detecting the subsequent annihilation of the antineutron, which produces a distinct multi-pion vertex signature of several emitted light hadrons. For a 400 kt $\cdot$ years exposure, DUNE will be sensitive to a lower limit on the bound neutron lifetime of 6.45×10^{32} years.

In summary, DUNE will conduct a multi-channel search for BNV, placing stringent lower limits on the lifetimes of neutron-antineutron mixing and nucleon decays, thereby constraining or ruling out various theoretical GUT frameworks.

3.3.2 Detection of Supernova Neutrinos

Beyond its accelerator-based objectives, DUNE will serve as a highly sensitive observatory for low-energy astrophysical neutrinos originating from core-collapse supernovae within our galaxy and its satellites [9, 75]. Detecting a Core-Collapse Supernova Neutrino Burst (SNB) provides a real-time probe into the extreme environment of stellar collapse, yielding data spanning a several-tens-of-seconds-long profile across all neutrino flavors [75]. Scientific interest in these events has flourished since Supernova 1987A (SN1987A) in the Large Magellanic Cloud, where the detection of a few tens of $\bar{\nu}_e$ events provided baseline validation for the core-collapse model and set foundational constraints on exotic physics [76, 77]. Capturing high-statistics neutrino profiles from a future nearby supernova is expected to offer unprecedented insights into the explosion mechanism, nucleosynthesis, shock-wave propagation, and the nature of the compact remnant, such as a neutron star or black hole [75].

Given the rarity of these events, occurring roughly twice per century in the Milky Way, it is vital that the experimental physics community maintains continuous, multi-flavor sensitivity to capture a complete burst history [75]. While next-generation water Cherenkov and liquid scintillator detectors, such as Hyper-Kamiokande and JUNO, are exceptionally sensitive to the $\bar{\nu}_e$ flux through inverse β decay, DUNE offers a uniquely complementary capability. Owing to its LAr medium, DUNE's dominant interaction channel is electron neutrino charged-current absorption on argon nuclei:

$$\nu_e + {}^{40}\mathrm{Ar} \rightarrow e^- + {}^{40}\mathrm{K}^*, \tag{3.1}$$

which is accompanied by distinct nuclear de-excitation γ ray signatures [75]. This makes DUNE uniquely sensitive to the ν_e component of the flux, allowing it to probe critical early-time features like the initial "neutronization burst" that occurs within the first 50 ms of the core collapse [75].

For a standard core-collapse event at a benchmark distance of 10 kpc (just beyond the Galactic center), the 40 kt DUNE FD is projected to record from several hundred to several thousand total events, completely dominated by the ν_e charged-current channel [75]. Conversely, subdominant channels will also provide essential cross-checks: these include anti-electron neutrino charged-current absorption ($\bar{\nu}_e + {}^{40}\mathrm{Ar} \rightarrow e^+ + {}^{40}\mathrm{Cl}^*$) and neutral-current elastic scattering on electrons ($\nu_x + e^- \rightarrow \nu_x + e^-$) across all active flavors [75]. DUNE's fine-grained LAr calorimetry can resolve these events over a dynamic energy range down to approximately 5 MeV [75]. This high-resolution sensitivity allows DUNE to track time-dependent flavor transitions caused by matter-induced Mikheyev-Smirnov-Wolfenstein (MSW) or collective neutrino-neutrino oscillation effects, or trace the formation of a black hole via a sudden termination of the neutrino flux [75].

To ensure that no galactic burst is missed, DUNE employs highly efficient triggering systems utilizing real-time data selection chains across both charge and light collection systems [75]. By exploiting the tight temporal coincidence of multiple low-energy events against dominant radiological backgrounds (such as the β decay of ^{39}Ar), an energy-weighted multiplicity trigger scheme can achieve nearly 100% trigger efficiency out to the edge of the Milky Way ($\sim$20 kpc) while keeping fake trigger rates below one per month [75]. This system is robust enough to maintain

a $\sim$70% burst detection efficiency even for more distant events yielding as few as 10 total interactions, such as a core collapse inside the Large Magellanic Cloud [75]. DUNE will feed these real-time burst metrics into the upgraded SuperNova Early Warning System (SNEWS 2.0) network to guide global multi-messenger follow-up campaigns [75].

DUNE also possesses vital standalone directional reconstruction capabilities to facilitate prompt multi-messenger astronomy alerts. While standard absorption interactions lose directional information due to quasi-isotropy, neutrino-electron elastic scattering preserves the forward trajectory of the incoming neutrino [75]. Utilizing advanced tracking algorithms alongside techniques such as "brems flipping" to reconstruct primary electron tracks and secondary bremsstrahlung photons, DUNE can cleanly isolate these elastic scattering events. For a standard supernova localized at a distance of 10 kpc, the DUNE FD can achieve a standalone pointing resolution of approximately 4.3° at 68% coverage under realistic interaction-channel misidentification conditions [78].

Finally, event timing in DUNE provides an exceptional handle on fast stellar dynamics. During the intense first 50 ms of a 10 kpc supernova, the mean time interval between successive neutrino interactions spans merely 0.5 to 1.7 ms [75]. While the TPC drift time at an electric field strength of 500 V/cm provides a time resolution of 0.6 ms, scintillation light detection yields sub-microsecond timing resolution (< 1 µs) [75]. This fine temporal precision guarantees that DUNE's measurement of the burst profile is bounded solely by statistical arrival rates rather than hardware limitations [75]. This capability allows the experiment to readily resolve fast astrophysical phenomena, such as the 10 ms oscillations driven by Standing Accretion Shock Instabilities (SASI) or the narrow 1–2 ms neutrino "trapping notch" that immediately precedes the core bounce [75]. When combined with data from water Cherenkov and liquid scintillator facilities, DUNE's unique ν_e dataset will prove instrumental in resolving the global flavor composition and pinning down neutrino mass ordering effects via matter-induced resonant conversions deep within the stellar envelope.

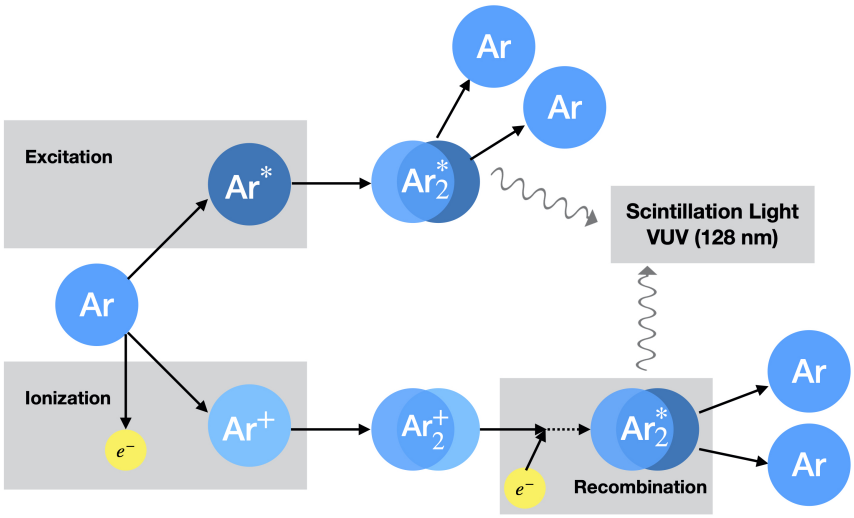

Figure 3.4: Schematic representation of the microphysical pathways governing scintillation and ionization processes following particle interactions in LAr. The dashed vector represents omitted intermediate non-radiative relaxation steps.

3.4 Liquid Argon Time Projection Chambers (LArTPCs)

3.4.1 Microphysics of Scintillation and Ionization in LAr

The operation of a LArTPC relies on two primary microphysical mechanisms initiated by ionizing radiation: atomic excitation and atomic ionization. Both pathways ultimately generate VUV scintillation photons characterized by a spectrum narrow-peaked at 128 nm [79], but only the ionization pathway yields the free electrons that form the primary TPC tracking signals. These fundamental initial states are classified as follows:

1. **Direct Excitation:** An incident charged particle excites an argon atom, forming a neutral excited state (Ar*).

2. **Ionization:** An incident charged particle ionizes an argon atom, producing a free electron and a positive ion pair ($\mathrm{Ar}^{+} + e^{-}$).

The specific evolutionary pathways of these initial states are illustrated schematically in Figure 3.4. Both mechanisms converge on the synthesis of an excited molecular dimer, or excimer state, denoted as $\mathrm{Ar}_2{}^{*}$, which acts as the precursor to scintillation emission [80].

The direct excitation pathway involves relatively few intermediate transitions. Following the formation of Ar^*, the excited atom rapidly binds with a neighboring ground-state argon atom to form the excimer:

$$\mathrm{Ar}^* + \mathrm{Ar} \rightarrow \mathrm{Ar_2}^*. \tag{3.2}$$

The excimer then undergoes radiative de-excitation to the dissociative ground state [81, 82]:

$$\mathrm{Ar_2}^* \rightarrow 2\mathrm{Ar} + \gamma, \tag{3.3}$$

where γ represents the VUV scintillation photon.

The reaction kinematics following primary ionization encompass a broader cascade of intermediate processes. The initial ionization step is expressed as:

$$\mathrm{Ar} \rightarrow \mathrm{Ar}^+ + e^-. \tag{3.4}$$

The mean energy expended to produce a single electron-ion pair in LAr (commonly denoted as the W-value) is experimentally characterized at 23.6 eV [83]. It should be noted that while individual ionization events can occur down to the first ionization potential of argon ($\sim$ 15.76 eV), the higher average W-value accounts for competing energy allocation into atomic excitation and sub-threshold electronic fluctuations. Following ionization, the positive argon ion undergoes rapid self-trapping with a neighboring ground-state atom to form a localized molecular ion:

$$\mathrm{Ar}^+ + \mathrm{Ar} \rightarrow \mathrm{Ar_2}^+. \tag{3.5}$$

In the absence of an external electric field, the thermalized ionization electron is highly likely to undergo recombination with the molecular ion core. If recombination occurs, it initiates a dissociative recombination sequence that culminates in the production of the same $\mathrm{Ar_2}^*$ excimer state populated by direct excitation:

$$\mathrm{Ar_2}^{+} + e^{-} \rightarrow \mathrm{Ar}^{**} + \mathrm{Ar} \tag{3.6}$$

$$\mathrm{Ar}^{**} \rightarrow \mathrm{Ar}^{*} \tag{3.7}$$

$$\mathrm{Ar}^{*} + \mathrm{Ar} \rightarrow \mathrm{Ar_2}^{*}, \tag{3.8}$$

where Ar^{**} represents a highly excited atomic state that undergoes non-radiative relaxation to the lowest excited atomic state, Ar^{*}.

Analogous to the direct excitation pathway, the resulting $\mathrm{Ar_2}^{*}$ excimer returns to the dissociative ground state via Equation 3.3, releasing a 128 nm photon. Under zero-field conditions, approximately 57% of the total observable scintillation light is a product of this electron-ion recombination cascade, while the remaining 43% is driven by direct atomic excitation [80].

3.4.2 Operating Principles and Field Dependence

The LArTPC framework is designed to simultaneously exploit these dual signatures, charge transport and scintillation light, to achieve sub-millimeter tracking precision and high-fidelity calorimetric energy reconstruction. The fundamental operating principle of a standard single-phase LArTPC is illustrated in Figure 3.5. An applied, uniform electric field drifts the ionization cloud along linear trajectories toward a segmented anode plane where the charge is recorded.

The complementary nature of ionization charge and scintillation light yields as a function of this external electric field was definitively investigated by Kubota et al. [80]. Their measurements demonstrated that as the magnitude of the applied electric field increases, the total collected ionization charge increases monotonically, while the complementary scintillation light intensity decreases. This anti-correlation persists up to an electric field threshold of approximately 8–10 kV/cm, beyond which both the charge and light yields plateau. This asymptotic behavior signifies that at sufficiently high field strengths, electron-ion recombination is entirely suppressed. The persistence of a baseline scintillation signal in this high-field regime provides direct empirical confirmation that a significant fraction of the total light yield is generated via independent non-recombining pathways, specifically direct atomic excitation.

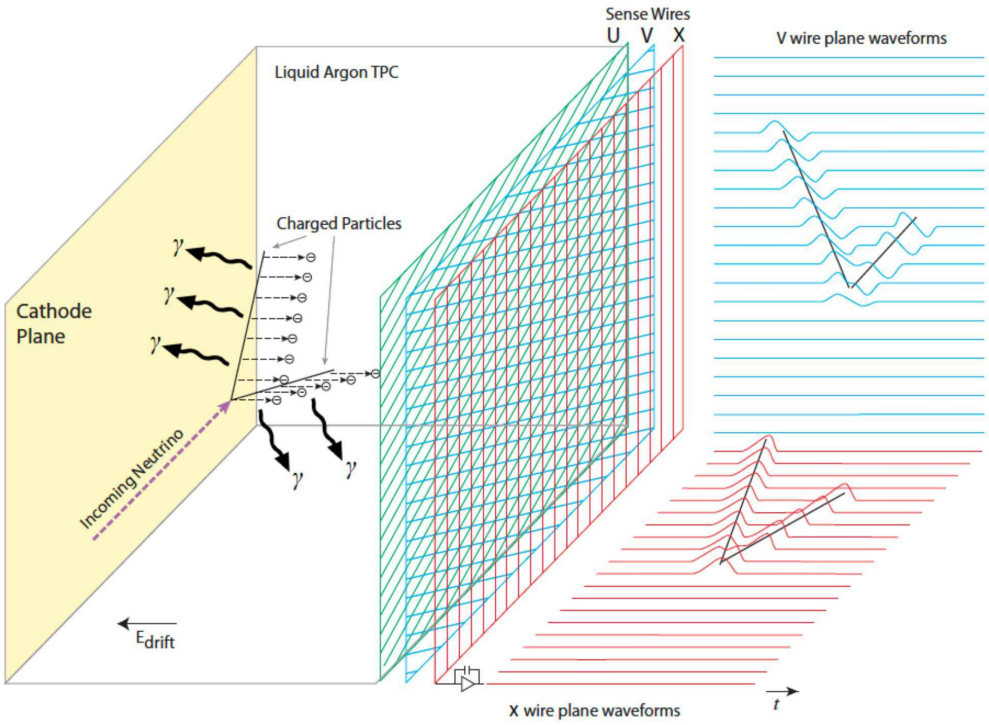

Figure 3.5: Illustration of a single-phase horizontal-drift LArTPC. Ionization electrons produced along a charged particle track are driven through a uniform electric field toward an anode wire plane for signal readout. Image retrieved from Ref. [9].

For instance, minimum-ionizing particles (MIPs), such as cosmic-ray muons, deposit approximately 55,000 ionization electrons per centimeter in a LArTPC operating at a nominal drift electric field strength of 500 V/cm. Concurrently, this optical pathway yields approximately 24,000 photons per MeV of deposited energy for a MIP at the same field strength [17]. Because the prompt scintillation light is detected on a nanosecond timescale, it provides the precise interaction timestamp (t_0) required to resolve the spatial coordinate along the drift axis.

Originally proposed for neutrino detection by Carlo Rubbia in 1977 [84], LArTPC technology has achieved operational maturity through generations of experiments, including ICARUS [66], MicroBooNE [67], the Short-Baseline Neutrino (SBN) Program [85], and DUNE [20]. Decades of R&D have yielded substantial advancements in industrial cryogenic scaling, multi-ton argon purification, ultra-low-noise cold electronics, and highly efficient VUV photon detection systems [86].

3.4.3 Electrostatic Configuration and Anode Readout Architecture

Establishing a highly uniform drift electric field across the active volume requires biasing parallel cathode and anode planes, structurally managed by a surrounding field cage. The field cage typically consists of equally spaced metallic profiles that envelop the active fiducial volume and span the drift length. Interconnected via a precision resistor chain, these electrodes drop the high-voltage potential linearly from the cathode to the anode, imposing a homogeneous electrostatic field that guarantees parallel, non-distorted electron drift trajectories.

Upon reaching the anode plane, the arriving ionization electrons induce currents or terminate on specialized readout electrodes to register the signal. In wire-based TPCs, the anode plane assembly historically consists of three parallel wire planes positioned in close proximity. The first two layers, designated as the U and V induction planes, are biased with potentials that allow drifting electrons to pass through the wire grids without being captured. As the electron cloud approaches and passes these wires, it generates a characteristic bipolar induced current waveform. The final layer, designated as the collection (W) plane, is biased with an attractive potential that forces the drifting electrons to terminate directly on the wires, producing a unipolar current pulse whose integrated area is strictly proportional to the total collected ionization charge. Modern pixelated architectures replace these wire grids with continuous arrays of charge-collecting pixel anodes, offering unambiguous 2D spatial localization at the anode interface.

3.4.4 Three-Dimensional Spatial Reconstruction

The primary advantage of the LArTPC is its imaging capability, providing fully three-dimensional topological representations of complex particle interactions. Transverse to the electron drift direction, spatial coordinates are fixed by the geometrical position of the active wire or pixel channel at the anode. In wire-based systems, a single wire orientation resolves only a two-dimensional projection of the event trajectory; therefore, signals from multiple overlapping induction and collection planes, oriented at relative stereo angles, must be mathematically deconvolved to extract intersecting 2D wire positions.

Resolving the third spatial coordinate along the longitudinal drift axis relies on a precise time-of-flight mapping of the drifted ionization cloud. The physical drift distance traveled by the ionization electrons, Δx, is evaluated analytically via:

$$\Delta x = v_{\text{drift}} \cdot (t - t_0), \tag{3.9}$$

where v_{drift} is the electron drift velocity dictated by the local electric-field magnitude, t is the absolute timestamp recorded by the front-end digitizers when the signal registers at the anode, and t_0 is the true absolute interaction time.

While the charge readout electronics continuously record the arrival time t, the TPC is intrinsically blind to the absolute event time t_0. Consequently, the spatial position along the drift axis remains undetermined without an independent determination of this translation offset. For accelerator-based neutrino beams, t_0 can be synchronized with the external beam-spill timing signals provided by the accelerator master clock. For non-beam or cosmic-ray events, t_0 is derived from the detection of prompt scintillation light using high-speed optical systems, such as PMTs or SiPM-based detectors. Because VUV scintillation photons propagate through the LAr volume on a nanosecond timescale whereas ionization electrons drift at a typical velocity of roughly 1.6 mm/μs, with drift times $\mathcal{O}(\text{ms})$, the optical trigger determines the t_0 with sub-microsecond precision, enabling effective 3D localization of interactions.

3.4.5 Calorimetric Precision and Neutrino Flavor Identification

Although LArTPCs interact with and measure only the final-state charged particles traversing the medium, they serve as powerful instruments for investigating neutral particles, such as incoming neutrinos. This is achieved through the total kinematic reconstruction of the secondary charged particles produced when an incident neutrino undergoes a weak-force interaction with an target argon nucleus (^{40}Ar).

Long-baseline neutrino oscillation programs exploit this combination of spatial granularity and calorimetric resolution to perform robust neutrino flavor identification. For instance, a charged-

current muon neutrino (ν_μ) interaction produces a primary final-state muon (μ^-), which exhibits a characteristically clean, continuous, long track topology. In contrast, a charged-current electron neutrino (ν_e) interaction produces a primary electron (e^-), which rapidly undergoes bremsstrahlung and pair production to develop a complex electromagnetic shower topology. The capability of a LArTPC to cleanly differentiate these topologies, while separating electron-induced showers from neutral-current-induced photon conversion backgrounds (1 MIP vs. 2 MIP separation via dE/dx profiling at the track vertex), provides the experimental sensitivity necessary for precision measurements of the neutrino mass ordering and leptonic CP violation parameters.

3.5 LBNF Neutrino Beam

DUNE will utilize an intense neutrino source in the form of a high-intensity beam generated at the LBNF accelerator complex. Figure 3.6 illustrates the primary components and structural aspects of the LBNF neutrino beamline. The facility will leverage the PIP-II [87] upgrade, which will initially deliver 1.2 MW of proton beam power, with a built-in design option to scale up to multi-megawatt beam power in subsequent phases.

Protons accelerated during each beam spill strike a segmented graphite target, producing a hadronic shower consisting primarily of pion and kaon secondaries. These secondary particles are subsequently focused by magnetic focusing horns that can preferentially select mesons of a chosen electric charge sign (positive or negative). Following this focusing stage, the mesons traverse a long, helium-filled decay pipe where they decay into tertiary leptons, including a beam dominated by ν_μ or $\bar{\nu}_\mu$ depending on the selected magnetic horn polarity.

At the downstream end of the decay pipe, a massive beam absorber halts any remaining hadrons and primary protons while allowing the weakly interacting neutrinos to pass through unhindered. The resulting neutrino beam is directed through the Earth's crust toward the ND and FD complexes to provide the high-statistics baseline required for DUNE's neutrino oscillation measurements. Over the full projected data-taking lifetime of the experiment, the LBNF beamline is expected to deliver a cumulative exposure of approximately 1.2×10^{21} protons-on-target (POT).

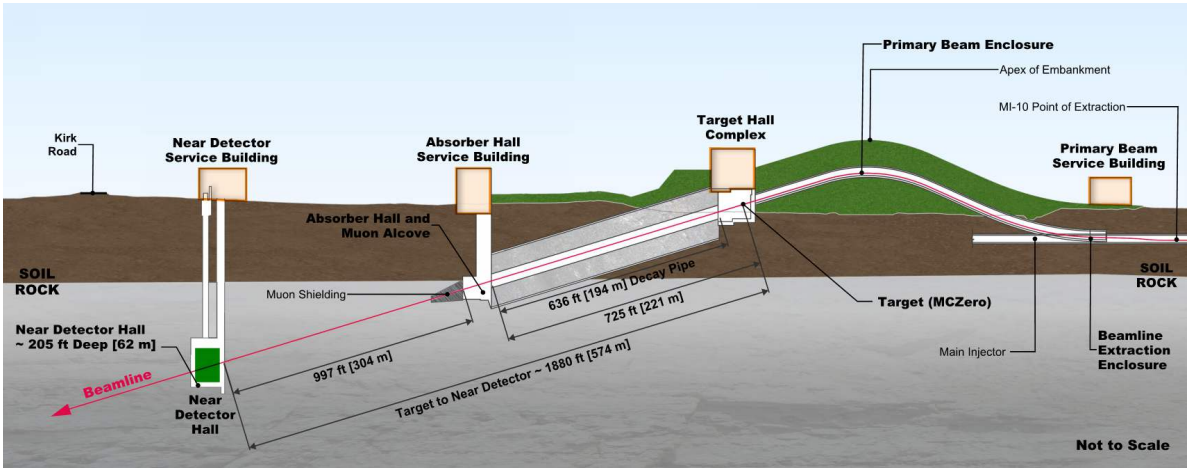

Figure 3.6: Illustration of the neutrino beamline and ND complex for the LBNF/DUNE project. Image retrieved from Ref. [9].

3.6 The DUNE Near Detector (ND)

3.6.1 The Purpose of the ND

The ND will be located 574 m from the beam source to monitor the beam parameters and measure, prior to the onset of neutrino oscillations, the unoscillated neutrino interaction rate as well as the muon and electron (anti)neutrino energy spectra [65]. The ND complex plays a crucial role in constraining systematic uncertainties for downstream oscillation measurements and providing essential empirical data to anchor the neutrino interaction model. To achieve highly sensitive measurements of the CP-violating phase δ_{CP} within a reasonable experimental timeframe, DUNE must constrain systematic uncertainties across as many independent parameters as possible [9,65].

The reconstructed energy spectrum is simultaneously influenced by the neutrino-argon cross section, the initial neutrino flux, and the specific detector response. Consequently, the ND must be capable of independently isolating these effects to provide the empirical inputs needed to model them accurately. Because the neutrino flux profile differs inherently between the ND and the FD due to geometric line-of-sight differences and the subsequent impact of neutrino oscillations, the simulation frameworks must rely on robust data to account for these variations.

Furthermore, the ND must measure neutrino interactions with significantly greater precision than the FD and be capable of handling the severe pile-up of multiple overlapping neutrino inter-

actions within a single beam spill, amounting to roughly 60 ν interactions per 1.2 MW beam spill. Beyond its role in the oscillation program, the ND will host a dedicated high-precision physics program of its own, measuring Standard Model cross sections, the weak mixing angle, and nuclear parton distribution functions. It will simultaneously be used to search for BSM phenomena, including light dark matter, dark photons, axion-like particles, and heavy neutral leptons [65].

3.6.2 DUNE Liquid Argon Near Detector (ND-LAr)

The DUNE ND-LAr detector is a core component of the ND complex tasked with measuring the unoscillated beam neutrino flux and flavor spectra [65]. ND-LAr utilizes the same LAr detection medium as the FD to inherently cancel out major cross-section and detector-response systematic uncertainties. Because ND-LAr will be exposed to an exceptionally high rate of neutrino interactions, it must feature rapid signal differentiation to disentangle dozens of concurrent neutrino interactions within each individual beam spill. Monolithic LArTPCs like the DUNE FD cannot handle such high interaction rates, primarily due to spatial ambiguities inherent in wire-based charge readouts.

To mitigate this pile-up challenge, ND-LAr is subdivided into 70 optically isolated TPC volumes equipped with a pixelated charge readout. This pixelated architecture enables native 3D event reconstruction, which eliminates wire-matching ambiguities and assists in cleanly distinguishing individual neutrino interactions. The short drift distance of 50 cm also relaxes requirements on LAr purity, as the ionization electrons travel a shorter distance before detection. Each TPC is lined with photon detectors to capture prompt scintillation light enabling drift coordinate reconstruction. The modularity and optical isolation of these TPC units drastically simplify the matching of scintillation light signals to their corresponding charge ionization tracks for individual neutrino interactions.

Figure 3.7 illustrates the modular architecture and the structural detector components of an ND-LAr LArTPC module. Each individual module houses two optically isolated TPCs within a $1 \times 1 \times 3$ m^3 volume. The complete ND-LAr detector will comprise a 7×5 array of these modular units.

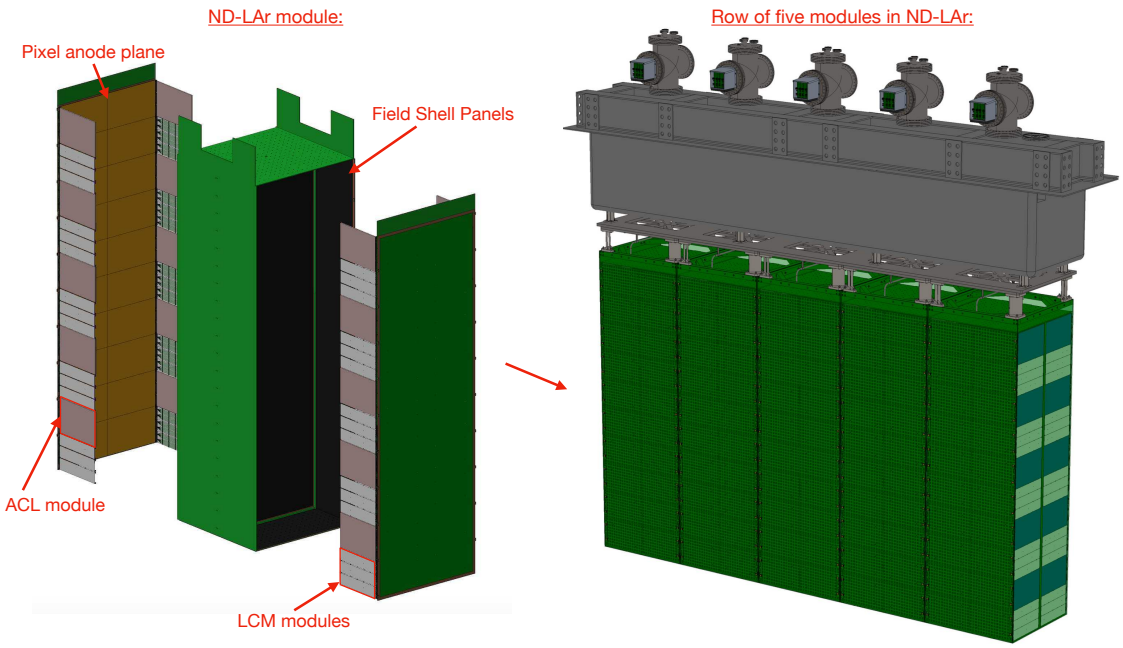

Figure 3.7: (Left) Illustration of an ND-LAr LArTPC module. (Right) Illustration of a single row of five modules in ND-LAr. Images retrieved from Ref. [65].

The modules leverage the ArgonCube technology [88] and utilize the LArPix [89] application-specific integrated circuit (ASIC) architecture for charge detection. LArPix utilizes a dense array of charge-sensitive pixel pads that directly collect the drifting ionization charge. Each individual pixel acts as an independent readout channel digitized via a localized channel-by-channel threshold. The low data volumes produced by this localized pixelated architecture enable continuous, dead-time-free readout. To maximize light yield, the photon detectors are placed directly inside the field cage, covering approximately 25% of the inner TPC boundaries.

3.6.3 SAND, ND-GAr, and PRISM

The complete ND complex layout is illustrated in Figure 3.8. ND-LAr is complemented by other upstream and downstream sub-detectors that fulfill distinct roles within the complex. Positioned immediately downstream of ND-LAr is ND-GAr, a detector comprising a high-pressure gaseous argon TPC (HPgTPC) surrounded by an electromagnetic calorimeter (ECAL) and a 0.5 T magnet. ND-GAr is designed to measure the momenta and charge sign of high-energy particles exiting the upstream ND-LAr volume. Additionally, its low density allows it to reconstruct low-

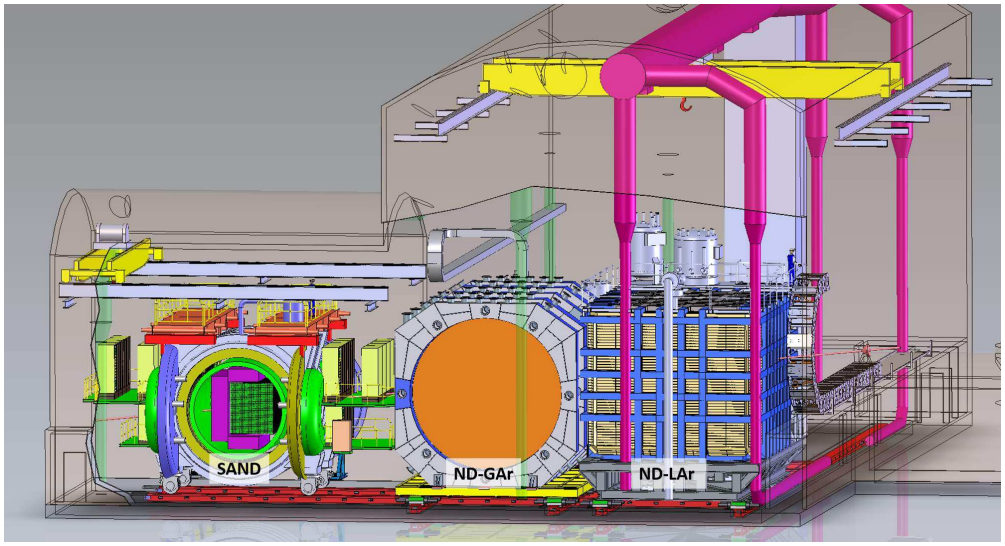

Figure 3.8: Illustration of the ND complex and detector components. The neutrino beam enters from the right side of the image. Image retrieved from Ref. [65].

energy final-state particles from neutrino interactions occurring directly within its gaseous volume at thresholds below what is achievable in liquid detectors, thereby providing vital constraints on neutrino-argon interaction models.

Due to missing energy from undetected neutral particles or misidentified final-state hadrons, the relationship between the visible final-state particles and the true incident neutrino energy is highly complex. The true neutrino energy spectra will differ substantially between the ND and the FD, creating a potential source of bias in neutrino oscillation parameter measurements if left unconstrained by data-driven methods.

To resolve this issue, the complex incorporates the DUNE Precision Reaction-Independent Spectrum Measurement (DUNE-PRISM) framework. DUNE-PRISM enables the motorized translation of ND-LAr and ND-GAr to various off-axis angles relative to the central neutrino beam axis. By leveraging the systematic decrease in mean neutrino energy with an increasing off-axis angle, DUNE-PRISM can construct custom, data-driven linear combinations of neutrino energy spectra to predict the FD event rate independently of interaction modeling assumptions.

While the neutrino beamline is designed for maximum operational stability, minor beam instabilities within tolerance limits must still be monitored continuously to minimize additional sys-

tematic errors. Because ND-LAr and ND-GAr will periodically move off-axis during PRISM data-taking, a dedicated on-axis instrument is required to provide uninterrupted beam monitoring. This role is fulfilled by the System for On-Axis Neutrino Detection (SAND). SAND reuses the superconducting magnet and iron yoke calorimeters from the KLOE experiment [90], reconfiguring the inner tracking volume with a dedicated target and tracking system. A primary capability of SAND is to provide continuous, high-statistics standalone measurements of the on-axis neutrino interaction rate and energy spectrum to verify beam alignment. All detector components of the ND complex are detailed comprehensively in the DUNE ND Conceptual Design Report [65].

3.7 DUNE ND-LAr Technology and Prototyping

3.7.1 Charge Readout System

The charge readout system (CRS) for ND-LAr, responsible for detecting drifted ionization electrons, utilizes LArPix. LArPix is instrumented directly on the ND-LAr anode planes and consists of a dense array of charge-sensitive pixels (gold-coated conductive pads) fabricated on a printed circuit board (PCB), which are digitized and read out by localized ASICs. The anodes are constructed out of a 2×10 array of individual pixel "tiles" that collectively span the full 1×3 m^2 area of each anode plane.

The LAr-facing side of each pixel tile features an 128×80 grid of 10,240 pixels with a 3.72 mm pitch, while the reverse side is instrumented with a 10×16 array of 160 ASICs (both shown in Figure 3.9). A block diagram of the LArPix electronics is provided in Figure 3.10. A summary of the LArPix operating principles is outlined here, while a more detailed discussion can be found in Ref. [23].

Each ASIC includes a front-end amplifier and a 10-bit analog-to-digital converter (ADC) dedicated to each individual pixel. When drifting ionization charge approaches and terminates on a pixel, it induces a current described by Ramo's Theorem for a pixel at a fixed bias voltage. A charge-sensitive amplifier (CSA) integrates this induced current, outputting the resulting signal to a buffer. The output of the buffer is directed concurrently to a discriminator and the ADC. If

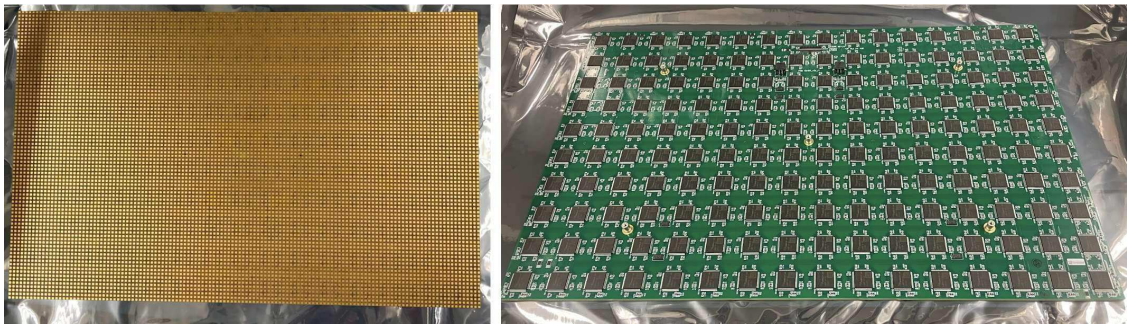

Figure 3.9: (Left) LArPix anode pixel tile for ND-LAr measuring 47.55 cm by 29.69 cm with 10,240 pixel pads at 3.72 mm pitch. (Right) LArPix pixel tile with 160 ASICs to digitize and read out signals produced on the pixels.

the signal provided to the discriminator exceeds a configurable, channel-dependent threshold, the digitization of the pixel in question is initiated via a "self-trigger."

Upon exceeding the threshold, a configurable integration timer is set, typically at 1.8 µs, allowing additional current to contribute to the total accumulated signal. After this time delay passes, the CSA output is sampled by the ADC to generate a data packet, where the recorded ADC value is directly proportional to the total charge detected by that pixel. The data packet also contains complementary information, such as the absolute timestamp at which the discriminator was triggered (discretized into 0.1 µs time ticks) and pixel-identifying coordinates used to map the physical position of the interaction in (x,y,z) space.

These packets are subsequently transmitted from the ASIC to the data acquisition system (DAQ) via a warm controller card designated as the Pixel Array Control, Monitor, and Network Card (PACMAN). The PACMAN board provides power, digital control, and data I/O for up to 10 anode tiles, interfacing directly with the DUNE timing system via a 62.5 MHz clock to ensure accurate synchronization with the rest of the detector sub-systems. It mounts to a feedthrough on the cryostat flange, connecting to each LArPix anode tile via internal cabling.

Additionally, the PACMAN board interfaces with an Ethernet switch using standard RJ-45 Ethernet cables, which route the data to the central DAQ via optical fiber. These connections facilitate bidirectional communication, allowing physics data to stream from the ASICs to the DAQ and enabling configuration commands to be sent from the DAQ to the anode tiles. A PACMAN

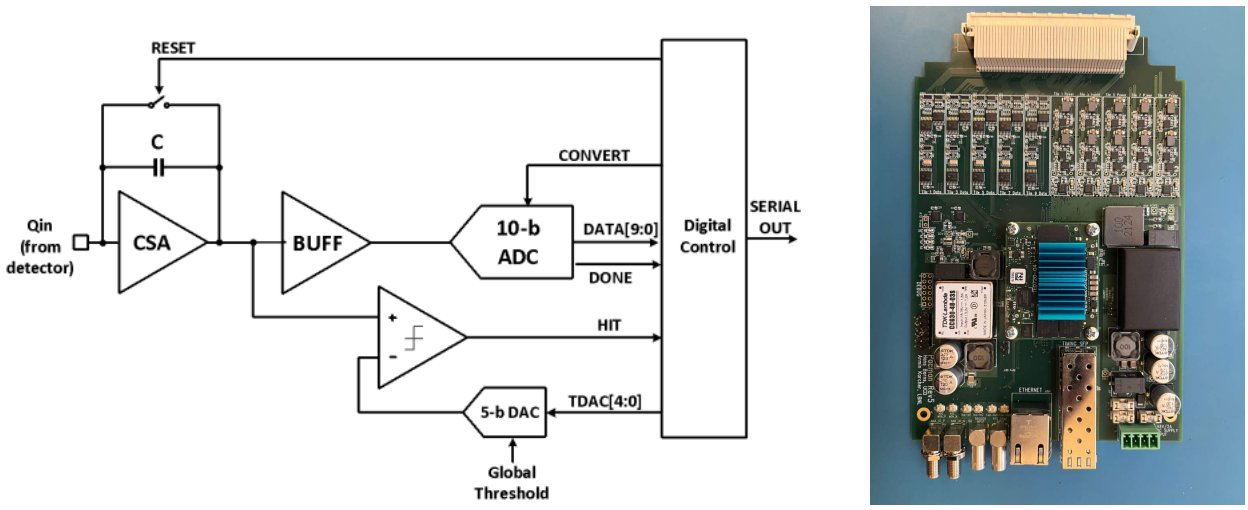

Figure 3.10: (Left) A block diagram of the LArPix charge readout system for ND-LAr. (Right) A picture of the PACMAN warm controller board.

board capable of reading out 10 pixel tiles is shown in Figure 3.10. Saving only sparse packet data dramatically reduces the bandwidth and storage overhead compared to traditional, monolithic LArTPC wire-based charge readouts that save full digitized waveforms without zero suppression. This efficient design enables continuous readout of the CRS, guaranteeing that rare or interesting activity within the TPCs is captured.

To calculate the physical amount of charge detected by an individual pixel, the following equation is applied:

$$Q = (K - P) \cdot G \cdot \frac{V_{\text{ref}} - V_{\text{cm}}}{2^N}, \tag{3.10}$$

where K represents the total ADC counts from the data packet, P is the channel-specific pedestal value in ADCs, and G is the gain of the CSA. For the analysis presented in this thesis, G is taken to be an average value of 0.221 ke^-/mV based on measurements of a sample of LArPix ASICs operating in LAr (where 1 ke^- = 1,000 electrons). The least significant bit (LSB) for an N-bit ADC is expressed as $(V_{\text{ref}} - V_{\text{cm}})/2^N$ in units of mV/ADC, where V_{ref} and V_{cm} are LArPix configuration parameters defining the maximum and minimum ADC input signals, respectively. The 2^N term represents the maximum possible dynamic range step for an N-bit ADC. While $N = 10$ for the final ND-LAr design, earlier LArPix iterations, such as those used in certain prototypes like the 2x2 Demonstrator, utilized $N = 8$. The average channel-specific pedestal is

extracted from dedicated, specially-configured calibration runs that randomly sample the baseline ADC value on a given channel repeatedly over a designated time frame.

3.7.2 Light Readout System

Figure 3.7 shows the schematic layout for an ND-LAr module alongside its two constituent photon detector technologies. Both systems are designed to convert the 128 nm VUV LAr scintillation light to a visible wavelength that is more efficiently detected by commercial SiPMs. The SiPMs used are Hamamatsu S13360-6050CS sensors, which feature a 6×6 mm^2 active area with a pixel pitch of 50 μm. Because these sensors are optimized for the detection of green light, it is necessary to up-convert the deep UV scintillation light using wavelength-shifting (WLS) materials.

The Light Collection Module (LCM) utilizes a bundle of optical fibers coated in tetraphenyl butadiene (TPB), which shifts the wavelength of incoming photons from 128 nm to approximately 430 nm. A fraction of the photons entering the fibers are trapped via total internal reflection. These trapped photons are further shifted up to roughly 490 nm as they propagate along the fibers, which wrap around and terminate at an active SiPM surface. Each individual LCM contains two SiPMs, with three vertically stacked LCMs treated as a single unified readout element in the detector.

The ArCLight (ACL) module similarly leverages TPB to provide an initial wavelength-shifting stage, subsequently utilizing a bulk WLS substrate as a secondary shifting medium. The bulk WLS component consists of a $300 \times 500 \times 10$ mm^3 plate composed of EJ-280, produced by Eljen Technology. A dichroic film coats the bulk WLS plate, trapping the photons shifted by the second stage within the substrate. Six SiPMs align along one edge of the ACL module to detect the trapped, shifted scintillation photons.

The wavelength-shifting mechanisms for both detector types are visualized in Figure 3.11. One ACL module matches the physical footprint of approximately three stacked LCM modules. The ACL and LCM units are stacked in an alternating vertical pattern along the sides of the TPC facing perpendicular to the drift direction, as illustrated in Figure 3.7.

A specialized summing ADC aggregates the signals from every group of six SiPMs to form

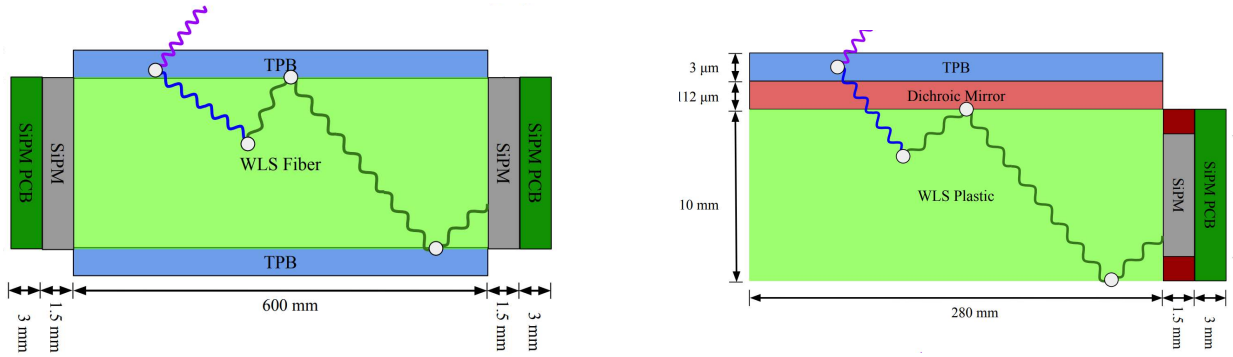

Figure 3.11: (Left) A diagram showing the operating principle of a LCM photon detector module. (Right) A diagram showing the operating principle of an ACL photon detector module.

"sum channels," which are utilized to drive the self-triggering logic of the light readout system (LRS) within each TPC module. If a sum channel in an ND-LAr module surpasses a programmed threshold, all SiPMs within that module are digitized by the system ADCs and written to disk. While the CRS operates independently from the LRS, timing information is strictly synchronized across both systems to enable charge-light matching. Alternatively, for accelerator beam spills, ND-LAr can leverage an external beam trigger signal to prompt LRS digitization. Conversely, out-of-time events from cosmic rays or radioactive decays within the LAr volume must be recorded utilizing the self-triggering operational mode.

3.7.3 Early ND-LAr Prototyping

An extensive prototyping program for ND-LAr has been underway throughout the 2020s. Prototyping relies first on component-level testing conducted mainly via small-scale test stands, and second on integrated system testing of all combined detector components within an operational LArTPC. The primary objective of these integrated tests is to verify that all sub-systems can operate concurrently without interference. Furthermore, these tests demonstrate the viability of the large-scale production and assembly methodologies required to construct the final ND-LAr detector.

The program initiated with small-scale integrated tests using "SingleCube Demonstrators." Each SingleCube consists of a single readout plane instrumented with a 70×70 array of pix-

els covering an area of approximately 30×30 cm^2, a miniature field cage for drift field uniformity, a high-voltage cathode, a single photon detector unit, and an active drift length of roughly 30 cm. The SingleCube demonstrators successfully validated the integrated performance of the underlying ND-LAr technologies at a compact scale. These decentralized tests were accelerated during the COVID-19 pandemic, which prevented international travel to the University of Bern where central integrated testing was hosted. Consequently, multiple SingleCube demonstrators were dispatched to collaborating institutions worldwide, including CSU, allowing institutions to independently evaluate and mature the technology.

3.7.4 The 2x2 Demonstrator

The next phase of the prototyping pipeline involved the construction, commissioning, and operation of the 2x2 Demonstrator (or "2x2"). This apparatus was designed to execute large-scale integration testing of ND-LAr technology within a LAr environment and to conduct a technical evaluation of modular LArTPC performance under realistic beam neutrino interactions. The 2x2 Demonstrator is located in the MINOS underground hall at Fermilab, directly exposed to the Neutrinos at the Main Injector (NuMI) beamline. Situated approximately 100 m underground, the facility benefits from a cosmic-ray background that is reduced by roughly two orders of magnitude relative to the surface.

To track through-going muons and provide a downstream muon spectrometer, scintillator tracking planes repurposed from the MINERvA experiment [91] are installed upstream and downstream of the TPC array. The 2x2 Demonstrator configuration features a 2×2 array of optically isolated LArTPC modules, each containing two independent TPC volumes. Each individual module is instrumented with two pixelated LArPix [89] anode planes (one per TPC), a central shared cathode plane, and light detectors covering the side boundaries. The modules are numbered sequentially according to their production order: Module-0, Module-1, Module-2, and Module-3.

The anode-to-cathode distance spans approximately 30 cm in each TPC, which is shorter than the nominal 50 cm drift length specified for the final ND-LAr design. A cathode high-voltage bias

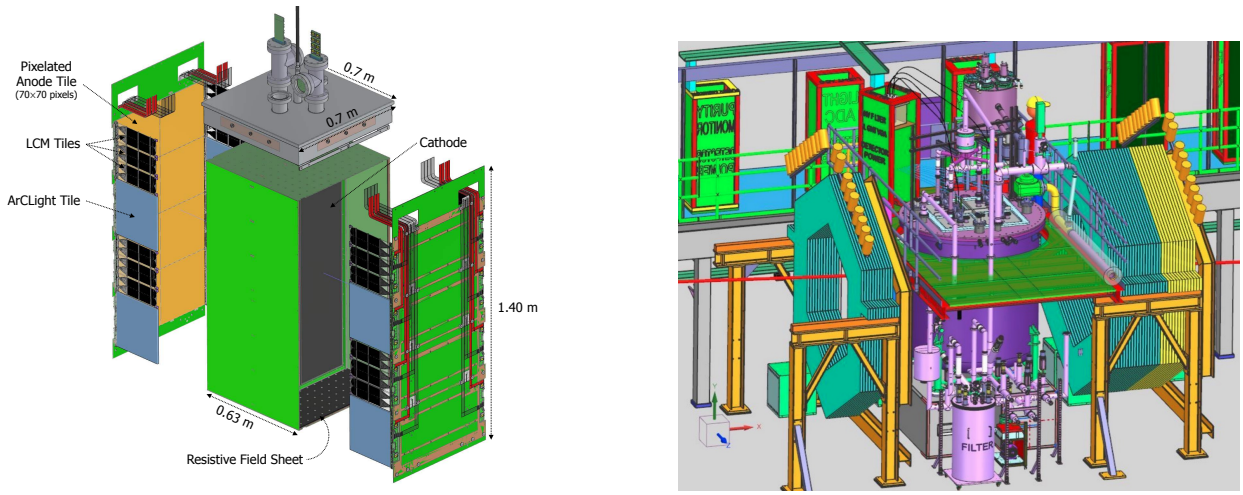

Figure 3.12: (Left) A diagram of a 2x2 Demonstrator LArTPC module [23]. (Right) A 3D model of the 2x2 Demonstrator.

of approximately -15 kV generates the nominal drift electric-field strength of 500 V/cm. Rather than a traditional wire or metal profile cage, a resistive field shell composed of a carbon-loaded Kapton film (DR8) laminated to a G10 substrate is deployed to enforce a uniform drift field [92]. The transverse boundaries of the TPCs are lined with alternating ACL [93] and LCM [94] modules to detect scintillation light. Pictures of each type of photon detector for the 2x2 Demonstrator, alongside a LArPix tile, are shown in Figure 3.14.

Figure 3.12 provides a technical schematic of an individual 2x2 Demonstrator LArTPC module alongside a structural model of the entire detector installation within the MINOS hall. Additionally, Figure 3.13 displays a picture of the interior layout of a 2x2 Demonstrator module and the subsequent mechanical installation of the array inside the 2x2 Demonstrator cryostat.

At the time of writing this thesis, the 2x2 has completed two major data-taking campaigns, spanning 2024 and 2025. The initial run was performed during active NuMI beam operations, while the secondary run was executed during a beam-shutdown window. The primary objective of the shutdown run was to evaluate the cryogenic system in preparation for a high-intensity neutrino beam run scheduled for 2026, and to validate various calibration sources, including the ^{220}Rn injection technique evaluated in this work.

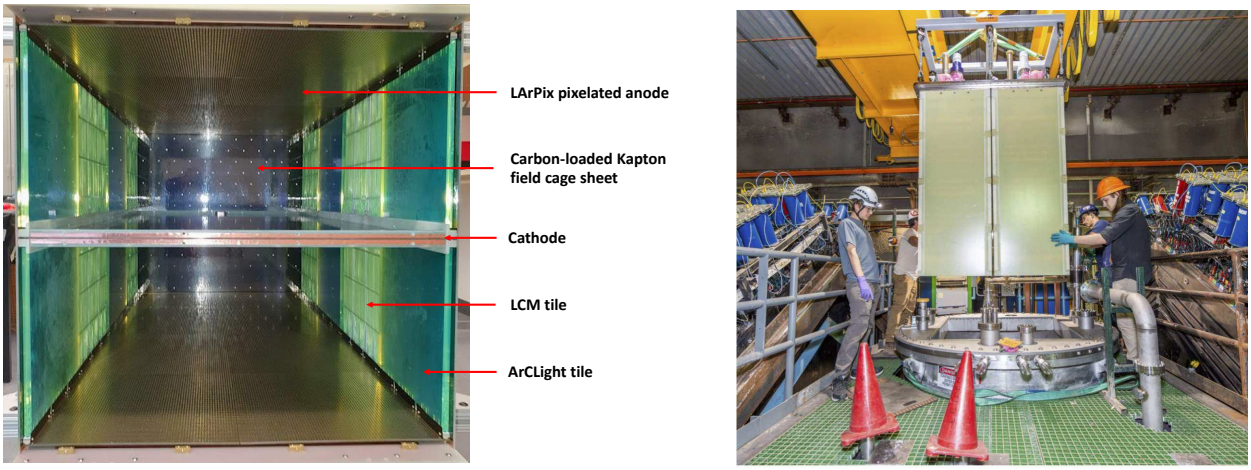

Figure 3.13: (Left) A picture of the inside of a 2x2 Demonstrator LArTPC module, retrieved from Ref. [23]. (Right) A picture of the 2x2 Demonstrator LArTPC modules being installed into the cryostat at Fermilab.

3.7.5 Full-Scale Demonstrator (FSD)

Because the 2x2 LArTPC modules are smaller than the final operational specifications, they do not represent the ultimate form-factor for the ND-LAr detector. The third milestone in the ND-LAr prototyping pipeline involves the construction, commissioning, and operation of a full-scale ND-LAr module, designated as the Full-Scale Demonstrator (FSD). The FSD replicates the precise final dimensions of an ND-LAr module, measuring $1 \times 1 \times 3$ m^3 with a full active drift length of 50 cm. A structural diagram of the FSD architecture, which is identical to the final ND-LAr design, is shown in Figure 3.7. A picture of the inside of the FSD TPC and the completed FSD module before insertion into the cryostat is shown in Figure 3.15.

The FSD shares an overarching purpose with the 2x2 program, aiming to conduct integrated system testing of full-scale ND-LAr components in preparation for upcoming production phases. The FSD operates within a cylindrical cryostat measuring 1.5 m in diameter and 4 m in height, supported by a high-purity LAr cryogenic infrastructure at the University of Bern Laboratory for High Energy Physics (LHEP). The detector successfully executed its operations in late 2024.

Unlike the 2x2, the FSD is not evaluated within an active neutrino beamline; instead, it relies on cosmic-ray muon interactions to benchmark detector performance and verify that the system satisfies baseline technical design specifications, such as drift electric-field uniformity and ionization

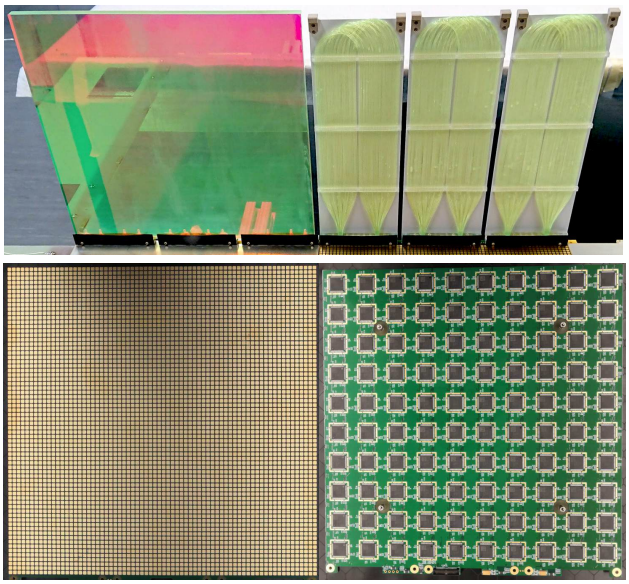

Figure 3.14: (Top) Picture of an ACL light detector module and three LCM light detector modules that comprise the light readout system for the 2x2 Demonstrator. (Bottom) Picture of both sides of a LArPix pixelated anode tile for the 2x2 Demonstrator TPC modules (left shows pixels, right shows ASICs).

electron drift velocity. Additional primary objectives of the FSD program include evaluating and refining component manufacturing and Quality Assurance/Quality Control (QA/QC) validation protocols, testing the full module assembly workflow to establish standard operational procedures for the upcoming Module Assembly and Test Facility (MATF), and prototyping the LAr cryogenic sub-systems with an emphasis on thermal management and chemical purity constraints.

The engineering results of this technical demonstration are currently being finalized for publication of the DUNE ND Technical Design Report. The FSD represents the final validation step for the ND-LAr technological layout and will be directly succeeded by the construction and testing of five full-scale production modules at the MATF facility in the 2026/2027 timeframe.

3.8 The DUNE Far Detector

3.8.1 Overview

The DUNE FD complex will be located at SURF in Lead, South Dakota, situated approximately 1.5 km underground. The full deployment plan comprises four instrumented LArTPC modules, each containing an active mass of over 10 kt. To maximize the scientific yield and hedge

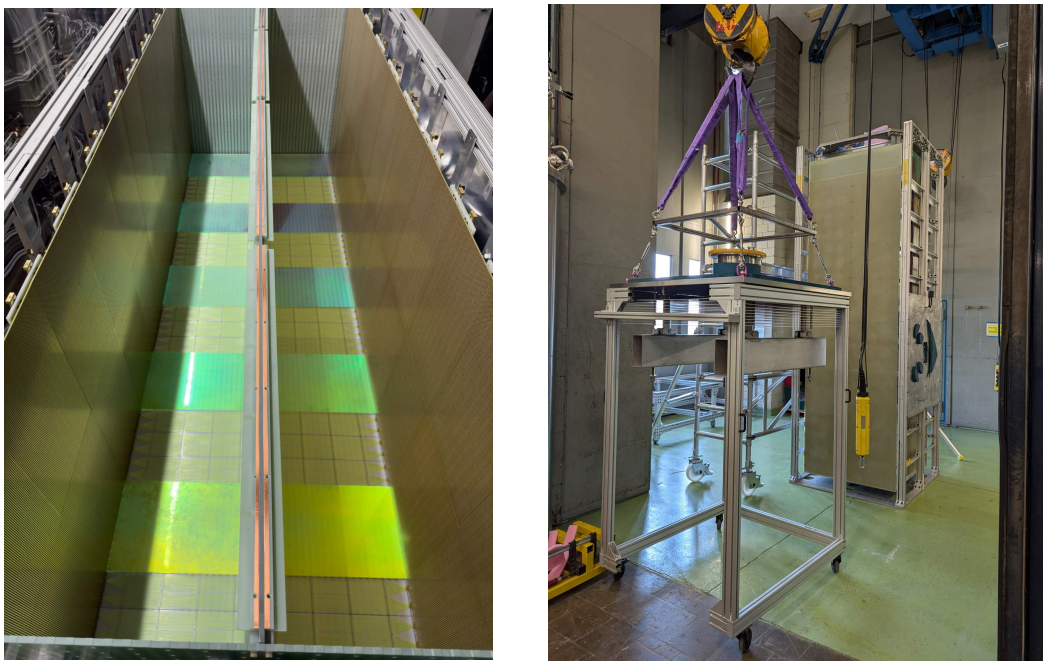

Figure 3.15: (Left) The inside of the FSD showing the pixel planes, cathode, and light detectors. (Right) Completed FSD module staged before inserting into the cryostat.

technical risks, DUNE implements two distinct LArTPC design paradigms for Phase I. The first operational module to be constructed will utilize the Vertical Drift (VD) configuration, while the second module will deploy the Horizontal Drift (HD) design. Figure 3.16 shows a schematic of a full-scale FD cryostat.

3.8.2 Vertical Drift Far Detector Design

The VD design optimizes the active volume footprint by establishing a vertically oriented electric-field vector across a single cathode plane suspended at the mid-height of the detector. This configuration splits the active volume into upper and lower drift regions, drifting ionization electrons vertically upward or downward toward respective anodes positioned at the ceiling and floor of the cryostat. The VD detector technology are visualized and illustrated in Figure 3.17.

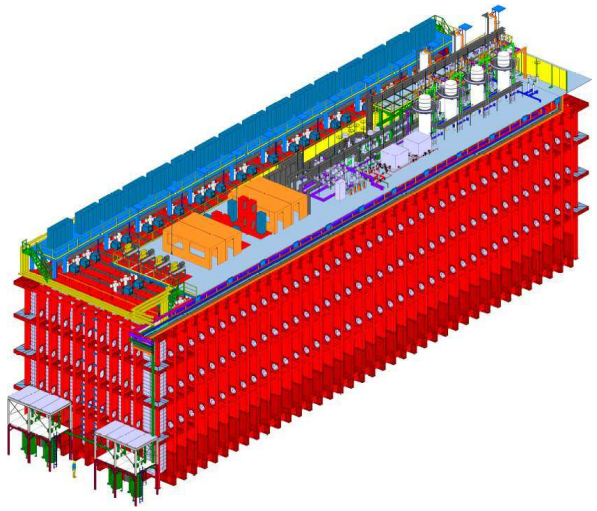

Figure 3.16: Schematic of a 65.8 m (L) by 18.9 m (W) by 17.8 m (H) outer-dimension cryostat that houses a 10 kt FD module.

Charge Readout System

The VD design replaces traditional wire-based frame assemblies with Charge Readout Planes (CRPs). These CRPs leverage multi-layer PCB technologies, where conductive traces are etched onto large, rigid perforated sheets to form the ground, induction, and collection planes. This approach significantly simplifies the structural engineering and tensioning requirements inherent to fine wire matrices while preserving high spatial sampling performance.

To establish the drift field, the central cathode plane is biased at approximately -300 kV, generating a uniform 500 V/cm drift electric field across both the upper and lower active LAr volumes. Uniformity of this field is maintained via a surrounding field cage along the lateral boundaries.

Photon Detection System

Because LAr scintillates brightly at a VUV wavelength of approximately 128 nm associated with particle interactions in the detector, a high-efficiency photon detection system (PDS) is vital to establish the precise event interaction time (t_0). The VD design utilizes photon detectors based on the advanced X-ARAPUCA concept [95], which improves upon geometric light collection limits via spectral wavelength shifting and optical trapping.

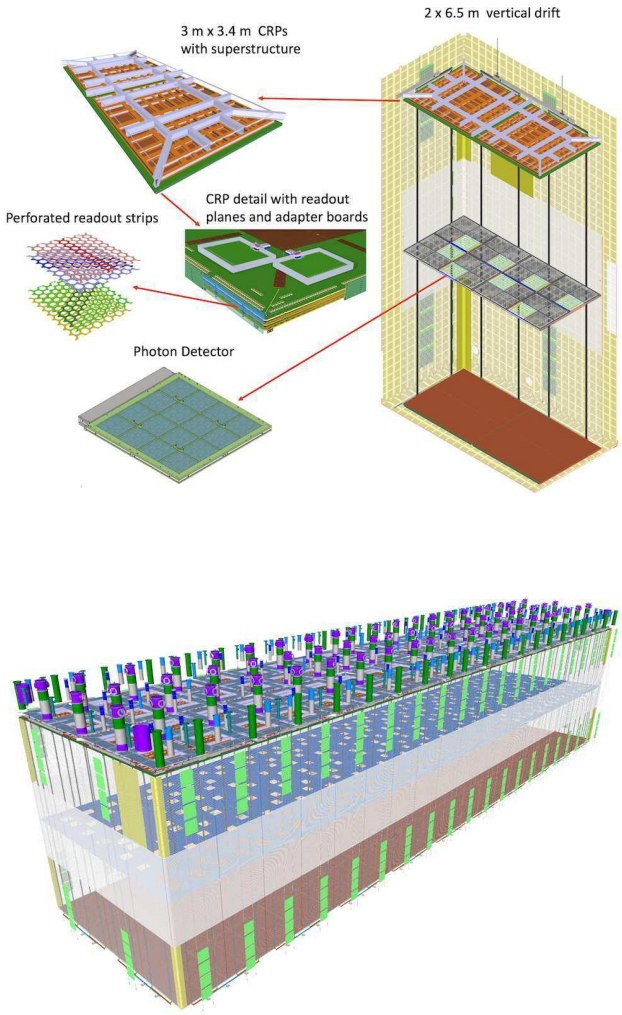

Figure 3.17: (Top) Illustration of the DUNE FD VD LArTPC concept utilizing top and bottom charge readout planes and a central cathode plane. (Bottom) A perspective view of the FD2-VD detector. Images retrieved from Ref. [22].

As shown schematically in Figure 3.18, incident 128 nm photons first encounter a dichroic filter window coated with p-Terphenyl (pTP). The pTP shifts the VUV light to an intermediate wavelength of roughly 350 nm, matching the high-transmission window of the underlying dichroic filter. Once inside the cell, the photons hit a secondary internal WLS plate that shifts them further to approximately 430 nm. Because the dichroic filter is opaque to wavelengths above its cutoff, the 430 nm light cannot escape back through the entry window, so they are reflected from the surface of the dichroic filter, trapping the photons inside of the X-ARAPUCA volume. Highly reflective Vikuity™ foils line the interior cavity walls to allow the photons to reflect inside the cavity until

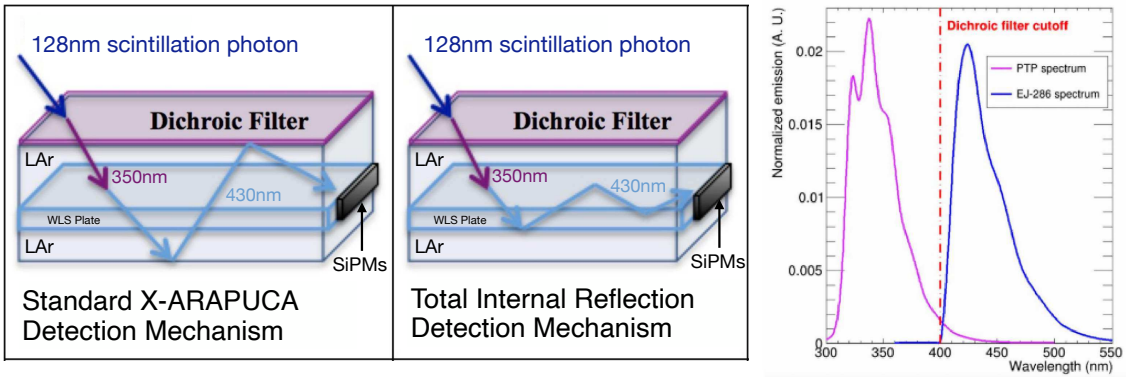

Figure 3.18: Scintillation emission spectra of the WLS materials used in the X-ARAPUCA design, denoting the DUNE FD X-ARAPUCAs' dichroic filter cutoff wavelength. Image retrieved from Ref. [96].

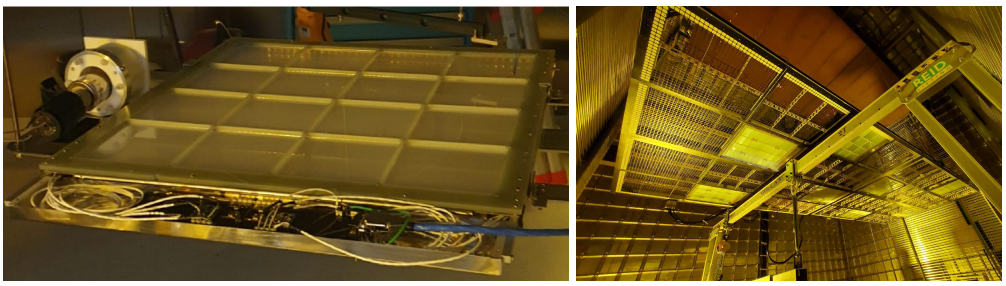

Figure 3.19: Picture of the inside of ProtoDUNE-VD looking down inside the active cryostat volume during construction.

they are detected by SiPMs mounted along the internal edges. The WLS plate also serves as a light guide, utilizing total internal reflection to route trapped photons toward the photodetector boundaries. This design feature is the primary reason why the X-ARAPUCA has an improved photon detection efficiency in comparison to the original ARAPUCA design [97]. Figure 3.18 details the emission profiles of the WLS materials used in the X-ARAPUCA design.

The VD PDS employs large X-ARAPUCA modules measuring 60×60 cm^2, fitted with 16 dichroic filters per side and 160 edge-mounted SiPMs. These modules are installed at two distinct operational locations: along the outer cryostat side walls behind the field cage and directly inside the central cathode plane structures. The cathode-mounted units are double-sided, enabling concurrent light collection from both the upper and lower drift volumes. Figure 3.19 shows pictures of a VD-style PDS module and the cathode inside ProtoDUNE-VD.

As the central cathode operates at an extremely high voltage nearing -300 kV, conventional

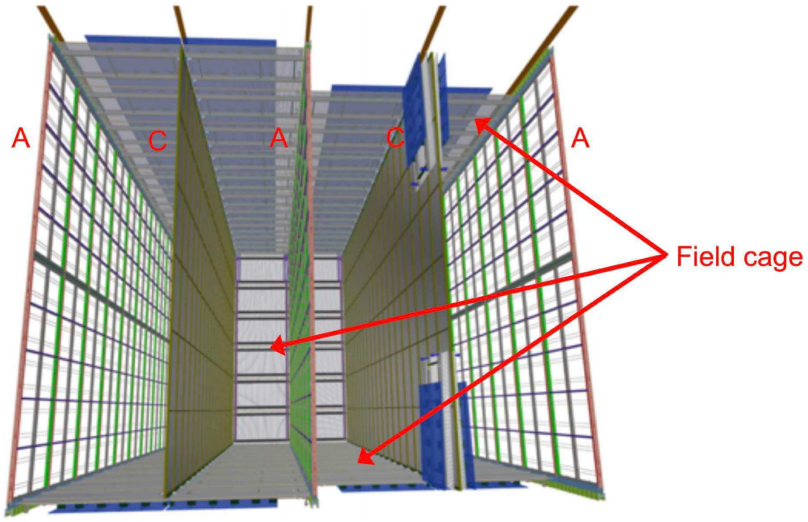

Figure 3.20: Illustration of the DUNE FD Horizontal Drift LArTPC architecture, where "A" denotes Anode Plane Assemblies (APAs) and "C" denotes Cathode Plane Assemblies (CPAs). Image retrieved from Ref. [9].

copper infrastructure cannot be used for power or signal routing due to severe dielectric breakdown and noise injection risks. Instead, power-over-fiber (PoF) and signal-over-fiber (SoF) technology is used [98, 99]. Warm-side drive lasers transmit optical power over non-conductive fiber lines to gallium arsenide (GaAs) Photovoltaic Power Converters (PPCs) embedded in the cathode, which generate the necessary voltage rails for the active cold electronics and SiPM biasing. For data return, analog signals from the SiPMs modulate cold-rated Fabry-Pérot laser diodes, passing optical data back to warm-side InGaAs photodiodes. This ensures that the VD PDS is fully electronically isolated, leading to lower overall noise levels and a significantly lower risk of sparking.

3.8.3 Horizontal Drift Far Detector Design

The Horizontal Drift (HD) design utilizes a segmented architecture where alternating vertical planes of anodes and cathodes establish a horizontal electric field across the active LAr medium. A schematic of the HD LArTPC technology is shown in Figure 3.20. The following describes the operating principles and design of the charge and light readout systems in the HD FD.

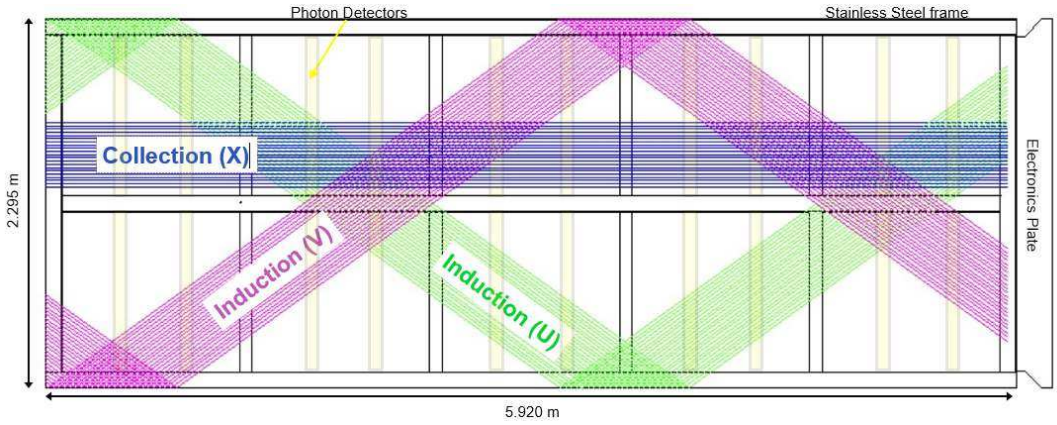

Figure 3.21: The detailed structural layout of an APA relevant for the horizontal-drift FD, highlighting the wire plane orientations and embedded photon detectors. Only a partial section of each wire plane is shown for clarity. Image retrieved from Ref. [9].

Charge Readout System

In the HD design, the anode planes are constructed from Anode Plane Assemblies (APAs), which are modular 2.3 m wide by 6 m high stainless steel structural frames. As shown in the detailed structure in Figure 3.21, each APA is instrumented with four distinct planes of copper-beryllium wires spaced at a nominal pitch of approximately 5 mm.

- **Ground Plane (G):** The outermost plane acts as an electrostatic shield.
- **Induction Planes (U and V):** The subsequent two wire planes are biased to maximize electron transparency, allowing drifting ionization electrons to pass through unimpeded. As the electron clouds pass by these wires, they induce bipolar signals. Because the U and V planes are oriented at ± 35.7 degrees relative to the vertical axis, matching these induction signals provides the necessary spatial correlations to reconstruct the 2D position of the ionization in the plane transverse to the drift direction by tracing patterns across the wire planes.
- **Collection Plane (X):** The final, vertically oriented wire plane is biased to fully collect the arriving ionization electrons. This produces unipolar signals on the wires that can be integrated to measure the total ionization charge.

To establish the drift electric field in the HD detector, Cathode Plane Assemblies (CPAs) are positioned parallel to the APAs and biased at -180 kV, generating a uniform 500 V/cm drift electric

field across the twin LAr active volumes. Uniformity of this field is maintained via a surrounding field cage along the lateral boundaries.

Photon Detection System

The HD module similarly deploys the X-ARAPUCA optical capture mechanics described in Section 3.8.2. However, the physical integration and packaging differ significantly from the VD module. In the HD module, the PDS modules are directly embedded within the structural frames of the APAs, occupying the spaces between wire planes. Each HD PDS module spans nearly the full width of an APA and comprises four independent structural units termed "supercells." A supercell is a single X-ARAPUCA housing measuring 488×100 mm^2, featuring six dichroic filter windows and a single WLS plate. Figure 3.22 shows a schematic of an HD PDS module and a picture of an APA in ProtoDUNE-HD.

Each supercell contains 48 SiPMs. To minimize channel counts, sets of six SiPMs are passively ganged together into eight distinct sub-groups. These eight channels are then combined via an active ganged cold amplifier circuit into a single differential analog signal per supercell. Signal processing is managed externally by warm readout front-end electronics designated as DAPHNE, which provides high-speed amplification, digitization, and direct data streaming to the central DAQ system.

3.8.4 Prototyping at the CERN Neutrino Platform

To validate the structural integrity, manufacturing tolerances, and physical performance of these FD charge and photon readout technologies, an intensive international prototyping program was established at the CERN Neutrino Platform near Geneva, Switzerland [20–22]. Funded and constructed by a global collaboration of universities and national laboratories, these large-scale testbeds, named "ProtoDUNEs," have provided essential engineering baselines over the last decade.

The prototyping validation sequence progressed through several distinct phases:

- **ProtoDUNE-SP (Single-Phase):** Operated from 2018 to 2020, this module represented the

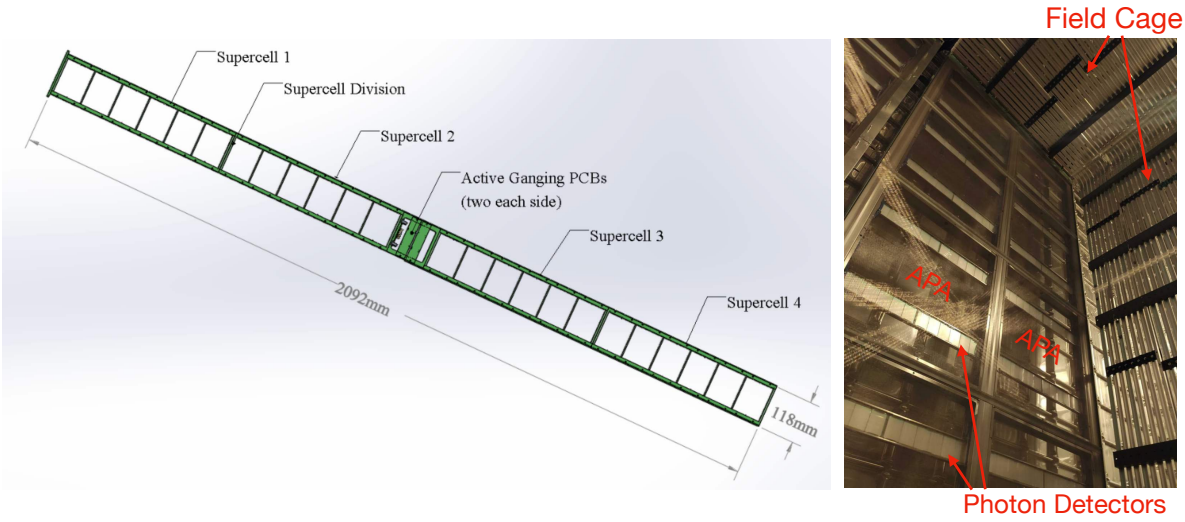

Figure 3.22: Detailed schematic drawing of a DUNE FD PDS module assembly showing four constituent supercells. Image retrieved from Ref. [96].

largest monolithic LArTPC constructed at that time, boasting an internal cryostat volume of $11.4 \times 10.8 \times 11.4$ m^3 and active TPC tracking dimensions of 7.2 m $\times$ 6.0 m $\times$ 6.9 m [9, 19, 20]. ProtoDUNE-SP successfully validated the full mechanical scale of the HD FD LArTPC design. There were three APAs per anode plane within this specific prototype layout. Exposure to a dedicated H4 test beam delivering pions, kaons, protons, muons, and electrons with well-defined momenta ranging from 0.3 GeV/c to 7.0 GeV/c allowed researchers to enable assessment of the detector technology's performance, development of calibration procedures relevant for the DUNE FD, and auxiliary cross section measurements of value to the future DUNE physics program. The comprehensive structural features and core physics performance data derived from this run cycle are thoroughly documented in Refs. [19, 20].

- **ProtoDUNE-DP (Dual-Phase) to ProtoDUNE-VD:** Concurrent with early single-phase testing, the dual-phase concept was evaluated using ProtoDUNE-DP [21]. This design drifted ionization electrons through a LAr volume up into an overhead gaseous argon layer to amplify the signal via gas electron multipliers before collection. However, rapid advancements in LAr purification and ultra-low-noise cold electronics demonstrated that satisfactory

signal-to-noise ratios could be achieved without complex gaseous amplification systems. Consequently, the collaboration pivoted to the VD single-phase concept [22], culminating in the construction and successful operation of ProtoDUNE-VD during the 2025–2026 run cycle. This prototype verified the mechanical integrity of large-area PCB-based CRPs alongside the fiber-optic PoF and SoF cathode readout configurations.

- **ProtoDUNE-HD (Horizontal Drift Run 2):** Following the conclusion of the initial single-phase run, the HD infrastructure was upgraded to ProtoDUNE-HD and operated in 2024. This installation cycle incorporated major technological improvements designed to match the final production specifications of the second FD module. In particular, it featured optimized X-ARAPUCA supercell configurations (see Figure 3.22), revised passive-to-active ganging routing, a streamlined configuration of two APAs per anode plane, and production-grade implementations of the DAPHNE warm electronics interface.

Chapter 4

LArTPC Calibrations and Radioactive Decays

4.1 LArTPC Calibrations

4.1.1 Overview

Detector calibration is a fundamental prerequisite for the successful execution of the DUNE physics program. To ensure the precise and accurate reconstruction of incident particle energies and spatial coordinates within LArTPCs, systematic calibration techniques must be deployed to continuously map and verify detector response parameters. Specifically, establishing an accurate absolute and relative energy scale is critical. Beyond total energy reconstruction, characterization of the differential charge deposition per unit length (dQ/dx) and differential energy deposition per unit length (dE/dx) serves as a primary metric for particle identification and calorimetry [16].

4.1.2 Particle Energy Reconstruction

Converting raw digitized charge data, whether from wire-based or pixel-based TPC architectures, into reconstructed charge and energy quantities requires comprehensive corrections for both electronic and detector-induced effects. First, the readout electronics gain must be characterized on a per-channel basis to translate measured voltages into the absolute number of collected ionization electrons. This electronics gain calibration is typically achieved either via electronic charge injection into individual front-end channels or by exploiting standard-candle tracks from cosmic-ray muons with well-modeled energy deposition profiles. To ultimately reconstruct the total kinetic energy of the incident charged particles, and by extension the incident neutrino energy, additional corrections must be applied to account for charge attenuation from electronegative impurities (detailed in Section 4.1.3) and the electron-ion recombination fraction (discussed below).

Accurate calorimetric reconstruction further requires a detailed accounting of electron-ion recombination, the process by which a fraction of the initial ionization electrons recombine with

argon ions before drifting. The microphysics involved in the scintillation emission process is described in Section 3.4.1. Because direct in-situ measurement of recombination on an event-by-event basis is experimentally unfeasible, analytical and empirical models are typically employed.

Historically, the Birks model [100] and the modified Box model [101] have served as the standard empirical frameworks for LArTPCs, describing the recombination experienced by a variety of particles in LAr. While these formulations successfully describe recombination effects for high-energy track-like particles, they break down when applied to low-energy electronic and nuclear recoils. This limitation arises because track-based frameworks do not capture the discrete microphysics that dominate for low-energy interactions in LAr. This deficiency motivated the development of the Noble Element Simulation Technique (NEST) [102, 103], a model semi-empirically tuned to a diverse suite of low-energy recombination measurements across a wide range of electric field strengths and incident particle energies, developed specifically for both liquid xenon [102] and liquid argon [103]. NEST provides well-characterized ionization-electron and scintillation-photon yields for electron-recoil, nuclear-recoil, and alpha events. This model is used for the analysis in this thesis (Chapters 5 and 7) to accurately simulate β and γ interactions in the 2x2 Demonstrator ND-LAr prototype.

4.1.3 Electron Lifetime Calibration

In addition to electronics gain, the attenuation of the primary ionization signal due to drift across the TPC must be meticulously corrected. Electronegative impurities within the bulk LAr, primarily oxygen and water, capture drifting ionization electrons, thereby reducing the total charge arriving at the anode planes as a function of drift distance [82, 104, 105]. While DUNE implements a high-capacity LAr purification and recirculation system designed to minimize this attenuation [9], sub-percent level corrections remain imperative for precision neutrino calorimetry.

The rate of change of the moving ionization charge due to attachment processes can be parameterized as follows [106, 107]:

$$\frac{dQ}{dt} = -k_A \cdot n \cdot Q, \tag{4.1}$$

where k_A represents the electron attachment rate constant specific to the impurity species, n is the volumetric concentration of that impurity, and Q is the instantaneous ionization charge during drift toward the anode plane. Integrating this expression yields the total survival charge as a function of drift time, $Q(t_{\text{drift}})$:

$$Q(t_{\text{drift}}) = Q(0) \cdot e^{-k_A \cdot n \cdot t_{\text{drift}}}, \tag{4.2}$$

where t_{drift} is the total transit time of the ionization electrons from the interaction point to the anode. To quantify the collective impact of these electronegative impurities without requiring explicit knowledge of individual attachment rate constants or distinct chemical concentrations, the bulk detector purity is conventionally expressed via the electron lifetime, τ, defined as:

$$\tau \equiv \frac{1}{k_A \cdot n}. \tag{4.3}$$

Substituting Equation 4.3 into Equation 4.2 yields the standard attenuation formula:

$$Q(t_{\text{drift}}) = Q(0) \cdot e^{-t_{\text{drift}}/\tau}. \tag{4.4}$$

By encapsulating multiple unconstrained impurity sources into a single macroscopic parameter, the electron lifetime serves as a convenient metric for monitoring and calibrating the liquid argon purity.

Utilizing Equation 4.4, the true initial ionization charge can be recovered provided two experimental inputs are known: the individual event drift time and the overall electron lifetime. For non-beam datasets, such as those where cosmic-ray data are collected, the drift time is determined by calculating the temporal separation between the prompt scintillation light trigger and the subsequent charge arrival at the anode plane. Alternatively, for accelerator-based physics experiments such as DUNE, the interaction time can be inferred directly from the high-precision synchronization signal provided externally by the accelerator.

In surface-level or shallow-depth LArTPCs, τ can be spatially and temporally mapped across

the active volume utilizing high-statistics data samples of cosmic-ray muon tracks traversing the TPC. By correlating the track trajectory with the precise interaction time (t_0) determined from the detection of prompt scintillation light, the signal attenuation can be robustly characterized and corrected. Alternatively, the electron lifetime can be measured without prompt scintillation light by leveraging "anode-to-cathode" cosmic-ray muon tracks that fully span the drift distance. In this configuration, the charge-readout timing alone provides the relative time intervals between individual track segments, allowing the drift time to be resolved independently of the absolute interaction time.

However, deep-underground detectors, such as the DUNE Far Detector, must perform electron-lifetime calibrations despite an extremely low cosmic-ray flux that precludes fine-grained tracking statistics. To compensate, dedicated purity monitors (detailed in Sections 4.1.5 and 6.3.4) can be deployed to provide real-time, independent measurements of the bulk LAr purity. A primary drawback of purity monitors, however, is that they measure the electron lifetime at a localized position outside the active TPC volume, leaving them unable to capture spatial variations or convective non-uniformities within the bulk detector. Consequently, mapping the internal electron lifetime across the fiducial volume requires alternative methodologies, motivating the investigation into the utilization of naturally occurring and artificial sources of radioactivity for LArTPC calibration.

4.1.4 Electric Field Calibration

Beyond energy-scale corrections, independent spatial calibrations are required to ensure the high-fidelity reconstruction of interaction vertices. This spatial mapping relies primarily on characterizing local electric field (E-field) distortions and the associated electron drift velocity. Localized E-field non-uniformities cause drifting ionization electrons to deviate from their nominal parallel trajectories, introducing systematic biases into the reconstructed spatial coordinates. Furthermore, variations in the local E-field magnitude alter the electron drift velocity from its expected value, directly degrading the precision of the longitudinal drift coordinate reconstruction. To mitigate these systematic shifts, it is critical to implement a data-driven correction map that accounts for

these localized field non-uniformities.

The longitudinal drift coordinate x in a LArTPC can be calculated using the following formula:

$$x = x_{\text{anode}} \pm v_{\text{drift}} \cdot \Delta t, \tag{4.5}$$

where x_{anode} is the x-coordinate of the anode plane, v_{drift} is the electron drift velocity, and Δt is the precise time interval between the occurrence of the interaction and when the ionization electrons are detected at the anode. The drift coordinate requires explicit tracking because, unlike the transverse coordinates localized directly by physical wire or pixel channel positions, it relies on an integrated time-of-flight. The drift velocity of electrons changes with the value of the E-field and LAr temperature, and is well known at the nominal DUNE E-field strength of 500 V/cm.

A common source of E-field distortions in LArTPCs near the surface is space charge. Cosmic-ray-induced activity is abundant at the surface and produces a large number of free argon ions, which slowly drift to the cathode and build up to levels that can cause significant E-field distortions. For instance, ProtoDUNE-SP for DUNE [108] and MicroBooNE [109] had to address these space charge effects because they were located at the surface. To handle this, a space charge calibration is performed to establish a data-driven map of the electric field distortions, enabling correction of the effect. On the other hand, the DUNE ND-LAr detector would not experience a significant space-charge effect, due to its underground location as well as it's reduced TPC size compared to typical large LArTPCs.

For underground detectors, such as the DUNE FD and ND-LAr, this effect is negligible due to the much lower cosmic-ray flux. However, other sources of distortions can also exist; for instance, certain detector materials could charge up; or a failed element of the field cage could introduce E-field distortions. As a result, it is still critical to monitor the FD LArTPCs for E-field non-uniformities. Given the low cosmic-ray flux, this motivates the investigation of using alternative techniques that utilize low-energy radioactive sources to perform an E-field calibration.

4.1.5 Common Calibration Techniques and Sources

Cosmic-Ray Muons

For detectors located near the surface, cosmic rays are a source of minimum-ionizing particles (MIPs), which are useful for LArTPC detector calibration. Cosmic-ray muons often create clean, straight tracks in a LArTPC and have a known stopping power of about 2.1 MeV/cm. They are useful for measuring electron lifetime by fitting the observed dQ/dx vs drift time with the exponential function from Equation 4.4. They can also be used to measure E-field distortions in a LArTPC by measuring the deviations of the tracks from the expected trajectories. Channel-by-channel electronics gain variations can also be measured using cosmic-ray muons utilizing their known energy deposition characteristics in LAr. However, as already mentioned, the low cosmic-ray flux deep underground limits the effectiveness of calibrations utilizing cosmic-ray muons.

Purity Monitors

A common device utilized alongside LArTPCs is a purity monitor. A purity monitor is a dedicated device used to measure electron lifetime measurements in LAr, described in further detail in Section 6.3.4. Because they can be employed in both deep-underground and on-surface environments, they are very useful as a method for monitoring LAr purity in LArTPCs. However, purity monitors only measure electron lifetime at a single location, which is an average measurement across the length of the device, not within the LArTPC itself. This means that electron lifetime measurements from purity monitors would not be ideal for calibrating out ionization electron attenuation, since LAr purity may differ inside and outside the TPC, and it may vary across the detector volume. Therefore, it is still important to measure the electron lifetime at various points within the TPC to ensure charge attenuation is correctly calibrated. Nevertheless, purity monitors remain very useful for detecting sudden changes in electron lifetime, especially when the LArTPC cannot operate, such as when the detector is being filled with LAr.

4.2 Calibrations Utilizing Sources of Radioactivity

4.2.1 Using Ambient Radioactivity for Calibrations

There are sources of naturally occurring radioactivity that can be detected in a LArTPC. In liquid argon retrieved from the atmosphere (as used in DUNE), ^{39}Ar β decays, which have an endpoint energy of about 565 keV, occur at a rate of about 1 Bq/kg [110]. For one of the DUNE FD modules, the total fiducial mass will be about 10 kt, leading to around ten million ^{39}Ar β decays occurring in the fiducial volume per second. These decays have a well-known spectrum and occur at a high rate, making them an excellent candidate for a calibration source.

However, there are some reasons why using ^{39}Ar β decays for calibration would be challenging. For one, given their low energies and the substantial size of typical LArTPCs, detecting them using photon detectors mounted at TPC boundaries would be challenging. This difficulty therefore limits the capability of a LArTPC to reconstruct the drift coordinate of low-energy decays. On the other hand, light-tagging events in close proximity to photon detectors may be viable in certain cases, offering potential applications for calibration. Even with only 2D reconstruction, ^{39}Ar β decays may still be useful, such as for channel-by-channel gain measurements.

Another issue is that other radiological sources can produce signals in a LArTPC that overlap in energy with signals from ^{39}Ar β decays. For instance, ^{85}Kr β decays are expected to occur in the LAr bulk at a lower rate than ^{39}Ar [110], but with a similar spectrum and slightly higher endpoint energy (687 keV). The structural materials in or surrounding a LArTPC could also be sources of radioactivity. Most materials contain trace amounts of ^{232}Th and ^{238}U which produce a variety of α, β, and γ decays. For example, ^{208}Tl is a significant contributor toward the end of the thorium chain and emits a characteristic 2.614 MeV gamma ray, which was notably observed in the MicroBooNE detector coming from G10 fiberglass structural support struts [111]. Materials outside the detector can also be sources of background, such as γ rays from the ambient rock in underground caverns. While deep underground deployment drastically reduces the cosmic-ray background, rock radioactivity creates a flux of high-energy gammas and neutrons that must be managed. These external backgrounds can be significantly reduced with active and passive shielding layouts. For example,

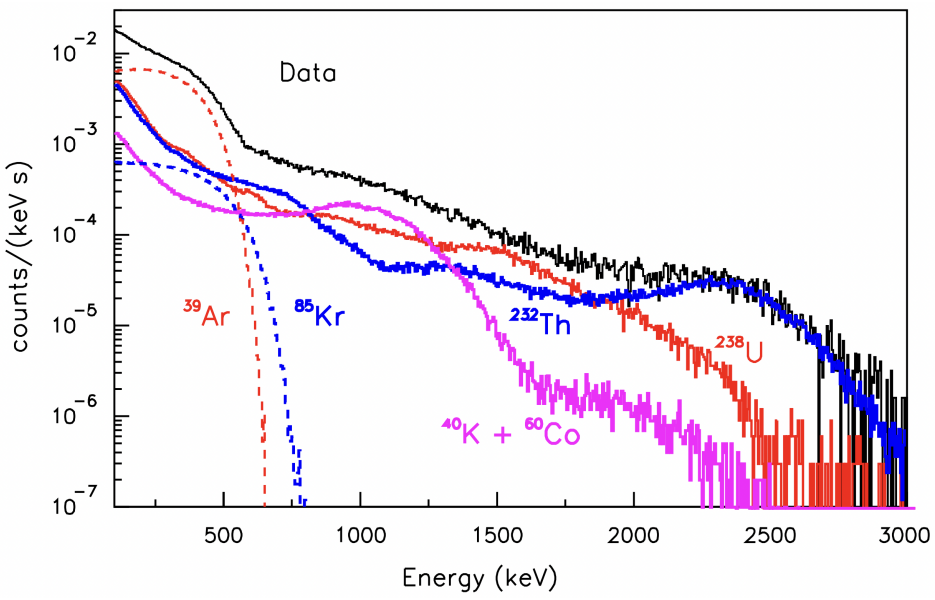

Figure 4.1: Spectra observed in the WARP LArTPC from multiple radiogenic sources, namely from β decay sources in the LAr and external γ sources. Image retrieved from Ref. [110]. Note that the detector was surrounded by 10 cm of lead shielding.

the WARP collaboration published a detailed study of radioactive backgrounds in an underground LArTPC [110], providing useful insight into the kinds of radioactive backgrounds that may be present in the DUNE LArTPCs.

4.2.2 Injected Calibration Sources: ^{220}Rn vs ^{222}Rn

To overcome the limitations of ambient radioactivity and provide uniform, controllable low-energy calibrations throughout the active volume, the intentional introduction of gaseous radioisotopes into noble element detectors has been widely investigated. Extensive previous research has established the utility of injecting radon isotopes, for instance ^{220}Rn (also called "thoron") into liquid xenon dark matter detectors such as XENONnT [112, 113] and liquid argon systems for ^{222}Rn such as MicroBooNE [114]. There appears to be very little research into ^{220}Rn injection for LArTPCs, which motivated the investigations into the technique described in this thesis.

The ^{222}Rn isotope has a relatively long half-life of 3.8 days, allowing it to distribute uniformly via diffusion and fluid convection currents across massive volumes. However, its decay chain includes the long-lived daughter isotope ^{210}Pb ($t_{1/2}$ = 22.3 years), which could plate out onto

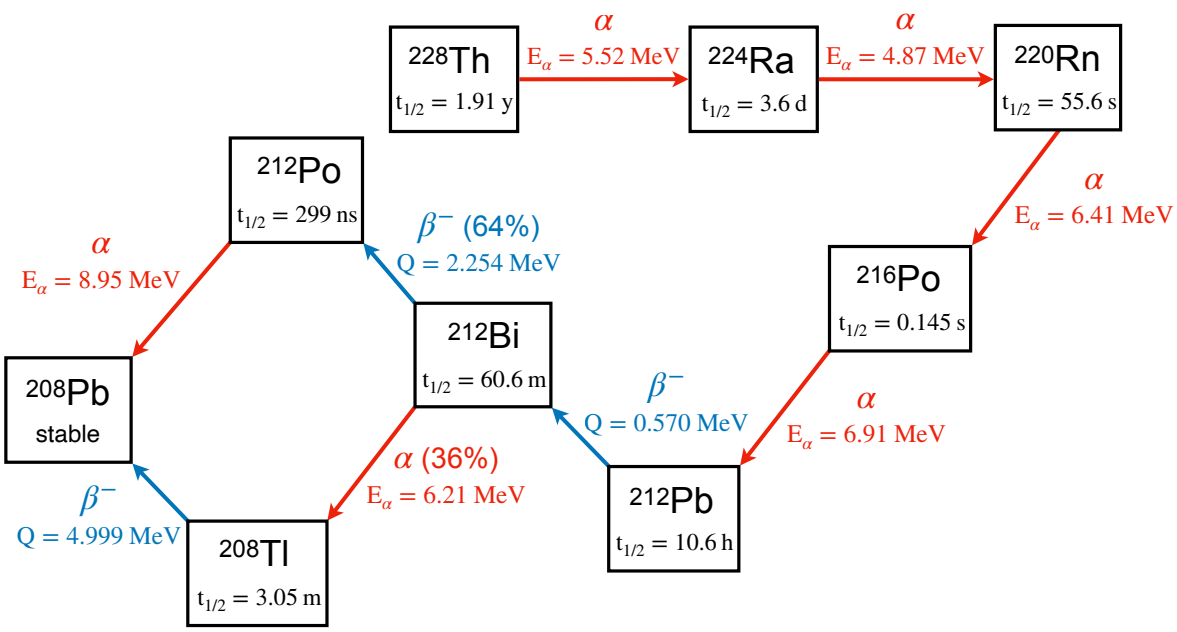

Figure 4.2: The α and β decays associated with the ^{228}Th decay chain. The longest half life in the chain is 10.6 h, with everything decaying eventually to the stable ^{208}Pb after several days.

TPC surfaces, creating permanent, long-lived internal backgrounds that would hinder subsequent high-sensitivity low-energy physics searches.

In contrast, the ^{220}Rn isotope ($t_{1/2} = 55.6$ s) and its immediate daughter ^{216}Po ($t_{1/2} = 0.145$ s) decay away rapidly. The longest-lived intermediate daughter in the ^{220}Rn decay chain is ^{212}Pb, which possesses a manageable half-life of only 10.6 h. This ensures that the entire injected activity completely decays away to negligible levels within around a week post-injection, introducing no long-term radiological contamination into the fiducial volume. It also potentially supports effective mixing throughout the active volume of a LArTPC given its relatively long lifetime. For this reason, ^{220}Rn provides a significant advantage over ^{222}Rn for large-scale LArTPCs.

4.2.3 Calibration Applications of the ^{220}Rn Decay Chain

The deliberate injection of ^{220}Rn into a LArTPC provides distinct topological, temporal, and calorimetric signatures highly valuable for comprehensive detector calibration. Rather than relying on a single isolated transition, this technique leverages specific daughter cascades to untangle independent detector response parameters across both the charge and light readout subsystems.

BiPo Tagging

One of the primary calibration channels exploited from an injected thoron system is the delayed coincidence between ^{212}Bi and ^{212}Po (commonly denoted as a "BiPo" pair). In this sequence, ^{212}Bi undergoes a β decay (endpoint energy of 2.25 MeV) directly to the highly unstable alpha emitter ^{212}Po ($t_{1/2} \approx 294$ ns). Because of the relatively short half-life of the intermediate ^{212}Po state, their respective signals may temporally overlap within both the charge and light readout windows of the TPC.

This distinct pairing is uniquely powerful due to the highly complementary microphysics of the two decay modes. The ^{212}Po α decay generates a large, localized deposition of kinetic energy (8.78 MeV) characterized by a high local stopping power (dE/dx). This dense ionization track experiences extreme recombination and significant quenching within the electronic channel, which drastically limits its free ionization charge yield. Concurrently, this high recombination rate produces an exceptionally large prompt scintillation light emission with an elevated prompt-to-late intensity ratio, which enables utilizing pulse shape discrimination to separate different particle populations. Conversely, the preceding ^{212}Bi β decay behaves as a minimum-ionizing electronic recoil, producing a continuous, low-dE/dx ionization track that avoids severe quenching and deposits a significant, readily detectable amount of free ionization charge.

This combination provides an ideal mechanism for absolute interaction time (t_0) tagging. The massive scintillation flash from the α decay serves as an effective way to determine the t_0 for the β decay. This optical trigger uniquely tags the spatial vertex and time of the overlapping β decay, bypassing the traditional blind-spot where β signals are difficult to successfully tag with photon detectors in LArTPCs. This absolute 3D localization enables a suite of core calibrations described in previous sections, including precise measurements of electron lifetime, measurements of E-field and ionization electron drift speed, and energy scale calibrations. Additionally, the high light yield from the ^{212}Po α decay presents an effective way to measure absolute and relative light yield variations across a large LArTPC.

RnPo Cascades and Convective Flow Mapping

Beyond localized microphysics calibrations, earlier steps in the thoron chain, specifically the sequential α decays of ^{220}Rn ($t_{1/2} = 55.6$ s) and ^{216}Po ($t_{1/2} = 0.145$ s), often referred to as "RnPo" pairs, offer a powerful mechanism for tracking fluid dynamics in LArTPCs. Because ^{216}Po is short-lived but long enough to be carried by bulk fluid transport, mapping the spatial displacement between the initial ^{220}Rn α vertex and the subsequent ^{216}Po α vertex allows for the direct assessment of internal LAr convective flow patterns and macroscopic mixing behavior.

However, executing these RnPo flow measurements introduces severe experimental challenges in DUNE-scale LArTPCs. Because both steps in this cascade are pure alpha decays, they suffer from the previously described heavy ionization quenching. The resulting sub-MeV equivalent electronic charge yields generate signals that frequently fall below the nominal front-end electronic noise thresholds of long-drift wire or pixel channels. While dark matter experiments like XENON100 have successfully leveraged similar spatial matching techniques in liquid xenon [113] due to lower thresholds and compact geometries, isolating these pure α-α pairs in large-scale neutrino detectors requires exceptionally low-noise charge readouts or highly sophisticated optical tracking networks capable of isolating localized photon clusters without corresponding TPC charge tracks. However, tagging the two-alpha topology using photon detectors would be more more achievable, given the bright scintillation flashes produced by α decays. Therefore, they may still have some utility for assessing LAr flow patterns without charge information, but to a lesser extent.

Chapter 5

Low-Energy Calibrations Using Ambient Radioactivity in the 2x2 Demonstrator

5.1 Cryogenic System

A schematic of the 2x2 Demonstrator system is shown in Figure 5.1. Section 3.7.4 provides a detailed overview of the 2x2 system, while this section provides an overview of the 2x2 cryogenic system. Liquid argon (LAr) from the supply dewars is first passed through a series of purification filters to remove electronegative contaminants, primarily oxygen and water. The filtered LAr is then transported via transfer lines to the top of the cryostat, where it is distributed through a manifold that splits into four separate fill lines, one for each module. This configuration delivers purified argon directly into the active volume, promoting uniform filling across the detector.

The active LAr volume has a mass of approximately 2.4 metric tons, while the total LAr volume corresponds to approximately 8 metric tons when the cryostat is full. In addition to the fill system, the cryostat is equipped with a continuous recirculation and recondensation loop. Boil-off argon gas from the ullage region beneath the cryostat lid is extracted and re-condensed using a cryocooler before being returned to the liquid phase. A Barber Nichols pump located near the bottom of the cryostat drives liquid recirculation through the external filtration system, after which purified LAr is reintroduced into the detector volume. This closed-loop system maintains long-term argon purity during operation.

5.2 Low-Energy Reconstruction

5.2.1 Charge Reconstruction Methodology

The 2x2 Demonstrator charge readout system was described in Section 3.7.1. The reconstruction of charge activity is performed using data collected by the 2x2 Demonstrator by grouping

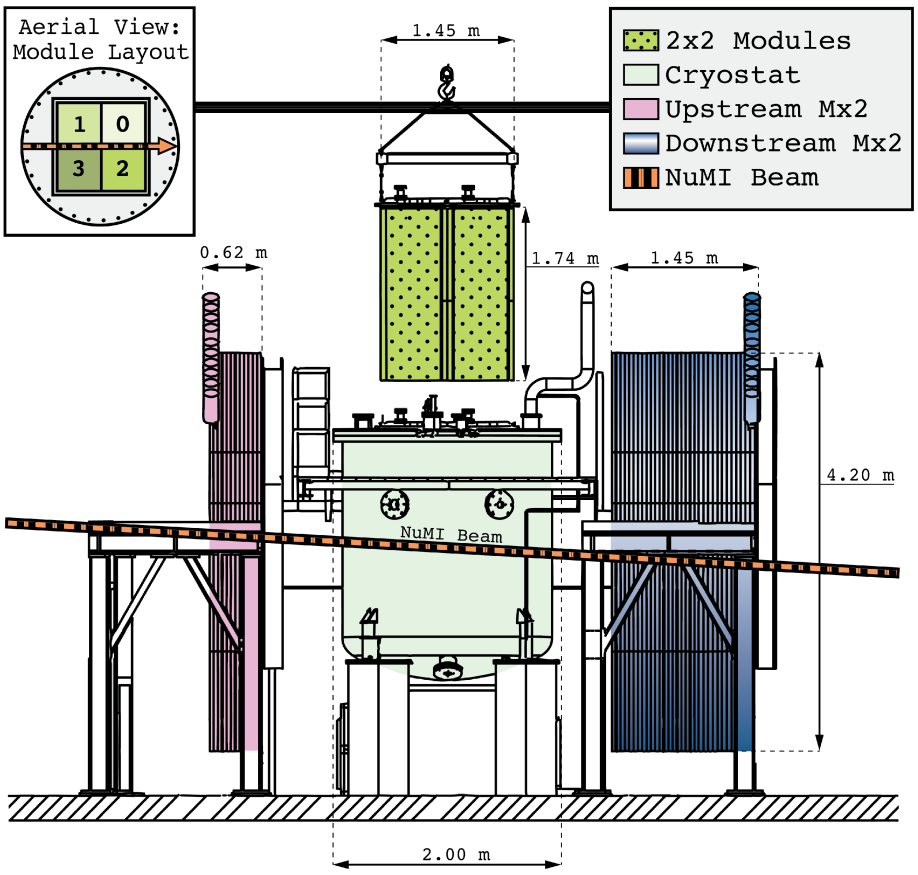

Figure 5.1: Schematic of the 2x2 Demonstrator in the NuMI neutrino beam at Fermilab. MINERvA modules, which are downstream and upstream of the 2x2 LArTPCs and used for tracking through-going muons, is not used for the analysis presented here. Image retrieved from Ref. [115].

spatially and temporally proximate pixel digitizations (commonly referred to as "hits") using the density-based clustering algorithm DBSCAN [116]. To facilitate three-dimensional clustering, hit timestamps are mapped to a spatial coordinate along the drift axis by multiplying the drift time by the expected electron drift speed in LAr. At the nominal DUNE drift electric-field strength of 500 V/cm, this speed is approximately 0.16 cm/µs. The DBSCAN algorithm is configured with a nearest-neighbor distance parameter of 2 cm, corresponding to approximately four to five pixels, and a minimum cluster size of one hit. This distance should be sufficiently large to group together hits from individual low-energy electron tracks. While a minimum cluster size of 1 increases susceptibility to noise hits, this threshold is necessary to maximize the reconstruction efficiency for low-energy interactions, which can frequently manifest as single-pixel hits. The clustering process

is applied to the whole dataset prior to matching to light activity.

The total charge of a cluster is defined as the sum of its constituent hit charges. A detailed description of the calibration and conversion of raw signals to total cluster charge in kiloelectrons (ke^-) is provided in Section 3.7.1. To reject backgrounds from cosmic-ray interactions, the selection restricts the sample to clusters containing ten or fewer hits. Furthermore, low-energy clusters closely correlated in time with track-like activity are mitigated by applying additional constraints on local cluster multiplicity and size; this topological filter is implemented following the charge-light matching procedure detailed in Section 5.2.2.

5.2.2 Charge-Light Matching Methodology

Charge-light matching is performed by correlating synchronized timestamps across the charge readout system (CRS) and light readout system (LRS). Because the pulse-per-second (PPS) signal is distributed simultaneously to both systems to reset their local clocks, the relative time elapsed since the preceding PPS pulse serves as the primary metric for cross-system event matching. To resolve ambiguities across different seconds, a Unix timestamp with second-level precision is used in tandem. At the nominal drift E-field strength of 500 V/cm and a maximum drift distance of approximately 30 cm within the 2x2 Demonstrator TPCs, the maximum electron drift time, and thus the physical bound on the time separation between a light trigger and its associated charge signal, is roughly 190 µs.

For the analysis presented in this work, a PPS timestamp matching window spanning from -100 µs to $+300$ µs relative to the light trigger time is enforced. This window is sufficiently wide to encompass charge deposition occurring near the cathode, while incorporating a 100 µs buffer region on either side to facilitate sideband evaluation and purity estimation. Concurrently, matched events are required to fall within the same one-second Unix timestamp interval. Because the maximum physical time difference between corresponding signals is well under one millisecond, the probability of a true coincidental pair straddling a second boundary is minimal, resulting in negligible efficiency loss.

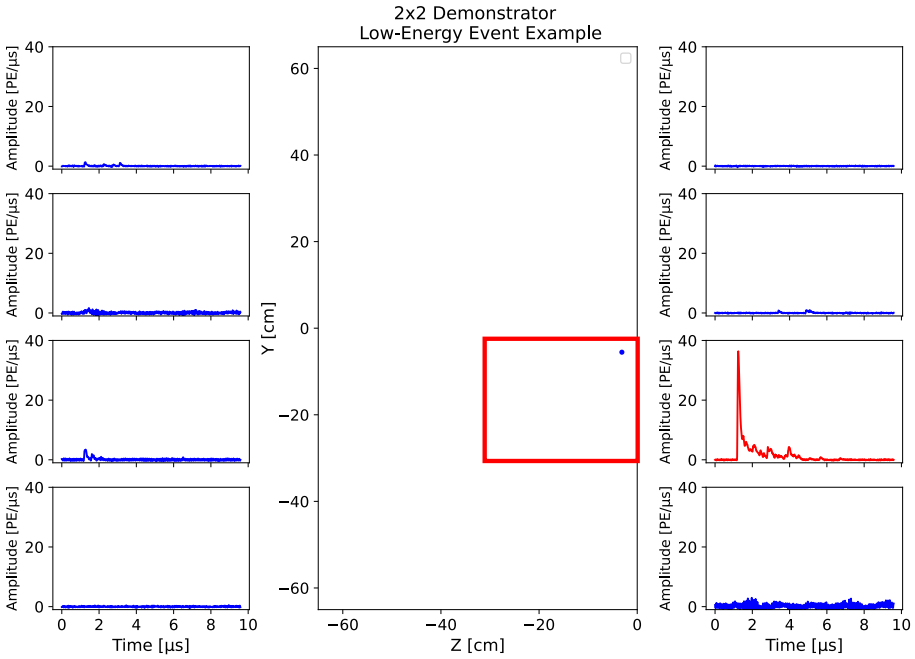

Figure 5.2: An example visualization of the optical proximity requirement for improving the selection purity of charge-light matching. The center panel shows the position of the reconstructed light-matched charge cluster. Left/Right panels show LRS waveforms for each light detector tile in the same TPC as the cluster. Red waveforms indicate optical hits found using a simple threshold on the waveforms in a given LRS trigger. The red box indicates the physical matching window around the light detector center.

As will be demonstrated in Section 5.3, temporal matching alone is insufficient to provide a high-purity sample of low-energy events; a spatial charge-light proximity requirement must be introduced to optimize selection purity (defined as the fraction of correctly matched charge-light pairs). Achieving high purity is essential for robust calibration applications. This topological filter requires that the reconstructed charge activity be physically localized near the specific light detector tile that registered an optical hit.

Here, an optical hit is defined as a summed waveform exceeding a predefined amplitude threshold. For the data collected at the University of Bern, an absolute gain calibration across all modules was unavailable to establish a conversion factor in units of ADC/PE (where PE refers to number of photoelectrons). For this reason, the hit-finder threshold for these runs was pinned to the hardware self-trigger threshold. Conversely, for the 2x2 Demonstrator dataset, an absolute LRS gain calibration was successfully performed across all modules, enabling the conversion of raw waveform amplitudes to PE. This permitted a standardized hit-finder threshold of 10 PE to be applied. This

value was selected to sit safely below the typical hardware trigger thresholds of individual light detectors while remaining sufficient to capture the primary signal driving any given LRS trigger. Figure 5.2 illustrates the spatial proximity requirement, which utilizes a rectangular selection window centered on the physical geometry of the light detector tile (corresponding to either one ACL or three LCMs) to isolate nearby charge activity.

5.3 Charge-Light Matching Selection and Performance

Figure 5.3 displays the reconstructed drift time (t_{drift}), defined as the difference between the cluster charge timestamp and the light trigger timestamp, for charge-light-matched events. To compute t_{drift}, the charge timestamp is defined as the midpoint of the constituent hit timestamps within a given cluster. A prominent event excess is observed between approximately 0 and 200 µs. This interval is designated as the "signal region" and corresponds to clusters reconstructed within the physical boundaries of the active TPC volume. Outside of this window, the drift time distribution remains flat, a characteristic signature of accidental matches between uncorrelated charge and light signals. The application of the optical proximity selection significantly suppresses these sideband regions while preserving a large sample of events within the signal region, qualitatively demonstrating that the spatial cut enhances the purity of the charge-light matching.

To quantify this improvement, a data-driven method has been developed to measure the selection purity using the pre-trigger sideband ($t_{\text{drift}} < 0$). The sideband is integrated over a baseline interval of $t_{\text{drift}} \in (-100\ \mu\text{s}, -50\ \mu\text{s})$ and subsequently scaled to the width of the signal region ($t_{\text{drift}} \in [0, 200\ \mu\text{s}]$). This scaling provides an estimate of the combinatorial background from accidental matches falling within the signal window. Subtracting this background estimate from the total integral of the signal region yields the net number of correctly matched signal events. The selection purity fraction is then defined as the ratio of the estimated signal counts to the total integrated counts within the signal region. Figure 5.4 shows the estimated selection purity evaluated as a function of cluster charge. Enforcing the optical proximity selection yields a purity exceeding 95% across nearly all charge values for all 2x2 Demonstrator LArTPC modules, fulfilling a critical

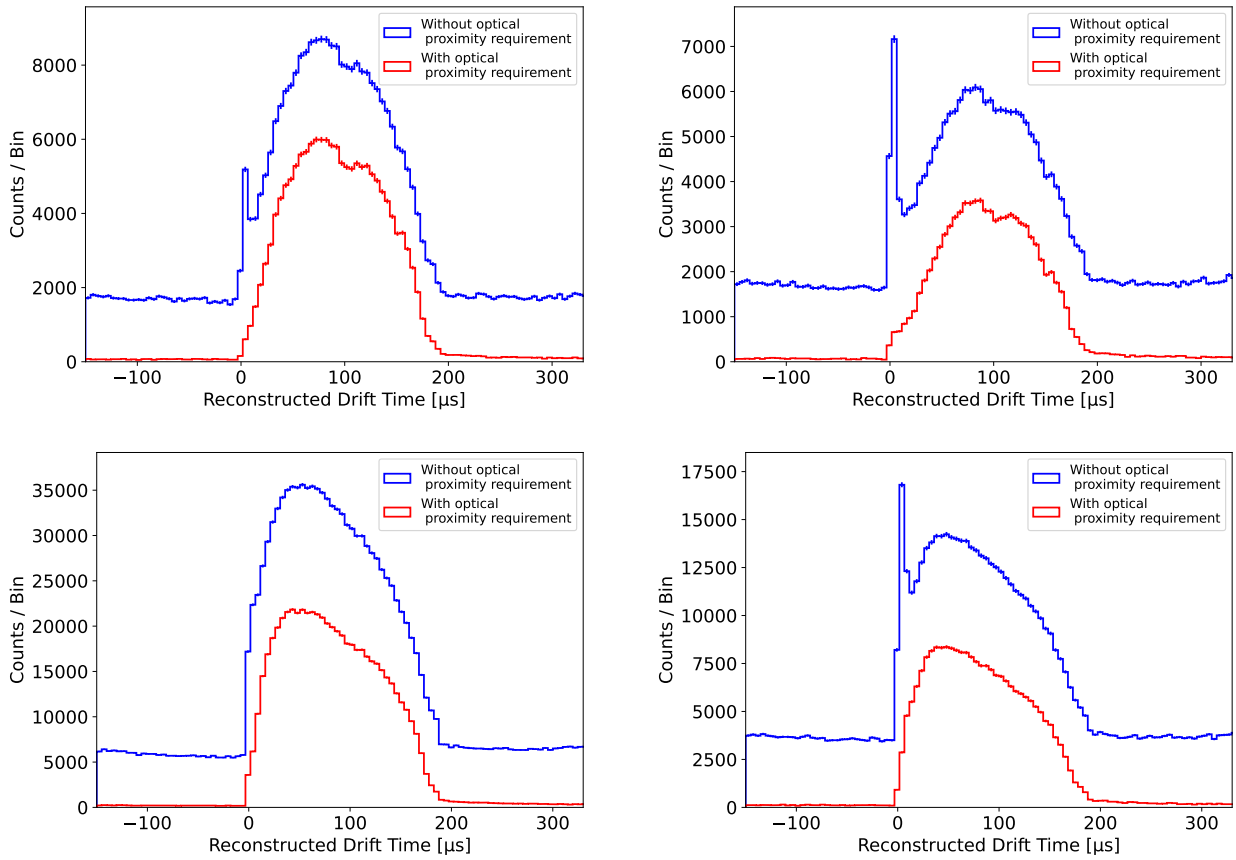

Figure 5.3: Reconstructed drift time of low-energy charge-light-matched charge clusters in Module-0 (top left), Module-1 (top right), Module-2 (bottom left), and Module-3 (bottom right) using data taken at the University of Bern.

prerequisite for utilizing low-energy activity in detector calibrations.

The spatial distributions in the y-z plane associated with charge-light-matched clusters obtained during the University of Bern tests are shown in Figures 5.5–5.8 for Module-0 through Module-3, respectively. The majority of charge clusters are reconstructed in close physical proximity to the light detectors ($z \approx \pm 31$ cm), reflecting the geometric acceptance and the higher photon detection probability for localized interactions. Localized excesses or "hot spots" are visible near the corners of the TPCs; these are likely caused by corner-crossing cosmic-ray muons that deposit small, isolated charge clusters capable of passing the hit-multiplicity cuts. To minimize this contamination, charge clusters located near the physical corners of the TPC are fiducially removed. Notably, this cosmic-ray background is significantly less pronounced in the 2x2 Demonstrator dataset due to a reduction of approximately two orders of magnitude in cosmic-ray flux, associated with the

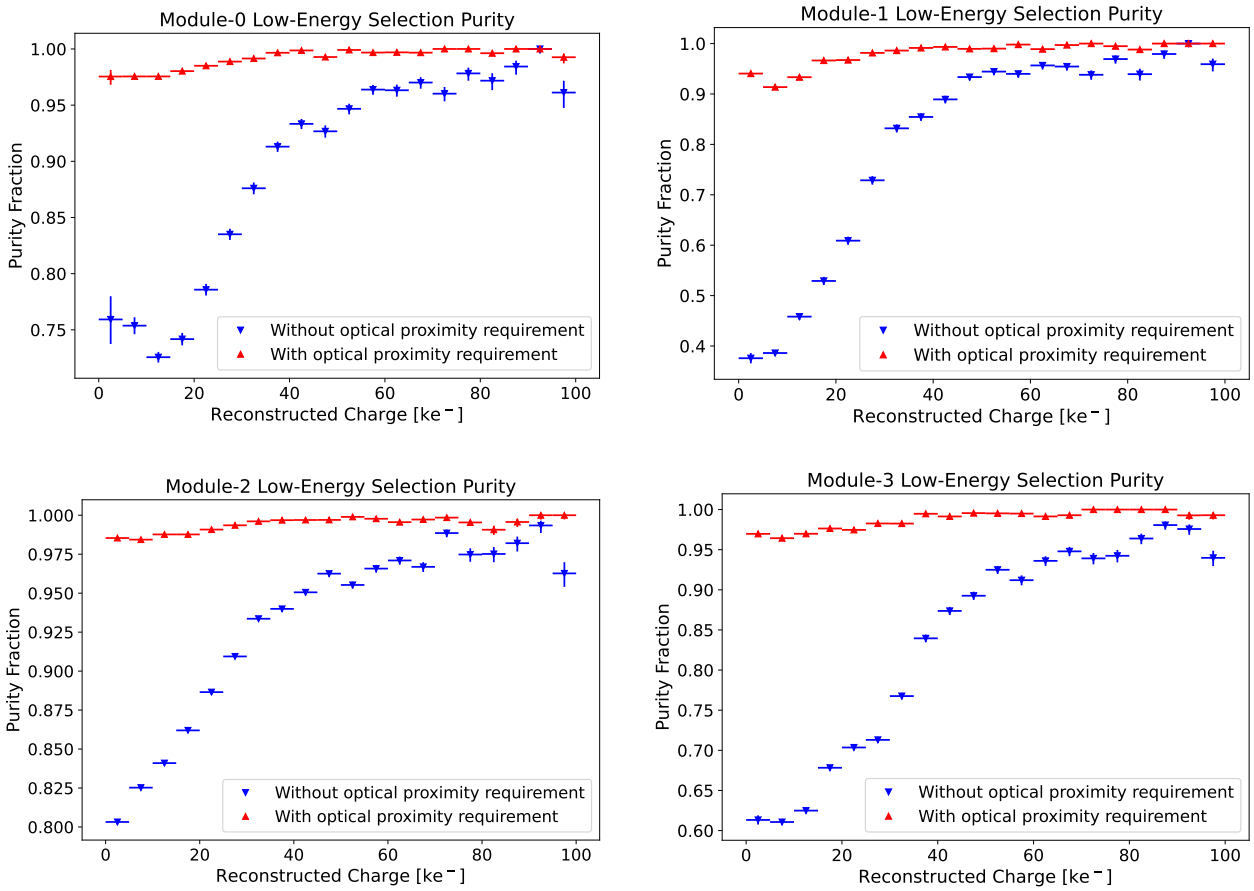

Figure 5.4: Selection purity of low-energy charge-light matched events in the single module tests at the University of Bern.

detector being located underground.

Furthermore, systematic module-to-module variations are observed in the total amount of reconstructed activity. These discrepancies are primarily driven by differences in the hardware and operational configurations of the LRS during the prototyping phase at the University of Bern, where SiPM bias settings and trigger thresholds were systematically varied to evaluate detector performance. Consequently, variations in event rates across modules do not inherently imply non-uniformities in intrinsic photon detection efficiency (PDE), except for the known higher PDE for LCMs compared to ACLs. A notable exception is Module-0, the first built prototype, which featured early-generation light detectors (primarily ACLs) with a lower PDE than those deployed in subsequent modules [23]. This performance baseline motivated the hardware optimizations imple-

mented in later iterations. Crucially, Modules 0 and 1 were operated with higher LRS thresholds than Modules 2 and 3, explaining why the latter two modules reconstructed significantly more activity further away from the light detectors. Ultimately, the Bern datasets demonstrate that minimizing LRS thresholds significantly increases the low-energy event yield, indicating that the primary constraint on low-energy reconstruction efficiency is dictated by the hardware configuration rather than light yield limitations.

Data collected from the 2x2 Demonstrator LArTPC modules during testing at the University of Bern are utilized to demonstrate the matching procedure and evaluate its effectiveness. The low-threshold LRS configurations deployed in the Bern tests enabled reconstructing more low-energy activity further from the light detectors. In effect, this helps to facilitate an evaluation of electric-field uniformity using low-energy activity (as detailed in Section 5.6.1) by ensuring that sufficient statistics are available to measure reconstructed spatial offsets at multiple locations within the TPC. The fundamental mechanics and performance of the charge-light matching process remain highly consistent between the Bern datasets and the 2x2 Demonstrator data, despite the higher LRS thresholds implemented in the latter. These elevated thresholds in the 2x2 Demonstrator were partially motivated by a strategy to mitigate bandwidth constraints associated with reading out all modules simultaneously upon an LRS trigger. However, subsequent findings from the ^{220}Rn injection test (Chapter 7) established that trigger rates could be increased by at least an order of magnitude without encountering data-acquisition bottlenecks. Crucially, the final selection purity for charge-light matched low-energy clusters within the 2x2 Demonstrator remains above 95% across all modules, as illustrated in Figure 5.9.

The data epoch analyzed for ambient radioactivity calibrations spans from October 10, 2025, to November 22, 2025, coinciding with the second data-taking run of the 2x2 Demonstrator. Certain data intervals within this period were excluded from the analysis due to non-optimal CRS or LRS hardware configurations, interleaved purity monitor runs that interrupted nominal data acquisition, or specialized calibration runs employing non-standard readout schemes. Nevertheless, the remaining dataset provides ample statistics to perform robust detector calibrations utilizing low-

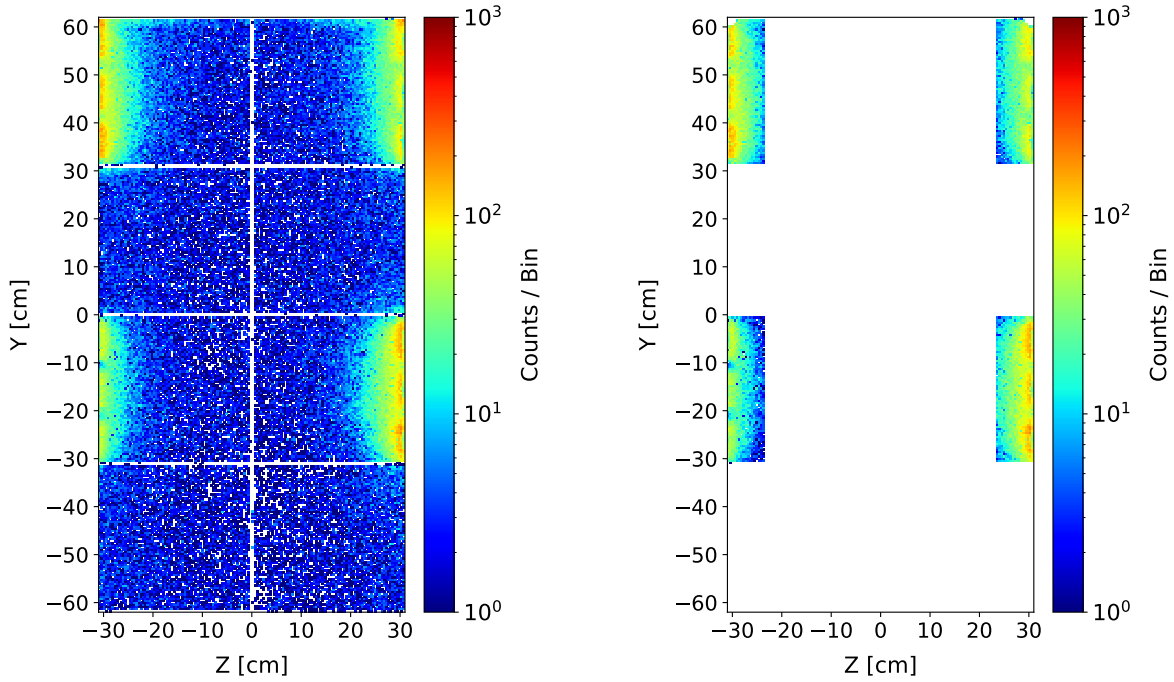

Figure 5.5: Spatial distributions of charge-light matched low-energy charge clusters in the single module test at the University of Bern with Module-0 before (left) and after (right) applying optical proximity selection.

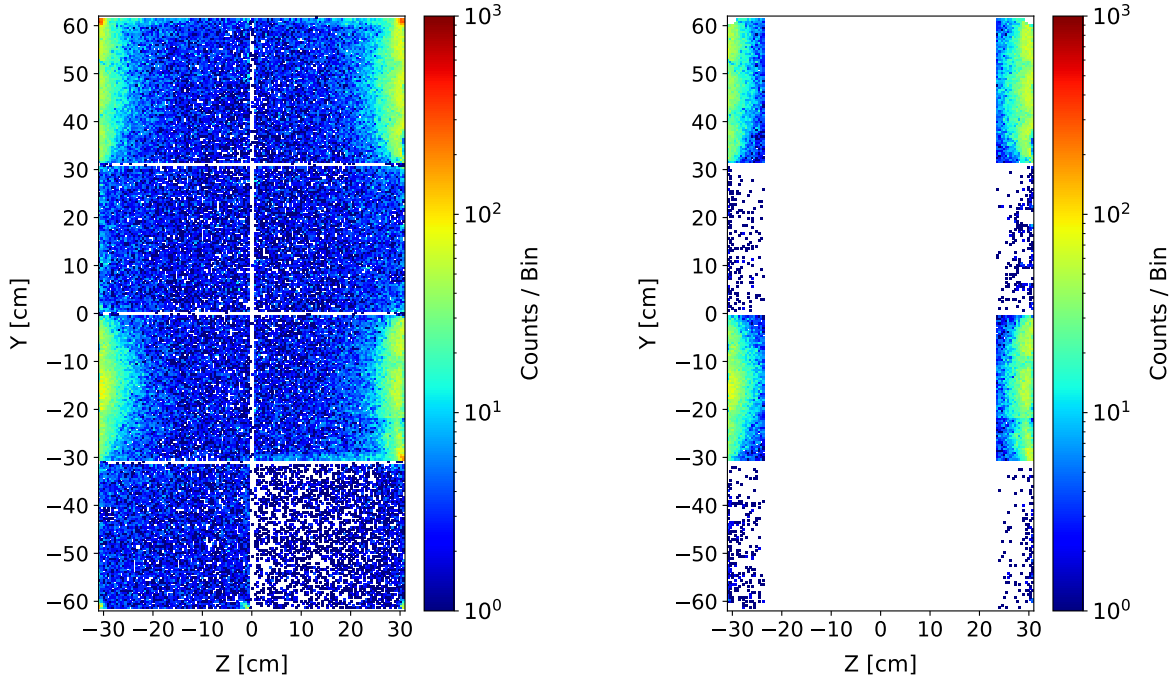

Figure 5.6: Spatial distributions of charge-light matched low-energy charge clusters in the single module test at the University of Bern with Module-1 before (left) and after (right) applying optical proximity selection.

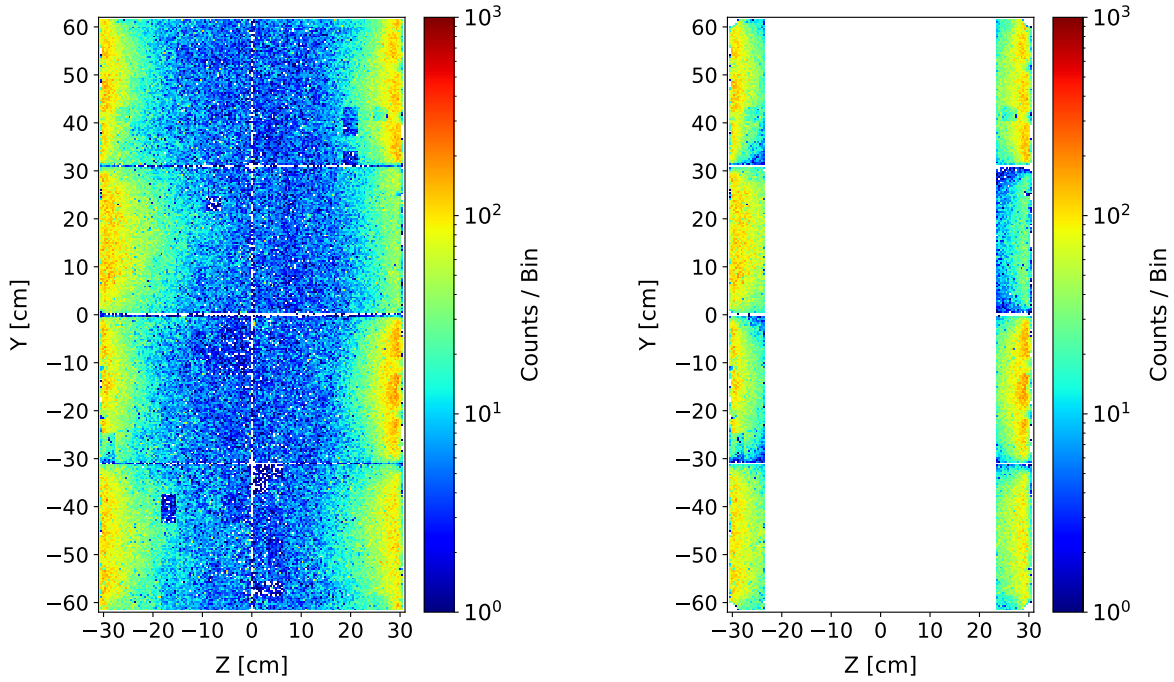

Figure 5.7: Spatial distributions of charge-light matched low-energy charge clusters in the single module test at the University of Bern with Module-2 before (left) and after (right) applying optical proximity selection.

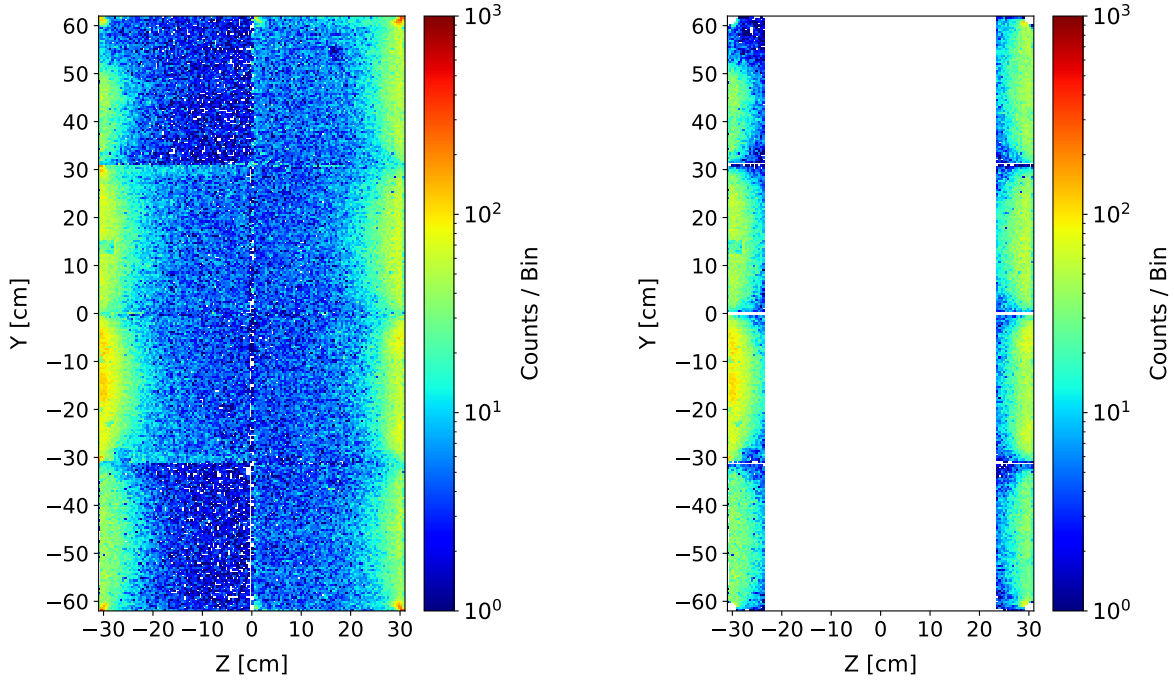

Figure 5.8: Spatial distributions of charge-light matched low-energy charge clusters in the single module test at the University of Bern with Module-3 before (left) and after (right) applying optical proximity selection.

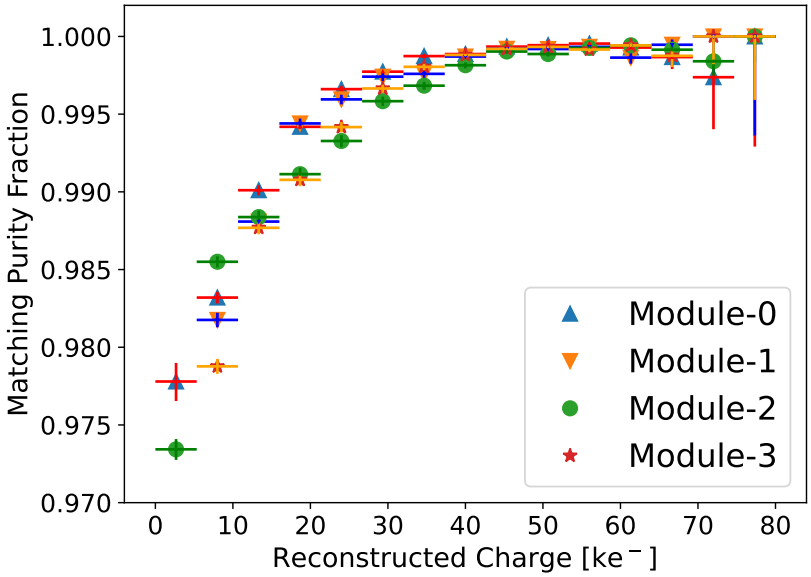

Figure 5.9: Selection purity for charge-light-matched low-energy charge clusters using data collected by the 2x2 Demonstrator at Fermilab (with optical proximity selection enforced).

energy activity, such as charge scale calibration and electron lifetime measurements, leveraging the low-background environment of the underground 2x2 Demonstrator deployment site at Fermilab.

5.4 Charge Scale Calibration

5.4.1 Reconstructed Charge Spectrum of Low-Energy Activity

A charge scale calibration provides a way to verify the alignment of data and MC simulation by anchoring the detector response to a well-understood reference signal (e.g. from cosmic-ray muons or radiogenic backgrounds). By fitting the experimental data with a model that incorporates a detailed simulation of the reference signal, an overall charge scale correction factor can be extracted. Beyond providing a simple normalization correction, this calibration offers a diagnostic tool to evaluate whether the fundamental detector response is properly modeled and whether underlying physical processes are correctly captured in simulation. Aligning the simulated spectrum with the experimental data often necessitates an additional resolution smearing component to account for unmodeled electronics noise, recombination fluctuations, or localized channel-to-channel gain variations. Ultimately, a charge scale calibration establishes a robust method for validating

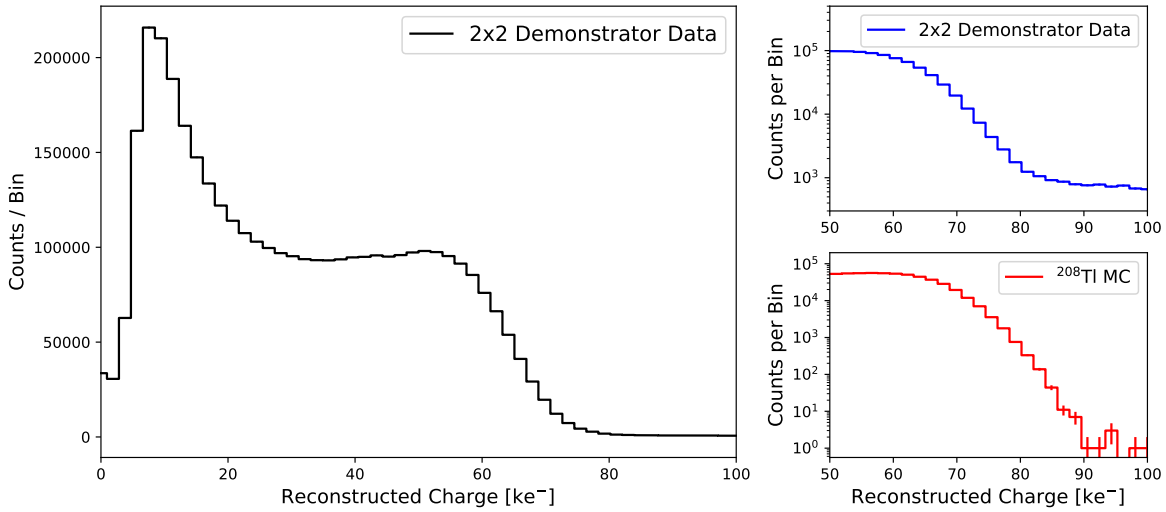

Figure 5.10: Reconstructed low-energy charge spectrum for charge-light-matched topological blips in the 2x2 Demonstrator, compared to a simulated spectrum of ^{208}Tl γ decays distributed uniformly throughout the TPC active volume. The charge spectra associated with both data and MC simulation incorporate electron lifetime corrections.

the modeling of ionization charge and detector response, serving as an essential prerequisite for accurate energy reconstruction.

In practice, isolating a clean reference signal can be challenging when multiple overlapping and indistinguishable background sources contribute to the observed spectrum. By investigating the reconstructed charge spectrum of charge-light-matched low-energy clusters, a candidate charge range is identified at high reconstructed charges. Figure 5.10 shows the charge spectrum using the 2x2 Demonstrator dataset in addition to the simulated spectrum of ^{208}Tl γ decays. A prominent spectral shoulder is visible above approximately 50 ke$^-$ in both the experimental data and the MC simulation, strongly suggesting that the event rate in this charge regime is dominated by signals originating from ^{208}Tl γ decays.

To further understand this observed feature, it is important to discuss the underlying γ interaction processes that can occur. Photons propagating through matter primarily undergo three distinct interactions that produce detectable ionization electrons in a LArTPC: the photoelectric effect, Compton scattering, and pair production [117]. The photoelectric effect involves the complete absorption of an incident photon by a bound atomic electron, resulting in the ejection of a pho-

toelectron. Because this process transfers nearly the entire energy of the primary photon to the kinetic energy of the outgoing electron, the ideal photoelectron energy distribution manifests as a delta function centered at the primary photon energy.

The second primary interaction is Compton scattering, in which an incident photon scatters off a quasi-free atomic electron at a deflection angle θ relative to its initial trajectory. At small scattering angles ($\theta \approx 0$), the photon retains nearly all its initial energy, transferring minimal kinetic energy to the target electron. Conversely, a maximum energy transfer occurs when the photon undergoes a catastrophic back-scattering event at $\theta = \pi$. The kinetic energy of the recoiling Compton electron (E_{e^-}) depends on the scattering angle θ, the incident photon energy $h\nu$, and the electron rest mass m_0 (0.511 MeV/c^2), governed by the kinematic relation [117]:

$$E_{e^-} = h\nu \left(\frac{\frac{h\nu}{m_0c^2}(1 - \cos\theta)}{1 + \frac{h\nu}{m_0c^2}(1 - \cos\theta)} \right). \tag{5.1}$$

Evaluating this expression at maximum back-scattering ($\theta = \pi$) for the 2.615 MeV line of ^{208}Tl yields a maximum recoil electron energy of approximately 2.382 MeV. This absolute kinematic boundary produces a distinct, sharp truncation in the continuous electron recoil spectrum, conventionally referred to as the "Compton edge." Because all intermediate scattering angles are physically allowed, the resulting Compton spectrum is continuous from zero up to this maximum threshold.

The third prominent photon interaction mechanism is pair production, wherein the incident photon is entirely annihilated within the intense nuclear Coulomb field to create an electron-positron pair. This process exhibits a strict kinematic threshold of $2m_0c^2 = 1.022\ \mathrm{MeV}$ to generate the invariant masses of the pair. The remaining kinetic energy is shared between the two particles according to this formula:

$$E_{e^-} + E_{e^+} = h\nu - 2m_0c^2. \tag{5.2}$$

As a result, the energy deposition from the pair resembles a spike at the energy of the photon. In an actual detector, the photon can lose energy by other means (e.g. by Compton scattering) before

pair production happens, meaning that the spike can occur at lower energies. In effect, this means the spectrum recorded by a detector could appear continuous from pair production in practice (similarly for the photoelectric effect).

The pair-produced positron eventually decelerates and annihilates with an ambient electron, generating a pair of 0.511 MeV photons. If both annihilation photons escape the active volume of the TPC without undergoing secondary interactions, their missing energy shifts the reconstructed event energy downward, producing a discrete feature known as the "double escape peak" at 1.593 MeV for a 2.615 MeV primary γ ray. If only a single annihilation photon escapes, a corresponding "single escape peak" appears at 2.104 MeV. The presence of these localized peak structures serves as an unambiguous topological signature of pair production.

Cross sections compiled in the Particle Data Group (PDG) review [16] dictate that at the relevant 2.615 MeV energy scale in LAr, Compton scattering remains the heavily dominant interaction mechanism over both the photoelectric effect and pair production. Consequently, the Compton edge profile forms the definitive structure driving the high-energy boundary of the ^{208}Tl spectrum inside a LArTPC. In the experimental data and MC simulation distributions shown in Figure 5.10, this high-charge boundary does not appear as an idealized, sharp drop-off; rather, it is noticeably smeared by detector effects, including electronics noise and recombination fluctuations.

To systematically verify that the observed high-charge feature is dominated by the ^{208}Tl Compton edge, a formal numerical fit can be performed directly on this region using the MC simulation templates. The comprehensive fit model must incorporate a component to account for residual cosmogenic muon-induced backgrounds, which, while strongly suppressed due to the 2x2 Demonstrator's underground location, remain a non-zero contribution. To robustly accommodate potential mis-modeling of the detector resolution within the simulation framework, a Gaussian smearing kernel is applied to the MC template during the fitting procedure. The final composite fit model thus evaluates the convoluted MC template alongside a background function parameterized to capture the cosmogenic continuum. The overall quality of this fit provides an explicit, quantitative test of whether the high-charge spectrum is indeed driven by environmental ^{208}Tl γ rays, while simulta-

neously extracting the absolute charge scale calibration factor necessary to precisely align the data and MC simulation energy scales.

5.4.2 Fitting Procedure

To extract the absolute charge scale calibration factor for the 2x2 Demonstrator, a composite analytical model is fit to the experimental data. Prior to executing the fit, attenuation of detected charge along the drift axis due to ionization loss from electronegative impurities must be corrected. For simulated events, this charge-attenuation correction is straightforward because the exact electron lifetime and the event drift time are known a priori. For the experimental data, a continuous history of the electron lifetime was established via a series of purity monitor measurements performed throughout the second data-taking run of the 2x2 Demonstrator. The discrete electron lifetime measurements are interpolated using a cubic spline fit to assign an instantaneous electron lifetime value to each charge cluster based on its absolute timestamp. The corresponding charge attenuation is then corrected on an event-by-event basis utilizing this interpolated electron lifetime and the drift time derived from the charge-light matching procedure. The purity monitor time-series data and the accompanying cubic spline fit are shown in Figure 5.11; the comprehensive spectrum previously presented in Figure 5.10 incorporates this correction.

The MC simulation chain utilizes `edep-sim`, a GEANT4 wrapper, to generate primary interactions from 2.615 MeV γ decays distributed uniformly throughout the TPC active volume. The resulting localized energy depositions are subsequently processed through `larnd-sim`, a dedicated software package that simulates the full TPC detector response, including electron-ion recombination, charge diffusion and attenuation along the drift axis, and the LArPix pixelated electronics response [118]. Recombination effects are modeled via the Noble Element Simulation Technique (NEST) framework [17], which provides an empirical parameterization benchmarked against an extensive suite of electron-recoil recombination measurements across varied energies and drift electric field strengths.

Because the MC simulation accounts for only the radiogenic isotope and omits a cosmogenic

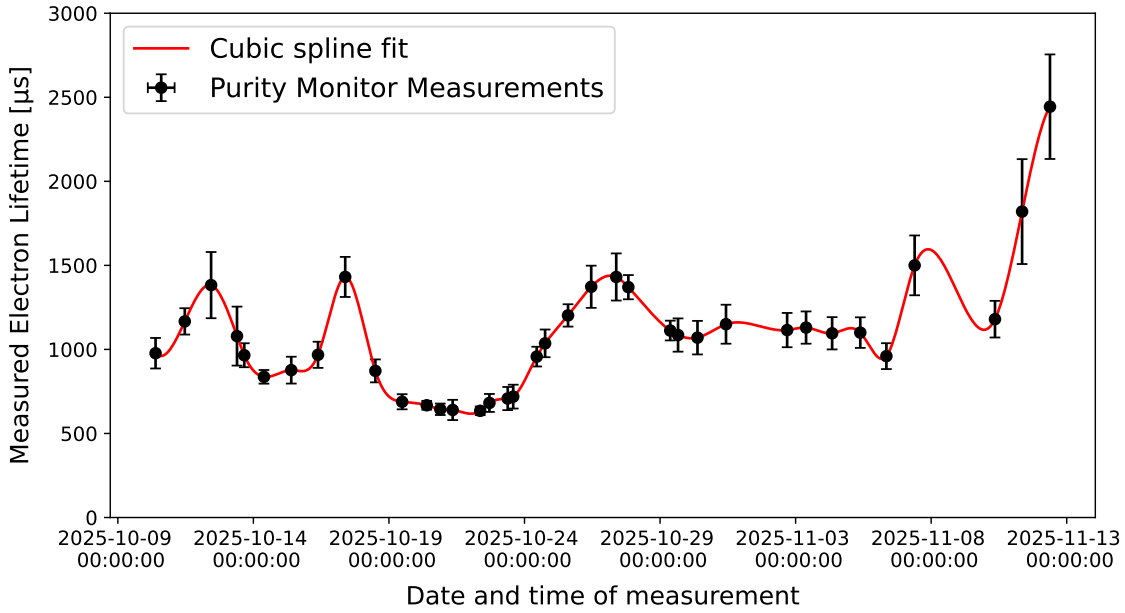

Figure 5.11: Purity monitor measurements recorded during the second data-taking run of the 2x2 Demonstrator, overlaid with the interpolating cubic spline fit. Each data point denotes the mean and standard deviation of eight consecutive measurements acquired at eight-minute intervals.

component, the fit model must account for the continuous background spectrum produced by residual cosmic-ray muons. A decreasing exponential function is chosen to account for the expected charge spectrum from the low-energy cosmic-ray-induced activity, which is expected to be dominated by electromagnetic secondaries, primarily δ rays [16].

The target fitting window is restricted to a charge range of 50–100 ke$^-$, which encapsulates the complete Compton edge feature alongside a portion of the high-charge continuum from cosmic-ray backgrounds. Restricting the lower bound to 50 ke$^-$ protects the fit from low-energy systematic biases introduced by unmodeled, sub-MeV radiogenic backgrounds that contribute to the spectrum. The composite function $f(Q)$ used to perform the optimization is defined as:

$$f(Q) = \left[N \cdot \text{Template}\left(\frac{Q}{s}\right) \otimes G(Q, \sigma_1, \sigma_2, F) \right] + \left[A \cdot e^{B \cdot Q} + C \right], \tag{5.3}$$

where N represents an overall normalization factor scaling the simulated ^{208}Tl template to the data statistics, Q is cluster charge in ke$^-$, and s is the absolute charge scale calibration factor that linearly adjusts the template along the charge axis. The background is parameterized by an expo-

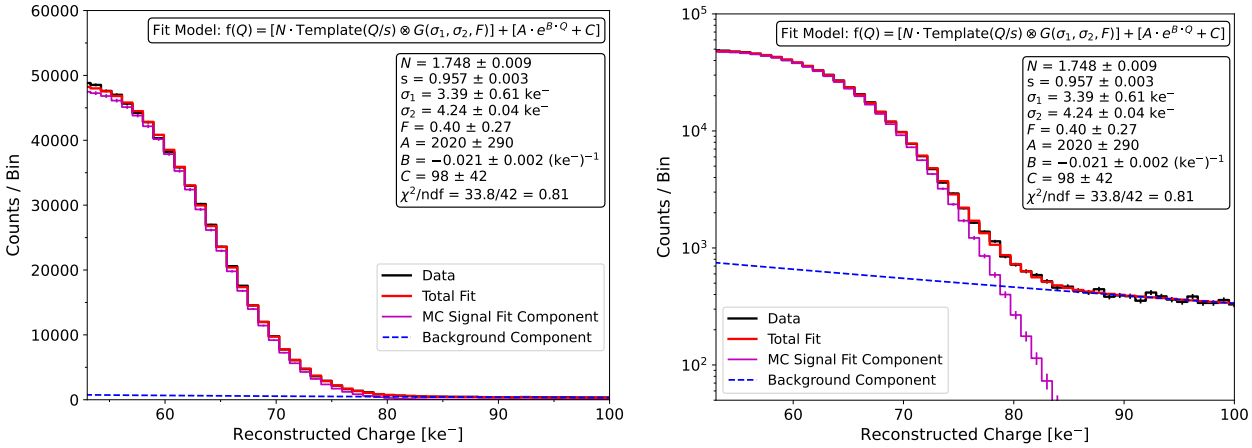

Figure 5.12: The result of the fit to the 2x2 Demonstrator low-energy charge spectrum using a model based on a ^{208}Tl γ decay MC simulation template, smearing kernel, and background model.

nential function with free parameters A, B, and C. To accommodate unmodeled detector smearing in the MC simulation, the template is convolved with a resolution kernel, G. A double-Gaussian formulation was selected for the kernel to capture multiple independent noise components, a choice motivated by historical characterizations showing that a small fraction of LArPix ASICs exhibit elevated electronic noise profiles relative to the nominal baseline channels [23]. The composite kernel G is defined as a linear combination of two Gaussian distributions:

$$G(Q, \sigma_1, \sigma_2, F) = (1 - F) \cdot G(Q, \sigma_1) + F \cdot G(Q, \sigma_2), \tag{5.4}$$

where σ_1 and σ_2 represent the standard deviations of the core and tail noise distributions, respectively, and the fraction F parametrizes the relative weight of the two smearing contributions.

5.4.3 Fit Results

The results of the optimization procedure, along with the extracted parameters, are illustrated in Figure 5.12. The absolute charge scale calibration factor is determined to be 0.957 ± 0.003 (stat), indicating that a 4.3% reduction in the simulated charge values is required to align the simulation with the experimental data. The reduced chi-squared value, $\chi^2/\text{ndf} = 33.8/42$, is close to unity, confirming that the composite model provides an excellent description of the physical data spec-

Table 5.1: Summary of uncertainties in the 2x2 Demonstrator charge scale measurement.

Parameter	**Value and Uncertainty**	**Measurement Deviation [%]**
Bin width	0.941 ± 0.094 ke$^-$	0.03
Fit start	53 ± 2 ke$^-$	0.01
Fit end	100 ± 2 ke$^-$	0.01
Electron lifetime	$\tau \pm \sigma_\tau$	0.70
Longitudinal Diffusion	$D_L = 4.01 \pm 0.18$ cm^2/s	0.01
Transverse Diffusion	$D_T = 7.57 \pm 0.11$ cm^2/s	0.08
Total Systematic		**0.71**
Statistical		**0.30**
Total Uncertainty		**0.77**

trum within the fitting window. Here, the number of degrees of freedom (ndf) is defined as the number of active bins minus the number of unconstrained fit parameters.

Table 5.1 itemizes the dominant sources of systematic uncertainty evaluated for this measurement. Variations in the histogram bin width and nominal boundaries of the fit range yield negligible contributions to the total uncertainty. The leading systematic uncertainty stems from the electron lifetime correction applied to mitigate charge attenuation in the data. This impact is quantified by shifting the purity monitor measurements to their upper and lower uncertainty bounds before calculating the cubic spline, and subsequently repeating the entire fitting procedure. The systematic uncertainty introduced by charge diffusion is evaluated by executing the fit against alternative MC simulation templates generated by varying the longitudinal and transverse diffusion coefficients within their reported uncertainties. Summing all individual contributions in quadrature yields a final measured charge scale reduction in MC simulation of $(4.3 \pm 0.3\ \text{(stat)} \pm 0.7\ \text{(syst)})\%$. This result demonstrates that the simulated charge response within the 2x2 Demonstrator MC simulation framework accurately matches the experimental data at the several-percent level.

5.5 Electron Lifetime Measurements

5.5.1 Measurement Methodology

The 2x2 Demonstrator provides an ideal testbed to evaluate the utility of ambient, low-energy radiogenic activity for in-situ electron lifetime measurements. As demonstrated in Section 5.3, the ability to tag low-energy interactions via the LRS is a critical requirement for this calibration. The underlying principle of the measurement relies on tracking the systematic attenuation of the reconstructed ionization charge as a function of drift time for a source with well-characterized spectral features. To implement this technique using charge-light-matched clusters, the fitting routine targets the charge regime dominated by ^{208}Tl γ decays ($Q > 50$ ke$^-$), as established in Section 5.4. Isolating a single-source energy window simplifies the calibration workflow and streamlines the validation of the method using MC simulation templates, as it avoids complex multi-isotope radiogenic simulations.

Furthermore, anchoring the calibration to a region dominated by a kinematic Compton edge significantly mitigates spectral shaping artifacts induced by spatial non-uniformities in the LRS PDE. In a standard continuous spectrum, higher-energy electrons produce more scintillation photons and are inherently more likely to satisfy the LRS trigger conditions than lower-energy electrons. Consequently, in localized regions characterized by degraded light collection efficiency, the reconstructed charge spectrum can suffer a geometric bias toward higher energies. However, the Compton edge feature arises fundamentally from a mono-energetic physical bound, the maximum kinetic energy transferable to an atomic electron during a 180° photon back-scatter, which is subsequently convolved with the detector response. Fluctuations near the Compton edge are driven by electronic noise and reconstruction smearing rather than variations in the electron energy. Therefore, the shape of the Compton edge spectral feature remains robust against threshold-driven light selection effects, establishing it as an optimal kinematic feature for assessing charge attenuation associated with finite electron lifetime in the detector.

Independent, high-frequency electron lifetime measurements were recorded by a purity monitor throughout the second data-taking run of the 2x2 Demonstrator (see Figure 5.11), providing

a set of reference measurements to validate the radiogenic calibration technique. Extracting the electron lifetime requires a statistical comparison of the charge distributions reconstructed across distinct drift-time intervals. The statistical stability of various estimation techniques was thoroughly evaluated using a bootstrapping approach, which involves drawing N data points from the charge-light-matched low-energy dataset with replacement, where N represents the total sample size.

Through successive bootstrap resamplings, it was discovered that directly comparing binned charge histograms yields highly non-Gaussian electron lifetime measurement distributions. This feature is primarily attributed to the discrete nature of the digital charge packets combined with uncalibrated pixel-to-pixel electronics gain variations across the CRS. To circumvent these systematic limitations, an unbinned statistical approach was developed to compare the drift-dependent charge distributions. Specifically, Kernel Density Estimation (KDE) was implemented to reconstruct the continuous probability density function (PDF) of the underlying charge spectrum without introducing arbitrary bin boundaries [119].

The KDE is constructed by centering a kernel function, selected as a Gaussian distribution in this work, at the charge value of each discrete data point, with the characteristic width dictated by a smoothing parameter known as the bandwidth. The individual localized kernels are subsequently summed and normalized to generate a continuous, smooth PDF. Figure 5.13 illustrates the analytical formulation of the KDE alongside an example application demonstrating the reconstruction of a underlying distribution from a sparse point set. For this analysis, a bandwidth of 0.1 ke^{-} was selected. The optimality of this value was verified by confirming that the resulting continuous PDF accurately maps the fine-grained structure of the underlying data without over-smoothing; the extracted spectral shapes exhibited strong stability under localized variations of the bandwidth parameter.

The electron lifetime measurement procedure is executed as follows. First, the reconstructed charge values are partitioned into discrete drift-time bins across a predefined spatial window along the drift axis, and the continuous KDE, $\hat{f}(q)$, is generated for the charge values within each bin.

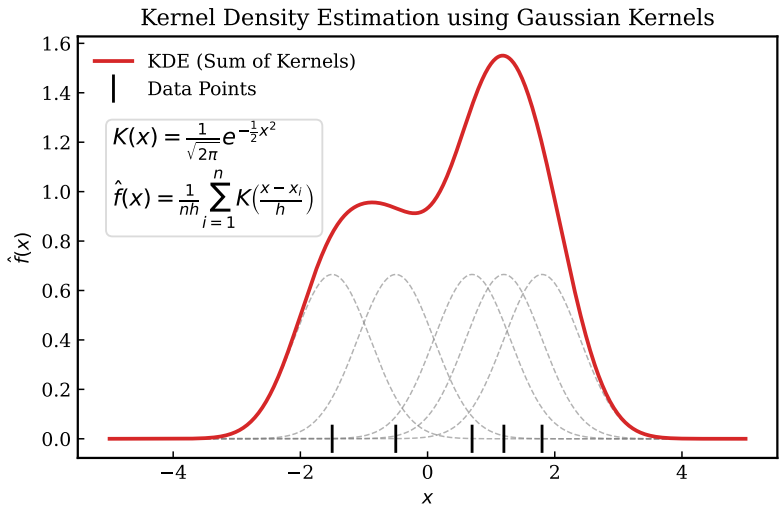

Figure 5.13: An example of using KDEs to determine the PDF associated with a set of discrete data points. The KDE is made up of a sum of individual Gaussian kernels, each with a width determined by the bandwidth parameter h.

The drift-time bin situated closest to the anode ($t_{\text{drift}} = 0$) within this selection window is designated as the reference distribution, $\hat{f}^0(q)$. For an arbitrary electron lifetime test value, τ, the charge distribution within the i-th drift-time bin is dynamically scaled by an exponential correction factor, $e^{(t_i - t_0)/\tau}$. This transformation explicitly compensates for the physical charge attenuation sustained by ionization clusters in bin i during their net transit time, $t_i - t_0$, where t_i and t_0 represent the mean drift times of the target bin and reference bin, respectively.

To quantify the statistical divergence between the corrected distributions, the integrated squared difference is computed between the reference distribution and each scaled bin within the defined charge window. Summing these individual variations across all active bins yields the function known as the total Integrated Squared Error (ISE):

$$\text{ISE}(\tau) = \sum_{i=1}^{n} \int_{q_1}^{q_2} \left[\hat{f}^0(q) - \hat{f}^i \left(q \cdot e^{(t_i - t_0)/\tau} \right) \cdot e^{-(t_i - t_0)/\tau} \right]^2 dq, \tag{5.5}$$

where q_1 and q_2 define the boundaries of the charge window used in the measurement. The optimal electron lifetime is extracted by numerically minimizing the ISE, which identifies the global value that best aligns the attenuated spectra across the entire drift distance.

The primary hyperparameters dictating this measurement are the selected drift-time window and the target charge integration boundaries. To guarantee a sample heavily dominated by the ^{208}Tl Compton edge feature while rejecting low-energy radiogenic contamination, a conservative lower threshold of $q_1 = 53$ ke$^-$ is enforced. The boundary q_2 is set sufficently high to capture the entire Compton edge (above 70 ke$^-$. The full lifetime estimation workflow is repeated over 1,000 bootstrap pseudo-experiments to robustly estimate the statistical uncertainty.

5.5.2 Applying the Measurement to MC Simulation

The ^{208}Tl MC simulation samples are generated in the same manner as described in Section 5.4 for the charge scale calibration. Samples with electron lifetimes ranging from 500 µs to 3000 µs are evaluated. For the MC simulation samples, the drift time is calculated in the same way as the data, except that the efficiency for creating an LRS trigger is set to be 100%. The charge values are grouped into drift-time bins to enable comparison of charge spectra at different distances from the anode.

Figure 5.14 illustrates the segmented charge spectra across the various drift-time bins alongside their corresponding continuous KDE smoothings for a MC sample generated with an intrinsic lifetime of $\tau = 1000$ µs. In addition, a scan of τ vs ISE - ISE$_{\text{min}}$ is shown which exhibits a well-defined global minimum, indicating the electron lifetime value that best models the charge attenuation within the sample. The statistical variation of this extracted electron lifetime value across the bootstrap ensemble is also shown in Figure 5.14, demonstrating an approximately Gaussian distribution. The 68% central confidence interval evaluated from the bootstrapped τ distribution is subsequently used to define the 1σ statistical uncertainty on the electron lifetime measurement.

As illustrated in Figure 5.14, the reconstructed τ value is considerably lower than the true simulated value. A similar systematic underestimation is consistently observed across the broad range of simulated electron lifetimes shown in Figure 5.15. Diagnostic studies revealed that by disabling ionization electron diffusion within the simulation framework, the majority of this bias is eliminated. The longitudinal and transverse diffusion coefficients utilized in these simulations,

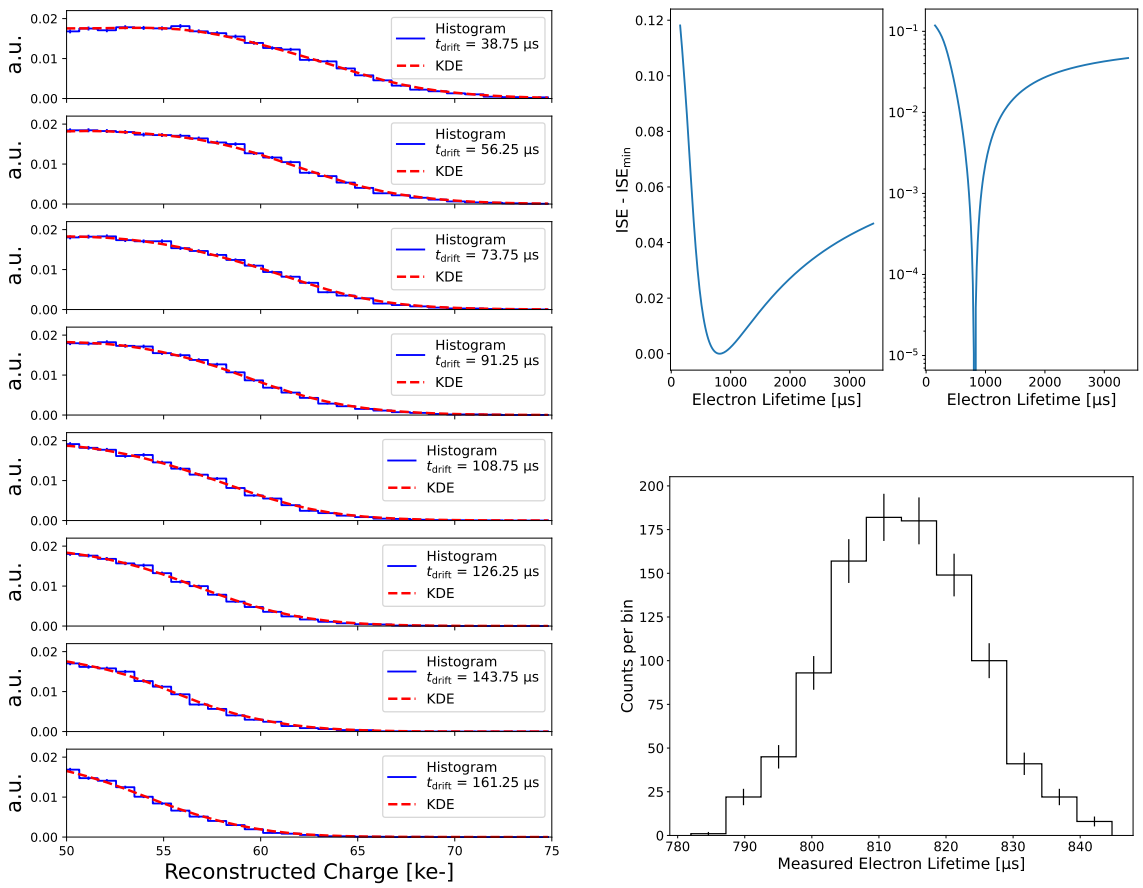

Figure 5.14: (Left) The charge distributions at different drift-time bins, (Top Right) a ISE-ISE$_{\text{min}}$ distribution for a range of τ values, and (Bottom Right) distribution of measured τ after bootstrapping (^{208}Tl simulation with $\tau = 1000$ µs). In this case, the measured electron lifetime unestimated the true value, which is accounted for in the bias correction factor.

$D_L = 4.01$ cm^2/s and $D_T = 7.57$ cm^2/s, respectively, were determined via conversation with ICARUS collaborators about a recent ICARUS measurement, performed near the 2x2 Demonstrator's nominal electric field of 500 V/cm.

As an ionization electron cloud drifts along the TPC electric field, thermal diffusion causes it to expand both longitudinally, extending the temporal duration of the signal, and transversely, broadening its spatial footprint across the pixel plane. The anode plane consists of highly segmented pixels, each functioning as an independent, self-triggering readout channel. Because transverse diffusion dilates the physical extent of the electron cloud, the total ionization charge is distributed over a larger number of individual pixels. Consequently, peripheral channels receiving these dif-

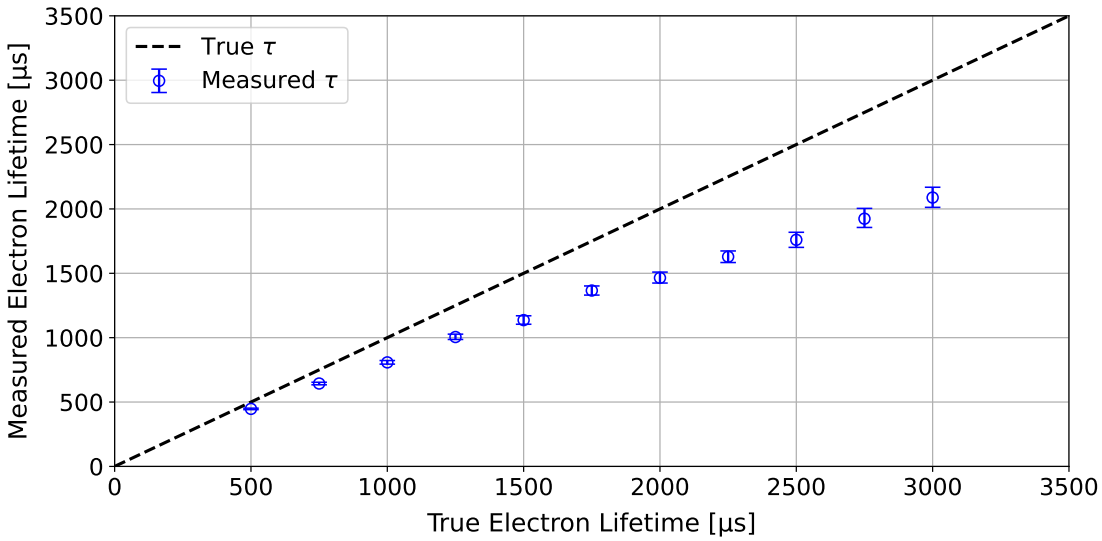

Figure 5.15: Electron lifetime measurements using 208Tl MC simulation samples with a variety of simulated electron lifetime values. A bias correction is extracted from these results and applied to data measurements.

fused signals may fail to surpass their localized hardware self-trigger thresholds, resulting in an unrecovered charge loss for the event.

Because the magnitude of this diffusion-induced charge dispersion scales with the drift distance, it manifests as a systematic, drift-dependent charge deficit that mimics the exponential attenuation of the electron lifetime. This spatial convolution artificially shifts the extracted electron lifetime to shorter values. While this thresholding effect could be partially mitigated by lowering the hardware trigger thresholds during detector operations, an empirical correction must be derived to compensate for this bias in existing datasets. The resulting correction function, along with its associated statistical uncertainties, is presented in Figure 5.16. This correction factor is designed to robustly account for the diffusion-induced threshold bias as well as any sub-dominant detector effects that systematically distort the electron lifetime measurement.

5.5.3 Measurements with 2x2 Demonstrator Data

Polynomial functions are used to fit the electron lifetime correction factor along with its upper and lower uncertainty bounds, as presented in Figure 5.16. Electron lifetime measurements using 2x2 Demonstrator data can be corrected using this parameterization in order to obtain a cor-

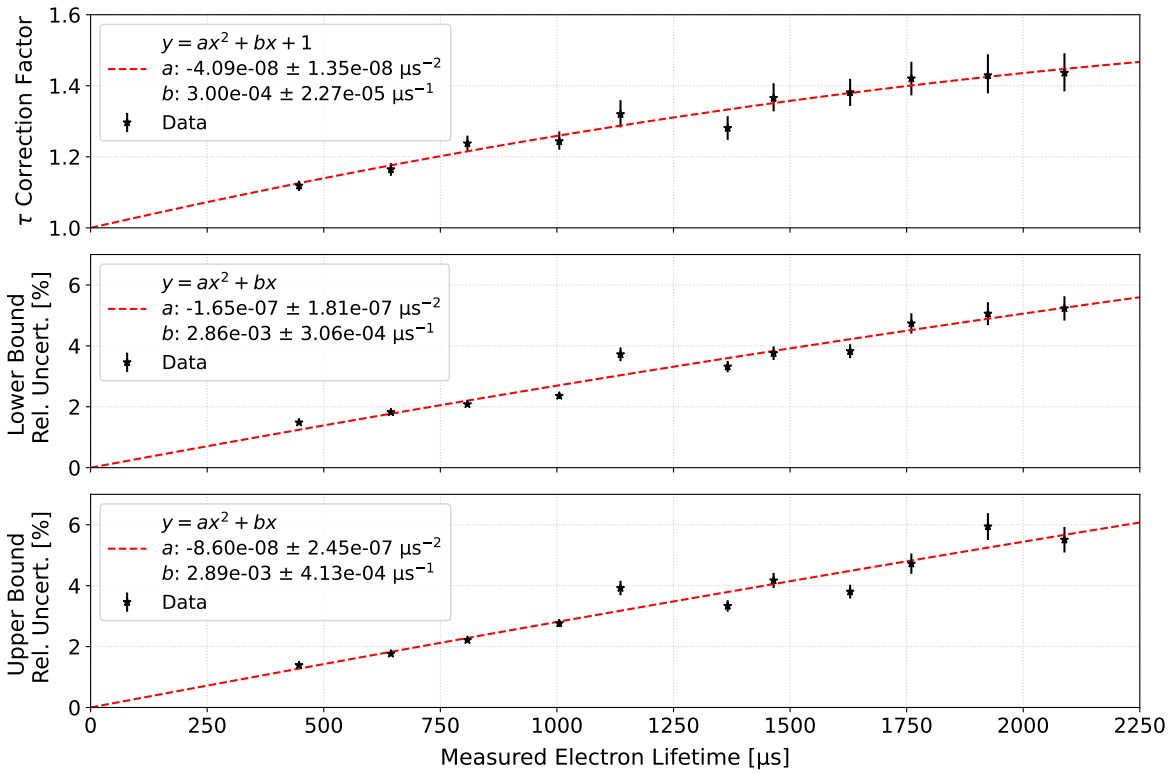

Figure 5.16: Empirical electron lifetime correction factor (and associated uncertainties) derived using MC simulation as a function of measured electron lifetime. This multiplicative correction factor is used to adjust electron lifetime measurements using 2x2 Demonstrator data.

rected, bias-free electron lifetime estimate with associated systematic uncertainty; the uncorrected measurement obtained from the methodology presented in Section 5.5.2 is used as input ("x" in Figure 5.16). Figure 5.17 displays the electron lifetime measurements extracted from 2x2 Demonstrator data, both prior to and following the measurement bias correction, overlaid with the independent purity monitor measurements recorded during the same time period. Significant variations in the electron lifetime are observable throughout this period, driven primarily by transient cryogenic operations. Notably, the radiogenic low-energy lifetime measurements closely replicate the trends resolved by the purity monitor. This strong correlation demonstrates that high-granularity purity tracking can be achieved on an hourly timescale utilizing ambient low-energy activity under this specific LRS threshold configuration. Consequently, this technique provides a passive mechanism to continuously monitor variations in LAr purity, offering a real-time diagnostic to alert system operators to underlying cryogenic anomalies.

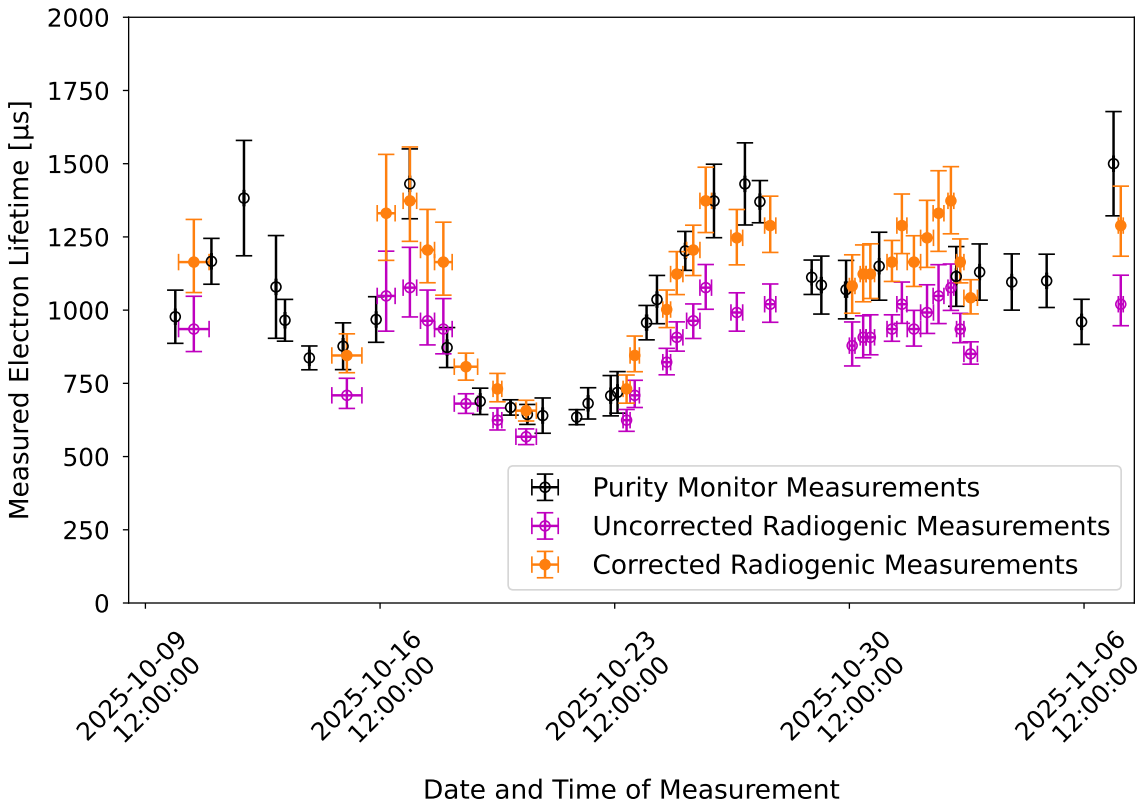

Figure 5.17: Radiogenic low-energy electron lifetime measurements acquired using 2x2 Demonstrator data, shown before and after the empirical measurement bias correction, compared alongside independent purity monitor measurements.

Due to the operation of a high-intensity xenon flash lamp, the CRS and LRS can not collect physics-quality data when the purity monitor is operating. The flash lamp introduces substantial optical backgrounds that saturate the LRS photon detectors, while simultaneously generating severe electromagnetic interference that degrades the signal-to-noise ratio across both the charge and light readouts. For this reason, each purity monitor measurement introduces at least one hour of detector dead time, during which physics-quality TPC data collection must be entirely suspended. Under nominal beam-target operations, such a situation is highly unfavorable given the critical importance of maximizing integrated beam exposure. However, because the second data-taking run was decoupled from the accelerator operations and focused heavily on characterizing the cryogenic purification performance, frequent purity monitor measurements were acceptable.

This operational constraint underscores why in-situ radiogenic electron lifetime measurements

provide a highly advantageous alternative to traditional monitoring methods. Ultimately, the results presented here establish that low-energy radiogenic calibrations align well with specialized purity monitors while delivering a continuous monitoring framework that is entirely non-disruptive to nominal TPC data-taking. This approach successfully tracks purity fluctuations on an hourly timescale, provides a vital and independent cross-check to existing instrumentation, and successfully validates the viability of ambient radiogenic calibrations for LArTPCs in deep underground environments.

5.6 Monitoring Electric Field Distortions

5.6.1 Electric Field Distortion Measurement Methodology

Data were acquired using each 2x2 Demonstrator LArTPC module at the University of Bern Laboratory for High Energy Physics (LHEP) after construction. These data were collected using relatively low LRS thresholds, making the data particularly suitable for evaluating E-field measurements using low-energy activity. The performance of the low-energy reconstruction in these tests was shown in Section 5.3, which demonstrated that low-energy events could be reconstructed with high selection purity in close proximity to the light detectors. Shown in Figures 5.18–5.23 are the spatial distributions of the reconstructed low-energy activity along the reconstructed drift coordinate and pixel-plane horizontal coordinate, near the TPC boundaries and light detectors.

In some modules, the distance between the outer edge of the event distribution and the physical TPC boundary increases at larger drift distances. This effect is most noticeable in Modules 1 and 2, but is less significant in the other modules. These features are consistent with transverse E-field distortions near the TPC boundaries, which push the drifting ionization electrons inward as they travel toward the anode. Because ionization electrons from larger drift distances spend more time traveling through the distorted E-field, they accumulate a larger spatial offset before they are detected at the anode. Specifically, the maximum spatial offsets are roughly 1–2 cm, which corresponds to about 3–7% of the total TPC width. In contrast, the spatial offsets in most other locations within the detector volume appear small or negligible, including Module-0 and Module-3

where the measurement is not directly applied. While the exact cause of these distortions is still under investigation, potential sources include non-uniformities in the resistive coating on the field shells or localized E-field distortions introduced by the light detector modules.

Two main quantities can be extracted from these boundary distributions: the spatial offset of the reconstructed activity and the transverse component of the E-field, the latter being one that can be calculated from the former. The spatial offset is valuable for understanding how E-field distortions physically shift the reconstructed interaction activity. The transverse E-field component is useful for directly comparing the distortion magnitude against the nominal E-field strength of 500 V/cm, and directly impacts the amount of electron-ion recombination that occurs. To extract these quantities, events are selected near the cathode, where the accumulated spatial offset is maximized. Specifically, the analysis targets a drift range of $|x| \in$ [5.0, 5.25] cm, effectively the distance from the cathode, which corresponds to a drift time window of about 1.6 µs. Events within this range are binned along the z direction, and a Crystal Ball function is fitted to the boundary edge of the distribution to extract the average spatial offset. To ensure sufficient statistics for a reliable fit, the data are integrated along the entire y axis. As a result, the extracted spatial offsets represent an average across both the vertical coordinate and the selected drift time interval. The results of these boundary fits are shown in Figure 5.20 for Module-1 and Figure 5.22 for Module-2. Due to the distortions appearing very small in Module-0 and Module-3, and other locations within the remaining Modules, an upper bound is determined for the remaining locations.

The spatial offset, Z_{offset}, is determined by subtracting the fitted mean parameter (μ) of the Crystal Ball function from the physical TPC boundary position. The total uncertainty on Z_{offset} is calculated by combining the statistical uncertainty on μ from the fit in quadrature with the pixel discretization uncertainty, defined as $d/\sqrt{12}$ where d is the pixel pitch. For the chosen drift slice, the measured average offsets span 1–2 cm. To calculate the average transverse E-field component from these offsets, the following relation is used:

$$\frac{v_T}{v_{\text{drift}}} = \frac{E_T}{E_{\text{drift}}}, \tag{5.6}$$

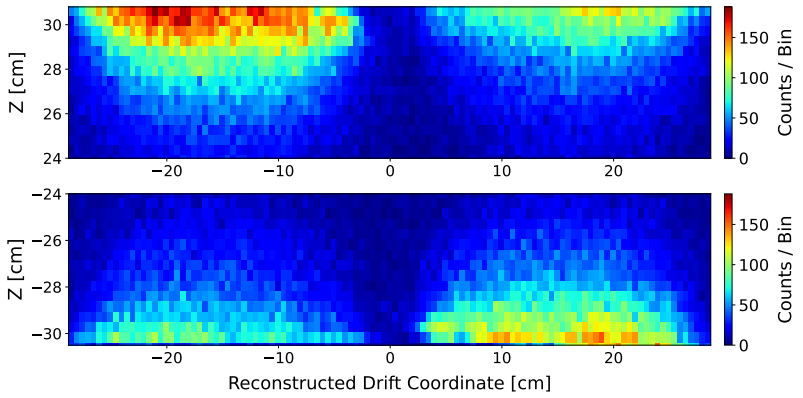

Figure 5.18: Spatial distribution of low-energy activity in the Module-0 LArTPC using data collected at the University of Bern. Events located close to the TPC boundaries and photon detectors are shown.

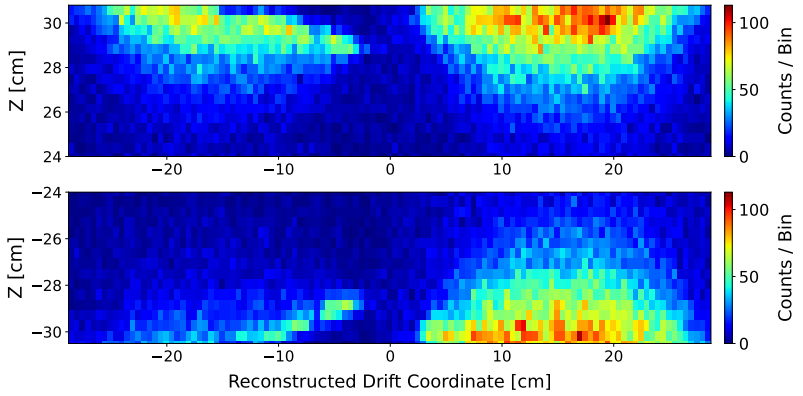

Figure 5.19: Spatial distribution of low-energy activity in the Module-1 LArTPC using data collected at the University of Bern. Events located close to the TPC boundaries and photon detectors are shown.

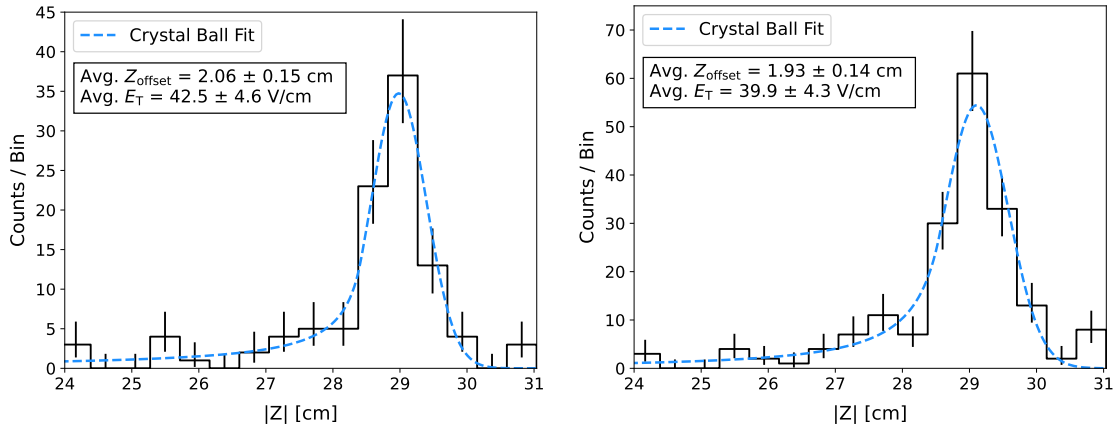

Figure 5.20: Measurements of average spatial offset and transverse E-field component in the top left quadrant of Module-1 (left) and bottom left quadrant of Module-1 (right).

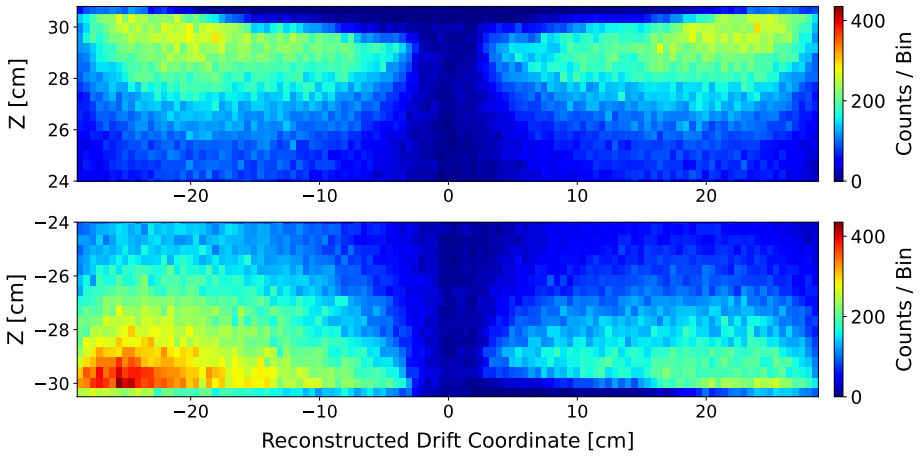

Figure 5.21: Spatial distribution of low-energy activity in the Module-2 LArTPC using data collected at the University of Bern. Events located close to the TPC boundaries and photon detectors are shown.

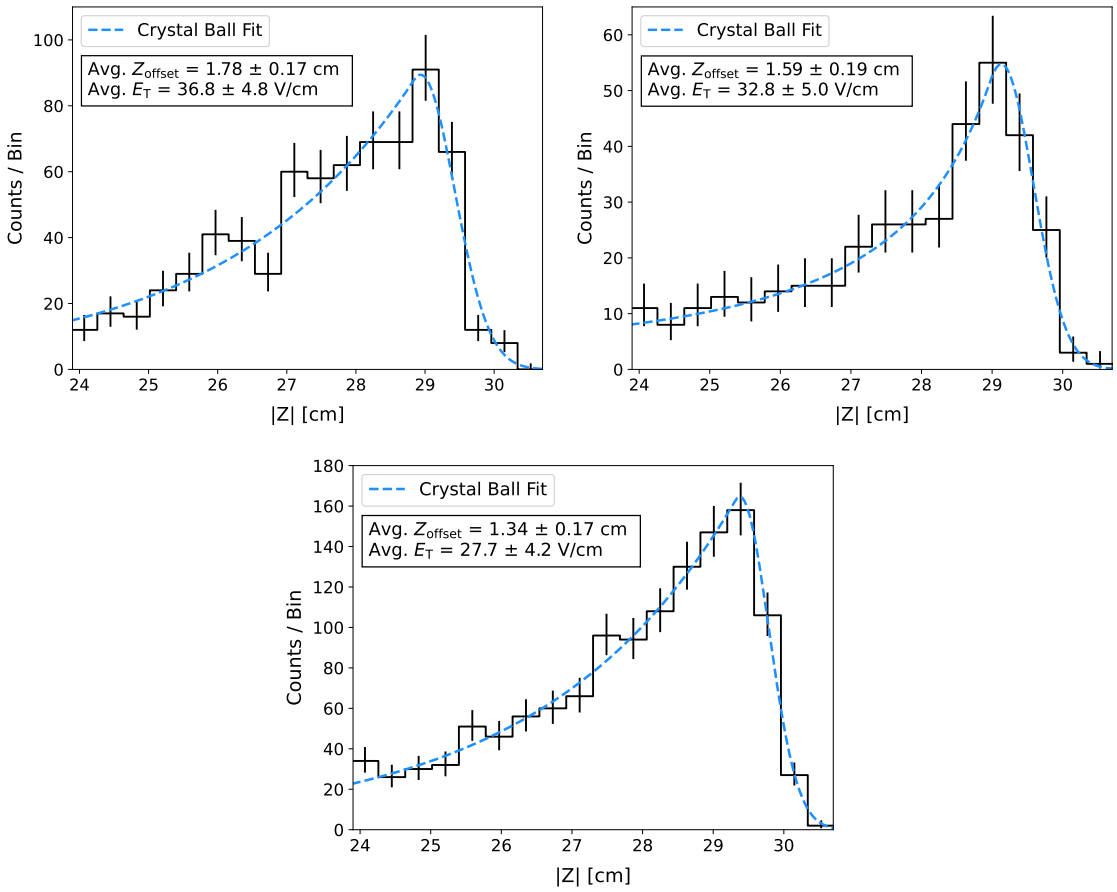

Figure 5.22: Measurements of average spatial offset and transverse E-field component in the top left qudrant of Module-2 (top left), top right quadrant of Module-2 (top right), and bottom right quadrant of Module-2 (bottom).

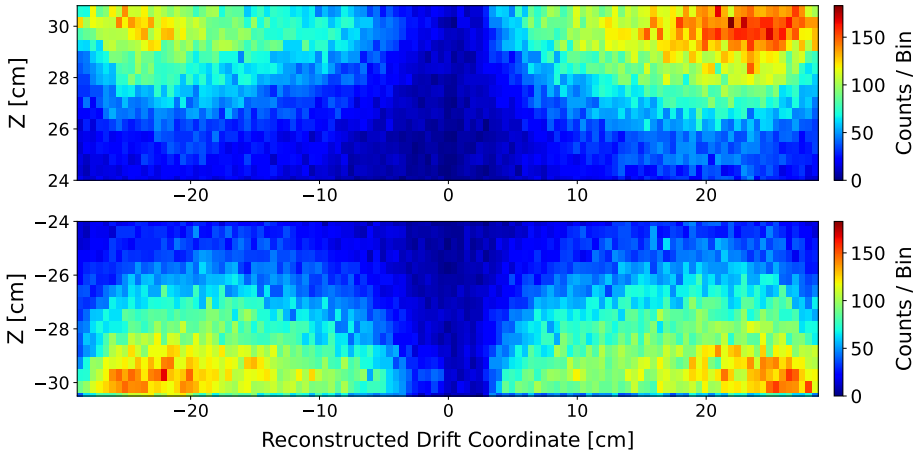

Figure 5.23: Spatial distribution of low-energy activity in the Module-3 LArTPC using data collected at the University of Bern. Events located close to the TPC boundaries and photon detectors are shown.

where v_T is the transverse drift velocity, v_{drift} is the longitudinal drift velocity, E_T is the transverse E-field component, and E_{drift} is the longitudinal E-field component. Because the total drift velocity vector is parallel to the total E-field vector, their components along the drift (x) and transverse (z) axes form similar triangles. Therefore, the ratio of the velocity components equals the ratio of the E-field components, as expressed in Equation 5.6.

Because the measured spatial offsets are small compared to the full drift width of the TPC, the E-field distortions are assumed to have a negligible impact on the overall longitudinal drift E-field component. Consequently, the longitudinal drift E-field component is approximated by its nominal value, $E_{\text{drift}} \approx 500$ V/cm. To determine the corresponding value of v_{drift}, the drift velocity parameterization as a function E-field strength developed by the MicroBooNE collaboration is used [120], which is based on a fit to ICARUS T600 data (during its run at LNGS) at a LAr temperature of 89 K [121]. Although the precise LAr temperature in the detectors may differ from 89 K, percent-level systematic shifts in v_{drift} have a negligible impact on the extracted E_T. Temperature variations within a realistic range of 87–89 K are therefore expected to have no significant effect on the final measurement. For this analysis, E_{drift} is fixed to the nominal 500 V/cm, and temperature-related velocity variations are treated as negligible.

With v_{drift} and E_{drift} established, only the transverse velocity component v_T must be determined.

Table 5.2: Average spatial offsets and transverse E-field components measured in the 2x2 Demonstrator LArTPC modules using data collected at the University of Bern. The polarity of z and x indicates the corresponding side of the module where the measurement was made. "Other Locations" refers to other modules or polarities not explicitly shown in this table.

Module	z, x **Polarity**	z **Offset**	**Transverse E-field**
Module-1	$+z, -x$	2.06 ± 0.15 cm	42.5 ± 4.6 V/cm
Module-1	$+z, -x$	1.93 ± 0.14 cm	39.9 ± 4.3 V/cm
Module-2	$+z, -x$	1.78 ± 0.17 cm	36.8 ± 4.8 V/cm
Module-2	$+z, +x$	1.59 ± 0.19 cm	32.8 ± 5.0 V/cm
Module-2	$-z, +x$	1.34 ± 0.17 cm	27.7 ± 4.2 V/cm
Other Locations	$\pm z, \pm x$	< 0.66 cm	< 15 V/cm

To find v_T, the measured Z_{offset} is divided by the average drift time of the events within the target selection range (X $\in$ [5.0, 5.25] cm). The standard deviation of these drift times is treated as a systematic uncertainty and propagated into the final calculation of E_T. The drift range window is kept narrow specifically to minimize this systematic effect; direct timing uncertainties from the readout electronics are small enough to be considered negligible. Once v_T is determined, Equation 5.6 yields the transverse E-field component E_T. These values are reported in Figures 5.20 and 5.22, as well as Table 5.2.

In locations where the transverse E-field distortions are very small, the event distribution tracks the physical boundary closely, causing the Crystal Ball fit to fail as it hits the edge of the physical TPC volume. For these regions, an upper bound on the transverse E-field is calculated instead. Because several pixel channels near the physical edges of the 2x2 Demonstrator TPCs are disabled, there is a lower limit to the smallest distortion this method can resolve. Furthermore, very small transverse offsets approach the intrinsic spatial resolution of the pixel pitch, making precise measurements difficult. Based on these geometric constraints, a conservative limit of 1.5 pixel pitches (about 0.66 cm) is used, which translates to an upper bound of approximately 15 V/cm for the transverse E-field component in the remaining, undistorted regions of the TPCs.

These results demonstrate that E-field distortions are present in multiple of the single modules tested at Bern. The distortions are inconsistent across the different modules, where some locations show little to no discernible E-field variation while others exhibit noticeable shifts. These specific

observations do not necessarily represent the expected E-field configurations in the final ND-LAr TPC modules, which will use a more traditional resistor-divider field cage design. However, because E-field distortions can still arise from alternative sources in ND-LAr, evaluating E-field uniformity remains highly important.

This analysis demonstrates that ambient low-energy activity provides a viable method to evaluate E-field uniformity in the final detector. Because these measurements are spatially averaged and primarily quantify distortions at the physical TPC boundaries, they do not provide a full 3D map to correct for localized spatial offsets. For this reason, this approach would not replace other track-based measurement methods, such as those utilizing cosmic-ray muons, but would instead serve as an independent cross-check and a continuous monitoring tool. Given that these boundary offsets are visually clear and straightforward to calculate, the distributions could be generated automatically during detector operations to provide real-time diagnostics of E-field uniformity. The primary technical constraint for implementing this method is ensuring that the LRS thresholds are set sufficiently low to detect low-energy interactions at a rate high enough to support regular measurements.

5.7 Observations from Charge Data

When using charge-light matching, low-energy clusters are mostly reconstructed close to the light detectors, especially when the LRS hardware thresholds are high. This localization limits the ability to evaluate the spatial distribution of low-energy backgrounds across the entire detector volume, which is necessary when identifying internal sources of radioactivity. Mapping this radioactivity throughout the detector could help pinpoint specific background sources, such as ^{208}Tl, which produces a distinct spectral signal in the 2x2 Demonstrator. Because the CRS provides continuous self-triggered readout independent of the light system, it can be used to analyze how low-energy activity is distributed across the entire pixel plane. One limitation is that the absolute drift coordinate cannot be determined without a light-matching timestamp (t_0), making it difficult to pinpoint source locations in 3D space. Nevertheless, analyzing the 2D projected spatial distri-

butions can still show whether there is an excess of activity near the boundaries of the TPC that are transverse to the ionization drift direction.

Figures 5.24–5.27 display the spatial distributions of low-energy activity in the 2x2 Demonstrator at Fermilab, using charge-only data from the CRS. Several features appear in these distributions that were not visible in the charge-light-matched sample. The most noticeable feature is an excess of reconstructed activity near the TPC edges, which steadily decreases toward the center of the TPC volumes. This pattern is consistent with a source of γ decays originating near the TPC boundaries. The $\mathcal{O}(1\ \text{MeV})$ photons from radioactive decays, such as the 2.615 MeV γ ray from ^{208}Tl, have characteristic interaction lengths of around 20 cm in LAr. As a result, the density of resulting ionization deposits is expected to decrease with distance from the boundaries, qualitatively matching the trend observed in the data. Alternatively, ^{222}Rn produced in the ^{238}U decay chain can emanate out of internal detector components and distribute itself throughout the liquid argon. The half-life of ^{222}Rn is relatively long, about 3.8 days, which provides ample time for the isotope to mix throughout the TPC active volume.

This falling event rate varies noticeably between individual TPCs and across the different modules. This variation could be caused by differences in the radioactive trace content within the detector components, such as field-shell panels coming from different manufacturing batches. Because formal radiological assays were not performed on these specific detector materials before assembly, it is difficult to identify the exact components driving the observed background rates. However, the spatial distributions strongly indicate that the activity originates from outside the active TPC volume, pointing to materials near the boundaries. A large amount of fiberglass G10 composite is used structurally in the 2x2 Demonstrator TPCs, making it a primary candidate for this background. The fact that the MicroBooNE collaboration observed a similar excess of ^{208}Tl decays near internal G10 structural struts supports this conclusion [111].

To further test the hypothesis that these γ decays originate near the detector edges, the reconstructed charge spectrum is compared across different regions of the TPC active volume. Figure 5.28 shows the charge spectrum extracted a few centimeters from the TPC boundaries com-

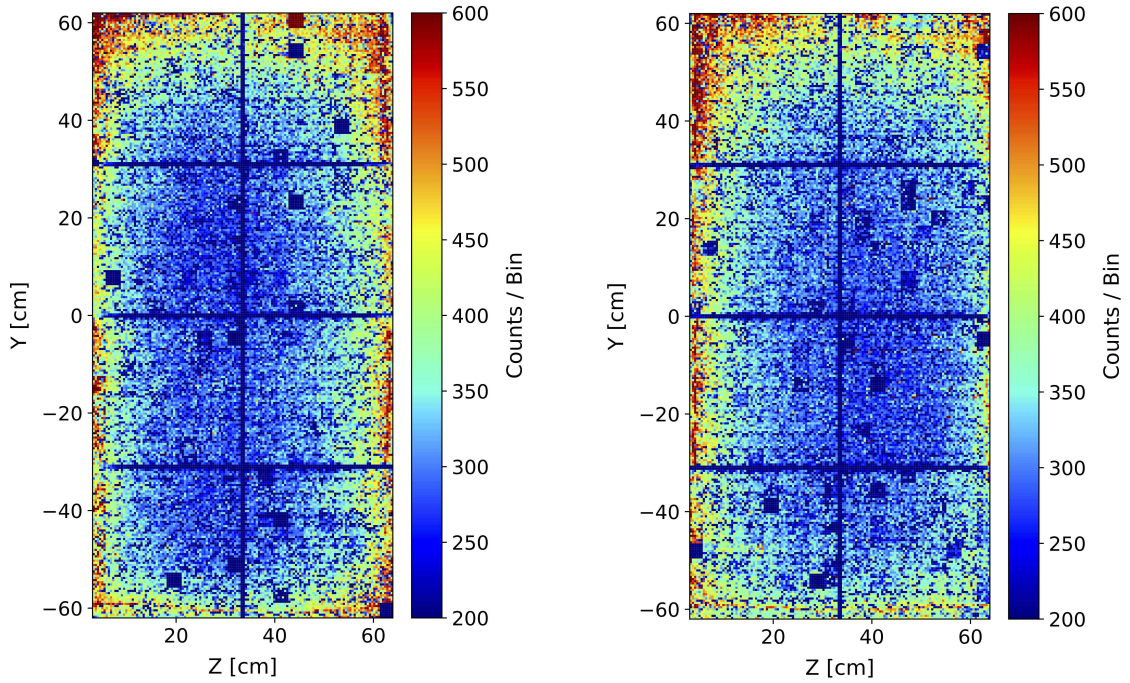

Figure 5.24: Spatial distribution of low-energy clusters in Module-0 of the 2x2 Demonstrator at Fermilab using only CRS data, without requiring charge clusters to be matched to a light flash in the LRS.

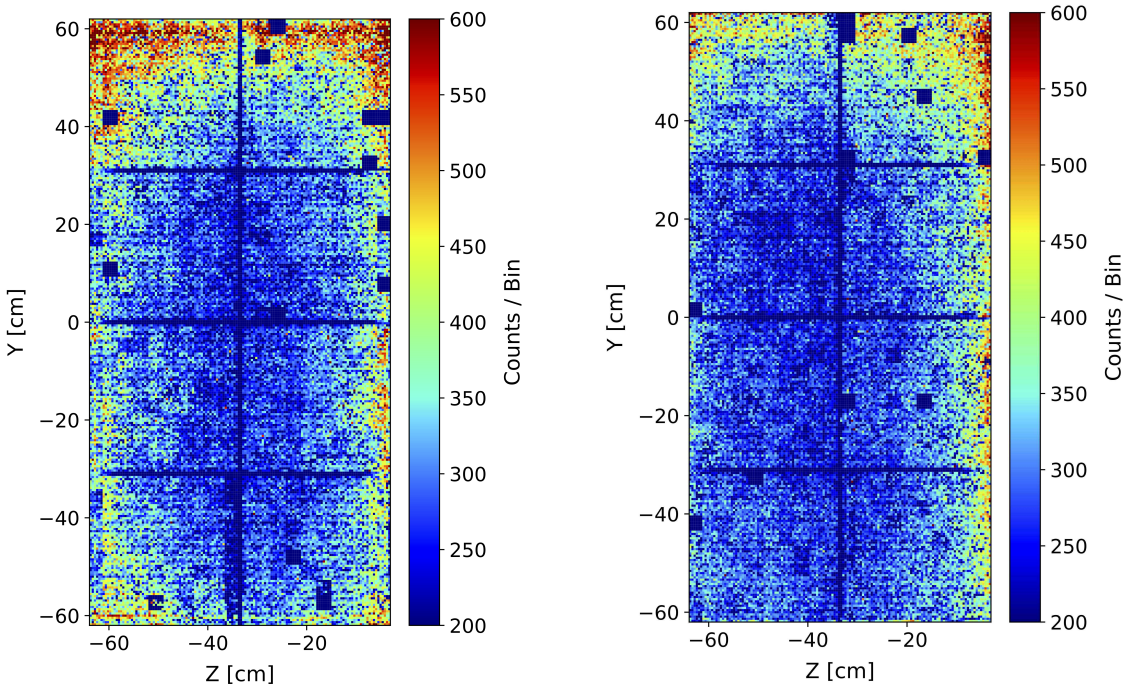

Figure 5.25: Spatial distribution of low-energy clusters in Module-1 of the 2x2 Demonstrator at Fermilab using only CRS data, without requiring charge clusters to be matched to a light flash in the LRS.

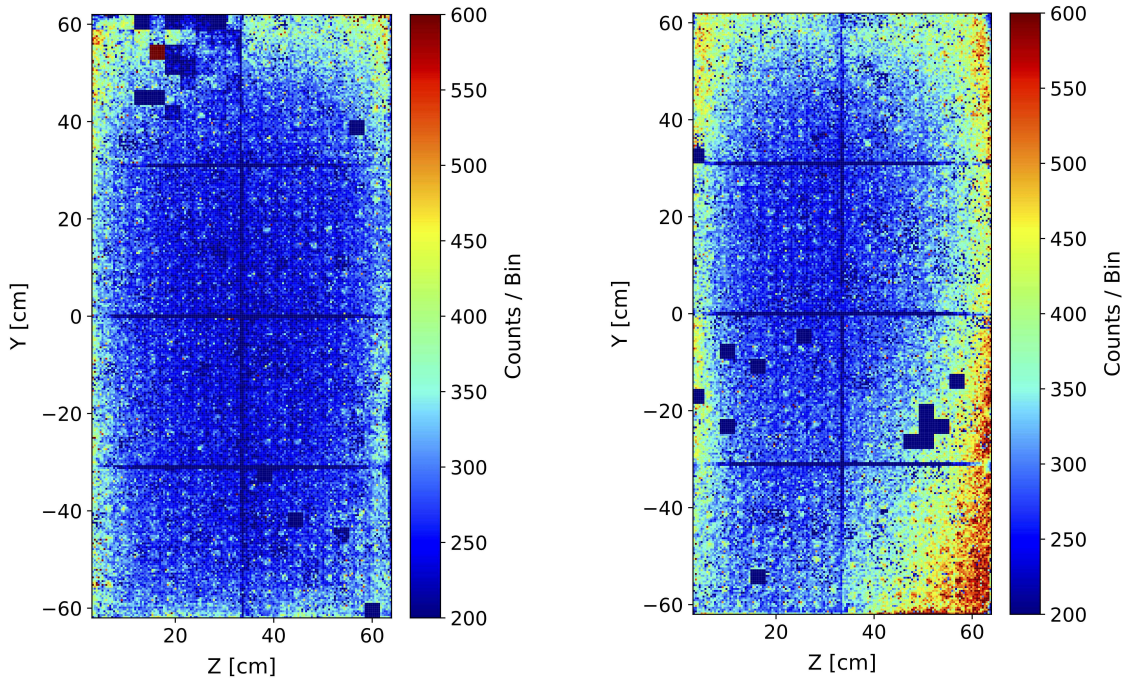

Figure 5.26: Spatial distribution of low-energy clusters in Module-2 of the 2x2 Demonstrator at Fermilab using only CRS data, without requiring charge clusters to be matched to a light flash in the LRS.

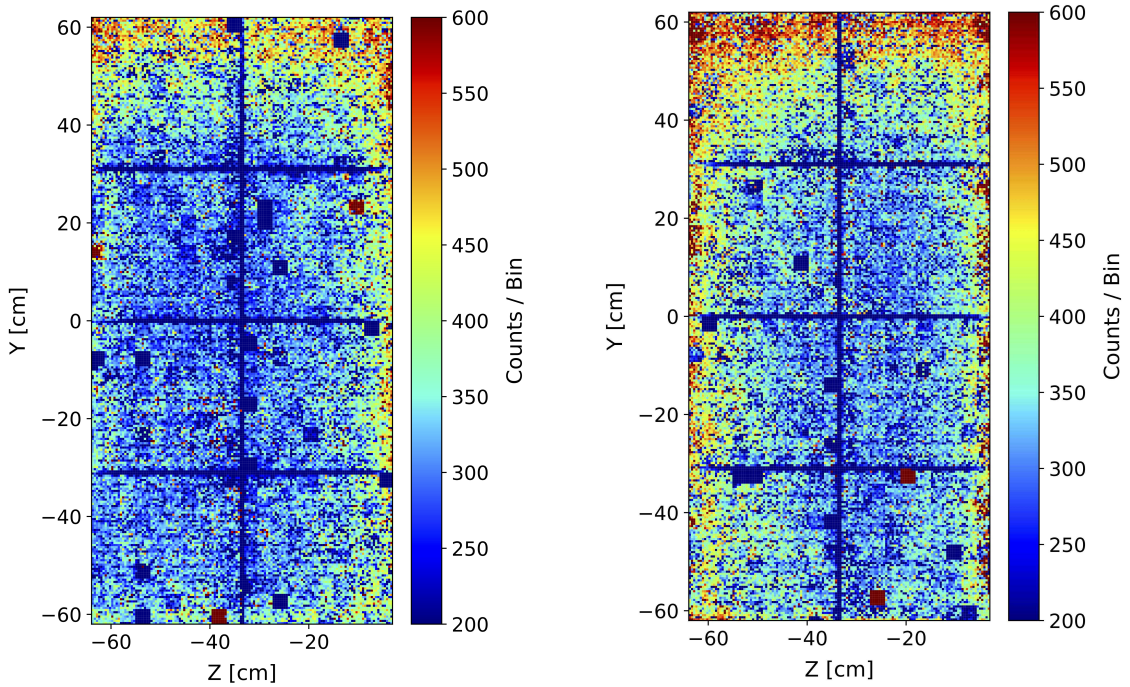

Figure 5.27: Spatial distribution of low-energy clusters in Module-3 of the 2x2 Demonstrator at Fermilab using only CRS data, without requiring charge clusters to be matched to a light flash in the LRS.

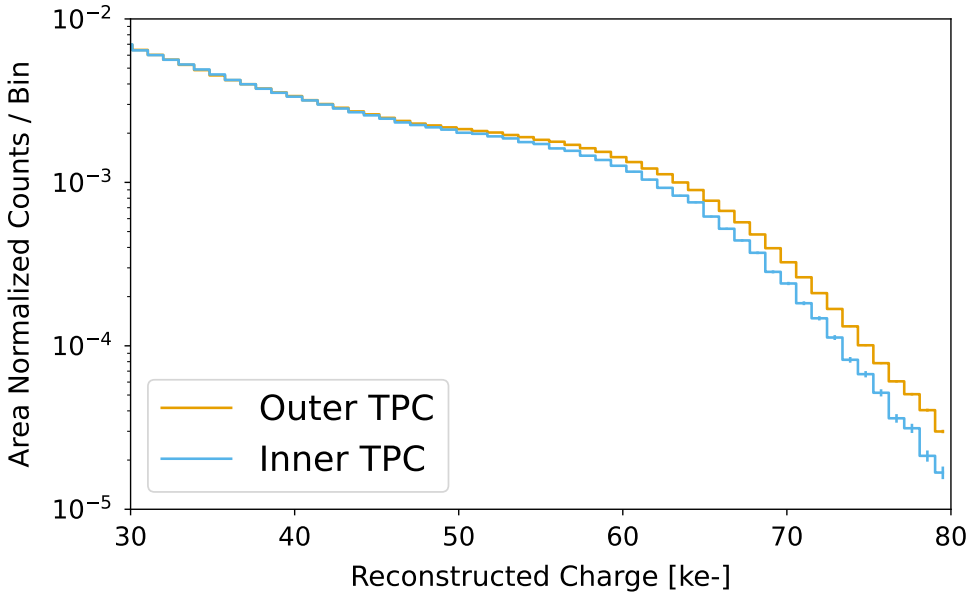

Figure 5.28: Comparison of the reconstructed charge spectrum at the outer edges of the TPC active volume vs. near the center of the TPC active volume.

pared against the spectrum from the innermost core of the TPC active volume. The characteristic ^{208}Tl Compton edge is visible in both regions but is significantly enhanced near the boundaries. The shift of the Compton edge toward lower charge values in the inner region can be explained by photons losing energy via intermediate scatters before reaching the center of the TPC active volume, which lowers the maximum energy that they can transfer to a single electron via Compton back-scattering. This spectral shift further indicates that the ^{208}Tl source is located near the TPC boundaries.

Chapter 6

Development of a LArTPC Test Stand for Evaluating the ^{220}Rn Injection Technique

6.1 Overview

This chapter describes the experimental configuration of the Colorado Argon Test Stand (CATS) which operates in Prof. Mike Mooney's laboratory at CSU. The facility consists of a large cryostat housing a cubic-foot-scale LArTPC equipped with both ionization charge and scintillation light readouts. The test stand also includes a dedicated LAr purification system, an isolated power and grounding network to achieve low-noise TPC performance, and a separate, smaller cryostat used for focused testing using a photon detector or a purity monitor. A primary goal of the CATS facility has been to demonstrate the viability of the ^{220}Rn injection method in LAr with the eventual goal of using the technique to calibrate the DUNE LArTPCs. Preliminary testing at CSU, described in Section 6.4.4, served as a proof-of-principle for the technique, which ultimately led to a full-scale test at the 2x2 Demonstrator (detailed in Chapter 7). The following sections describe the CATS hardware configuration, TPC charge-light matching performance, and results of ^{220}Rn injection tests in this setup.

6.2 Colorado Argon Test Stand

6.2.1 Cryogenic System

The CATS setup uses three separate cryostats. The primary cryostat (model CF-3642-F) houses the TPC and photon detector. It has a diameter of 96 cm and a height of 145 cm, and is positioned 15 cm above the floor to maintain proper electrical ground isolation. To reduce argon costs and filling time, a 170 L stainless steel internal bucket is placed inside this cryostat to minimize the total volume of LAr needed for testing. The TPC and bucket hang from a stainless steel crossbar

Figure 6.1: (Left) The TPC mounted on the CATS cryostat lid before installing the bucket. (Right) The TPC inside the internal bucket immediately before closing and sealing the cryostat.

connected to the bottom of the cryostat lid by four metallic rods. A pulley system is used to raise the lid to enable detector installation, after which the lid is lowered to seal the cryostat body. The assembled TPC installed on the cryostat lid is shown in Figure 6.1.

A second, smaller cryostat houses the LAr purification filters. The LAr filters utilize a copper-based catalyst and molecular sieve to remove electronegative impurities from the liquid argon, oxygen and water, respectively. The filters are discussed with further detail in Section 6.2.2. It connects to the LAr supply dewars on one end and to the primary cryostat on the other via vacuum-insulated cryogenic transfer pipes. A third cryostat (40 L) is used for fast-turnaround tests of ^{220}Rn injection using a photon detector. On separate occasions, this cryostat is used to perform tests of the LAr purity monitor to confirm that the LAr purification filters are working as intended. The full cryogenic layout for a standard TPC test is shown in Figure 6.2.

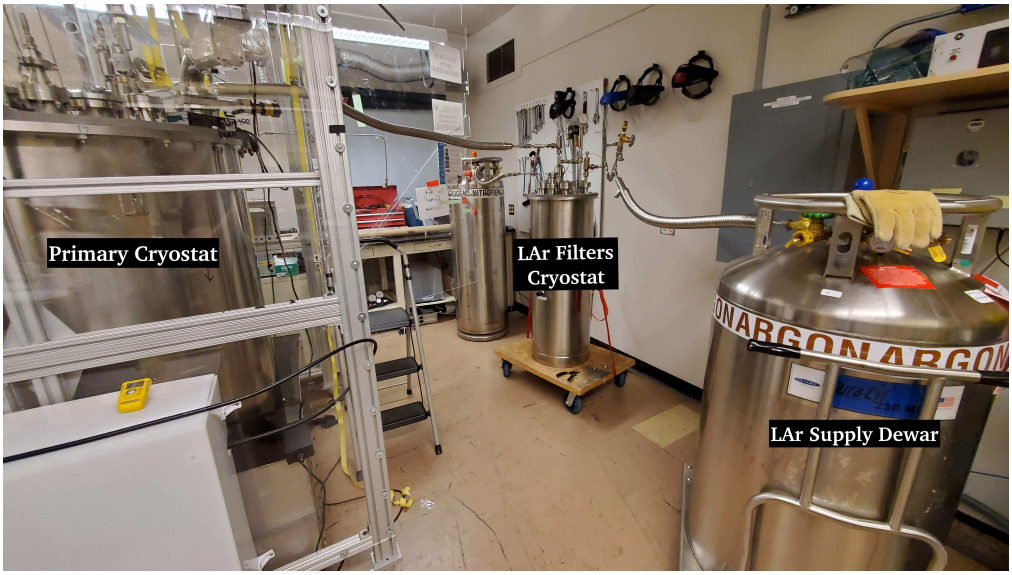

Figure 6.2: The cryogenic configuration in the CSU laboratory in preparation for a TPC test, showing the primary cryostat containing the TPC, the filter cryostat, and a LAr supply dewar. The 40 L cryostat is not shown here, as it is typically operated separately.

The TPC utilizes the LArPix pixelated charge readout system to detect drifted ionization electrons, which is the same technology used in the 2x2 Demonstrator [89] and described in Section 3.7.1. Power and data lines for the anode pixel tile pass through the cryostat lid via a Conflat feedthrough PCB, while the cathode high-voltage (HV) bias is supplied through a 20 kV SHV feedthrough. The liquid level is monitored using a capacitive level meter connected to an American Magnetics Model 1700 readout instrument, along with a series of Resistance Temperature Detector (RTD) sensors read out using an Arduino. The RTD sensors provide liquid level measurements at heights where the capacitive depth sensor is insensitive, and they indicate when LAr first begins to collect at the bottom of the bucket.

Liquid argon enters the system through a cryostat lid penetration into a fill tube that terminates near the bottom of the bucket. The liquid flow rate is kept relatively low to maximize the contact time with the filtration media, which is critical because the liquid argon is only filtered once before entering the primary cryostat. Cosmic-ray muon measurements using this TPC have demonstrated that electron lifetimes up to $\sim$ 1 ms are achievable in this setup. This purity level corresponds to an average charge attenuation of about 10% at the maximum drift distance of roughly 30 cm under the

nominal electric field strength of 500 V/cm. This performance is sufficient for detector testing; it ensures ionization signals remain easily detectable across the full drift volume, and the attenuation can be corrected offline once the electron lifetime is measured and the events are tagged by the light readout system.

For safety and pressure control, the cryostats are equipped with standard relief valves and burst disks to prevent dangerous over-pressure conditions. The primary cryostat sits inside an acrylic enclosure connected to an exhaust fan, which safely vents argon gas in the event of a leak. The laboratory is also monitored by stationary Oxygen Deficiency Hazard (ODH) sensors and personal ODH sensors worn by lab personnel, to alert lab personnel to low-oxygen levels.

The primary cryostat filling process typically takes between 12 and 16 hours and requires continuous monitoring. Operators track the progress using pressure gauges, the capacitive level sensor, and the RTD sensors. Inside the cryostat, an HV-rated cable connects the feedthrough to a tab at the top of the TPC cathode. HV is only applied when the LAr level completely submerges this tab to eliminate the risk of electrical sparking in the gas phase. As a result, the maximum duration of a data-taking run is limited by the argon boil-off rate. However, the cryostat can be topped off with additional LAr when the level drops too low.

The liquid argon is delivered in 230 L ultra-high-purity (UHP) dewars from a local Matheson vendor [122]. Completely filling the bucket requires two to three dewars, as a significant portion of the liquid boils away while cooling down the filters, transfer lines, TPC, and bucket assembly. The vendor certifies that this UHP argon contains less than around 0.5 ppm of oxygen and water impurities. To prevent the external piping from acting as an antenna that could introduce electronic noise into the charge and light readouts, ceramic isolation disks are installed in the transfer lines near both the filter and primary cryostats. Additionally, a particulate filter is installed at the outlet of the main LAr filter to trap small particles, which previous tests linked to HV instabilities.

Before filling with LAr, the cryostats undergo an evacuation and purge cycle to clear out atmospheric impurities. The vessels are evacuated using a combination of a mechanical roughing pump and an Edwards NEXT85D turbomolecular pump to remove air and allow the internal detector ma-

terials to outgas. Once the pressure drops to approximately 10^{-4} mbar, the pumps are isolated and the volume is back-filled with UHP argon gas to near atmospheric pressure. This pump-and-purge cycle is repeated at least five times, followed by a final continuous pump-down lasting at least 24 hours. This extended pumping phase allows internal components to fully outgas, preventing trace impurities from entering the liquid argon later.

The detector is then filled with filtered argon gas prior to the liquid transfer. The entire system is checked for vacuum integrity using a helium leak detector, which has successfully resolved multiple leaks to ensure low baseline pressures. System pressures are monitored using a combination of Inficon PSG550, Lesker KJL275803LL, and Omega PX309-015CG5V gauges, which together span the full pressure range of the system. These gauges are read out through a LabJack DAQ interface connected to a computer, and the pressure data are monitored and recorded using a custom LabVIEW application.

6.2.2 LAr Purification System

The LAr purification system consists of two primary components: an oxygen remover and a moisture absorber. The filter media are housed inside separate Conflat nipple assemblies within a dedicated cryostat (Figure 6.3). Oxygen removal is performed using an activated copper catalyst (GetterMax-133), while water removal is accomplished via a molecular sieve (RCI-DRI-13X). To prevent the filters from absorbing atmospheric impurities when not in use, they are kept under continuous pressure using UHP argon gas. During detector filling, LAr first flows through the molecular sieve to trap water, and then passes through the copper getter to capture oxygen. The system contains 4.8 kg of the copper catalyst and 1.9 kg of the molecular sieve, ensuring efficient single-pass filtration for the CATS cryogenic operations.

The copper getter works by reacting elemental copper with dissolved oxygen molecules to form copper oxide, removing the oxygen from the liquid argon. The catalyst must be activated before its first use and regenerated once it becomes saturated. This regeneration process is performed by flowing a mixture of hydrogen and argon gas (5% H_2 by volume) through the media. Throughout

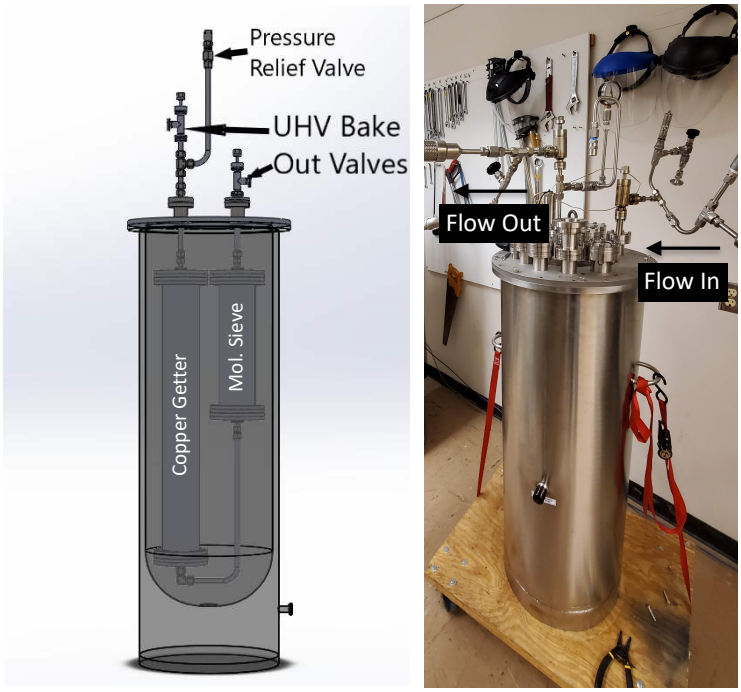

Figure 6.3: (Left) Schematic diagram of the LAr filters. (Right) Photograph of the LAr filter cryostat assembly in the CSU laboratory.

most of the regeneration cycle, the inlet gas consists of roughly 1% of the hydrogen mixture, with the remaining 99% supplied by UHP gaseous argon. The hydrogen drives a chemical reduction reaction:

$$\mathrm{CuO} + \mathrm{H_2} + \text{heat} \rightarrow \mathrm{Cu} + \mathrm{H_2O}, \tag{6.1}$$

which restores the solid copper media and releases water vapor. Because this reaction produces water, the copper getter must be placed downstream of the molecular sieve to prevent the sieve from absorbing the newly created moisture. The filter vessel is heated above 175°C to initiate the reduction reaction using an inline Omega gas heater and heat tape wrapped around the getter enclosure. Ideally, the copper media is maintained between 200°C and 215°C for optimal reaction efficiency. The molecular sieve is regenerated at lower temperatures because the process relies simply on evaporating the trapped water molecules. This sieve operates by trapping water molecules through electrostatic attraction to sodium cations within the media structure. Heating the material releases these bound water molecules, which are then swept out of the system.

Thermocouples monitor temperatures at the system inlet, the exhaust outlet, and various locations on both filter enclosures. The overall reaction rate and status are monitored by tracking the internal getter temperature and the exhaust temperature profile. Figure 6.4 shows the temperature trends during the initial phase of a filter regeneration cycle. A relay system, controlled by a computer program, automatically switches the heaters on and off to maintain the target temperature range, causing the small oscillations visible in Figure 6.4 that are correlated with the relay ON/OFF state. During the active reduction phase, short temperature spikes appear in the copper getter data; their frequency is one indicator of the reaction status.

To verify that the regeneration is complete, operators briefly increase the hydrogen concentration in the inlet gas. Because the reduction reaction is exothermic, a sudden temperature spike in the getter and exhaust indicates that unreacted copper oxide remains. If no temperature spike occurs, the reaction is considered complete. To ensure full media recovery, the final hour of the process is performed using the full 5%/95% hydrogen/argon gas mixture. The entire regeneration cycle, including the preheating phase, typically takes between 24 and 36 hours, depending on the total volume of LAr filtered since the previous cycle.

6.2.3 High-Voltage and Power Isolation Systems

The negative HV cathode bias voltage required to establish the TPC drift electric field is supplied by an SRS PS370 HV power supply. The power supply output connects to an external HV filter through an SHV cable rated for 20 kV. The HV filter is a passive RC circuit designed to attenuate voltage fluctuations to provide a highly stable baseline bias for the TPC. The output of the filter enclosure connects to the primary cryostat flange via a second 20 kV SHV cable. The physical filter enclosure is visible in the bottom-left corner of Figure 6.2.

The HV filter contains an internal resistance of 200 MΩ, while the internal TPC field cage resistance is 1530 MΩ. With the anode end of the field cage held at ground potential, this configuration creates an 11.6% voltage drop between the power supply and the cathode. Consequently, setting the power supply to -17.2 kV yields a cathode voltage of approximately -15.2 kV, which

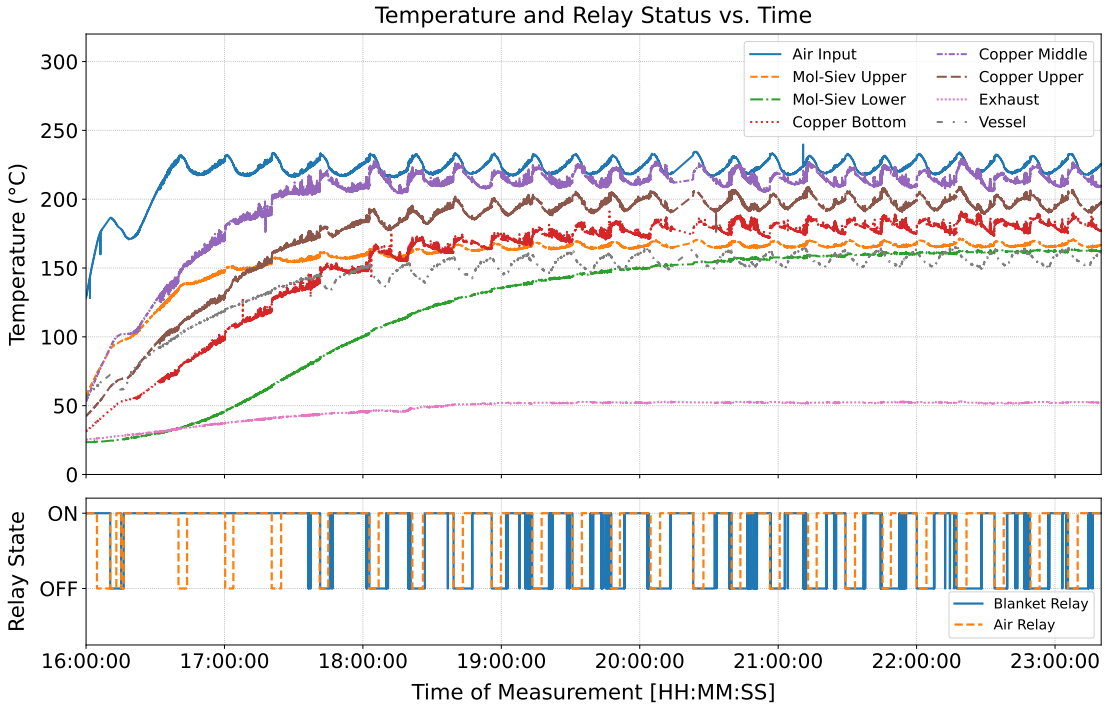

Figure 6.4: Recorded LAr filter temperatures during the beginning of a standard regeneration cycle.

provides the nominal 500 V/cm drift electric field strength. The power supply voltage is ramped up slowly to minimize the risk of sudden electrical discharges or current surges, and the output current is continuously logged to monitor field stability.

Standard building power and ground networks are often too unstable and electrically noisy for low-noise TPC operations. To address this, a dedicated isolation network was constructed to deliver an independent power source and ground. The design is based on a system developed at Fermilab for LArTPCs operating at Fermilab and CERN. Figure 6.5 shows the labeled components of this isolation system.

The primary components separating the system from the building grid are a double-shielded isolation transformer and a saturable inductor (see Figure 6.6). The transformer primary winding connects to the main building power panel, while its secondary winding feeds the isolated electronics panel. Two internal Faraday shields sit between the windings to shunt high-frequency electrical noise away from the isolated panel. This isolated panel features a dedicated ground bus bar and isolated outlets. The reference ground for the charge and light readout systems (referred to as "de-

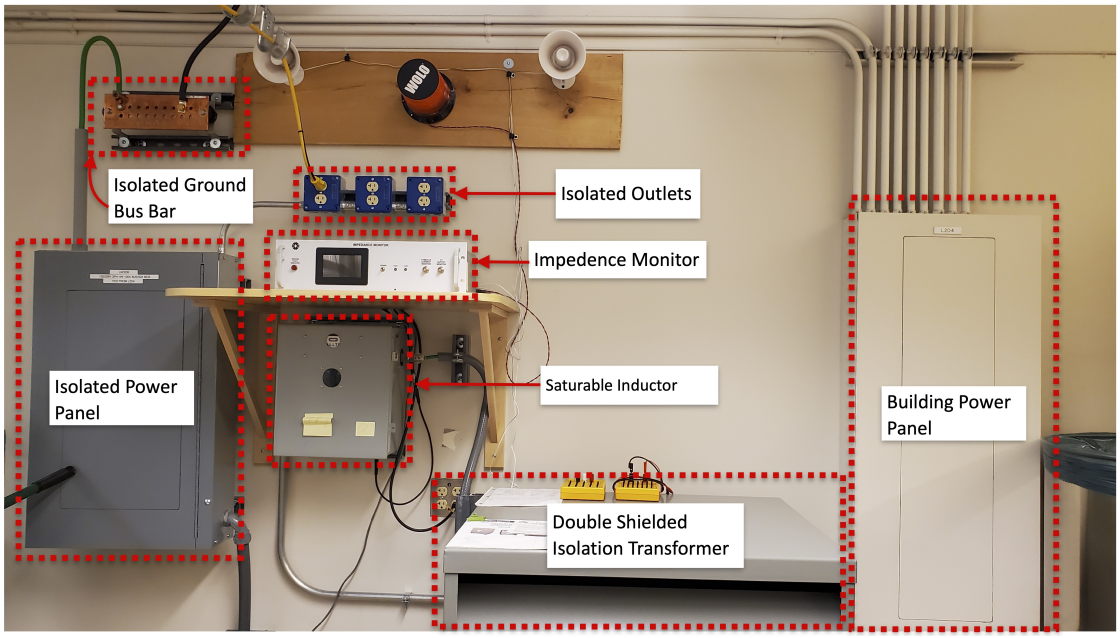

Figure 6.5: Annotated photograph of the power and ground isolation system for CATS.

tector ground") is drawn directly from this isolated bus bar via the heavy black cable shown in the top left corner of Figure 6.5. All power supplies and supporting electronics for the CATS facility plug directly into these isolated outlets.

The saturable inductor sits directly between the isolated detector ground and the main building ground. It presents a high impedance to low-level AC currents, suppressing ground-loop noise during normal operation. In the event of a severe fault that creates a large potential difference between the two grounds, the inductor saturates. This drops its impedance to a minimum under high current, allowing the fault current to safely pass and trip the primary circuit breaker.

A current transformer is also looped around the isolated ground line. Any stray current flowing through the ground line induces a signal in this transformer, which is routed to a Ground Impedance Monitor (GIZMO) module mounted above the saturable inductor chassis. Loaned by Fermilab, the GIZMO device continuously checks for accidental short-circuits between the isolated detector ground and the building ground. It works by sending a small AC "stimulus current" into the isolated ground loop. The returning current is measured by the current transformer and compared to the original signal. If the return current drops to zero, it indicates a direct short to

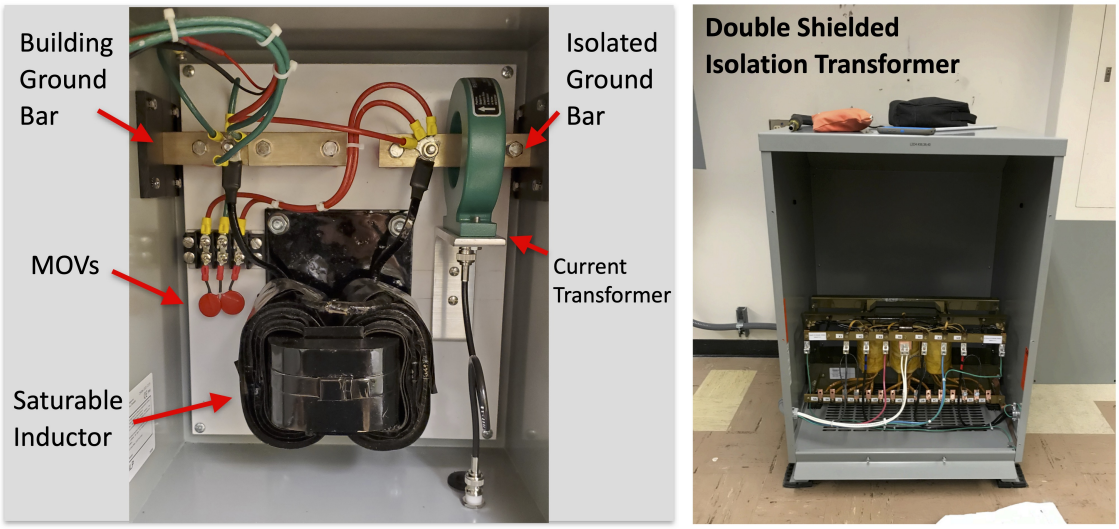

Figure 6.6: (Left) Internal components of the saturable inductor chassis. (Right) Exterior of the double-shielded isolation transformer unit.

the building ground, which immediately triggers an audio alarm. A partial return signal indicates a high-resistance short. This monitoring setup provides a straightforward way to verify that the detector is completely isolated from the building ground after installation and before operating any of the detectors, otherwise data quality could be negatively impacted. In addition, the system runs continuously during operations to alert personnel in the unlikely event of a short to building ground during data-taking.

6.2.4 Charge Readout System

The main components of the TPC, including the field cage, cathode, anode, and light detector, are shown in Figure 6.7. The cathode plane is biased to generate a uniform electric field that drifts ionization electrons from interacting charged particles toward the grounded anode plane. Operating at the nominal electric field of 500 V/cm, ionization electrons travel at a drift velocity of approximately 0.16 cm/μs. Given the total drift distance of about 30 cm, the maximum ionization electron drift time across the volume is roughly 190 μs.

The anode plane is instrumented with a LArPix [89] pixel tile provided by Lawrence Berkeley National Laboratory (LBNL). This is the identical technology developed for the ND-LAr charge readout system and deployed in the 2x2 Demonstrator. The tile consists of a 70 $\times$ 70 array of

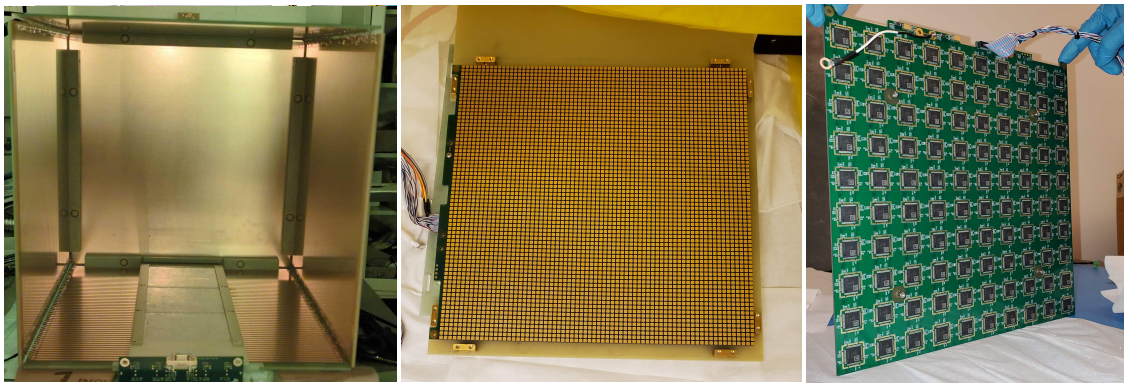

Figure 6.7: (Left) The CATS TPC field cage with the photon detector installed (anode removed for visualization purposes). (Middle) The pixelated side of the anode plane. (Right) The ASIC-instrumented side of the anode plane.

pixels with a 4.4 mm pitch. The foundational operating principles of the LArPix architecture are detailed in Section 3.7.1. The PACMAN controller card, also provided by LBNL, manages data acquisition along with power and data I/O for the pixel plane. It connects to the pixel tile via a dedicated feedthrough PCB on the cryostat lid and a 34-pin twisted-pair ribbon cable inside the vessel. Power is supplied to the card by an external 24 V DC power supply. The PACMAN communicates with the DAQ computer over a standard RJ-45 Ethernet connection, which transmits data packets to and from the pixel plane.

System monitoring and data acquisition are managed on the DAQ computer using Python scripts. Before collecting self-triggered physics data, several calibration steps are required. These include mapping and disabling noisy channels, setting individual channel trigger thresholds based on local pixel noise levels, and measuring baseline channel pedestals. Thresholds are optimized automatically on a channel-by-channel basis to the lowest stable values to minimize noise-induced data packets while maximizing sensitivity to true ionization signals. All of these steps are automated by dedicated Python scripts that are run on the DAQ computer.

6.2.5 Light Readout System

Reconstructing the absolute drift coordinate of an interaction requires a dedicated light readout system to detect LAr scintillation light and establish the event time (t_0). Because the CATS TPC

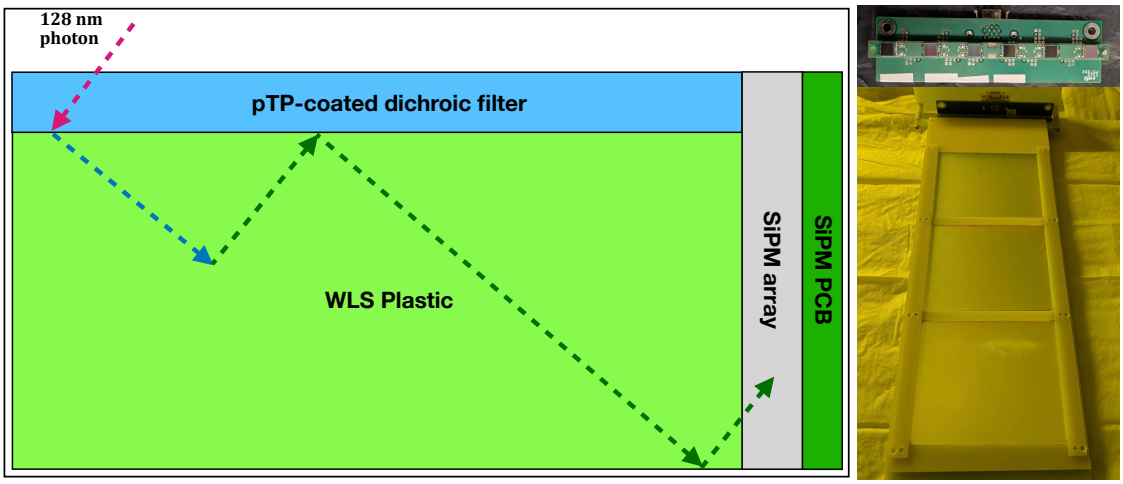

Figure 6.8: (Left) Operational concept of the CATS photon detector. (Right) The physical photon detector assembly (bottom) and the accompanying SiPM PCB motherboard (top).

is compact and surrounded by opaque field cage panels, the photon detector must be integrated directly inside the active volume. The system uses an X-ARAPUCA-style photon detector mounted along the bottom field cage panel, as shown in Figure 6.7.

The operating principle of the detector is illustrated in Figure 6.8, which is similar to the DUNE FD PDS (see Section 3.8.2). The device spans the entire length between the anode and cathode planes, with a width covering roughly one-third of the TPC. The interior of the device houses a wavelength-shifting (WLS) light-guide bar. The top acceptance window consists of dichroic filters coated with a p-Terphenyl (pTP) wavelength shifter. Incoming 128 nm scintillation photons are shifted by the pTP layer to approximately 350 nm and re-emitted isotropically.

Upon entering the WLS guide, scintillation photons are shifted a second time to roughly 430 nm. Because this second wavelength is above the cutoff frequency of the dichroic filter (around 400 nm), the filter acts as a mirror, trapping the photons inside the assembly. The internal walls of the detector casing are lined with 3M Vikuiti® reflective foils to prevent light loss. Six equally spaced Silicon Photomultipliers (SiPMs) are positioned at the anode end of the assembly, facing inward and optically coupled to the edge of the WLS guide. Photons trapped within the light guide via total internal reflection travel toward these sensors and are detected. Photographs of the complete photon detector and the SiPM array are displayed in Figure 6.8 (right).

The SiPMs (Hamamatsu S13360-6075PE) are soldered to a PCB motherboard containing only passive routing components for signal and bias lines. A cryogenic cable connects this motherboard to a DB15 feedthrough on the cryostat flange, which links to an external readout module and power supply on the warm side. A single DC power supply on the electronics rack provides the necessary operating bias voltage to all six SiPM channels simultaneously. Because the SiPMs are all biased using a single supply, it was important to down-select sensors which had similar breakdown voltage. This was done by selecting sensors with similar breakdown voltage in warm based on the vendor's measurements. The breakdown voltage of each sensor was then determined in LAr by scanning the bias voltage. To produce a sufficiently high overvoltage level for all six SiPMs to increase the photon detector gain, a bias voltage of 47 V was chosen, enabling 4–5 V over the breakdown voltage for all sensors.

A CAEN DT2740 digitizer directly records the raw signals from each SiPM without intermediate amplification. The digitizer is configured by a nearby computer via a USB 3.0 link using the CAEN WaveDump2 software package, which also manages the light data stream. During normal operations, the digitizer runs in a self-triggering mode with user-defined logic conditions. The data is saved in raw binary form to limit processing time and maximize achievable data rates. The raw data files are then automatically converted into CERN ROOT files by a dedicated decoder script for offline analysis.

Initial LAr tests used a basic logic condition that triggered whenever any single channel exceeded a set threshold. However, at low thresholds, the trigger rate increased dramatically due to localized 1–2 photoelectron (PE) signal pulses isolated to individual channels. Because these isolated signal pulses likely did not originate from physics events like cosmic-ray muons or radiogenic decays, a majority trigger logic was implemented for subsequent runs. This majority logic requires at least two of the six SiPM channels to exceed their thresholds within a tight coincidence window to validate a trigger. In subsequent tests, the thresholds were set to trigger on signals greater than two or three PE. Dedicated LED runs were performed to determine the single-PE gain for each SiPM, which was used to determine at what voltage level to trigger the digitizer readout.

When a valid light trigger occurs, the digitizer outputs a physical synchronization pulse. This pulse is routed to the LArPix PACMAN card through a LEMO cable, allowing the time of the light event to be registered in the charge data stream as a specialized external trigger packet. These hardware packets enable offline event matching between the independent charge and light systems. However, because the clocks in the LArPix and the CAEN digitizer systems operate on separate, independent internal oscillators, precise charge-light matching requires an offline clock synchronization workflow. The entire synchronization and charge-light matching workflow is described in Section 6.3.

6.3 Charge-Light Matching

6.3.1 Synchronization Procedure

Because the charge and light readout systems operate independently during LAr operations, they must be synchronized to achieve accurate charge-light event matching. However, the version of the PACMAN card used in CATS cannot accept an external clock signal, whereas the CAEN digitizer can. To work around this hardware limitation, a similar synchronization method as is used for the 2x2 Demonstrator was adopted.

To correct for relative clock drift, a pulse-per-second (PPS) signal is distributed simultaneously to both the PACMAN card and the CAEN digitizer, forcing both systems to reset their local timestamps at the exact same moment. This synchronization signal is generated by an SRS digital delay generator, which is controlled remotely by the DAQ computer.

While the PPS resets keep the clocks from drifting apart indefinitely, the fixed network or electronic delay between the two independent systems must still be determined to properly align the timestamps. Instead of measuring this delay directly, the PPS resets are initiated at a known time interval at the start of each data-taking run. The operational sequence follows these steps:

- A new charge data run is initialized automatically (typically every 10 min).
- The DAQ computer sends a command to turn off the PPS signal.

- The system waits for approximately 30 s while initial charge data begins to accumulate.
- The computer re-initiates the PPS signal simultaneously to both systems.

This re-initialization creates a distinct signature in the data: a large, continuously rising timestamp followed by a sharp drop back to zero. By identifying the first PPS reset after this re-initialization in both data streams, the offline matching code calculates and corrects for the constant electronic delay between the two systems. This alignment sequence is repeated at the start of every file to prevent long-term synchronization loss, which can occur if a PPS reset packet is occasionally missed by the LArPix readout.

One minor limitation of this approach is that the analysis software must process each data file from the very beginning to identify the first re-initialization event in that specific period. To avoid data complications and ensure alignment tracking remains straightforward, data-taking is stopped and restarted at regular intervals throughout a CATS run to keep individual data periods clearly separated.

Figure 6.9 displays the charge and light timestamps recorded over a single data-taking period. The charge timestamps correspond to external trigger packets embedded directly into the LArPix data stream whenever the light system fires. The light timestamps record the true trigger time of the CAEN digitizer. To align the entire dataset, the matching code processes the files chronologically, locates the first large timestamp spike, and introduces a time offset relative to that initial reset point. From there, each subsequent PPS reset is identified by a sharp drop in the timestamp value, which realigns the systems second by second. This loop continues until the next large timestamp spike is encountered, signaling the start of a new data file and prompting the code to re-establish the temporal alignment.

As a visual check, Figure 6.10 shows the aligned charge external triggers and light system self-triggers during a single one-second interval represented by vertical lines at their corrected relative timestamps. According to this plot, it appears the triggers in either system are lined up with each other in time. However, Figure 6.11 reveals that the absolute time difference between matched triggers can be as large as about 9 μs, with the light-trigger timestamps consistently lagging behind

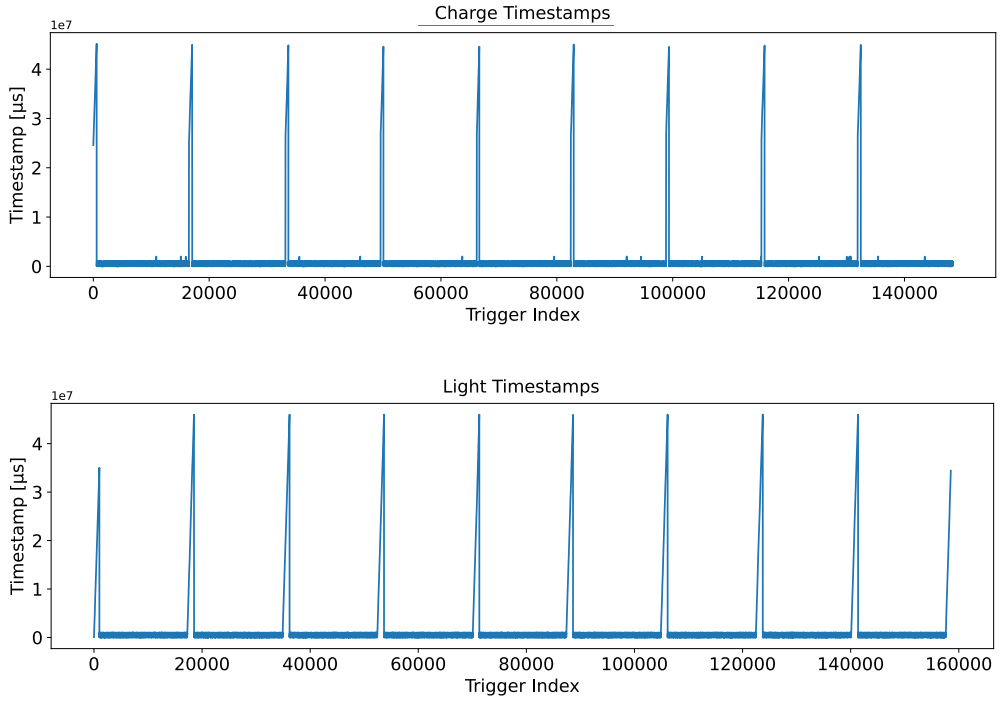

Figure 6.9: (Top) Charge trigger timestamps and (Bottom) light trigger timestamps during a single 90 min data-taking period. The large timestamp spikes correspond to the re-initialization of the PPS, marking the time at which the timestamps are re-aligned. Outside of these spikes, the timestamps oscillate between 0 and 10^6 µs due to the PPS resets.

the charge-trigger packets. This small, residual offset between simultaneous physical events may be caused by minor clock drift that accumulates within the 1 second windows between successive PPS resets.

Determining the exact event time is essential for calculating the drift distance inside the TPC, which is used to reconstruct the spatial coordinates of interactions and correct for charge attenuation from impurities. A timing error of 9 µs introduces an uncertainty of approximately 1.5 cm in the reconstructed drift coordinate, which represents about 5% of the maximum drift distance. While this spatial error is relatively small, it can be entirely eliminated by using the forwarded hardware trigger packet in the charge data stream. Because this external trigger timestamp is generated by the same 10 MHz master oscillator used to provide the timestamp for the pixel digitizations, it provides an inherently synchronized reference time that is unaffected by clock drift between the separate readout systems. The accuracy of these charge-trigger timestamps is limited

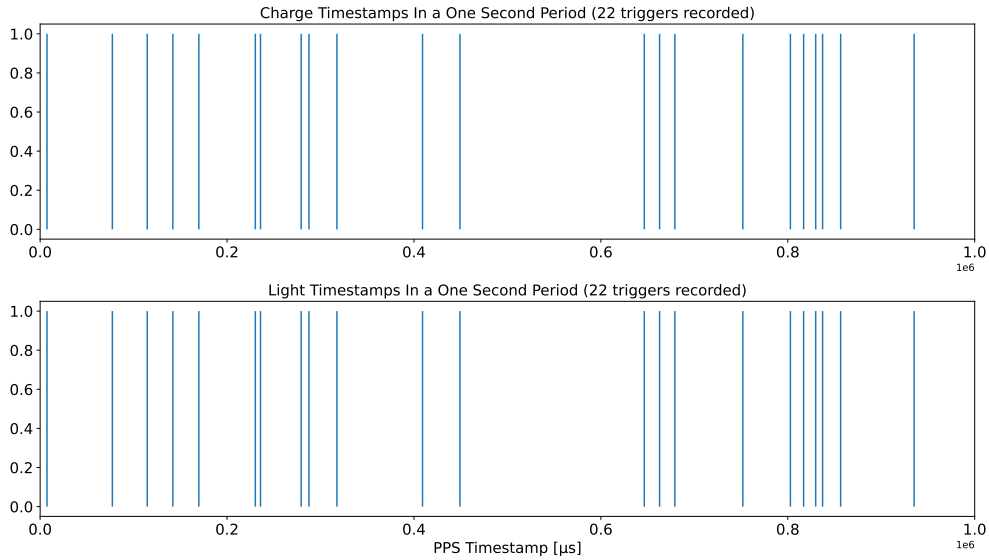

Figure 6.10: Visual alignment of charge and light triggers within a single one-second window, where each charge or light trigger signal is represented by a vertical line at its respective timestamp value.

only by the 100 ns sampling resolution of the LArPix clock. For this reason, the charge-trigger hardware packet is utilized as the primary time reference for drift coordinate reconstruction.

6.3.2 Light Readout Signal Deconvolution

The light detector signals recorded by the digitizer represent a convolution of the true scintillation light signal with the characteristic response of the readout electronics. Deconvolving these waveforms removes the electronic response, which is essential for determining precise event timing, resolving the true pulse shape, and calculating the absolute light yield from particle interactions inside the TPC.

The deconvolution method used for CATS relies on single-PE templates extracted from dedicated LED calibration runs. To build this template, the initial pulse regions of thousands of calibration waveforms are integrated and compiled into a histogram. This distribution exhibits distinct Gaussian peaks corresponding to discrete numbers of collected photoelectrons, including a baseline pedestal (0 PE) peak. The separation between these individual peaks and the width of the baseline pedestal peak provide a direct measure of the system's signal-to-noise ratio. By selecting and filtering waveforms with integrals falling under the 1 PE peak, an ensemble of single photon

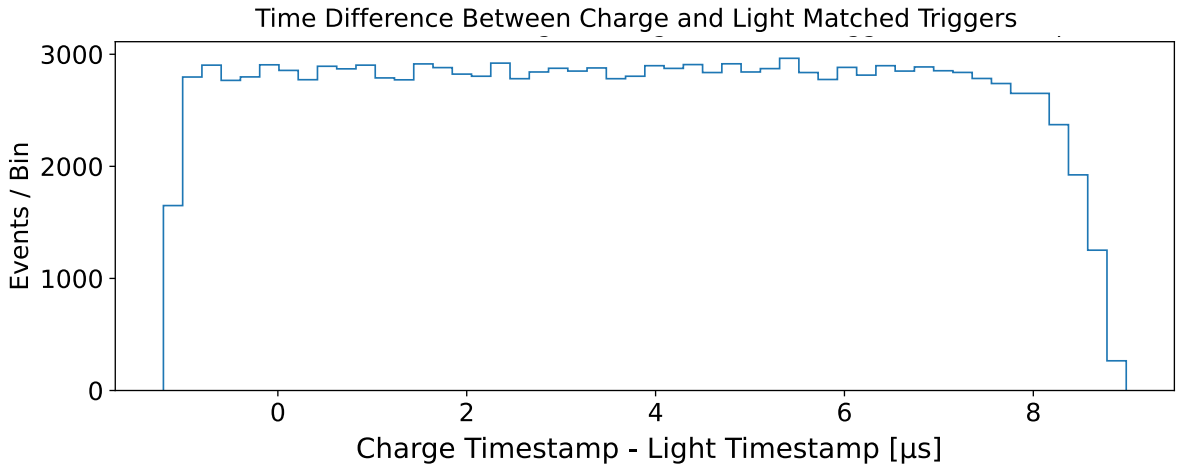

Figure 6.11: Absolute timestamp difference between matched charge and light triggers recorded during a single continuous data-taking run.

events is isolated and averaged to construct a clean single-PE impulse response template.

While an ideal multi-PE waveform would simply be a scaled copy of this single-PE template, individual photons arrive at the SiPM faces at different times throughout an event. To account for this temporal distribution, a frequency-domain deconvolution is used. According to the convolution theorem, the convolution of two time-domain signals is mathematically equivalent to the pointwise multiplication of their respective Fourier transforms in the frequency domain. Therefore, the electronics response can be isolated and removed by transforming the raw waveform into the frequency domain and dividing it by the frequency-domain template of the single-PE response. The full deconvolution workflow consists of the following steps:

1. Convert both the raw SiPM waveforms and the single-PE template into the frequency domain using a Fast Fourier Transform (FFT).

2. Divide the frequency spectrum of the raw waveform by the frequency spectrum of the template.

3. Apply a low-pass filter in the frequency domain to suppress high-frequency electronic noise.

4. Transform the filtered frequency spectrum back into the time domain using an Inverse Fast Fourier Transform (IFFT).

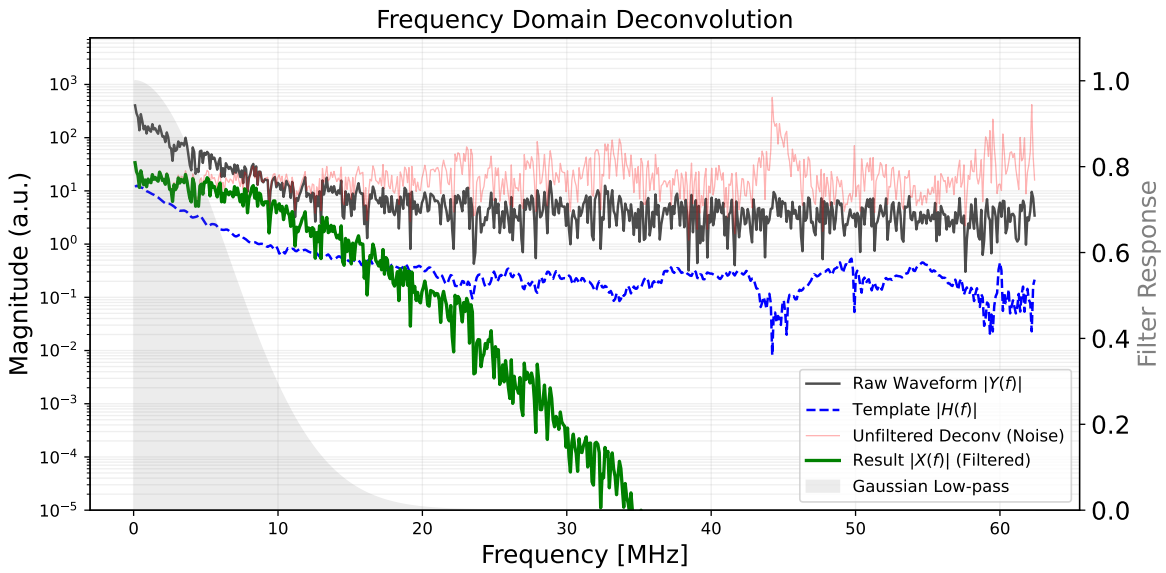

Figure 6.12: Visual representation of a calibration waveform and the single-PE template at successive stages of the deconvolution process in the frequency domain.

5. Perform a cross-correlation alignment to correct for minor phase shifts introduced during the deconvolution process, returning the processed pulse peaks to their original physical timestamps.

These processing steps are illustrated in Figure 6.12 for a single waveform event. The template used in this workflow is an average of many single-PE pulses and does not receive any upfront hardware filtering. After dividing the raw waveform spectrum by the template spectrum, a Gaussian low-pass filter with a cutoff frequency of 7.5 MHz is applied to remove high-frequency noise before performing the inverse transform.

Examples of raw and deconvolved time-domain signals are shown in Figure 6.13. The narrowing of the pulse widths indicates that the electronic response is effectively removed by this technique. The template introduces a small amount of low-amplitude baseline noise prior to the main pulse due to minor LED artifacts, but this effect has a negligible impact on the reconstructed signals.

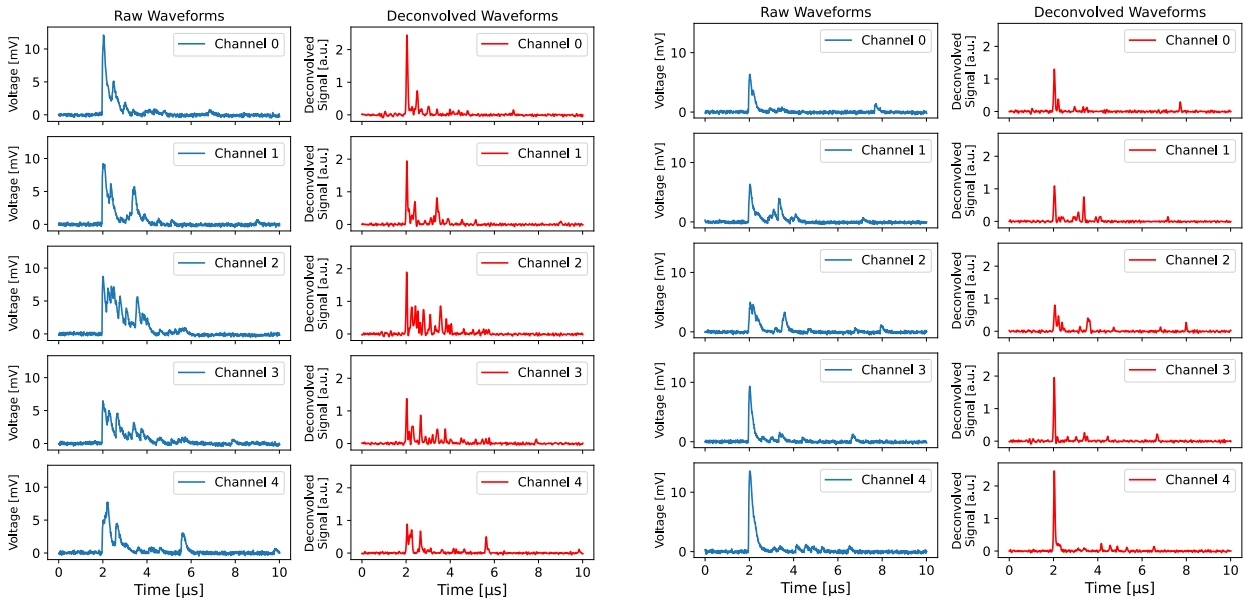

(a) Deconvolution example #1.

(b) Deconvolution example #2.

Figure 6.13: Examples of raw and deconvolved waveforms for individual physics events. The final SiPM channel is omitted because it was physically obscured inside the detector and recorded no signal.

6.3.3 Charge-Light-Matching Performance

Figure 6.14 displays the reconstructed drift time (t_{drift}) distribution alongside the x-coordinate versus t_{drift} distribution for low-energy charge clusters recorded in the CATS TPC. For this test, the HV was operated at a lower level such that the drift electric field strength was lower than the nominal 500 V/cm, which caused the maximum drift time to be larger than the previously stated 190 μs value. A clear excess of events is reconstructed between zero and the maximum drift time in the TPC of around 235 μs. The distributions show spatial features that align with the structural segmentation of the light detector's dichroic filter plates. Additionally, the event rate decreases as the distance from the SiPM array increases, with the lowest rates occurring farthest from the anode plane ($t_{\text{drift}} = 0$).

A spatial and charge selection filter is applied in Figure 6.14 to improve the charge-light matching selection purity, i.e. the size of the excess with respect to the sideband in Figure 6.14. In particular, the selection keeps charge clusters that are about 1.5 cm away from the light detector along the vertical axis, about 5 cm away from the center of the light detector along the horizontal pixel plane axis (x). In addition, clusters with a total charge less than 20 ke$^-$ are removed from the se-

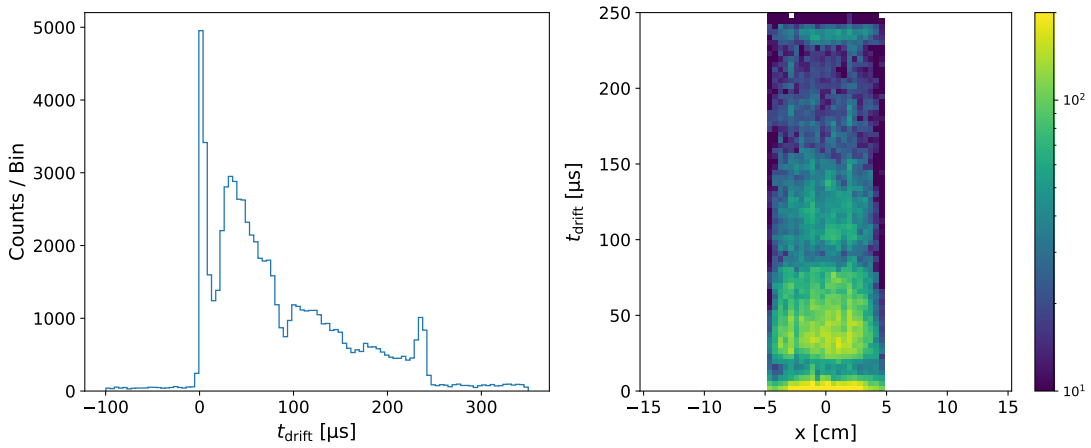

Figure 6.14: Charge-light-matched low-energy events in the CATS TPC in close proximity to the light detector. Note that the drift electric field was operated at a lower value than 500 V/cm for this test, such that the maximum drift time was around 235 µs.

lection. This procedure aligns with a similar one used in the 2x2 Demonstrator (see Section 5.3) to improve charge-light matching selection purity by requiring selected charge clusters to be nearby light detectors. The events producing the excess in Figure 6.14 are likely produced by a combination of ambient radiogenic decays and localized, point-like ionization depositions left by passing cosmic-ray muons.

One notable feature are be seen in Figure 6.14 which is a spike at low and high drift times. These events are likely due to cosmic-rays that clip the boundaries of the TPC producing short tracklets that pass the selection and mimic a low-energy event. A similar observation was made in the single module tests at the University of Bern (see Section 5.3). Additionally, less events appear to be reconstructed at large drift times, which is likely caused by lower photon detection efficiency as a function of the distance from the SiPMs (located near $t_{\text{drift}} = 0$). Scintillation photons arriving at the light detector surface at larger drift distance would have to reflect more times before reaching an SiPM compared to lower drift distances, such that they would more likely be absorbed by the detector materials. This may also explain why the peak at large drift times is smaller in amplitude compared to the peak at low drift times.

As a further demonstration of the low-energy reconstruction capability in the CATS TPC, Figure 6.15 shows examples of Michel electron candidates identified within the TPC volume. A

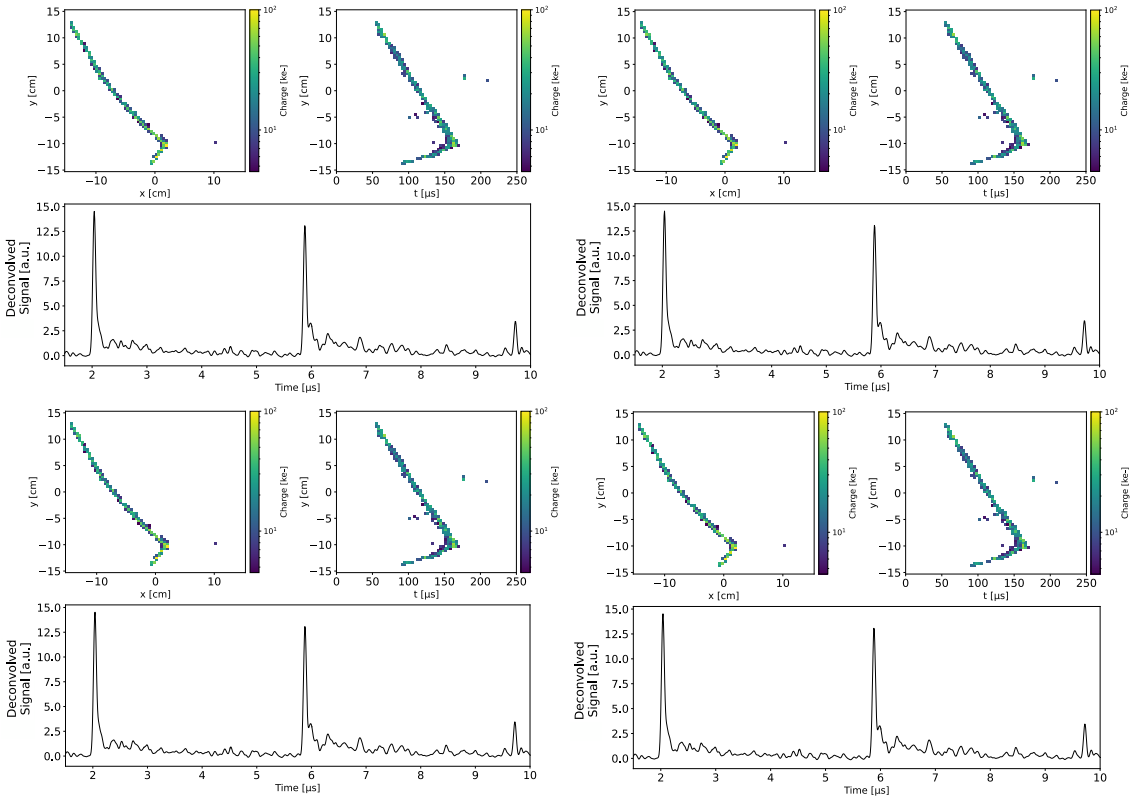

Figure 6.15: Charge activity and light waveforms for a select number of Michel electron candidates in the CATS TPC. The light waveforms are deconvolved and summed across the five unobscured SiPM channels. The time separation between light readout peaks appear consistent with the approximately 2.2 μs muon lifetime [123]. Note also that the drift electric field was operated at a lower value than 500 V/cm for this test, such that the maximum drift time was around 235 μs.

Michel electron refers to when a cosmic-ray muon loses all of its kinetic energy and stops inside a detector and eventually decays to produce an electron. This kind of event produces a very clear signature in a TPC, where the charge readout signals show a long track (the muon) that ends inside a detector with a short "tail" (from the electron). Additionally, the light readout will see two distinct peaks separated by upwards of multiple microseconds, due to the approximately 2.2 μs muon lifetime [123]. The four events displayed in Figure 6.15 have both of these features, which not only indicates these are likely true Michel electron events but also implies that the charge-light matching procedure is working effectively. An appropriate followup study would be to measure the muon lifetime for a large set of Michel electron candidates in CATS, which would be an even more robust way to demonstrate the charge-light matching effectiveness.

These results demonstrate that the offline charge-light matching framework developed for

CATS successfully enables the full 3D reconstruction of particle interactions inside the TPC. Future studies will focus on optimizing these event selection criteria and further characterizing the overall efficiency of the matching workflow.

6.3.4 Purity Monitoring System

A purity monitor is an independent diagnostic device that provides localized measurements of the electron lifetime outside of the main TPC volume. These devices are standard calibration tools across various LArTPC experiments, including ICARUS [66], MicroBooNE [67], and DUNE [9]. The purity monitor uses a TPC-like design consisting of a cathode, an anode, and a surrounding field cage. A photocathode is mounted directly on the cathode surface and emits electrons via the photoelectric effect when struck by ultraviolet (UV) light. An internal electric field drifts these photoelectrons toward the anode, where they are collected. As the electron cloud passes through the purity monitor, it generates distinct induction and collection signals at the cathode and anode. These signals are independently amplified and digitized. The LAr electron lifetime (τ) is calculated using the following relation [124]:

$$\tau = \frac{1}{\ln\left(Q_A/Q_C\right)}\left(t_2 + \frac{t_1 + t_3}{2}\right), \tag{6.2}$$

where the charge ratio is defined as:

$$\frac{Q_A}{Q_C} = \frac{g_C}{g_A}\frac{V^A_{\text{peak}}}{V^C_{\text{peak}}}. \tag{6.3}$$

In Equation 6.3, V^A_{peak} and V^C_{peak} represent the baseline-corrected peak amplitudes of the anode and cathode voltage waveforms. The factor $\frac{g_C}{g_A}$ is the ratio of the electronic gains of the two charge preamplifiers, which is calibrated to approximately 1 for the CATS purity monitor electronics.

A schematic diagram and a photograph of the CATS purity monitor device are shown in Figure 6.16. The assembly includes two wire grids fabricated by securing an electroformed copper wire mesh (Precision Eforming MC7) between stainless steel support rings [125]. These grids act

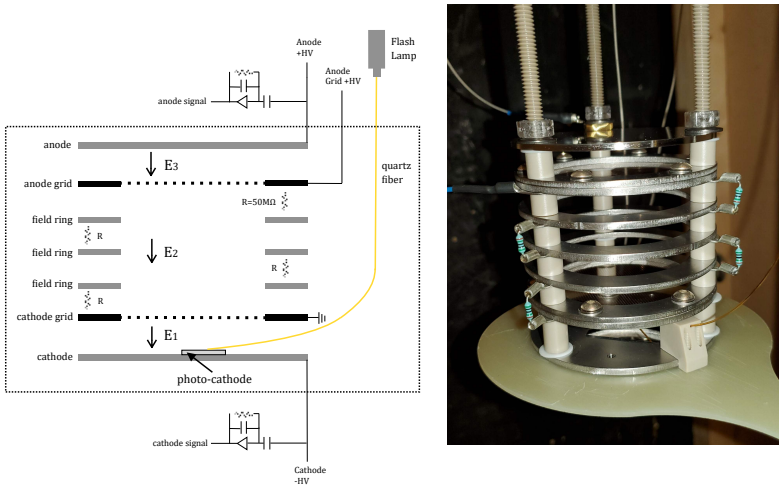

Figure 6.16: (Left) Schematic diagram of the CATS purity monitor components and electrical connections. (Right) Photograph of the assembled CATS purity monitor unit.

as Frisch grids to shield the anode and cathode from early induced currents caused by the drifting electrons before they pass the grid planes.

Three stainless steel field-shaping rings sit between the two grids. These rings are interconnected by 50 MΩ resistors to form a voltage divider that ensures field uniformity throughout the central drift region. The cathode, anode grid, and anode are biased independently by separate HV power supplies, while the cathode grid is tied directly to ground potential. The HV lines are filtered by custom low-pass filter PCBs to suppress high-frequency noise.

This biasing scheme establishes three distinct electric field regions (E_1, E_2, E_3) that can be adjusted independently. To ensure the drifting electron cloud passes transparently through the wire meshes without being collected prematurely by the grid wires, the electric fields must satisfy the following geometric transparency ratio:

$$\frac{E_2}{E_1} = \frac{E_3}{E_2} \geq 2.5. \tag{6.4}$$

This particular ratio was calculated as a safe lower limit for the particular geometry of the CATS purity monitor and its wire grids [125], based on the principles described in Ref. [126]. This lower limit is even higher than what was used for a similarly designed purity monitor in Ref. [124]

Figure 6.17: (Left) Photograph of the CATS purity monitor suspended from the cryostat lid. (Right) The dedicated 40 L CATS purity monitor cryostat assembly.

(which was 2 rather than 2.5), indicating that this requirement is likely sufficient to enable effective electron transparency through the wire grids.

When the purity monitor is being used, it is installed in 40 L cryostat, separate from the primary CATS TPC vessel, as shown in Figure 6.17. It hangs directly from the cryostat lid. Before LAr transfer, the vessel is sealed and undergoes a similar pump-and-purge cycle as the main TPC system described in Section 6.2.1. HV lines enter the cryostat through two SHV feedthroughs, with the anode channel routed through its own dedicated feedthrough. Inside the cryostat, HV-rated wires link the feedthroughs to the purity monitor electrodes; on the warm side, coaxial SHV cables connect the feedthroughs to the filtered power supplies. The cryostat lid is also instrumented with mechanical relief valves, a burst disk, a capacitive depth sensor, and an optical viewport to visually monitor the LAr level.

The photoelectrons are generated using a solid silver photocathode. Silver was selected over gold because earlier test stand characterizations demonstrated that it produced higher photoelec-

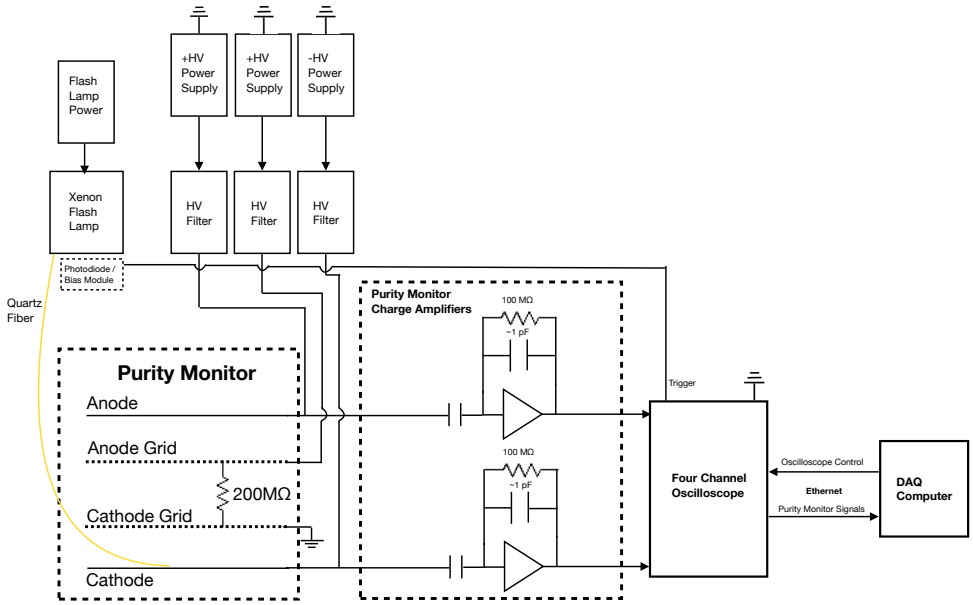

Figure 6.18: System block diagram of the purity monitor control, trigger isolation, and data acquisition networks.

tron yields under identical illumination conditions [124, 125]. Ultraviolet light is supplied by an external Hamamatsu L7685 xenon flash lamp housed in a shielded enclosure near the cryostat. The flash lamp chassis is tied directly to building ground to prevent its HV discharge pulses from inducing noise on the sensitive purity monitor electronics. A 600-µm UV-transmitting optical fiber routes the light from the flash lamp bulb to a fiber feedthrough on the cryostat lid. Inside the vessel, an internal fiber patch cord directs the UV light into the purity monitor housing, terminating at a 17° angle relative to the photocathode surface (see Figure 6.16). To completely electrically decouple the flash lamp from the purity monitor electronics, a ThorLabs photodiode module is used to trigger the data acquisition using light produced by the lamp. A secondary optical fiber collects a small fraction of the light from the flash lamp housing and routes it to the external photodiode. The resulting electrical pulse from the photodiode module triggers the readout electronics, decoupling the data acquisition system from the noisy flash lamp circuit. This triggering scheme, along with effective grounding of the flash lamp enclosure, enabled minimal noise induced on the purity monitor signals.

The raw current signals from the purity monitor are converted to voltage using custom charge-sensitive preamplifiers. These preamplifiers use an OP27G operational amplifier configured with a parallel resistor and a 1 pF feedback capacitor, where the small capacitance was chosen to produce a large overall charge gain. The output of each preamplifier is routed to a digital oscilloscope. The oscilloscope is triggered by the isolated photodiode signal, and multiple waveforms are averaged live to suppress random electronic noise and lower statistical uncertainties. A custom PC-based graphical user interface (GUI) controls the oscilloscope and automates data transfer over an Ethernet connection. A complete block diagram of the readout network is shown in Figure 6.18.

Representative signals from an early LAr test run are shown in Figure 6.19, yielding a measured electron lifetime of approximately 22 μs. This relatively low lifetime indicated that the filtration media had become saturated through extended use and required a full regeneration cycle before subsequent testing. This result demonstrated that the purity monitor as designed could effectively measure electron lifetime in liquid argon to support CATS operations. A useful followup application of the purity monitor would be to do dedicated testing of the impact of ^{220}Rn injection on LAr purity. Additionally, installing the purity monitor inside the same cryostat as the TPC would be advantageous to allow monitoring of LAr purity while filling the detector. This would be a useful capability because CATS currently does not use a LAr recirculation system, meaning that a purity monitor installed inside the cryostat could be used to identify low purity conditions before it would be too late to fix it.

6.4 Dedicated Tests of ^{220}Rn Injection

6.4.1 Injection Technique

To perform the injection, argon gas flows through a source container that holds the ^{228}Th isotope. The thorium source continuously emanates ^{220}Rn, and the gas stream carries the radon atoms out of the source housing, through piping, and into the cryostat that contains a LArTPC. Once inside the cryostat, the ^{220}Rn gas has the opportunity to mix into the LAr volume. Given sufficient mixing, the daughter isotopes of ^{220}Rn may enter the TPC active volume and decay to produce

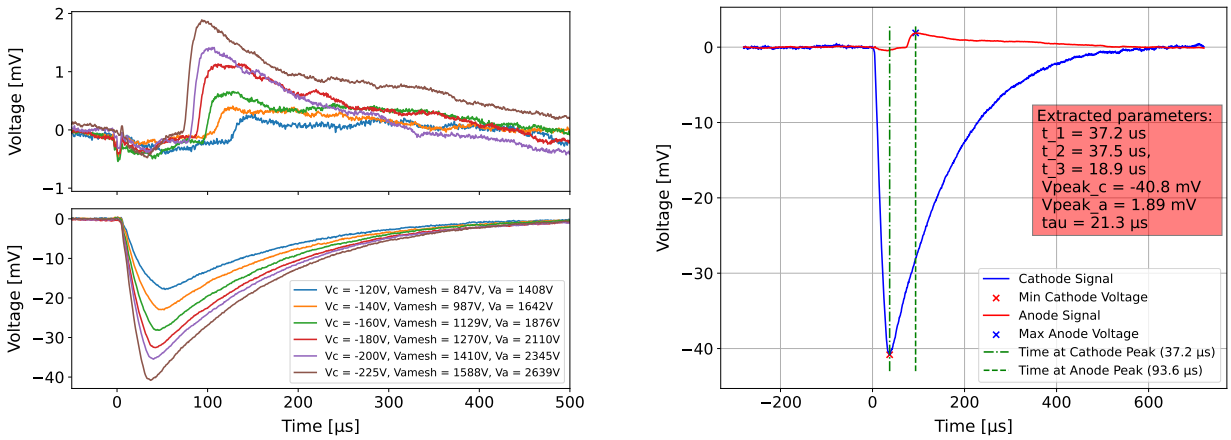

Figure 6.19: (Left) Averaged purity monitor waveforms under various electric field configurations. (Right) Electron lifetime measurement of $\sim$ 22 μs.

detectable signals. Minimizing dwell time of ^{220}Rn in the piping is important, due to the relatively short half life of 55.6 s, to ensure decays only occur once inside the cryostat.

6.4.2 ^{220}Rn Sources

The isotope of ^{228}Th, which emanates ^{220}Rn, can be directly obtained commercially in the form of a powder or dissolved in a solution. The NIDC (National Isotope Development Center) from the US Department of Energy is a prominent source in the US. Sometimes ^{228}Th is electroplated onto a metal film, which is then installed into a leak-tight enclosure [127]. Then, ^{220}Rn would emanate out of the surface of the film, and gas would be flowed across it to push the radioactive gas atoms out of the source enclosure. It is possible to obtain this kind of source commercially, such as from Eckert & Ziegler (E&Z) [128].

Another option is a flow-through source, which involves putting solid ^{228}Th (e.g. in powder form) within an enclosure without involving the electroplating process. In this case, since the material is loose, it is critical to use filters to prevent any ^{228}Th from escaping the source enclosure. This is important because ^{228}Th has a 1.91 y half-life, such that it would introduce a long-lived background into a detector. Flowing gas through the source would push ^{220}Rn atoms out of the source which can be directed into a detector. It is also important to avoid introducing impurities

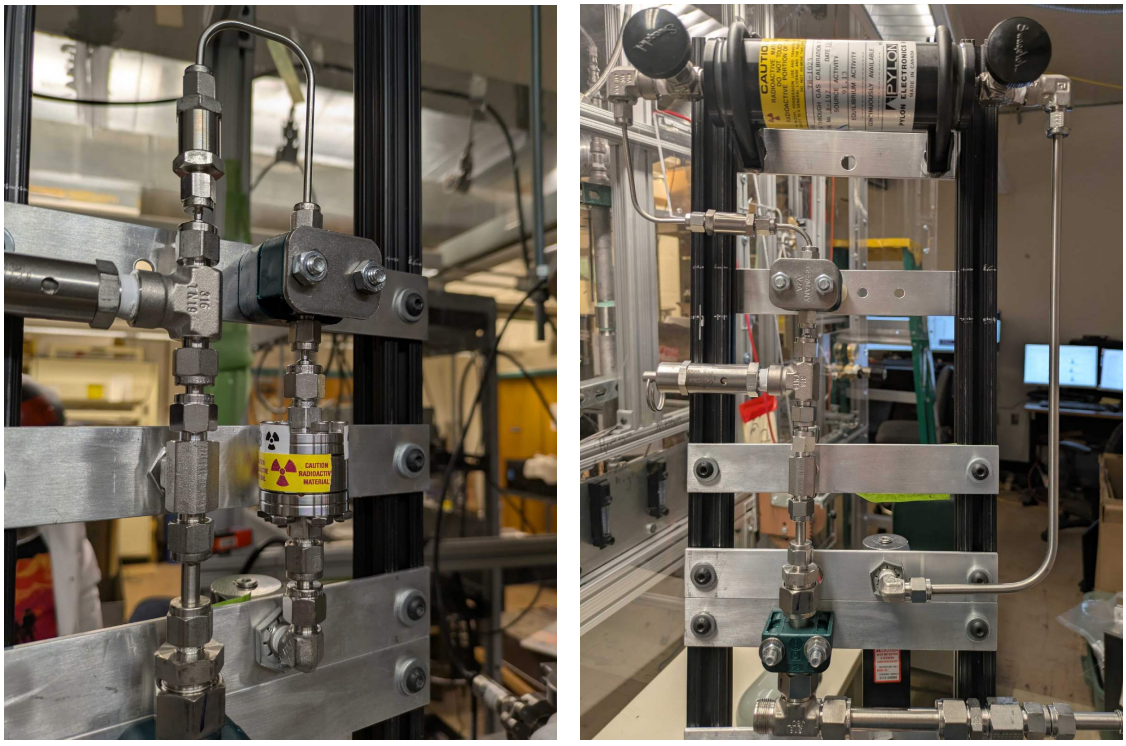

Figure 6.20: ^{220}Rn sources used for tests described in this thesis. (Left) Eckert & Ziegler electroplated source. (Right) Pylon Electronics flow-through source (TH-1025).

from the source into the LArTPC which may be trapped in the materials. To address this problem in CATS, the sources and transfer lines are purged with argon gas for at least 24 hours prior to an injection test.

In this thesis, results of tests with two ^{220}Rn sources are shown. The sources are shown in Figure 6.20 installed in the injection apparatus. Both sources were received by the CSU Radiation Safety department and tested for leaks. The first source is the AF-228-A2 [128] from E&Z, which has ^{228}Th oxide electroplated onto a platinum film. To allow for the removal of ^{220}Rn from the source, the manufacturer removed the window that normally covers the film. It had a ^{228}Th activity of about 30 kBq upon delivery in late 2023, but about 15 kBq in mid-2025. This source was the first source acquired for tests in Prof. Mike Mooney's lab at CSU. The injection apparatus was designed and constructed at CSU to allow high-purity argon gas to flow through the source. The source was tested in the CATS facility in 2024-2025. The second source is the Pylon Electronics Model TH-1025 [129]. It is a flow-through source with an activity of about 91 kBq upon delivery. Shown in

Figure 6.20, it reuses the same injection apparatus as the first source. The source enclosure is leak-tight and installed with a set of Swagelock valves, one on each side, to allow gas to flow through it at specific times. The interior of the enclosure is capped by two glass filters to prevent ^{228}Th or ^{224}Ra from escaping. By the time of the 2x2 Demonstrator tests with this source in November 2025, the ^{228}Th activity was about 80 kBq.

6.4.3 Experimental Setup

The ^{220}Rn injection technique was first evaluated at a small scale to provide preliminary validation before conducting tests inside a full LArTPC. For these initial runs, the light detector was suspended from the cryostat lid inside the 40 L vessel. Figure 6.21 shows the hardware configuration with labels highlighting key components.

The validation tests were conducted using unfiltered UHP-grade LAr. The liquid argon was intentionally not routed through the purification filters during filling because scintillation light signals are not significantly affected by the electronegative impurities that the filtration media targets. Bypassing the filtration system reduced the total volume of LAr required for each test cycle, enabling faster turn-around times between experimental runs.

The ^{220}Rn injection manifold was mounted directly to the cryostat and configured to introduce the radon-doped gas directly into the LAr inlet stream. Because the source was set up specifically for injecting into the LAr fill line, the injection only proceeded while the cryostat was filling with LAr. Light detector data acquisition began as soon as the SiPM array was submerged in liquid. Waiting for full submersion was necessary to ensure that the SiPM breakdown voltage, which is temperature-dependent, remained completely stable during data-taking.

The liquid level was continuously monitored using a combination of the capacitive depth meter and the RTD sensor array. The incoming LAr enters through the cryostat lid and is discharged near the bottom of the volume beneath a protective G10 plate. This plate acts as a diffuser to prevent the liquid from splashing onto the delicate dichroic filter windows during the initially turbulent filling phase.

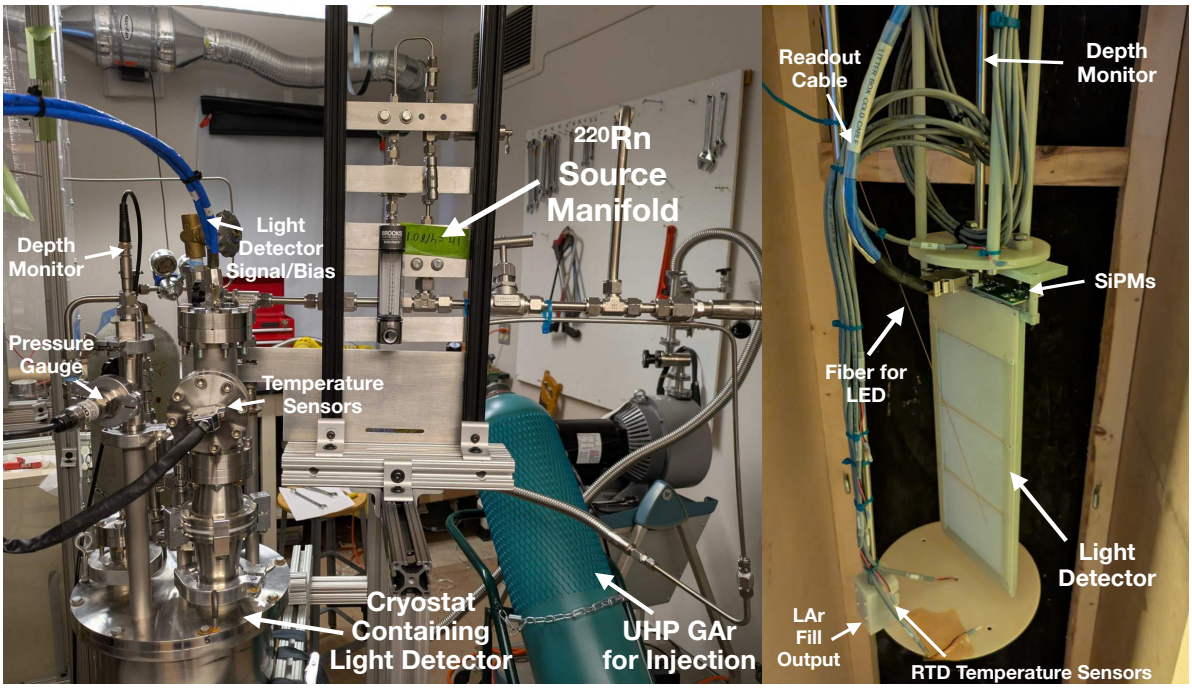

Figure 6.21: Experimental configuration for preliminary testing of ^{220}Rn inline injection using an isolated light detector in LAr.

6.4.4 ^{220}Rn Injection Performance

Figure 6.22 shows the light detector trigger rates during and immediately following ^{220}Rn injection for the two sources. Both datasets follow a matching trend: an initially elevated trigger rate during active injection, followed by a rapid exponential drop over several minutes once the source is isolated from the liquid argon volume. The high trigger rate during the active injection phase is consistent with the fast alpha decays of ^{220}Rn and its immediate daughter ^{216}Po, which have half-lives of approximately 55.6 s and 150 ms, respectively. To extract the effective half-life of this initial decay component, a single exponential function was fitted to the rate data. For the lower-activity E&Z source, the fit yields an effective half-life of $T_{1/2} = 56.4 \pm 0.9$ s, which is in excellent agreement with the true 55.6 s half-life of ^{220}Rn. This confirming result indicates that the elevated trigger rate is driven predominantly by ^{220}Rn α-decays, along with accompanying decays from ^{216}Po, which decays nearly instantaneously in comparison.

Applying the same single-exponential fit to the higher-activity Pylon source data yields an effective half-life of $T_{1/2} = 68.2 \pm 0.3$ s, which sits significantly higher than the true ^{220}Rn value.

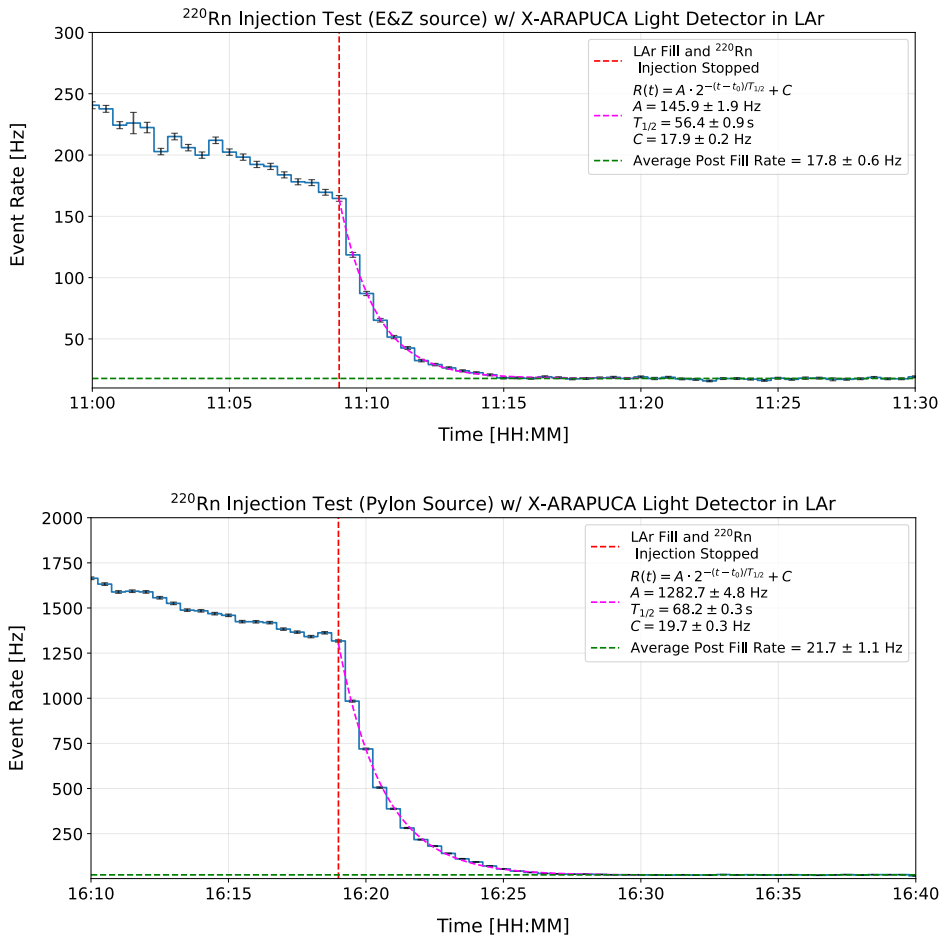

Figure 6.22: Recorded light detector trigger rates over time during and after inline ^{220}Rn injection using the $\sim$ 15 kBq Eckert & Ziegler source (top) and the $\sim$ 90 kBq Pylon source (bottom).

This upward shift is expected because the simplified single-exponential model does not account for the slow increasing rate from downstream decay products, specifically ^{212}Pb ($T_{1/2} = 10.6$ hours) and ^{212}Bi ($T_{1/2} = 60.5$ min). Because the injection window lasted less than a few hours, these longer-lived daughter isotopes did not have sufficient time to reach secular equilibrium with the parent chain. Consequently, a slow accumulation of un-equilibrated ^{212}Bi atoms inside the cryostat over time introduces a secondary, rising decay rate component post-injection, artificially shifting the single-exponential fit toward a longer apparent half-life. This effect is not visibly pronounced in the E&Z source data likely because a lower total amount of ^{220}Rn was introduced into the detector due to the lower original ^{228}Th activity in the source. As a result, the ^{212}Bi background component remains well below the ambient rate background such that it does not bias the primary half-life fit.

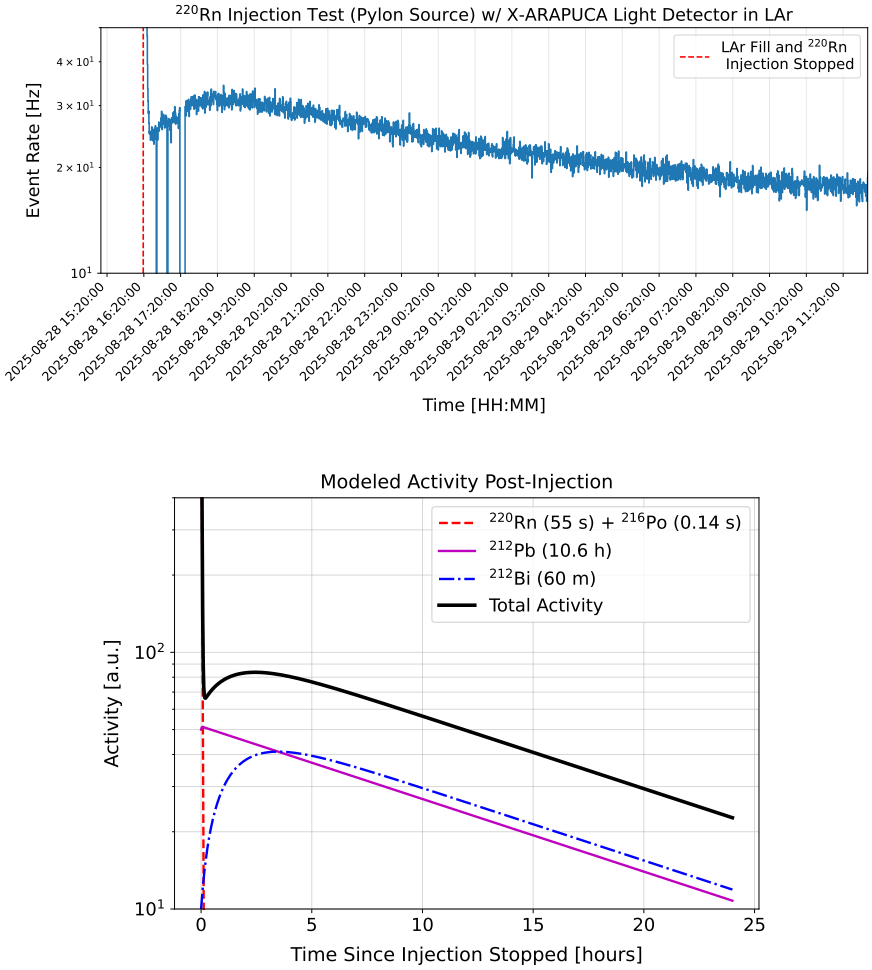

Figure 6.23: (Top) Zoomed view of the light trigger rates following Pylon source injection, revealing a distinct, slow increase long after ^{220}Rn stopped. (Bottom) Calculated isotope activity over time modeled using the Bateman equations.

Figure 6.23 (top) provides a focused view of the Pylon source dataset, revealing the subtle rise and subsequent long-term fall of the event rate after ^{220}Rn injection stopped. The time-dependent population dynamics of the isotope chain (^{220}Rn $\rightarrow$ ^{216}Po $\rightarrow$ ^{212}Pb $\rightarrow$ ^{212}Bi) are mathematically governed by a system of coupled, first-order linear ordinary differential equations. Originally derived by Harry Bateman [130], the rate of change of the i-th nuclide inventory is given by:

$$\frac{dN_i(t)}{dt} = \lambda_{i-1} N_{i-1}(t) - \lambda_i N_i(t), \tag{6.5}$$

where $N_i(t)$ represents the number of atoms of the i-th nuclide at time t, and $\lambda_i = \frac{\ln 2}{t_{1/2,i}}$ is its corresponding characteristic decay constant. The explicit initial activities are defined as $A_{\text{Rn}}(0) = A_{\text{Po}}(0) = A_0$, while the daughter activities satisfy $A_{\text{Pb}}(0) \ll A_0$ and $A_{\text{Bi}}(0) \ll A_0$.

Because the system was stopped before reaching full downstream equilibrium, the initial conditions are characterized by a high inventory of ^{220}Rn and ^{216}Po, achieving secular equilibrium almost instantly due to the 0.145-second half-life of ^{212}Po, and relatively low initial quantities of ^{212}Pb and ^{212}Bi. Given the vast differences in half-lives across the series ($\lambda_{\text{Po}} \gg \lambda_{\text{Rn}} \gg \lambda_{\text{Bi}} > \lambda_{\text{Pb}}$), the post-injection behavior simplifies conceptually into two distinct temporal phases. The first comes from the residual ^{220}Rn and ^{216}Po that clear out rapidly via exponential decay:

$$A_{\text{Rn}}(t) \approx A_{\text{Po}}(t) \propto e^{-\lambda_{\text{Rn}} t}. \tag{6.6}$$

Then, the system behavior is entirely dictated by the 10.64-hour half-life of the ^{212}Pb bottleneck. Because ^{212}Bi decays faster than its parent, it undergoes a transient grow-in period, rising to a peak before locking into a parallel, equilibrium trajectory with ^{212}Pb:

$$A_{\text{Bi}}(t) \propto \left(e^{-\lambda_{\text{Pb}} t} - e^{-\lambda_{\text{Bi}} t}\right). \tag{6.7}$$

The full versions of these equations (not shown for brevity) are solved numerically and the results are shown in Figure 6.23 (bottom). The model successfully reproduces the characteristic rise-and-fall profile observed in the data. The delayed peak represents the gradual increase in rate from ^{212}Bi as it transitions toward secular equilibrium with its parent, ^{212}Pb. While this simplified model omits higher-order detector physics, such as differing photon detection efficiencies for α versus β particles, plate-out of certain isotopes on detector surfaces, changing liquid volume over time, and gas-phase venting, the qualitative agreement shows that the delayed event profile is an expected consequence of the ^{220}Rn decay chain. This confirms why the single-exponential fit in Figure 6.22 yields an higher-than-expected half-life value for ^{220}Rn.

Overall, these results demonstrate that ^{220}Rn can be successfully injected into a volume of LAr

and identified using a photon detector. This success led directly to implementing the injection technique at the 2x2 Demonstrator in late 2025, described in Chapter 7. Additional tests have also been performed or are planned in CATS to further explore the technique for deployment in DUNE.

Chapter 7

Low-Energy Calibrations Using ^{220}Rn Injection in the 2x2 Demonstrator

7.1 ^{220}Rn Injection System

A dedicated ^{220}Rn injection test was carried out during the second data-taking run of the 2x2 Demonstrator in November 2025. This period followed the initial commissioning phase, which provided extensive data for low-energy calibration studies presented in Chapter 5. The same Pylon TH-1025 ^{220}Rn source previously characterized at CSU was deployed in the 2x2 Demonstrator setup, with an activity of approximately 80.2 kBq (2.17 μCi) at the time of use. The source and associated injection manifold were transported to Fermilab and installed shortly before the measurement campaign. The injection system was integrated into the existing cryogenic infrastructure at a location near the bottom of the cryostat, as illustrated in Figure 7.1. The argon carrier gas was supplied from a high-pressure LAr dewar and passed through dedicated Matheson O_2 and H_2O gas filters prior to entering the source to minimize contamination. The purified gaseous argon then was passed through a network of copper piping, through the source enclosure, and onward to the top of the cryostat, where it was introduced via a dipstick assembly. The total length of the gas transfer path from source to cryostat entrance was approximately 7.5 m, with a piping inner diameter of 0.25 inches.

The dipstick is a vertical insertion line that penetrates the cryostat flange and terminates below the LAr surface. It is typically used for ullage gas sampling but was repurposed for ^{220}Rn injection in this study. Its location lies between Module-1 and Module-3, as shown in Figure 7.2. To minimize the risk of contaminating or drawing vacuum on the source, all lines were purged with argon gas for 24 hours prior to injection, ensuring removal of residual impurities from the plumbing and source volume. After the purge period, two injection configurations were evaluated: direct

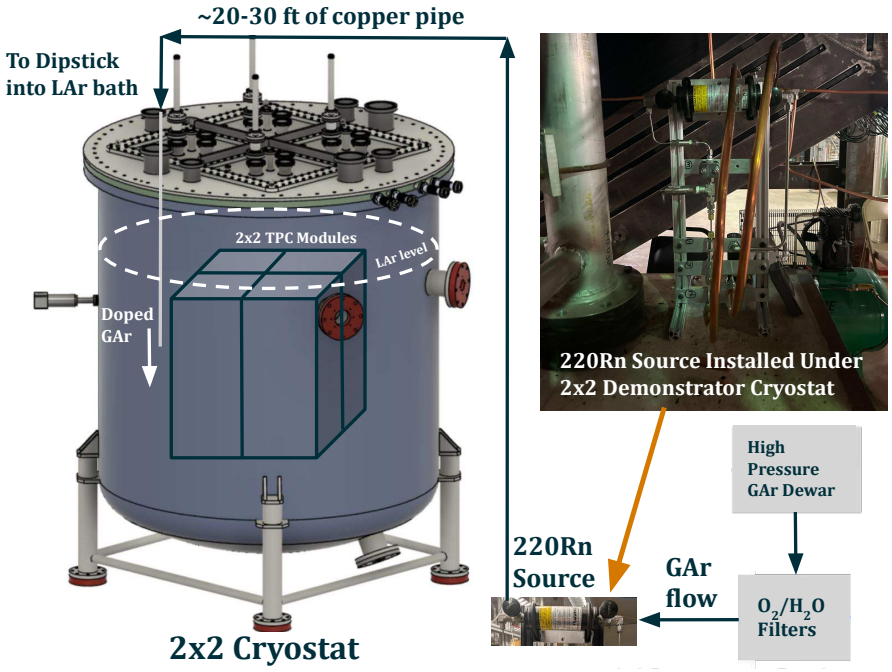

Figure 7.1: Illustration of the ^{220}Rn injection manifold as installed within the gas makeup line in the 2x2 Demonstrator.

injection into the LAr recirculation line (downstream of the purification filters), and injection via the dipstick. The recirculation-line configuration was found to be ineffective for delivering ^{220}Rn into the active volume, whereas the dipstick configuration resulted in a clearly observable and significantly more efficient injection into the detector.

7.2 Activity Introduced By ^{220}Rn Injection

7.2.1 Observations from the LRS Trigger Rates

General Trends

Figure 7.3 shows the LRS trigger rate before, during, and after the ^{220}Rn injection period. The initial 1–2 hours of the study were dedicated to testing an inline injection configuration. No measurable increase in trigger rate was observed during this phase. This is consistent with a low effective delivery efficiency of ^{220}Rn into the detector, likely driven by the combination of limited flow rate and the long transport path between the source and the LAr injection point. Under these conditions, a substantial fraction of ^{220}Rn atoms are expected to decay within the transfer lines

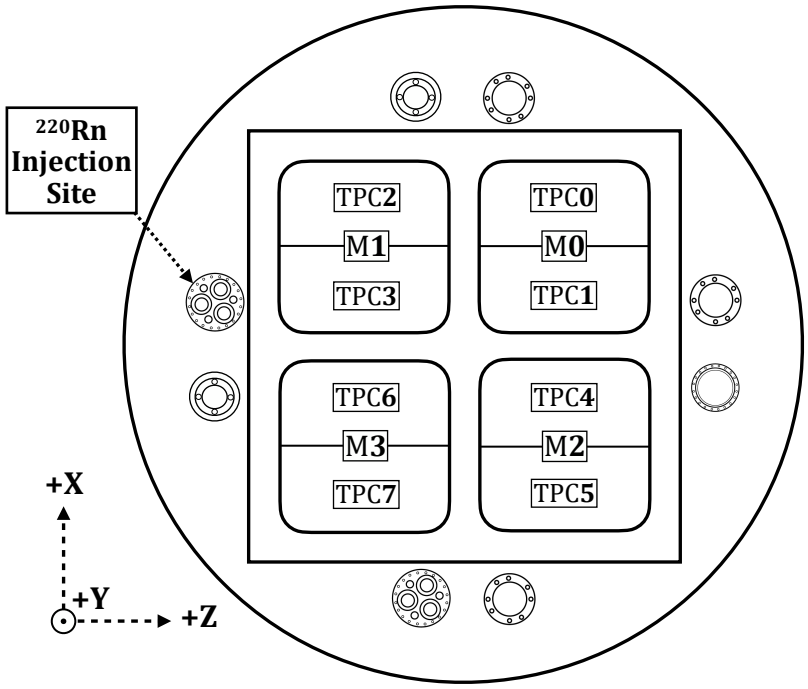

Figure 7.2: Top-down diagram of the 2x2 Demonstrator modules and TPCs, indicating the ^{220}Rn injection location.

prior to reaching the cryostat. In addition, charged daughter isotopes produced along the piping may plate out onto the inner surfaces of the gas lines, further reducing the fraction of activity entering the active volume. These effects are both strongly dependent on the residence time of the gas in the transfer system.

The transport efficiency could, in principle, be improved by reducing the total piping length, thereby decreasing the dwell time of radioactive species in the system. Alternatively, reducing the pipe cross-sectional area would increase the gas flow velocity for a given volumetric flow rate, also decreasing the residence time. However, under an incompressible approximation the flow velocity scales inversely with cross-sectional area, meaning that a smaller radius correspondingly increases linear gas velocity. In practice, neither modification was feasible within the constraints of the test, and higher flow rates were impossible due to pressure limitations of the source, likely exacerbated by back-pressure from the recirculation line.

Given these constraints, the injection strategy was switched to the dipstick configuration for the remainder of the study. This approach proved substantially more effective, producing a clear

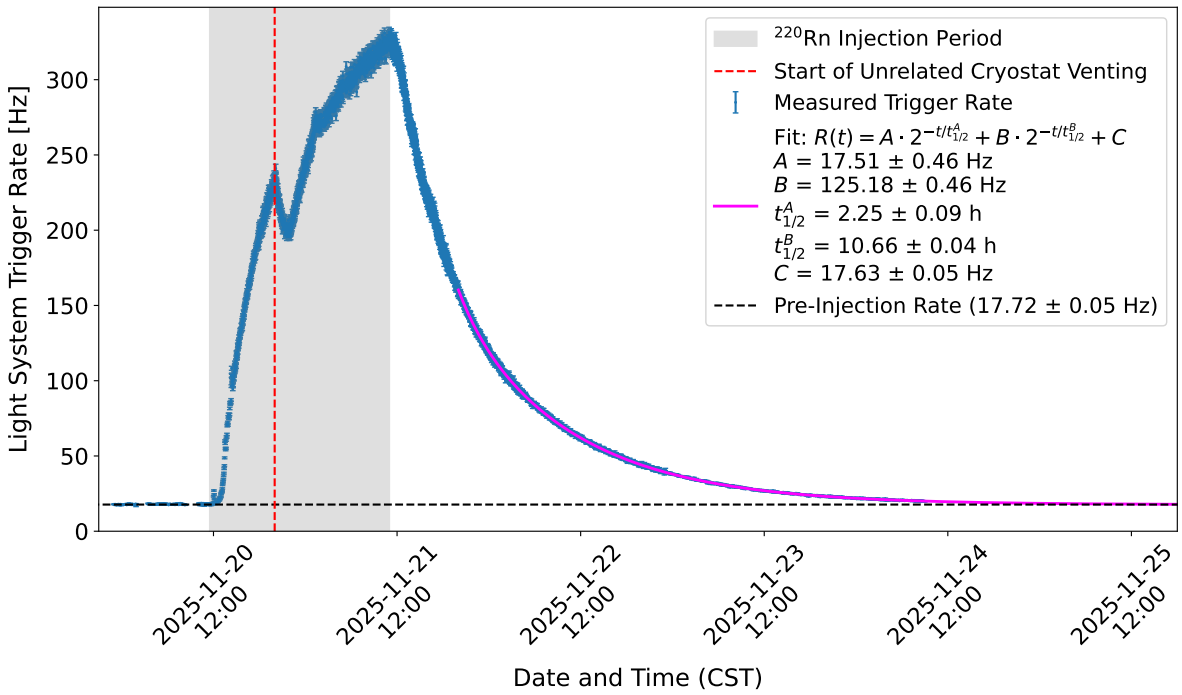

Figure 7.3: LRS trigger rates before, during, and after the ^{220}Rn injection period in the 2x2 Demonstrator.

and sustained increase in both CRS and LRS trigger rates over an approximately 22-hour injection period. The LRS trigger rates over the course of the injection, and long after the source was closed, are shown in Figure 7.3. Following the injection phase, data were collected for approximately three days to study the evolution of longer-lived components in the ^{220}Rn decay chain and their impact on the detector response.

Impact of Venting on LRS Trigger Rate

The ^{220}Rn injection began at 12:00 PM CST on November 20, 2025 and continued for approximately 22 hours. During this period, a prompt increase in LRS trigger rates was observed, indicating successful transport and mixing of radioactivity into the TPC active volume. The trigger rate rose from approximately 18 Hz, dominated by cosmic-ray interactions and ambient radioactive backgrounds, to roughly 325 Hz by the end of the injection period.

At approximately 8:00 PM CST on November 20, a cryostat venting event occurred to relieve internal pressure associated with unrelated cryogenic system behavior. During this interval, a transient reduction in trigger rate was observed. This behavior is consistent with the partial removal

of ^{220}Rn daughter isotopes from the ullage region via gas exhaust, followed by a corresponding decrease in the fraction of activity entering the active volume.

The injection was terminated at approximately 10:30 AM CST on November 21, at which point the gas flow was stopped and the source was closed. Following the injection period, a three-day observation phase was carried out in order to monitor the evolution of longer-lived components in the ^{220}Rn decay chain. The LRS trigger rates decayed in a manner broadly consistent with the 10.6-hour half-life of ^{212}Pb, reaching baseline levels after approximately five days (determined using the fit in Figure 7.3).

The venting event was determined to be unrelated to the injection procedure and was instead driven by a temporary loss of cooling performance in the cryogenic system, which led to increased boil-off and elevated internal pressure. The evolution of trigger rates, venting fraction, and cryostat pressure during the injection period is shown in Figure 7.4. The increase in venting rate appears to be correlated with a temporary suppression of LRS trigger activity, suggesting a reduction in the amount of ^{220}Rn decay chain activity reaching the TPC active volume.

From a physical standpoint, venting removes gas-phase argon from the ullage, thereby reducing the internal pressure of the cryostat. Any ^{220}Rn present in the gas phase are subject to the same removal mechanism. A plausible scenario is that a fraction of the injected ^{220}Rn does not immediately dissolve into the liquid argon, but instead remains in the ullage where it can be preferentially depleted by venting. In this picture, daughter isotopes produced from ^{220}Rn decay may subsequently migrate into the liquid phase on longer timescales, leading to a delayed appearance of activity in the TPC active volume. Venting would interrupt this transport pathway, producing an apparent reduction in observed activity within the detector. Further dedicated studies will be required to validate this interpretation and to resolve the interplay between the injection and cryogenic operations. If confirmed, this behavior would motivate consideration of alternative injection strategies, including direct injection into the ullage. Such an approach may be relevant for large-scale LArTPC detectors as gas-phase injection which could, in principle, improve mixing efficiency in large detectors due to faster transport dynamics in the gaseous phase compared to

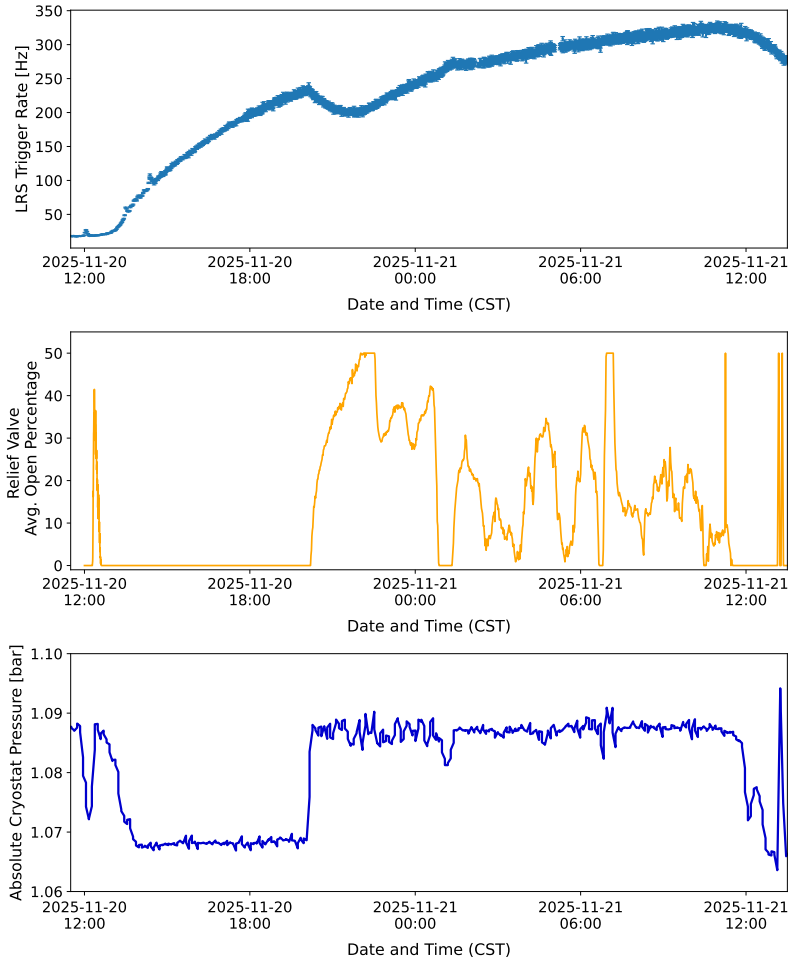

Figure 7.4: (Top) Light readout trigger rates, (Middle) cryostat relief valve open percentage, and (Bottom) internal cryostat pressure during the ^{220}Rn injection period in the 2x2 Demonstrator. The relief valve can turn on and off very quickly in response to pressure, so the graph shows an average of every two data points.

liquid argon.

Other Notable Observations

Besides the transient decrease in LRS trigger rate associated with cryostat venting, several additional features are evident in Figure 7.3. Following the end of the injection period, the trigger rate decrease can be described using a two-component exponential model. The first 12 hours of data (post-injection) are excluded from this fit due to concurrent cryogenic activities that could impact the extracted time constants. The later data are well described by two exponential components with characteristic timescales of approximately 2.2 hours and 10.6 hours, respectively. The shorter

timescale is consistent with removal processes such as LAr recirculation, which can transport and filter out charged or neutral daughter isotopes from the bulk liquid. This possibility is supported by the observation by the MicroBooNE collaboration of the absorption of ^{222}Rn decay products in LAr filter media [131]. The longer timescale is consistent with the decay of ^{212}Pb. Together, these components reflect the combined effects of radioactive decay and detector purification processes, and are consistent with expectations for a system in which LAr level and circulation remain approximately stable.

After closing off the source, the trigger rates do not immediately drop or return to baseline levels. This indicates that the observed activity is largely not caused by short-lived components of the decay chain, namely ^{220}Rn or ^{216}Po, whose half-lives are 55.6 s and 145 ms, respectively. Instead, this indicates the elevated trigger rates are induced by longer-lived daughter isotopes, such as the α decays of ^{212}Bi and ^{212}Po. Despite this observation, Section 7.5.1 describes the reconstruction performance of delayed coincidences of ^{220}Rn and ^{216}Po decays ("RnPo") which are used to make an assessment of LAr flow patterns in the 2x2 Demonstrator.

Another observation from Figure 7.3 is that the LRS trigger rate determined several days after injection (by extrapolating the exponential fit at future times) is about 17.66 ± 0.05 Hz, consistent with the pre-injection value of around 17.72 ± 0.05 Hz. This recovery is expected given the decay structure of the ^{220}Rn decay chain, in which the longest-lived relevant isotope is ^{212}Pb with a half-life of 10.6 hours, ensuring the entire chain decays to stable ^{208}Pb on the timescale of several days. This observation provides an important validation that ^{220}Rn injection does not introduce persistent long-lived backgrounds in the detector. In particular, long-lived backgrounds being introduced into the detector is disfavored as micron-scale particulate filters and an additional sintered filter are installed within the source manifold to prevent any of the powdered ^{228}Th from escaping the source. This result is significant in the context of the DUNE Far Detector, where long-lived radioactive backgrounds could impact low-energy rare-event searches (e.g. solar neutrino measurements).

7.2.2 Charge-Light-Matched Activity

Figures 7.5 and 7.6 show the spatial distributions of charge-light-matched low-energy activity projected onto the x–z plane, before and during the ^{220}Rn injection period, respectively. In addition, Figures 7.7, 7.9, 7.11, and 7.13 show the distributions before ^{220}Rn injection, while Figures 7.8, 7.10, 7.12, and 7.14 show them during ^{220}Rn injection, projected onto the y–z plane.

During ^{220}Rn injection, charge-light-matched low-energy activity is observed across a substantially larger fraction of the TPC active volume compared to the pre-injection period. This increase is consistent with the presence of BiPo decays in the detector, where the prompt scintillation signal from the ^{212}Po α decay provides a strong light signature that enables efficient tagging of the preceding ^{212}Bi β decay using the LRS.

Differences in observed event populations between modules are primarily attributable to variations in photon detection efficiency and LRS threshold settings. In particular, Module-0 exhibits reduced activity due to lower photon detection efficiency and higher effective LRS thresholds relative to the other modules. Module-1 also shows reduced rates compared to Modules 2 and 3, consistent with its comparatively higher LRS thresholds.

A prominent excess of activity is observed near the nominal cathode position ($x \approx \pm 35$ cm). This feature is consistent with previous observations in liquid argon and liquid xenon detectors, including XENON100, where ^{220}Rn progeny were found to plate out on cathode surfaces [113]. In this scenario, positively-charged daughter ions drift under the influence of the electric field and accumulate on the cathode, where they subsequently decay and produce a concentrated population of reconstructed signals. The cathode-associated activity is further investigated in Section 7.3.1, where measurements of the ionization electron drift speed and drift electric field variations are discussed.

7.2.3 Scintillation Light Waveform Shape

Scintillation light in LAr is produced through the formation and radiative decay of excited argon excimer states (detailed in Section 3.4.1). These excimers exist in either a short-lived singlet

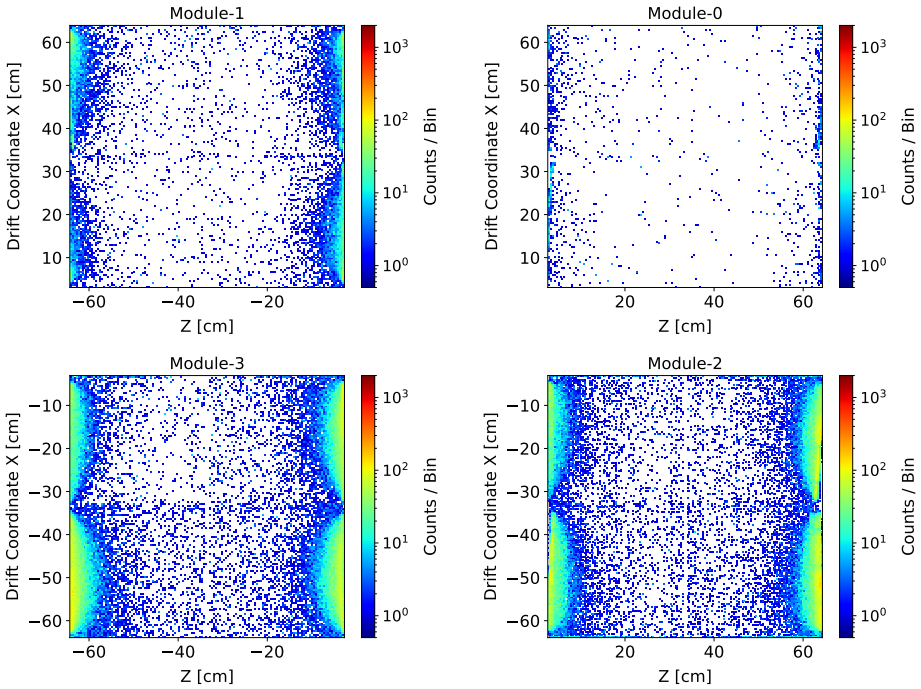

Figure 7.5: Reconstructed spatial distributions, in the 2x2 Demonstrator x–z plane, of charge-light-matched events before ^{220}Rn injection, over a 9.36 hour period.

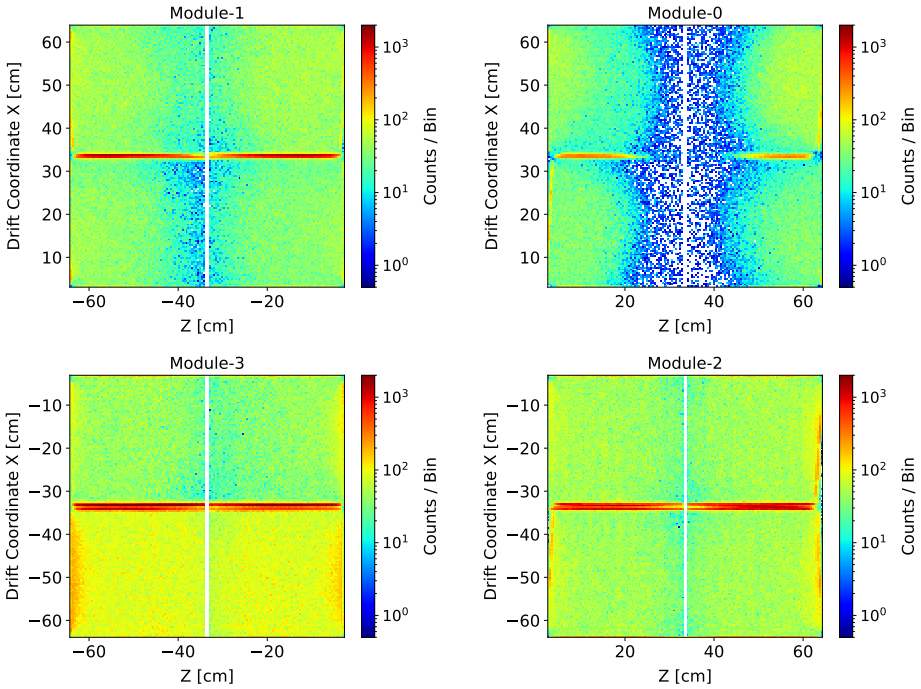

Figure 7.6: Reconstructed spatial distributions, in the 2x2 Demonstrator x–z plane, of charge-light-matched events during ^{220}Rn injection, over a 9.36 hour period.

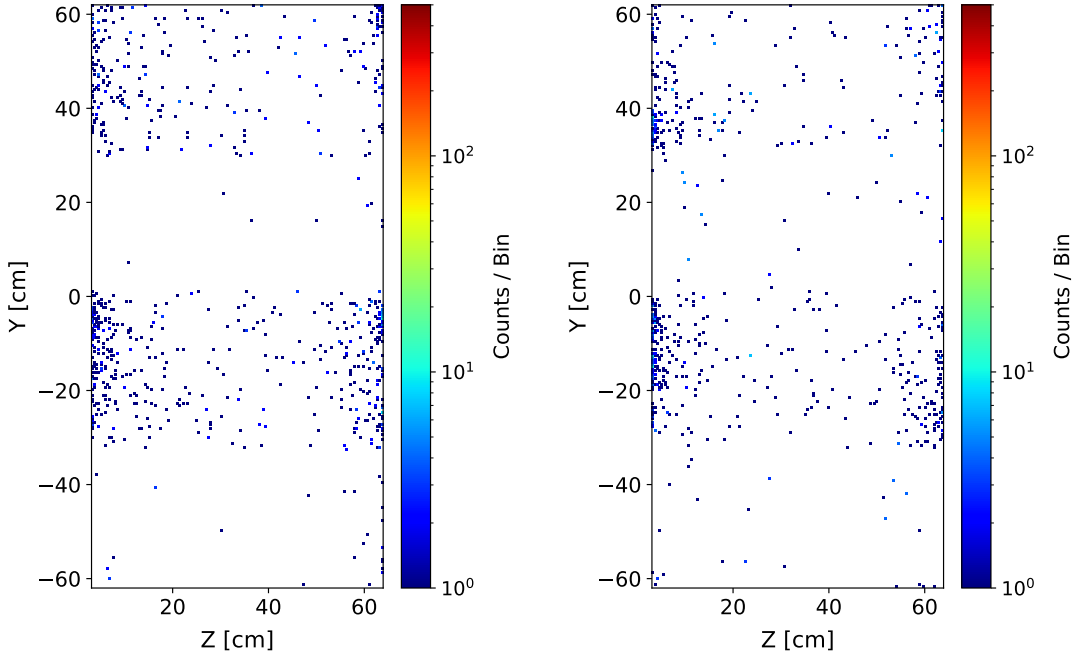

Figure 7.7: Reconstructed spatial distribution, in the 2x2 Demonstrator y–z plane, of charge-light-matched events before ^{220}Rn injection for (Left) Module-0 TPC-0 and (Right) Module-0 TPC-1.

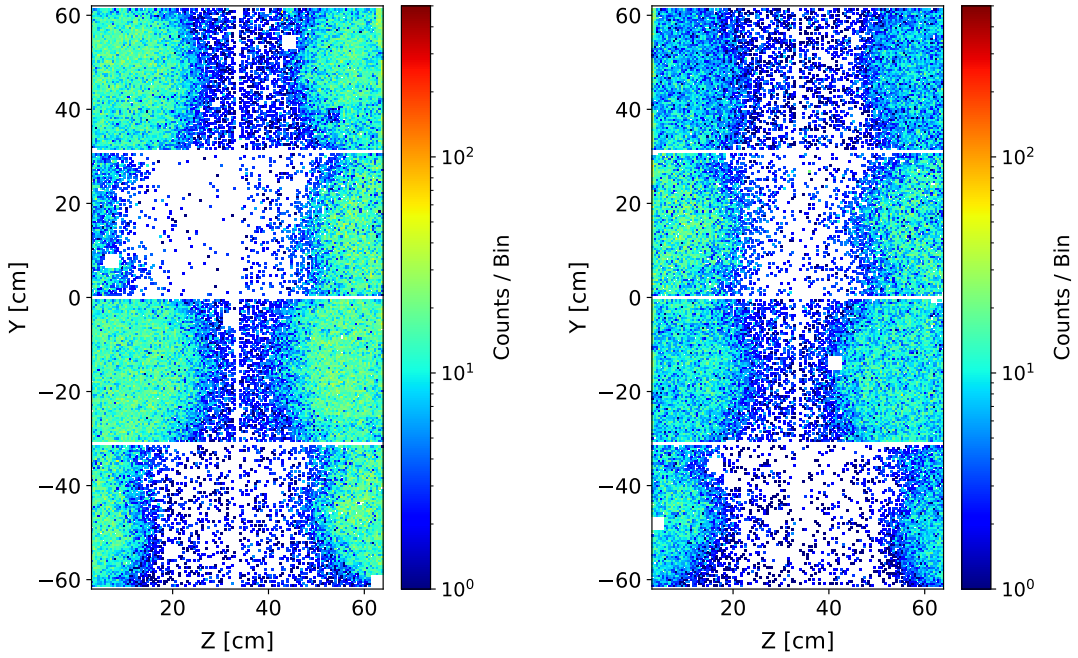

Figure 7.8: Reconstructed spatial distribution, in the 2x2 Demonstrator y–z plane, of charge-light-matched events during ^{220}Rn injection for (Left) Module-0 TPC-0 and (Right) Module-0 TPC-1.

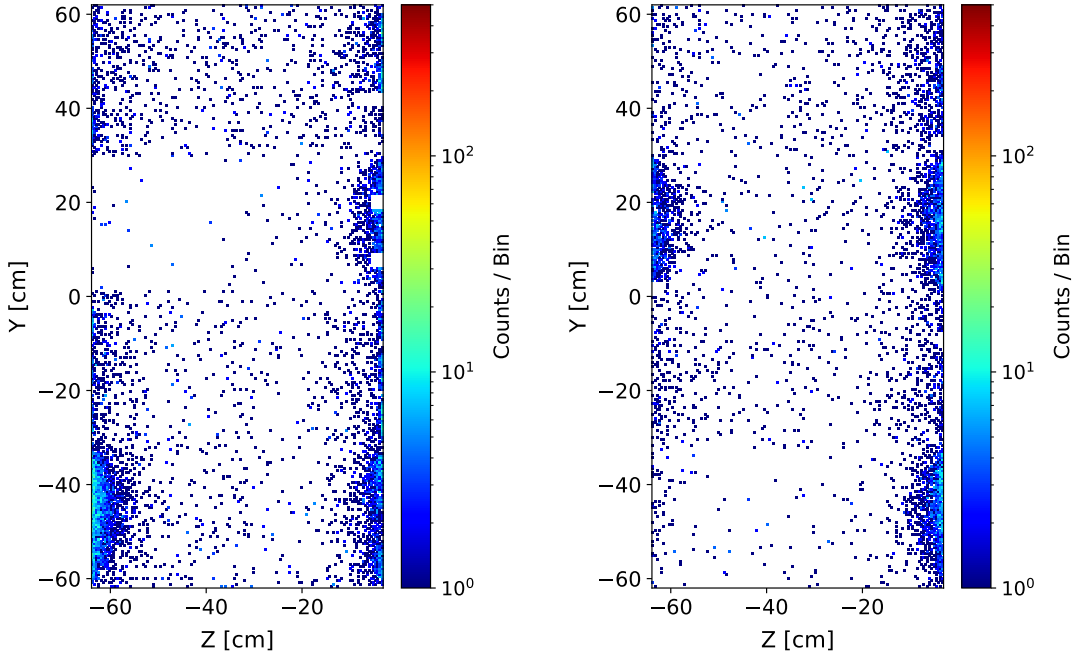

Figure 7.9: Reconstructed spatial distribution, in the 2x2 Demonstrator y–z plane, of charge-light-matched events before ^{220}Rn injection for (Left) Module-1 TPC-2 and (Right) Module-1 TPC-3.

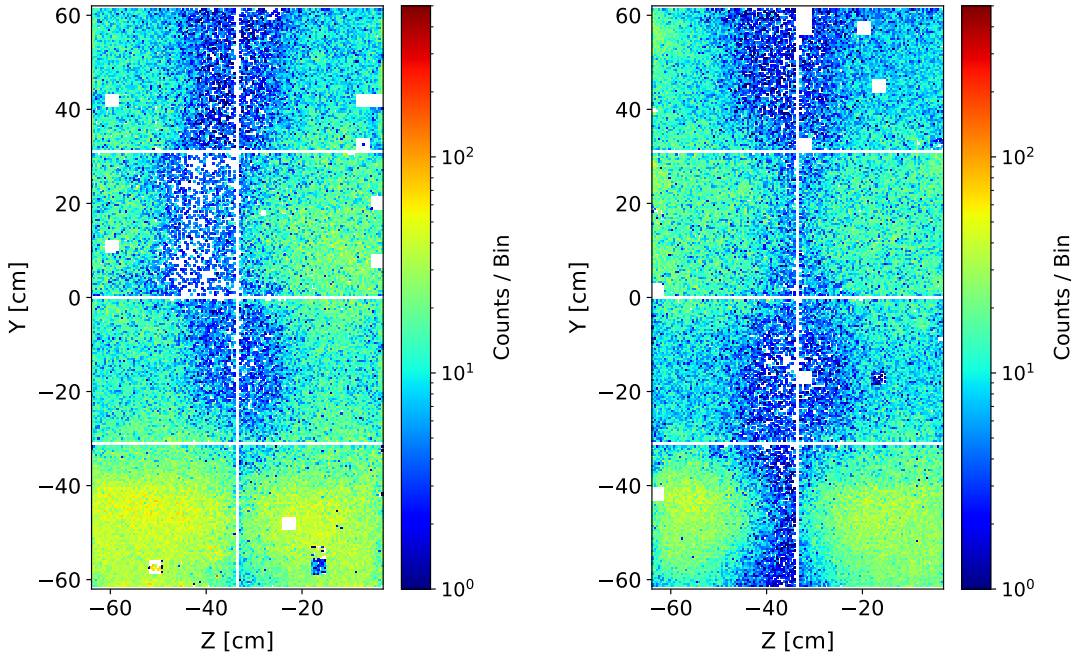

Figure 7.10: Reconstructed spatial distribution, in the 2x2 Demonstrator y–z plane, of charge-light-matched events during ^{220}Rn injection for (Left) Module-1 TPC-2 and (Right) Module-1 TPC-3 (over a 9.36 hour period).

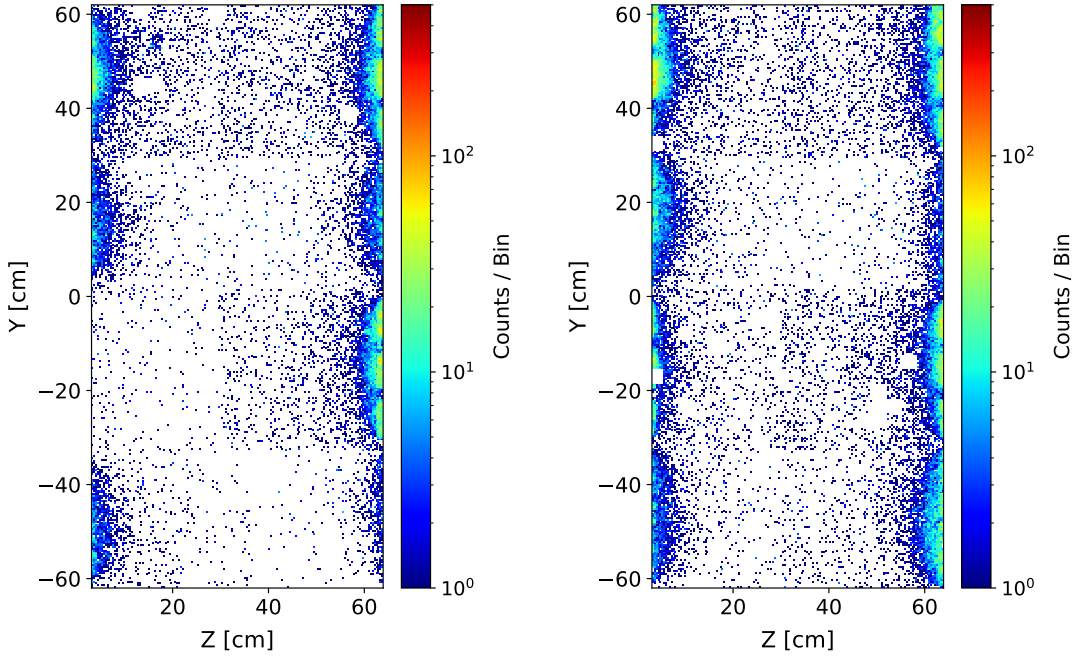

Figure 7.11: Reconstructed spatial distribution, in the 2x2 Demonstrator y–z plane, of charge-light-matched events before ^{220}Rn injection for (Left) Module-2 TPC-4 and (Right) Module-2 TPC-5 (over a 9.36 hour period).

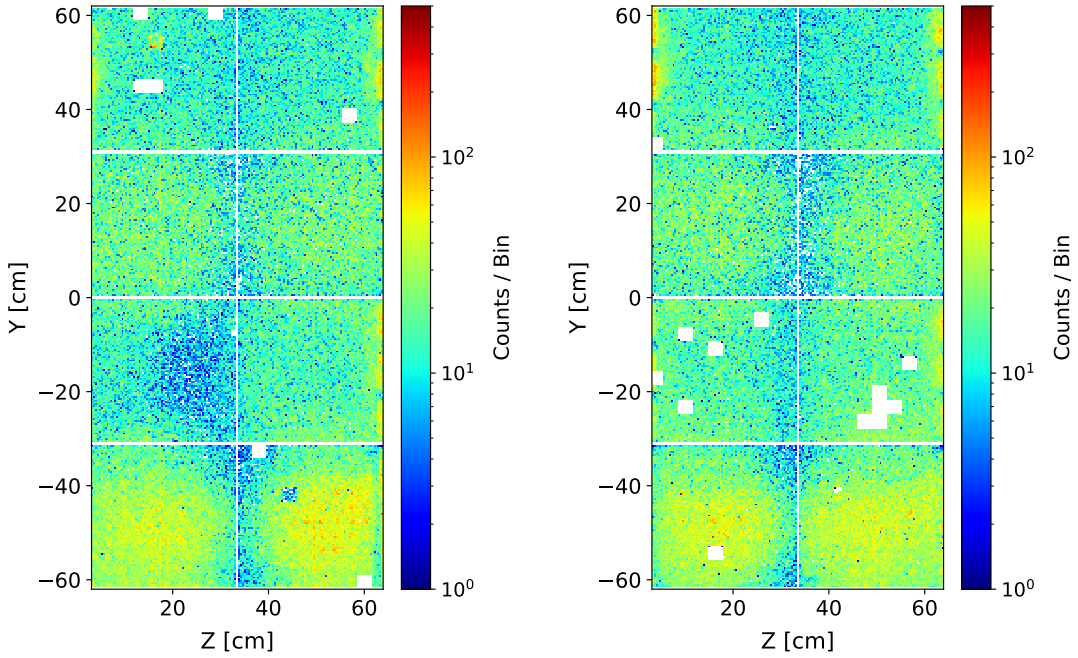

Figure 7.12: Reconstructed spatial distribution, in the 2x2 Demonstrator y–z plane, of charge-light-matched events during ^{220}Rn injection for (Left) Module-2 TPC-4 and (Right) Module-2 TPC-5 (over a 9.36 hour period).

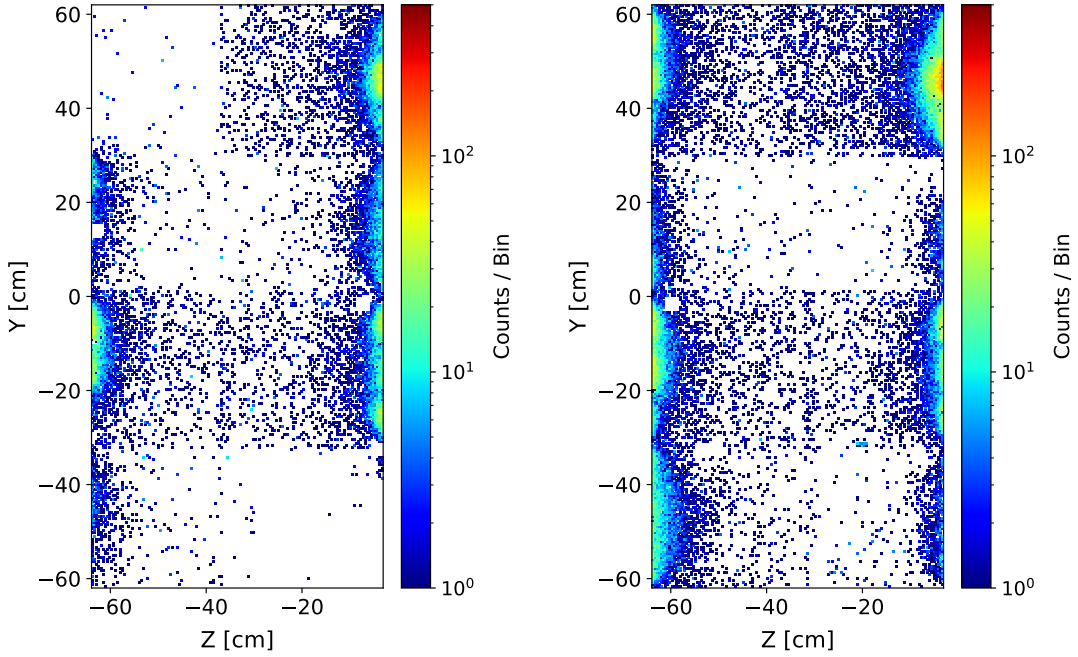

Figure 7.13: Reconstructed spatial distribution, in the 2x2 Demonstrator y–z plane, of charge-light-matched events before ^{220}Rn injection for (Left) Module-3 TPC-6 and (Right) Module-3 TPC-7 (over a 9.36 hour period).

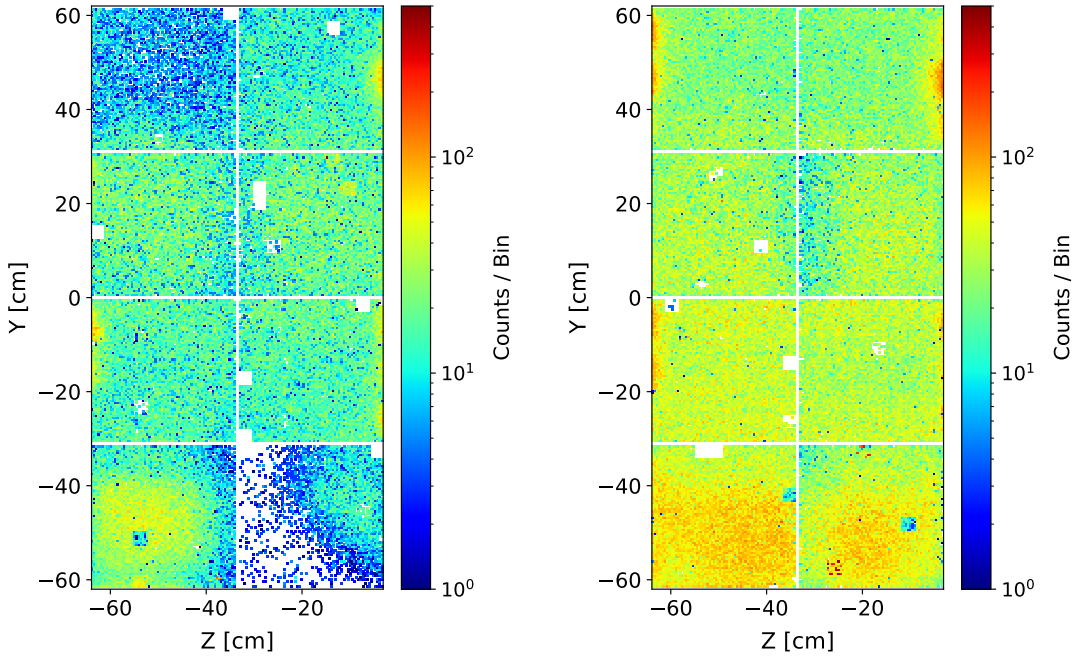

Figure 7.14: Reconstructed spatial distribution, in the 2x2 Demonstrator y–z plane, of charge-light-matched events during ^{220}Rn injection for (Left) Module-3 TPC-6 and (Right) Module-3 TPC-7 (over a 9.36 hour period).

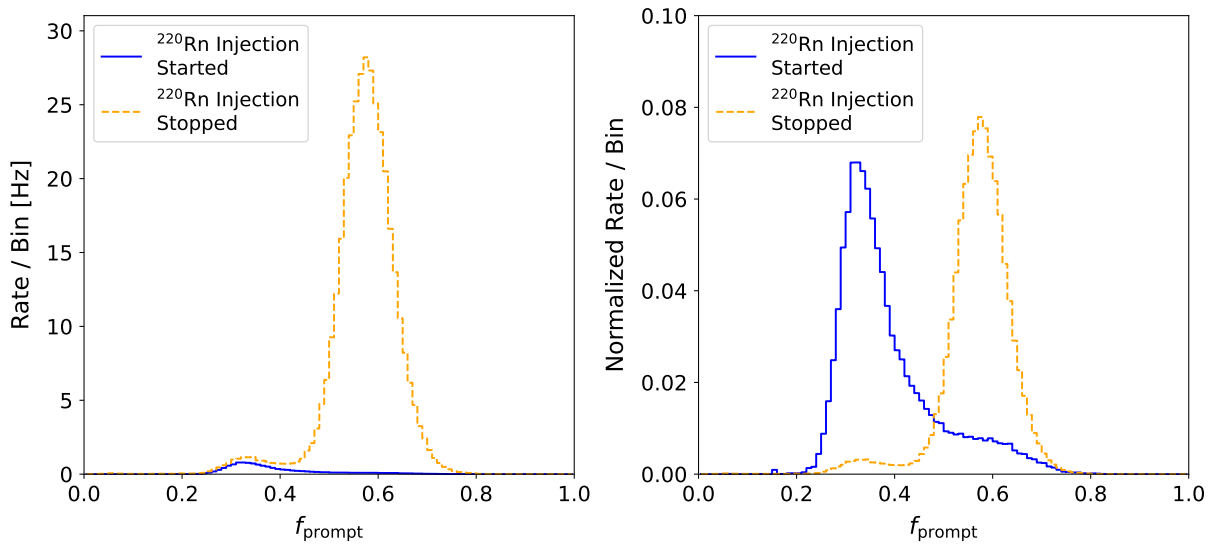

Figure 7.15: The f_{prompt} distribution for low-energy events at the start and end of the ^{220}Rn injection period. (Left) Non-normalized to show relative rate differences, (Right) Area normalized to show shape differences.

state or a long-lived triplet state, with decay times of approximately 6 ns and 1.6 µs, respectively. Because α particles produce much denser ionization tracks than electron recoils, they preferentially populate the singlet state, resulting in a larger fraction of prompt scintillation light and a reduced late-light component compared to β or γ interactions.

The difference in scintillation light waveform shape can be quantified using the f_{prompt} variable, defined in this work as the integral of the first 200 ns of the scintillation light waveform divided by the integral of the first 3200 ns of the pulse. The prompt integration window is made large enough to account for the electronics response of the 2x2 Demonstrator LRS that slows the rise of the pulse. Figure 7.15 shows the f_{prompt} distributions at the beginning and end of the ^{220}Rn injection period. Before the injection, most of the events have low-f_{prompt} (around 0.2–0.4), consistent with electron-recoil-like events. During the injection, a prominent high-f_{prompt} population appears, consistent with an increased rate of α decays in the detector. A smaller high-f_{prompt} population is also present before the injection begins, which is consistent with background events from naturally occurring α decays, such as those originating from the ^{232}Th and ^{238}U decay chains, and cosmogenic neutron scatters.

7.2.4 Pulse Shape Discrimination (PSD)

PSD exploits differences in scintillation light waveform shapes to distinguish between different classes of particle interactions. In LAr, PSD is particularly effective because particles with different ionization densities produce different singlet-to-triplet excimer ratios, leading to characteristically different time profiles of scintillation light. In the analysis presented in this work, PSD is used to discriminate between MIP-like events (e.g. β decays, cosmic-ray muons, and Compton-scattering photons from γ decays) and α decays.

A simple implementation of PSD uses the f_{prompt} variable by selecting events with low or high values to identify the different populations. However, as shown in Figure 7.15, the electron and α populations exhibit some overlap, making a single cut on f_{prompt} an imperfect discriminator. Improved separation can be achieved by fitting scintillation light waveform templates corresponding to the expected pulse shapes of different particle types and determining which template provides the best description of the observed waveform.

Figure 7.16 shows the pulse templates used in this analysis. The template for electron interactions is constructed from a high-purity sample of ^{208}Tl γ decays recorded during the pre-injection period used in the analysis described in Chapter 5. Because Compton-scattering photons produce electrons in LAr, these events provide a representative sample of the detector response to few-MeV electron recoils. Therefore, for brevity, this template is referred to as the β template for the rest of this chapter. The template is then normalized to the peak amplitude before being used in the fitting procedure. The α template is derived from the ^{220}Rn injection data by selecting double-peak events with sufficiently large peak-to-peak separation. Section 7.2.5 describes the double-peak-finding selection, which goes on to demonstrate that the resulting sample contains a high-purity sample of BiPo decays. To reduce background contributions, data from near the end of the ^{220}Rn injection period, where the LRS trigger rates were highest, are used. The individual waveforms are aligned and summed after replacing the region preceding the main pulse (which is located at approximately 2 μs) with zeros, and the resulting waveform is normalized to the peak amplitude to produce the average α pulse shape. This procedure is expected to yield a reliable α template be-

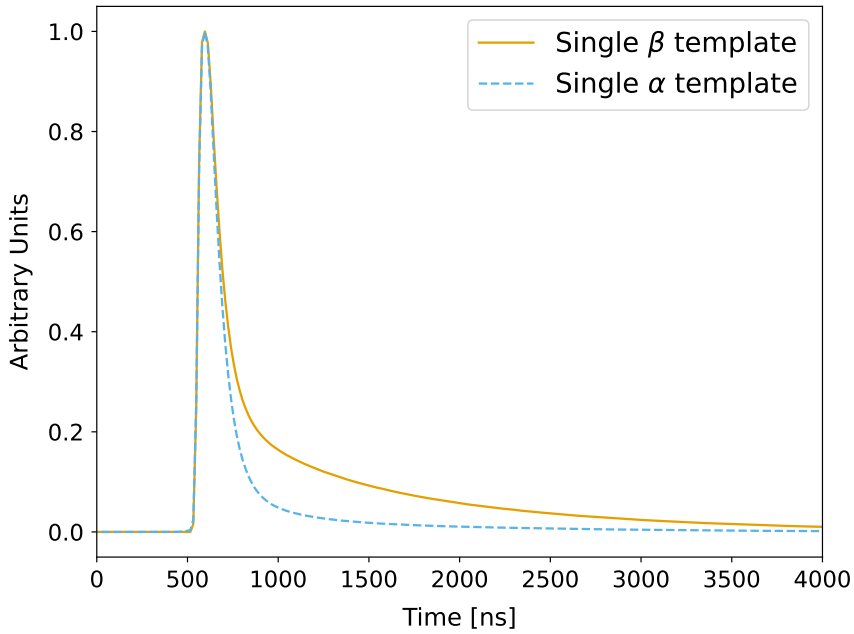

Figure 7.16: Templates of β and α scintillation light waveforms used for PSD, extracted from 2x2 Demonstrator data.

cause the overwhelming majority of events satisfying the LRS trigger conditions during the ^{220}Rn injection appear to originate from α decays, as evidenced by the prominent high-f_{prompt} population.

Candidates are classified as either β or α events by fitting the corresponding scintillation light waveform templates to the primary pulse in each waveform. The primary pulse is defined as the pulse that triggered the LRS, which occurs approximately 2 µs after the beginning of the waveform due to the configured pre-trigger window. The classification is determined by fitting both templates to the light waveform and selecting the template that yields the smallest χ^2.

BiPo candidates are identified from the subset of events classified as α candidates. An event is classified as a BiPo candidate if, in addition to the primary pulse being identified as an α, a secondary pulse is present before the primary pulse. Because the preceding β pulse is significantly smaller than the α pulse, it should have a small impact on the initial template classification, even when the two pulses partially overlap. Examples of single α, single β, and β+α (BiPo) candidates are shown in Figure 7.17.

The reconstructed drift-time distributions for each event class are shown in Figure 7.18. These distributions are used to measure the charge-light-matching selection purity using the sideband

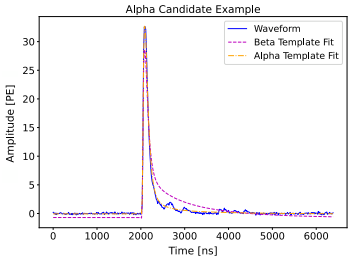

(a) Single α candidate example #1.

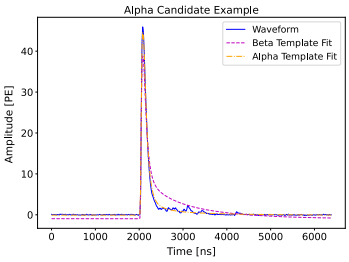

(b) Single α candidate example #2.

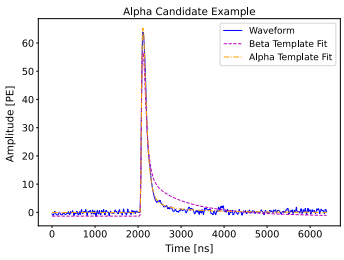

(c) Single α candidate example #3.

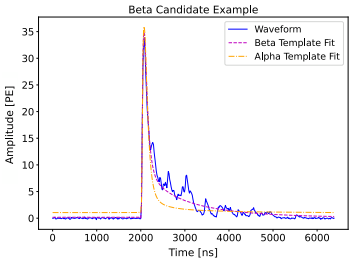

(d) Single β candidate example #1.

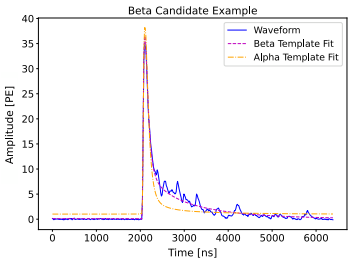

(e) Single β candidate example #2.

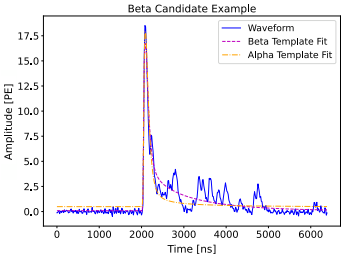

(f) Single β candidate example #3.

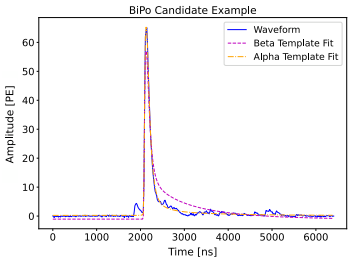

(g) β+α (BiPo) candidate example #1.

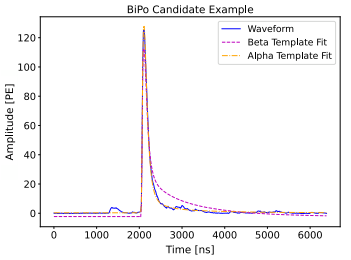

(h) β+α (BiPo) candidate example #2.

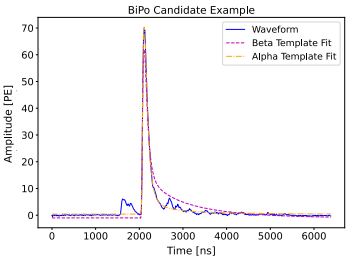

(i) β+α (BiPo) candidate example #3.

Figure 7.17: Examples of single α, single β, and β+α (BiPo) candidates found using the PSD template fit procedure applied to 2x2 Demonstrator data.

subtraction method described in Chapter 5. The resulting selection purities as a function of reconstructed charge are shown in Figure 7.19. The β candidates achieve selection purities approaching unity, whereas the α candidates exhibit substantially lower purity.

The reduced purity of the α selection is primarily due to contamination from ^{212}Bi α decays. Without requiring the characteristic two-pulse BiPo signature, the α selection contains both ^{212}Bi and ^{212}Po α decays. While the ^{212}Po α decay is preceded by a detectable β decay that produces an associated charge signal, the ^{212}Bi α decay produces too little ionization charge for the CRS to detect in most cases because of the detector thresholds. Consequently, these events contribute almost

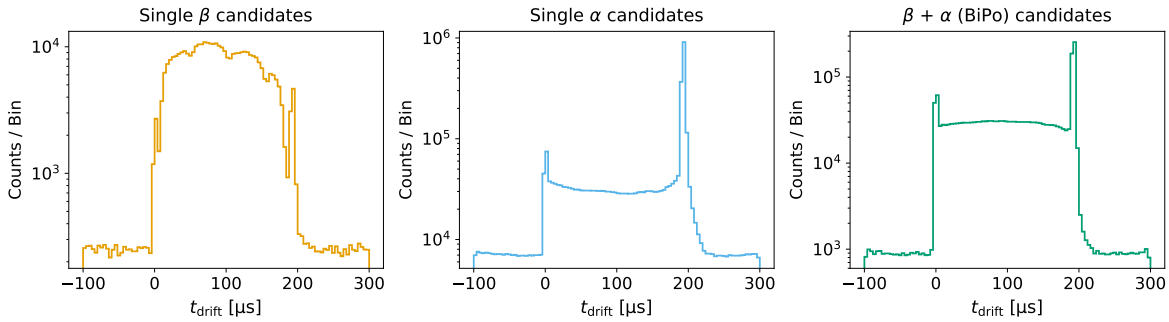

Figure 7.18: Reconstructed drift time for single β, single α, and β+α (BiPo) candidates determined using the PSD selection using 2x2 Demonstrator data.

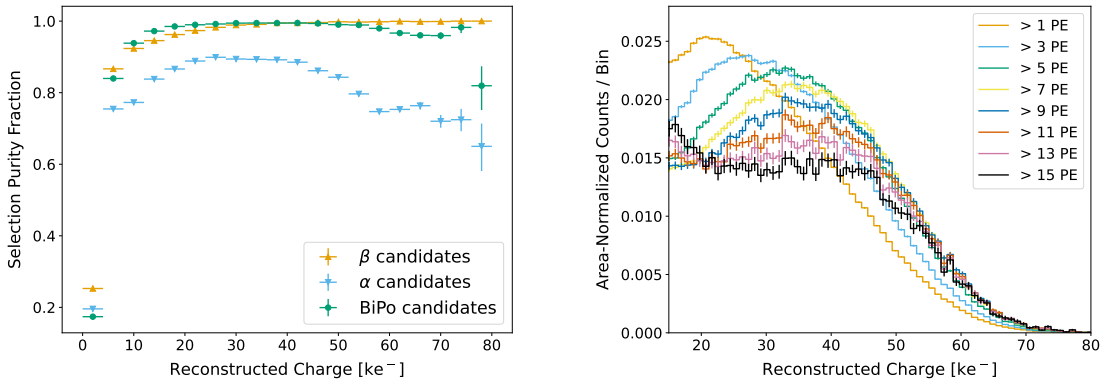

Figure 7.19: (Left) Charge-light-matching selection purity for single β candidates, single α candidates, and β+α (BiPo) candidates found using a template-based fitting algorithm to scintillation light waveforms. (Right) Charge spectrum of the β+α (BiPo) candidates as a function of the amplitude threshold on the β peak preceding the α peak (using 2x2 Demonstrator data).

exclusively accidental charge-light matches, reducing the overall selection purity. Requiring the presence of the preceding β pulse therefore provides an effective means of isolating true BiPo decays while suppressing contamination from isolated ^{212}Bi α decays. As shown in Figure 7.19, the resulting BiPo selection achieves a purity exceeding 95% for reconstructed charges above approximately 20 ke$^-$. For calibrations that require high-purity charge-light matching, such as electron lifetime measurements, it is important to restrict the analysis to the charge region where this high purity is achieved.

The f_{prompt} and reconstructed charge spectra for each event classification are shown in Figure 7.20. The α and β candidates form well-separated f_{prompt} populations, indicating that the template-fitting procedure provides good discrimination between the two event classes. As ex-

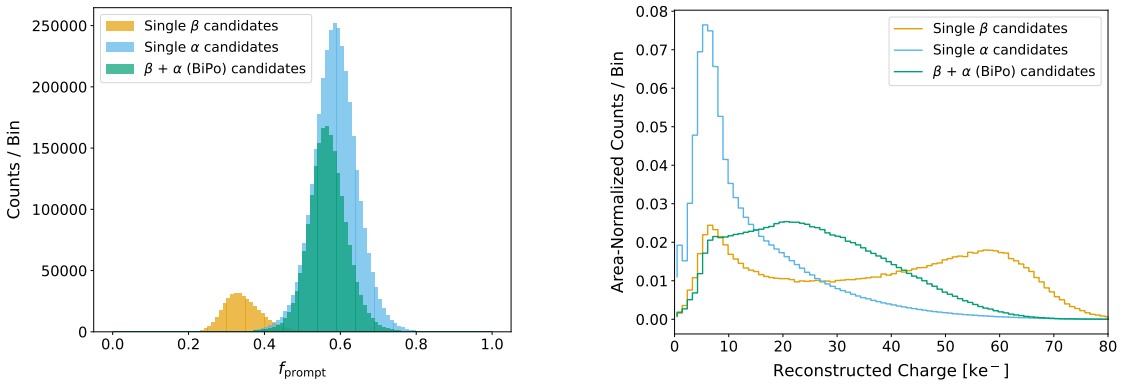

Figure 7.20: (Left) The f_{prompt} variable and (Right) reconstructed charge spectra for single β candidates, single α candidates, and β+α (BiPo) candidates found using a template-based fitting algorithm to scintillation light waveforms using 2x2 Demonstrator data.

pected, the BiPo candidates also populate the high-f_{prompt} region because they are selected from the α sample. However, their mean f_{prompt} is slightly lower than that of the full α sample. This is consistent with the presence of the preceding β pulse that partially overlaps with the α pulse, thereby reducing the measured prompt-light fraction.

The reconstructed charge spectra are also shown in Figure 7.20. The β candidate spectrum appears to closely resemble the spectrum obtained from ^{208}Tl γ decays, particularly at high reconstructed charge where the Compton edge is clearly visible. This population is expected to contain contributions from both naturally occurring ^{208}Tl and ^{208}Tl produced from the injected ^{220}Rn decay chain. In contrast, the α candidate spectrum is difficult to interpret because of its relatively low charge-light-matching selection purity. The BiPo spectrum exhibits a distinctly different shape and, in the region of high purity, is expected to be dominated by ^{212}Bi β decays.

To investigate the origin of the BiPo charge spectrum, Figure 7.19 shows the reconstructed charge spectrum for several selection thresholds on the amplitude of the preceding β pulse. The spectrum shifts toward higher reconstructed charge values as the minimum β pulse amplitude is increased. This behavior is expected because higher-energy β decays produce more scintillation light in LAr and are therefore more likely to satisfy a larger pulse amplitude requirement. This result indicates that requiring clusters with larger charge values could minimize this effect. More importantly, this implies that spatial variations on photon detection efficiency could impact the

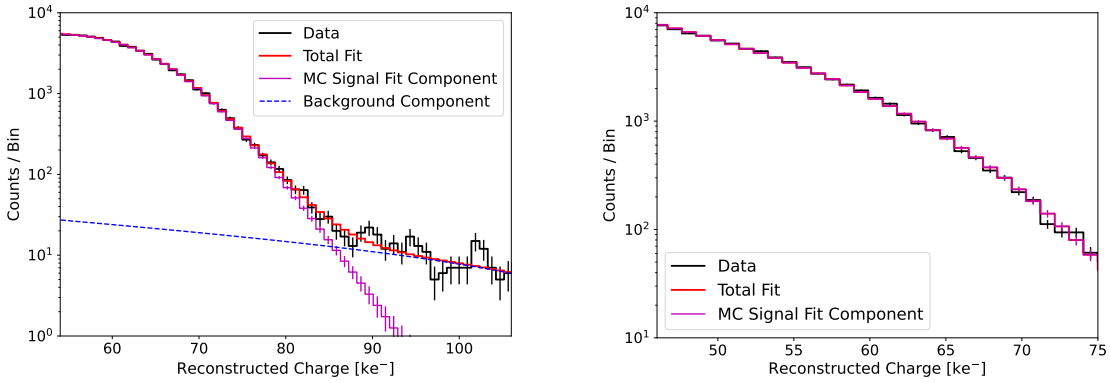

Figure 7.21: (Left) Fit to the β candidate charge spectrum in the 2x2 Demonstrator with $\chi^2/\mathrm{ndf} = 50.23/47$ and (Right) fit to the BiPo candidate charge spectrum in the 2x2 Demonstrator with $\chi^2/\mathrm{ndf} = 24.79/28$.

charge spectrum in a significant way. As a result, drift-dependent measurements, such as electron lifetime, should be restricted to higher reconstructed charge values.

As shown in Figure 7.21 (left), the corresponding charge spectrum for β candidates is fitted to the simulated ^{208}Tl 2.615 MeV γ ray Compton-scattering spectrum using the same method as Section 5.4. The fit describes the data well, indicating that the β candidates at high reconstructed charge values is dominated by ^{208}Tl γ decays. Figure 7.21 (right) shows the charge spectrum for BiPo candidates fitted with simulated ^{212}Bi β decays. The fit specifically excludes a background contribution, because cosmic-ray-induced backgrounds should be significantly suppressed by the BiPo selection procedure. This fit describes the data well, indicating that the BiPo candidate sample at high reconstructed charge values is dominated by ^{212}Bi β decays. These fitted regions are used for the electron lifetime measurements during the ^{220}Rn injection period, described in Section 7.4, because they should be minimally affected by light detection threshold and selection efficiency effects.

7.2.5 BiPo Double-Peak Identification

Identifying events containing a large α-like pulse preceded by a smaller β-like pulse provides a robust method for selecting BiPo decays in the ^{220}Rn data. This distinction is important because the inclusive α sample contains both isolated ^{212}Bi α decays and BiPo decays, with the latter providing the most useful sample for detector calibration. Because the half-life of ^{212}Po is approximately

300 ns, the β and α pulses are sufficiently separated to be individually resolved in a large fraction of BiPo decays. However, when the decay time is short, the two scintillation light pulses can partially or completely overlap on the waveform, reducing the efficiency with which BiPo events can be identified.

The finite readout window also affects the reconstruction. During the ^{220}Rn injection, the LRS was configured with a 2 µs pre-trigger window. Because the α pulse almost always triggers the LRS, any preceding β pulse occurring more than 2 µs earlier falls outside the recorded waveform. Although this has only a small impact on the overall BiPo selection efficiency, it complicates the measurement of the ^{212}Po half-life, which relies on the distribution of the time difference between the β and α pulses. In particular, a β pulse that lies outside the readout window can still contribute late scintillation light within the recorded waveform, potentially biasing the reconstructed time separation. To mitigate this effect, the selection requires that no significant signal is present before the reconstructed β pulse. The BiPo selection is applied to the sample of α candidates identified by the PSD template classification. Events are required to contain a pulse within the pre-trigger window in addition to the primary α pulse. This pulse is identified as the largest pulse in the pre-trigger region, and no earlier pulses are permitted on the waveform (otherwise the event is thrown out). These requirements suppress accidental activity while selecting the characteristic two-pulse topology expected from BiPo decays.

The timing of each pulse is determined by fitting its rising edge with a sigmoid function,

$$S(t) = \frac{a}{1 + e^{-\frac{(t-t_0)}{k}}} + b, \tag{7.1}$$

where a is the pulse amplitude, t_0 is the inflection point of the sigmoid (corresponding to the 50% rise time), k controls the steepness of the rising edge, and b is the vertical offset before the pulse. The fitted value of t_0 is taken as the pulse time. Figure 7.22 shows examples of the sigmoid timing fits applied to BiPo candidates in the 2x2 Demonstrator data.

The b parameter in the sigmoid fit accounts for the elevated waveform baseline that can result when the preceding β pulse partially overlaps the α pulse. Rather than defining the pulse time by

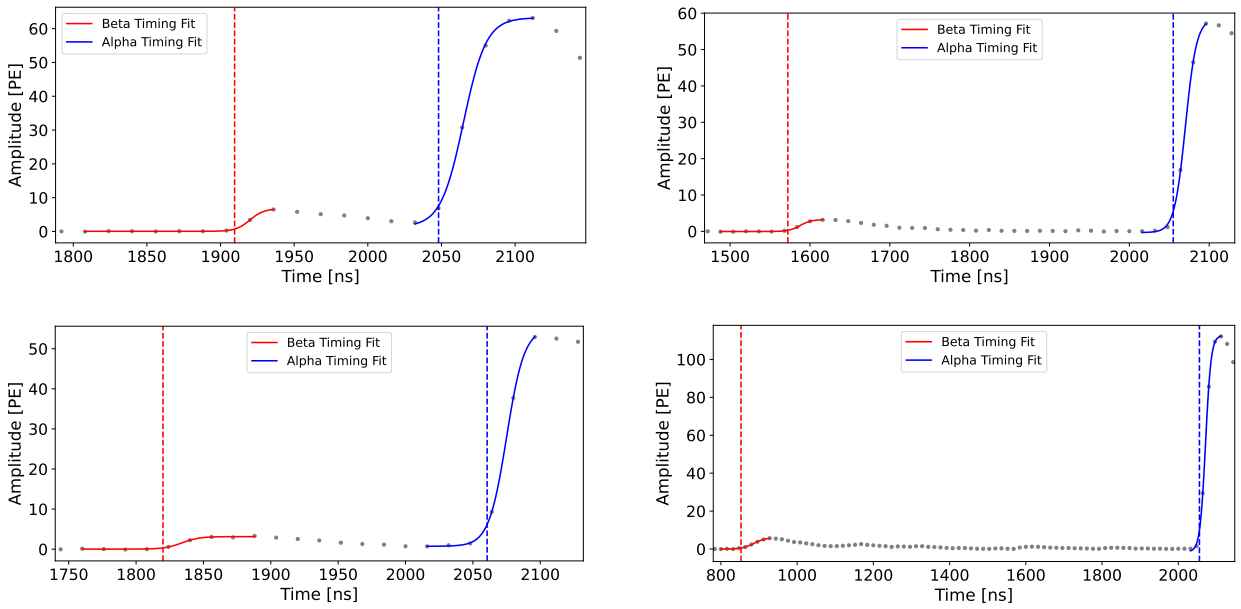

Figure 7.22: Examples of measuring the peak timing for BiPo candidates in the 2x2 Demonstrator using a sigmoid function.

the sigmoid inflection point, the timing of each pulse is taken to be the point at which the fitted sigmoid reaches 10% of its maximum amplitude. This time is given by the following formula:

$$t_{10} = t_0 + k \ln\left(\frac{0.1}{0.9}\right). \tag{7.2}$$

To suppress spurious peaks arising from electronic noise or baseline fluctuations, the fitted β pulse is required to satisfy a goodness-of-fit criterion based on the chi-squared distance,

$$\chi^2 = \frac{1}{2}\sum_i \frac{(S_i - W_i)^2}{S_i + W_i}, \tag{7.3}$$

where S_i and W_i denote the fitted sigmoid and measured waveform values, respectively. Events are required to satisfy $\chi^2 < 1$ for the β pulse. No corresponding requirement is applied to the α pulse because its substantially larger amplitude makes it readily distinguishable from noise and β decays. The resulting χ^2 distributions for the β and α pulses are shown in Figure 7.23.

The measured half-life, $t_{1/2}$, is determined by fitting the distribution of the reconstructed time difference, Δt, between the β and α pulses. The fit function is:

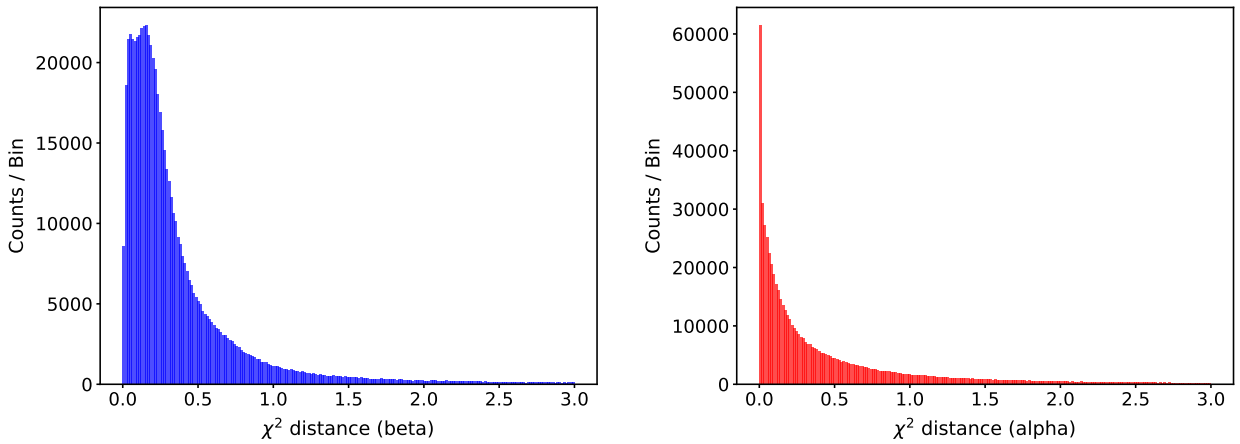

Figure 7.23: The χ^2 distance calculated between the fitted sigmoid and the waveform peak for both the β and α peaks found in BiPo candidate waveforms.

$$f(\Delta t) = A \cdot 2^{-\Delta t / t_{1/2}} + C, \tag{7.4}$$

where A is the normalization and C accounts for a constant background. The fit is performed only over a selected range of Δt to avoid regions where reconstruction biases may affect the measurement.

To evaluate the stability of the extracted half-life, scans are performed over both the lower and upper boundaries of the fit range, as shown in Figure 7.24 (top). The measured half-life is stable over a broad range of fit boundary values. The primary exception occurs when the lower fit boundary is placed at small Δt, where significant overlap between the β and α pulses degrades the timing reconstruction and biases the fitted half-life. The lower bound adopted for the nominal measurement, indicated by the vertical dashed line in Figure 7.24 (top left), is chosen to avoid this biased region while retaining sufficient statistical precision for the measurement. The upper fit boundary is selected to be as large as possible while remaining sufficiently below the 2 µs pre-trigger window. This minimizes the loss of events in which the preceding β pulse extends beyond the beginning of the recorded waveform, which would otherwise bias the reconstructed Δt distribution. The systematic uncertainty associated with the choice of fit range is estimated by varying both the lower and upper fit boundaries by $\pm$50 ns and taking the resulting variation in the fitted half-life as the corresponding systematic uncertainty.

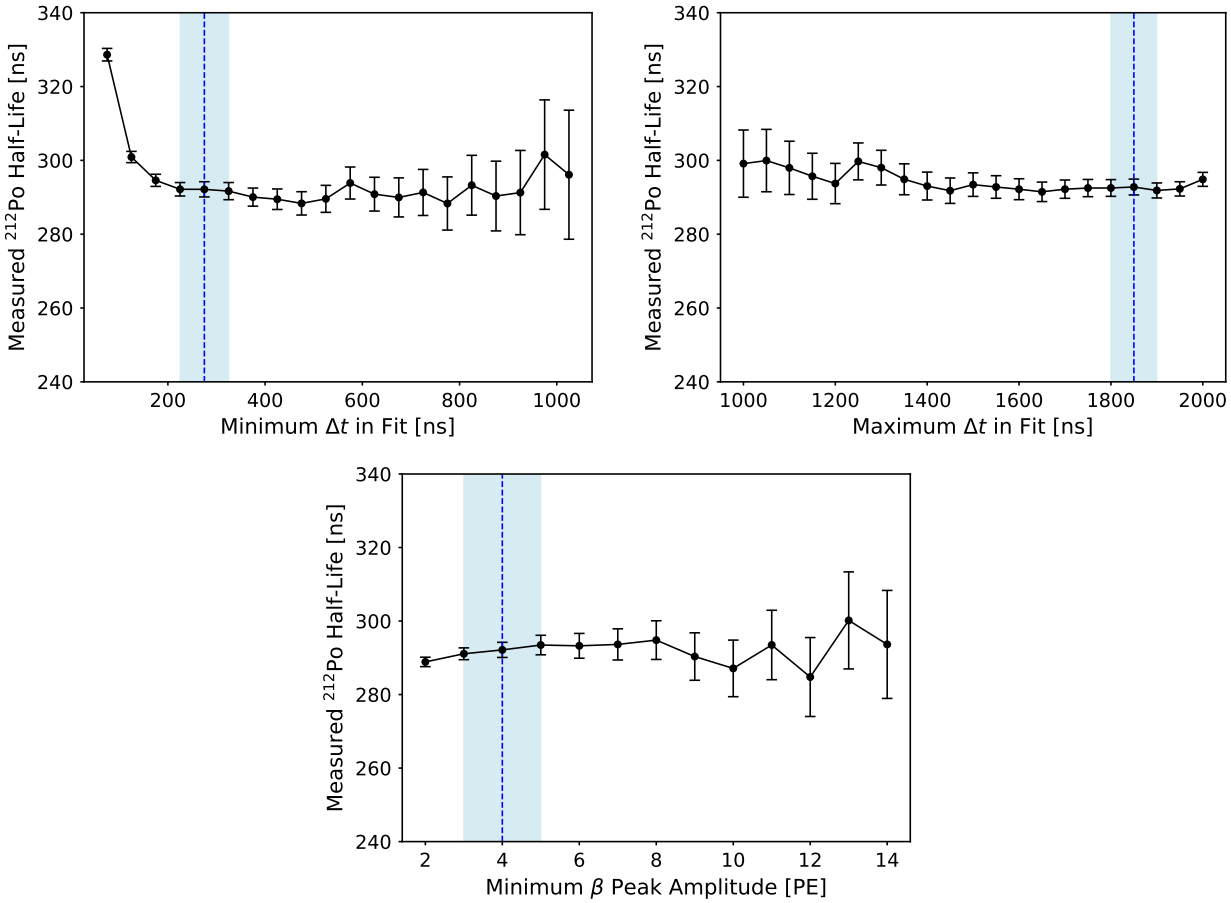

Figure 7.24: (Top Left) Scan of the lower bound on the fit range, (Top Right) scan of the upper bound on the fit range, and (Bottom) scan of the minimum β peak amplitude allowed in the fit, for the measured ^{212}Po half-life from BiPo candidates in the 2x2 Demonstrator. The vertical line indicates the chosen value for the measurement and the shaded region shows the range over which this parameter was varied to evaluate the contribution to the systematic uncertainty.

An additional source of systematic uncertainty is associated with the minimum amplitude required for the reconstructed β pulse. The dependence of the measured half-life on this requirement is shown in Figure 7.24 (bottom). The extracted half-life is stable over a broad range of minimum β pulse amplitudes, although a small bias toward shorter half-lives is observed for the lowest thresholds. This behavior is likely caused by the poorer signal-to-noise ratio of small β pulses, which degrades the timing reconstruction. The nominal minimum β pulse amplitude used in the analysis is 4 PE, and the corresponding systematic uncertainty is evaluated by varying this threshold by ± 1 PE. This variation accounts for any residual bias arising from the timing reconstruction of low-amplitude β pulses.

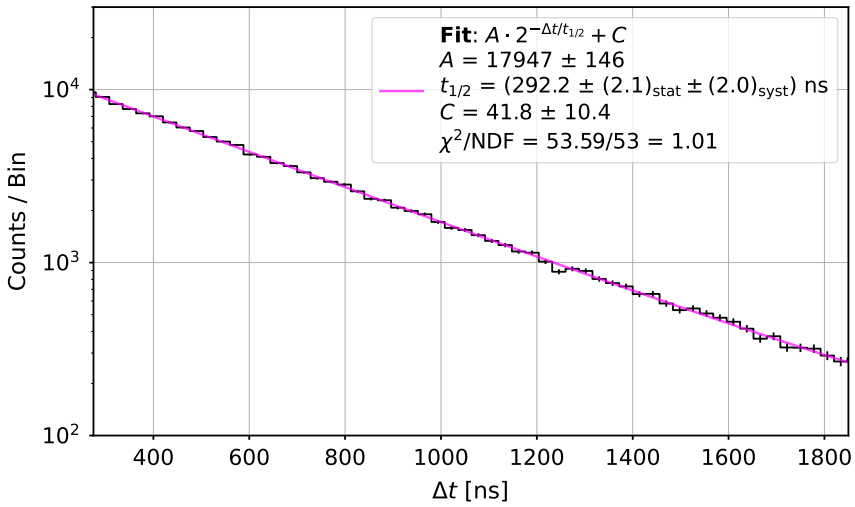

Figure 7.25: Distribution of Δt from BiPo candidate waveforms in the 2x2 Demonstrator and corresponding fit to measure ^{212}Po half life.

The reconstructed Δt distribution and the corresponding fit obtained using the nominal selection criteria ($\Delta t_{\text{min}} = 275$ ns, $\Delta t_{\text{max}} = 1850$ ns, and a minimum β pulse amplitude of 4 PE) are shown in Figure 7.25. Combining the statistical and systematic uncertainties in quadrature yields a measured half-life of

$$t_{1/2} = 292.2 \pm 2.9 \text{ ns}. \tag{7.5}$$

This result is compared with previous measurements in Figure 7.26. The measured half-life is consistent with the established values, providing strong evidence that the selected events originate from genuine BiPo decays in the ^{220}Rn decay chain. Previous measurements have been performed in liquid xenon dual-phase and scintillator-based detectors, where β decays produce much larger scintillation light signals than in the light readout systems of typical LArTPC neutrino detectors. Additionally, requiring a visible β pulse in the 2x2 Demonstrator substantially reduces the number of usable BiPo candidates, because of the relatively low light yield from β decays as opposed to the relatively large amount of scintillation light from α decays. This reduced efficiency is an inherent limitation of measurements in large LArTPCs that rely on resolving the β pulse from BiPo decays.

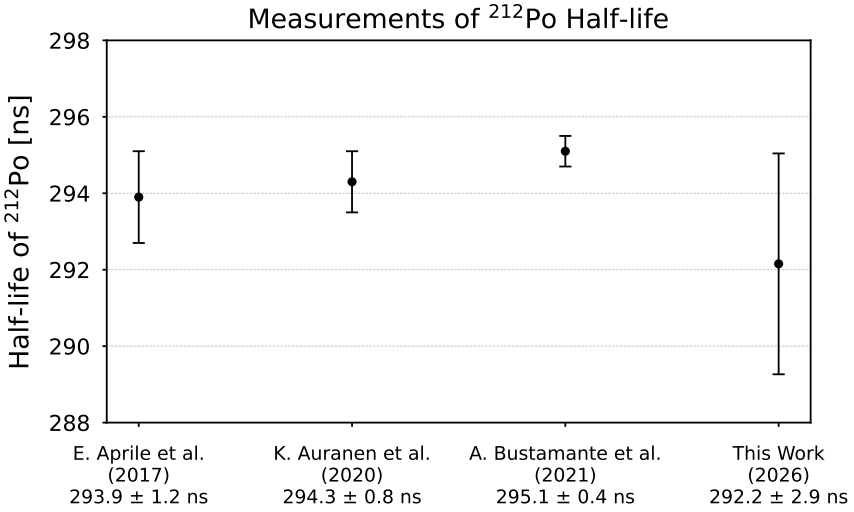

Figure 7.26: The ^{212}Po half-life measurement using 2x2 Demonstrator data compared to measurements produced by other experiments.

7.3 Drift Speed and Drift Field Measurements

7.3.1 ^{220}Rn Daughter Cathode Plate-Out

Previous studies, such as the XENON100 ^{220}Rn calibration [113], have shown that daughters of the ^{220}Rn decay chain can plate out onto detector electrodes. This behavior is expected because several α decays in the decay chain produce daughter isotopes that are initially ionized. Following an α decay, the daughter nucleus retains much of the parent atom's electron cloud and is therefore often negatively charged. However, the large energy released in the decay can also remove atomic electrons, leaving the daughter neutral or positively charged. Before these ions become neutralized, they can drift in the TPC electric field toward either the cathode or anode, depending on their charge. If they reach an electrode before neutralization, they may become attached to its surface and subsequently decay there, producing localized excess activity.

A similar plate-out process would occur in the 2x2 Demonstrator if a sufficient fraction of the ^{220}Rn daughters enters the TPC active volume. As shown previously, a significant excess of reconstructed activity is observed near the cathode (with a smaller excess near the anode). Approximately one-third of all LRS triggers originate from activity reconstructed near the cathode. In addition, the cathode activity is more concentrated toward the bottom of the TPC active volume,

which can be seen in the distributions projected along the y–z plane (Figures 7.10, 7.12, and 7.14). This observation can be explained by ^{220}Rn daughters entering the TPC active volume primarily through the openings in the lower field cage panel, which allows LAr to circulate into and out of the TPC. If ions enter near the bottom of the detector, they may drift toward the cathode before becoming uniformly mixed throughout the LAr volume. Such a process would naturally produce a higher density of plated-out activity near the lower portion of the cathode, qualitatively consistent with the observed spatial distributions.

7.3.2 Drift Speed Measurement Methodology

The concentration of ^{220}Rn daughter activity on the cathode provides a unique calibration sample for measuring the ionization electron drift speed. Assuming that these decays originate on the cathode surface, their position along the drift direction is known a priori. The measured drift time of the associated ionization therefore directly constrains the electron drift speed over the full drift distance of the TPC. Furthermore, if the plate-out is distributed across the cathode surface, the same technique can be extended to construct a two-dimensional map of the drift velocity, providing a probe of spatial variations in the electric field.

Figure 7.27 shows the reconstructed drift-time distribution for events associated with the cathode peak in one TPC. The TPC electronics response and longitudinal diffusion contribute to the width of the peak near the cathode. Furthermore, the distribution exhibits a pronounced tail extending several microseconds beyond the peak, likely caused by BiPo decays as the delayed ^{212}Po α decay can occur up to several microseconds after the preceding ^{212}Bi β decay. Depending on the particular low-energy event, the associated ionization may either extend the temporal width of the reconstructed charge cluster or be reconstructed as a nearby secondary charge packet, producing a non-Gaussian tail in the cathode timing distribution. To suppress this effect, events are required to have a reconstructed cluster temporal extent smaller than 2 µs. In addition, a requirement of reconstructed charge in the cluster greater than 12 ke$^-$ is applied to reject clusters dominated by isolated α interactions. Together, these selections remove most of the asymmetric tail while preserving the

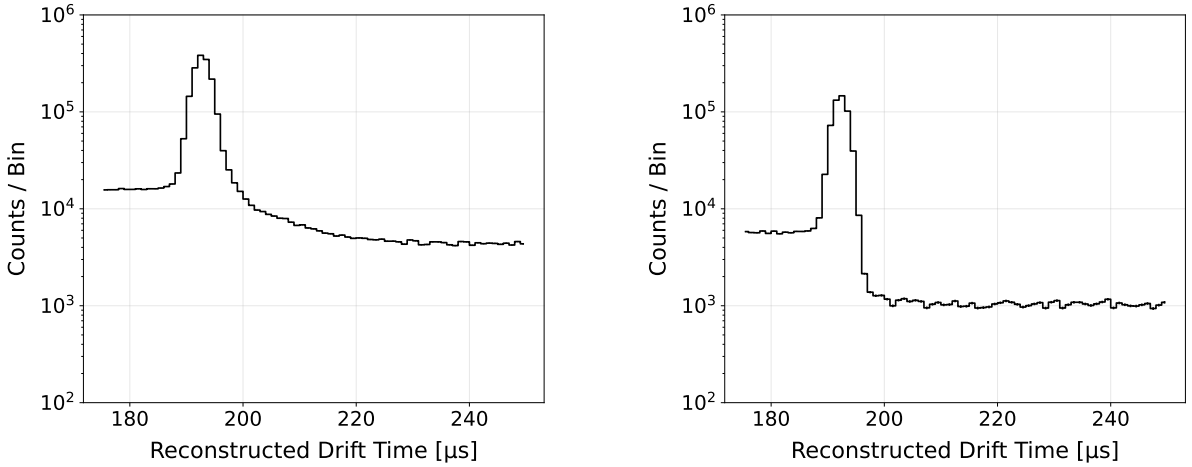

Figure 7.27: The reconstructed drift time of near-cathode activity during ^{220}Rn injection in the 2x2 Demonstrator for Module-3 TPC-7. (Left) No additional selection requirements. (Right) Selection requirements of drift time < 2 μs and Q > 12 ke$^-$, where Q is the reconstructed charge for low-energy clusters.

majority of the available statistics, resulting in the approximately Gaussian cathode-peak component shown in Figure 7.27. This simplified peak shape enables a more robust determination of the mean drift time between cathode and anode.

A model is constructed to extract the average drift time associated with cathode activity. The observed drift-time distribution consists of two primary components: a localized peak corresponding to activity originating on the cathode surface, and a more diffuse contribution from activity occurring within the TPC active volume. Both components are subject to broadening from ionization electron diffusion and the TPC electronics response, which are assumed to be well described by a common Gaussian smearing with standard deviation σ.

The bulk activity component transitions smoothly to the combinatorial background at the cathode drift time, t_c, which is modeled using a smeared step function. This contribution is combined with a Gaussian peak centered at t_c to describe the cathode-localized activity. The resulting fit function is:

$$\begin{aligned} f(t) = & A_{\text{peak}} \exp\left(-\frac{(t-t_c)^2}{2\sigma^2}\right) \\ & + \frac{A_{\text{bulk}} + m(t-t_c)}{2}\left[1 + \text{erf}\left(\frac{t_c - t}{\sqrt{2}\sigma}\right)\right] \\ & + B, \end{aligned} \tag{7.6}$$

where A_{peak} describes the amplitude of the cathode-localized peak, A_{bulk} and m parameterize the approximately linear bulk activity component, and B accounts for a constant combinatorial background.

The model is simultaneously sensitive to the cathode peak width and the transition of the bulk component to the background component, both of which share the same smearing parameter σ. The choice of using a shared σ is appropriate because both effects are dominated by the same longitudinal diffusion and TPC electronics response effects. Example fits for each TPC are shown in Figure 7.28. The primary parameter of interest is t_c, which provides a direct measurement of the mean drift time between cathode and anode.

7.3.3 Drift Speed Measurements by TPC

The reconstructed drift speed as a function of y and z position for each TPC is shown in Figures 7.29–7.32. To obtain these measurements, low-energy charge clusters matched to scintillation light signals are grouped into 3 cm spatial bins in both coordinates. In each bin, the cathode drift time is extracted using the fit procedure described in the previous section, provided that the bin contains at least 150 events to ensure sufficient statistical precision. The local drift speed is then computed as the ratio of the known drift distance, 30.272 cm, to the fitted cathode drift time, t_c. The resulting maps show spatial variations that are most pronounced near TPC and pixel-tile boundaries. In particular, an enhancement in the reconstructed drift speed is observed near the top-center region of the TPCs, which is consistent with the location of the high-voltage feedthrough and suggests local electric-field distortions in its vicinity.

In addition, the percent change of either the drift speed or the E-field magnitude can be calculated for different locations. The dependence of ionization electron drift speed on E-field magnitude established by MicroBooNE [120] is used to translate the drift speed measurements to an estimation of the local E-field magnitude in a given TPC; note that the small temperature dependence between MicroBooNE and the 2x2 Demonstrator will not impact the relative measurement that is discussed below. The ionization electron drift speed and E-field magnitude appear lower at

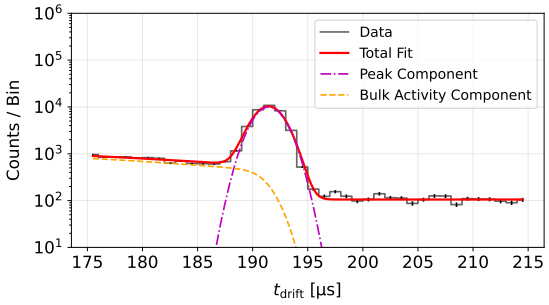

(a) Near-cathode activity in Module-0 TPC-0.

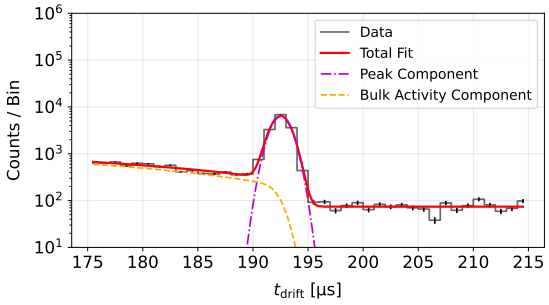

(b) Near-cathode activity in Module-0 TPC-1.

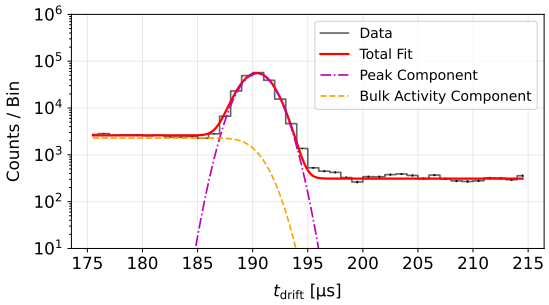

(c) Near-cathode activity in Module-1 TPC-2.

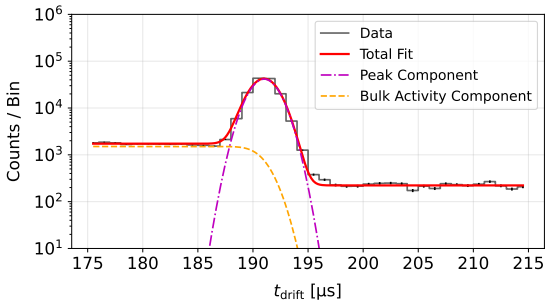

(d) Near-cathode activity in Module-1 TPC-3.

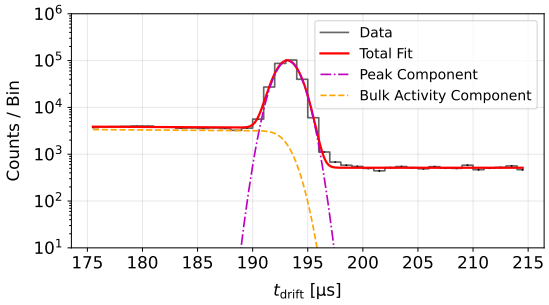

(e) Near-cathode activity in Module-2 TPC-4.

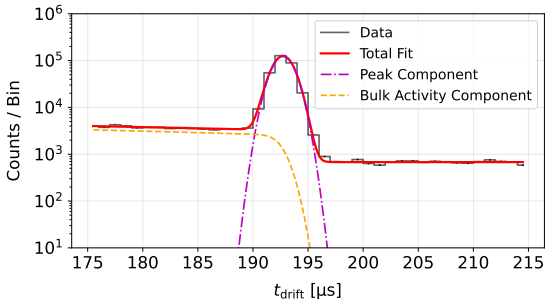

(f) Near-cathode activity in Module-2 TPC-5.

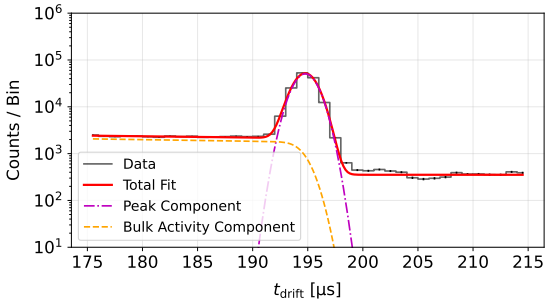

(g) Near-cathode activity in Module-3 TPC-6.

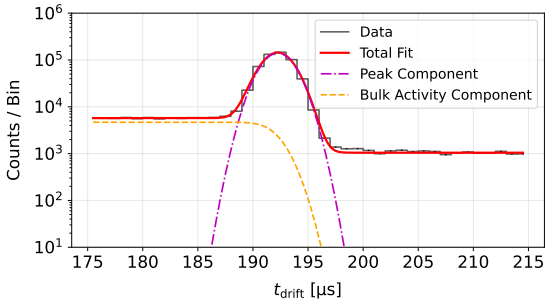

(h) Near-cathode activity in Module-3 TPC-7.

Figure 7.28: Reconstructed ionization electron drift time for activity near the expected cathode-to-anode drift time. The fit includes a bulk activity component, a Gaussian component to capture the peak time, and a combinatorial background component.

pixel tile boundaries by at least -1% and -2%, respectively. The largest variation by far is at the top center of each TPC, with up to $+3\%$ and $+7\%$, for variations in ionization electron drift speed and E-field magnitude, respectively (not visible with the restricted color scales). These larger variations near the top of the TPCs are likely to be associated with the high-voltage feedthroughs at the same locations. Overall, these results demonstrate that cathode plateout activity introduced by ^{220}Rn injection can be used to assess variations of average drift speed across different drift volumes, temperature-dependence of drift speed, and spatial variations of ionization electron drift speed and E-field magnitude, presenting a viable alternative to cosmic-ray muons for the DUNE LArTPCs.

To compare drift speeds between TPCs, regions susceptible to boundary effects are excluded. This is necessary because electric-field distortions near detector edges can bias the cathode timing. Events are therefore required to be more than 15 pixels away from pixel-tile boundaries (17 pixels for Module-2 due to smaller pixel pitch). The selected events are rebinned over this restricted area and the drift speed in each bin is evaluated using the same fitting procedure. For each TPC, the final drift speed is obtained from the distribution of per-bin measurements using the median value and central 68% confidence interval. The resulting per-bin distributions are shown in Figure 7.33, while the TPC-level measurements are summarized in Figure 7.34. Some variation is observed between modules and between TPCs within the same module. In general, the two TPCs within a given module show similar drift speed, with the exception of Module-3 where has a larger discrepancy.

A difference in field cage resistances provides a plausible explanation for module-to-module variations, as they can lead to small differences in the effective electric-field strength within each module by impacting the cathode bias voltage. However, such effects alone do not explain the TPC-to-TPC difference observed in Module-3, given that both TPCs share a common cathode high-voltage supply. The remaining possibilities are (i) residual measurement bias, (ii) localized electric-field distortions that differ between the two TPCs, or (iii) small differences in the effective drift distance. An independent drift speed measurement using cosmic-ray muons would help disambiguate between these possibilities.

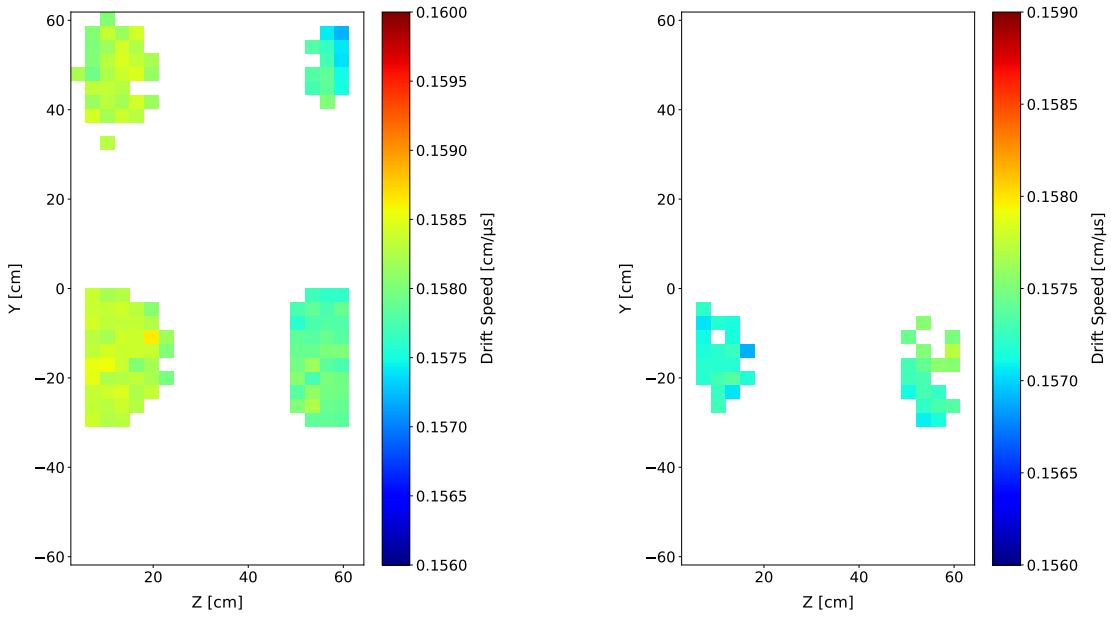

Figure 7.29: (Left) Measurements of ionization electron drift speed across Module-0 TPC-0 in the 2x2 Demonstrator. (Right) Measurements of drift speed across Module-0 TPC-1. White areas are regions with too few events ($<$ 150) to perform a measurement

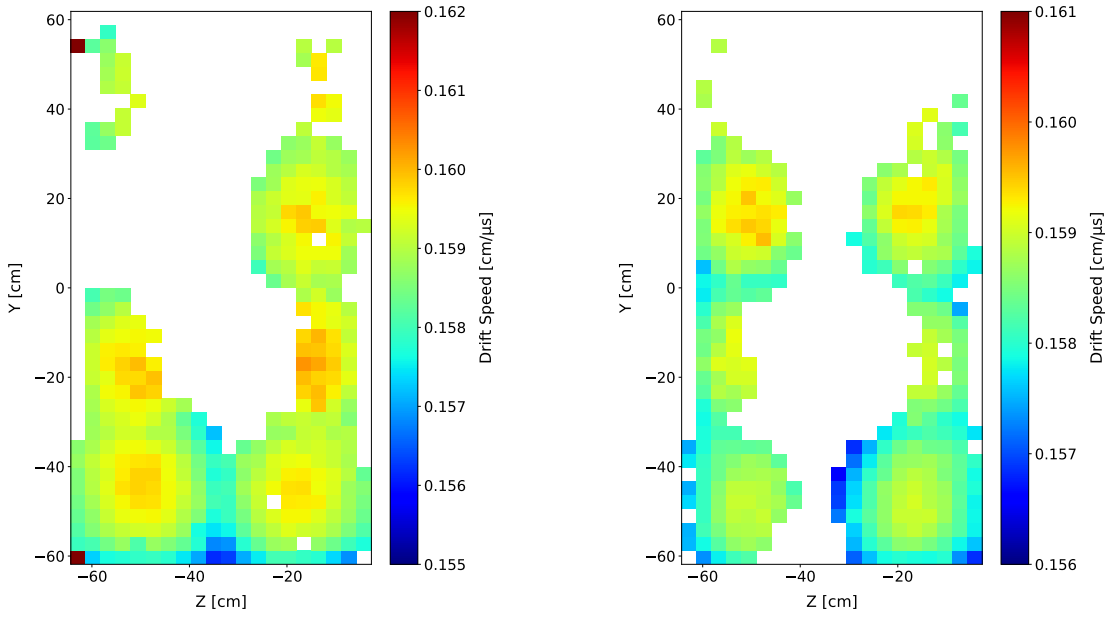

Figure 7.30: (Left) Measurements of ionization electron drift speed across Module-1 TPC-2 in the 2x2 Demonstrator. (Right) Measurements of drift speed across Module-1 TPC-3. White areas are regions with too few events ($<$ 150) to perform a measurement

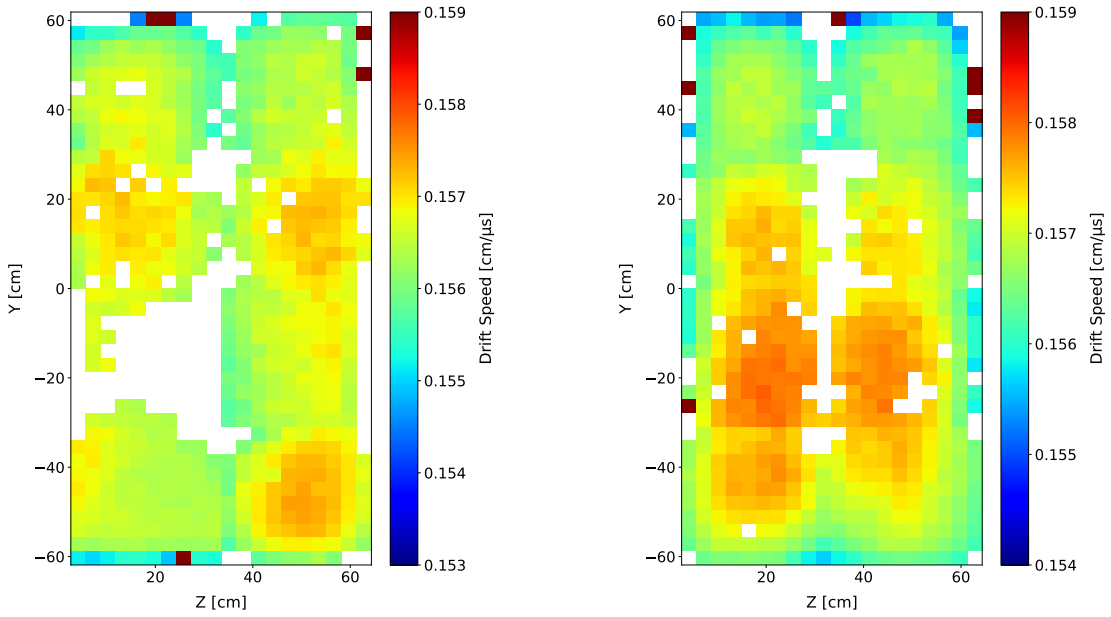

Figure 7.31: (Left) Measurements of ionization electron drift speed across Module-2 TPC-4 in the 2x2 Demonstrator. (Right) Measurements of drift speed across Module-2 TPC-5. White areas are regions with too few events ($<$ 150) to perform a measurement

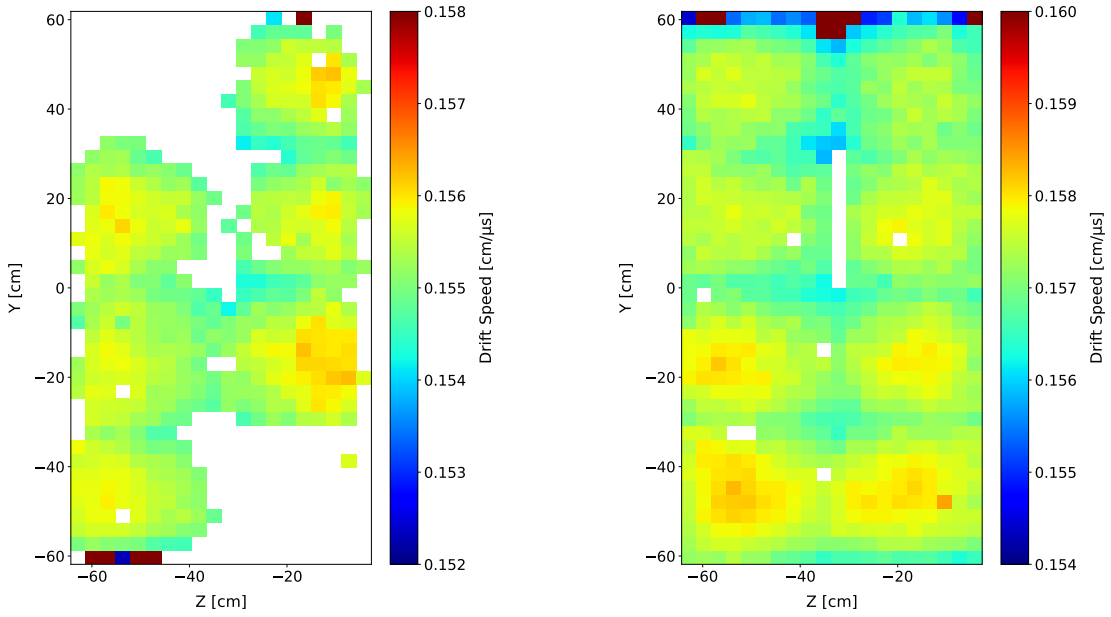

Figure 7.32: (Left) Measurements of ionization electron drift speed across Module-3 TPC-6 in the 2x2 Demonstrator. (Right) Measurements of ionization electron drift speed across Module-3 TPC-7. White areas are regions with too few events ($<$ 150) to perform a measurement.

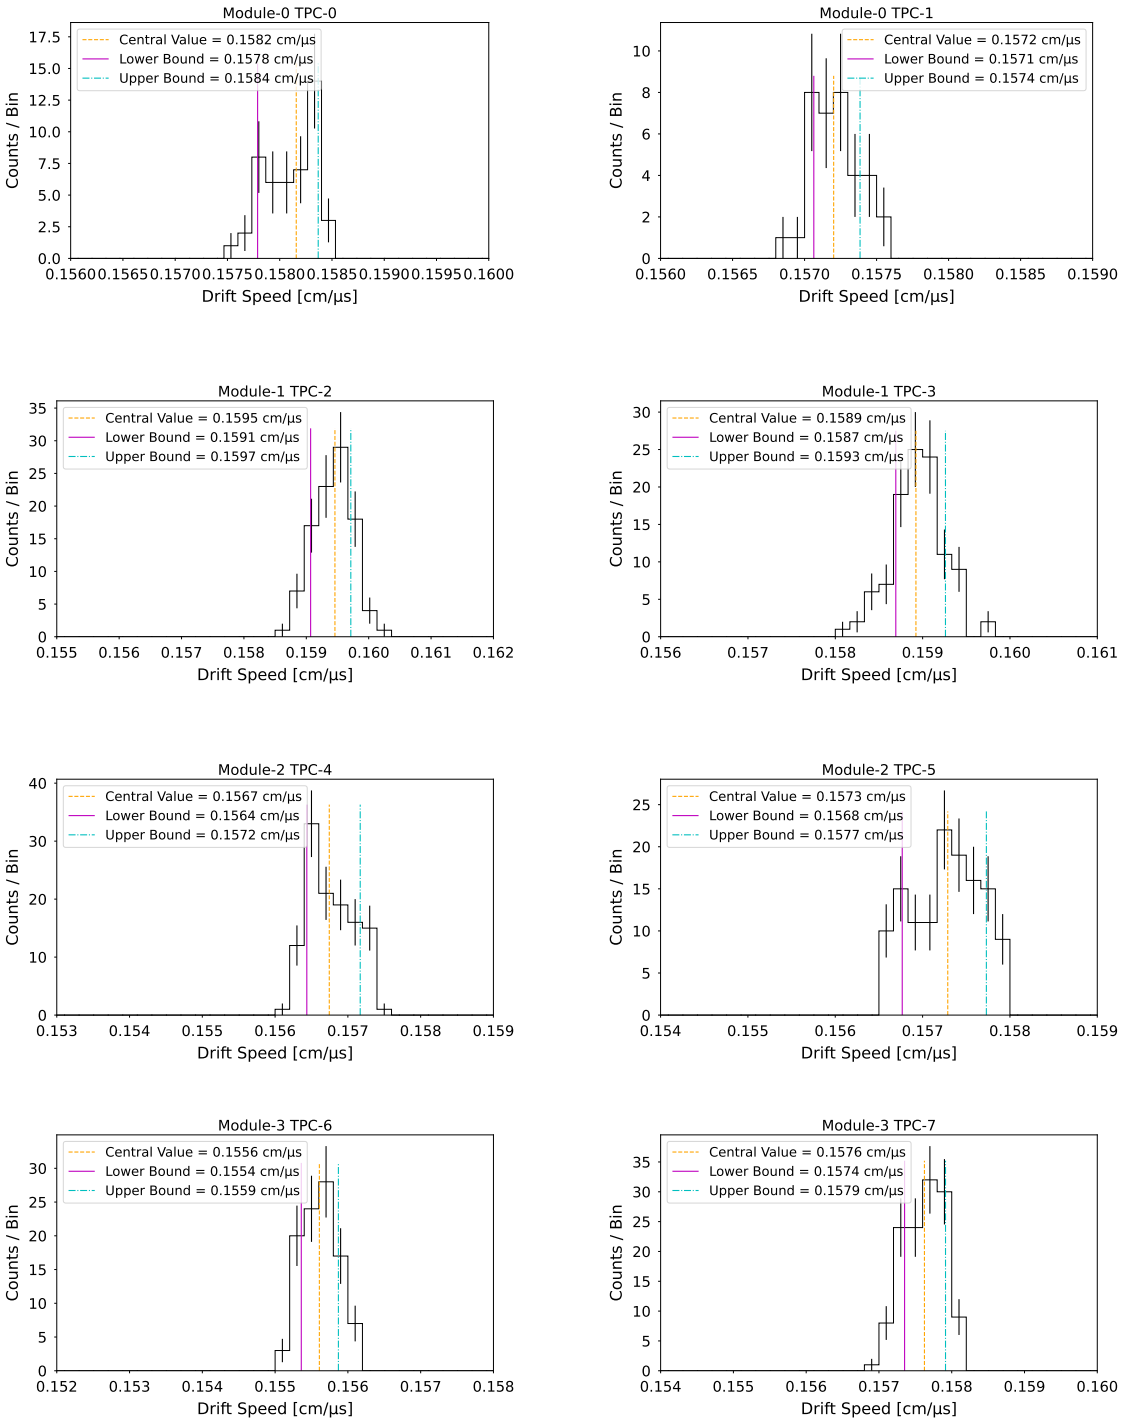

Figure 7.33: Individual ionization electron drift speed measurements in each TPC of the 2x2 Demonstrator using events more than 15 pixels (17 for Module-2) from the pixel tile edges.

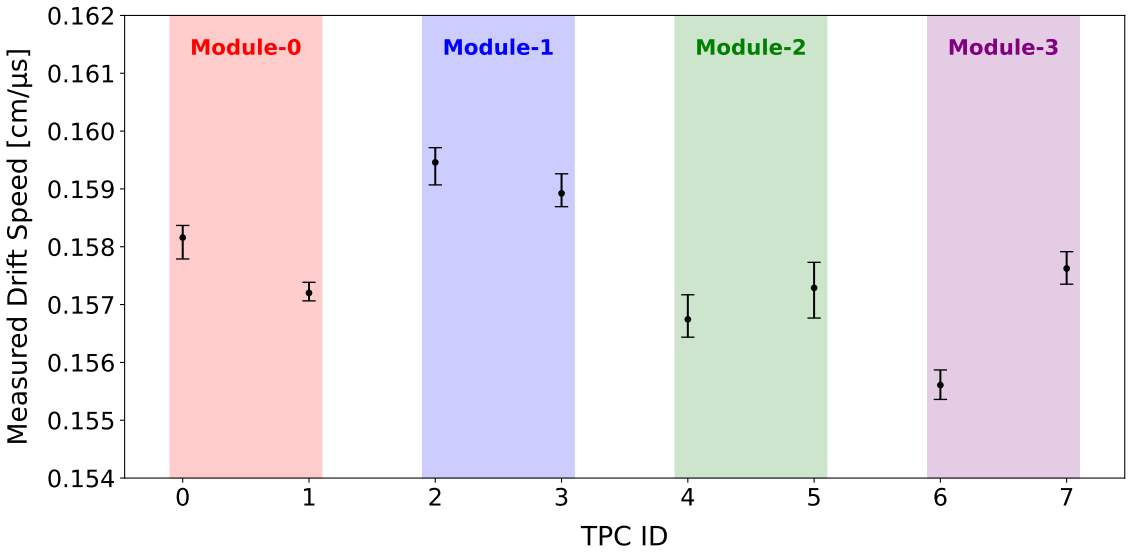

Figure 7.34: Ionization electron drift speed measurement in each TPC of the 2x2 Demonstrator using ^{220}Rn injection data.

7.3.4 Temperature-Dependence of Drift Speed

As shown in Figure 7.4, the cryostat internal pressure varied significantly over the course of the ^{220}Rn injection period. These variations were caused by unrelated cryogenic issues that led to a temporary reduction in cooling performance. Corresponding changes in pressure are expected to induce changes in the temperature of both the ullage gas and the liquid argon, as confirmed using the temperature sensors instrumented in the 2x2 Demonstrator, which showed an initial decrease in temperature followed by a subsequent recovery. Variations in temperature and pressure are expected to affect the ionization electron drift speed in LAr [132]. In general, lower (higher) LAr temperature leads to higher (lower) electron mobility and therefore higher (lower) drift speed. As a result, time-dependent changes in cryogenic conditions should manifest as corresponding variations in the reconstructed drift speed.

To investigate this effect, the drift speed is measured as a function of time in each TPC using the same fiducial selection described in the previous section to avoid TPC boundaries and pixel tile edges. The resulting time dependence is shown in Figure 7.35, where a gradual increase followed by a decrease in drift speed is observed, consistent with the evolution of the temperature over the same period of time.

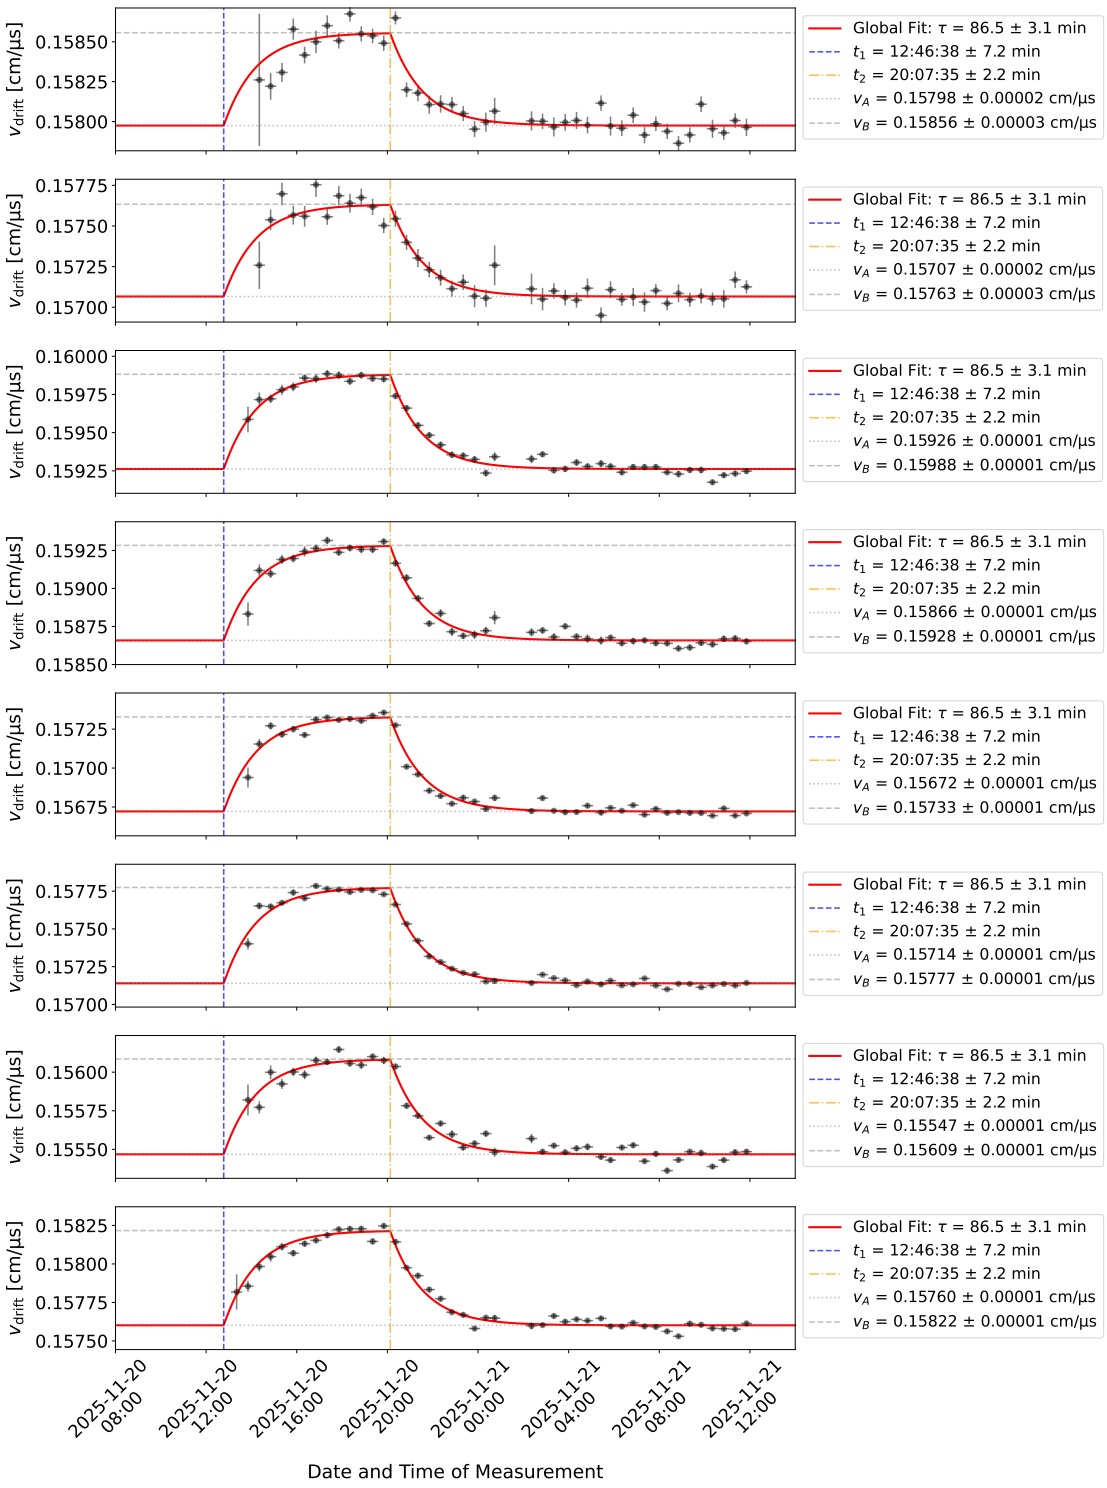

Figure 7.35: Time dependence of the ionization electron drift speed during a brief change in cryostat pressure, shown for each TPC of the 2x2 Demonstrator.

To quantify this behavior, the drift speed is modeled with a piecewise function describing two transitions in the operating conditions:

$$v_i(t) = \begin{cases} v_{\mathrm{A},i} & t < t_1 \\ v_{\mathrm{A},i} + (v_{\mathrm{B},i} - v_{\mathrm{A},i}) \left(1 - e^{-\frac{t-t_1}{\tau_p}}\right) & t_1 \leq t < t_2 \\ v_{\mathrm{A},i} + (v_i(t_2) - v_{\mathrm{A},i}) e^{-\frac{t-t_2}{\tau_p}} & t \geq t_2, \end{cases} \tag{7.7}$$

where i labels the TPC, $v_{\mathrm{A},i}$ is the baseline drift speed before the first transition and after full recovery, and $v_{\mathrm{B},i}$ is the drift speed during the intermediate state. The times t_1 and t_2 correspond to the times associated with the two observed pressure transitions, and τ_p describes the characteristic time scale over which the system responds to the pressure change. The value $v_i(t_2)$ denotes the drift speed at the end of the first transition.

A simultaneous fit to all TPCs is performed, in which the transition times and time constant are shared across all TPCs, while $v_{\mathrm{A},i}$ and $v_{\mathrm{B},i}$ are allowed to vary independently for each TPC. The fit results are shown in Figure 7.35. The extracted time constant (τ_p) is 86.5 $\pm$ 3.1 min and the fitted transition times t_1 and t_2 are consistent with the independently observed pressure changes.

The extracted time constant characterizes the timescale over which the drift speed responds to changes in cryogenic conditions. The observed variation in drift speed is at the level of approximately 0.3–0.4%. This corresponds to a variation in reconstructed drift time of roughly 0.5–1.0 µs over the full drift distance, and a shift in reconstructed position of approximately 0.1 cm. For comparison, in a meter-scale LArTPC, the same fractional variation would translate to timing shifts of roughly 40–50 µs and spatial shifts of several centimeters.

These results highlight the importance of stable cryogenic conditions for precise reconstruction in LArTPCs. The observed instability is attributed to intermittent reductions in LAr cooling power, an issue that was identified and is planned to be addressed prior to the next data-taking run of the 2x2 Demonstrator. The observed variation of drift speed with cryostat pressure allows a measurement of the relative dependence of ionization electron drift speed on temperature, expressed in

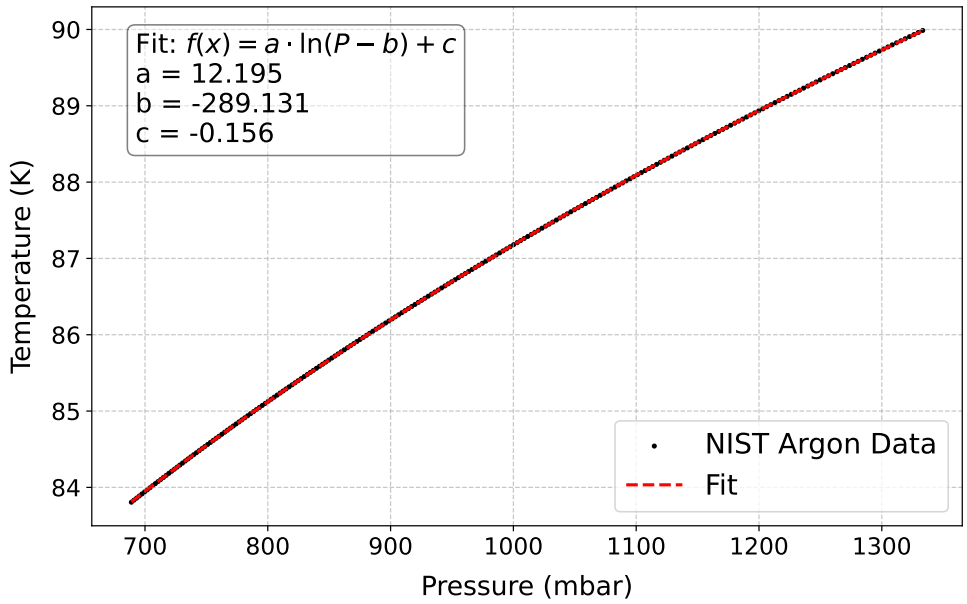

Figure 7.36: NIST pressure versus temperature data for argon [133].

units of $\%/\mathrm{K}$. This quantity was previously measured by Walkowiak [132] to be approximately $-1.97\%/\mathrm{K}$. While no explicit uncertainty is provided in that work, it serves as a useful benchmark for comparison with the 2x2 Demonstrator measurement. To extract this dependence, the fitted drift speeds $v_{A,i}$ and $v_{B,i}$, shown in Figure 7.35, are used to quantify the change in drift speed before and after the observed pressure transition. The fractional change in drift speed is combined with the corresponding change in temperature to obtain the variation in units of $\%/\mathrm{K}$.

Rather than relying directly on calibrated temperature sensor measurements, which were not available for this study, the temperature change is inferred from the measured cryostat pressure in the ullage. The conversion between pressure and temperature is obtained from NIST argon saturation properties [133], as shown in Figure 7.36. The average pressure immediately before and after the second pressure transition is used to determine the corresponding temperatures via this parameterization. The observed pressure change of approximately 19 mbar corresponds to a temperature variation of about 0.17 K, consistent with the trends observed using the 2x2 Demonstrator temperature sensors. The resulting temperature difference is combined with the measured change in drift speed to obtain a measurement in units of %/K.

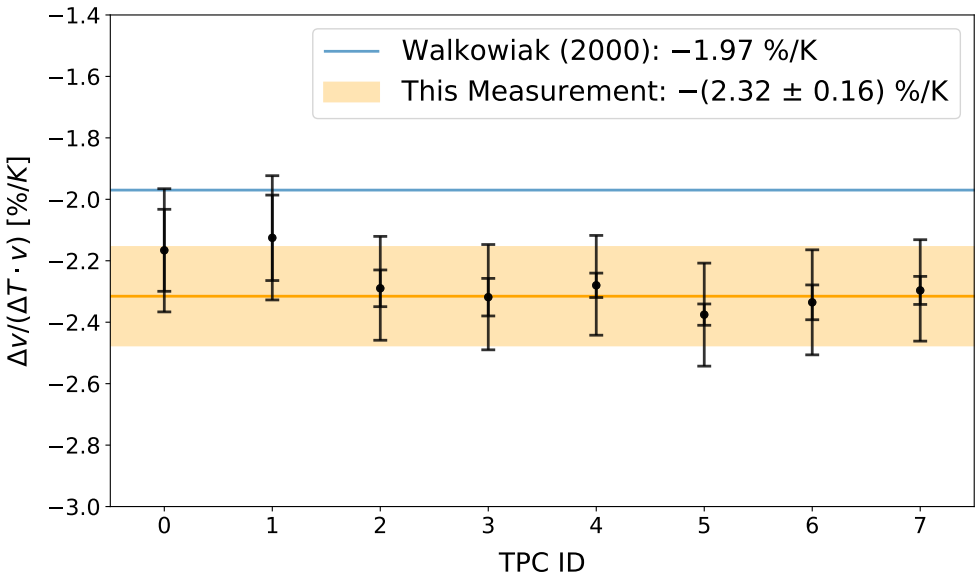

Figure 7.37: Measurement of ionization electron drift speed dependence on temperature (%/K) in the 2x2 Demonstrator.

Systematic uncertainties are evaluated from the stability of the pressure measurement. Specifically, the standard deviation of the pressure distributions before and after the transition is propagated through the NIST pressure-temperature conversion to obtain an uncertainty on the temperature difference. Statistical uncertainties arise from the uncertainties on $v_{A,i}$ and $v_{B,i}$ returned by the fit to the drift speed in each TPC as a function of time. These contributions to the measurement uncertainty are propagated independently to the final %/K values and combined in quadrature. To determine the uncertainty on the average across the TPC measurements, the fully correlated systematic uncertainties must be properly accounted for. This is done by quantifying the variation in the average upon varying the measurement by its systematic uncertainty bounds. This variation, in quadrature with the statistical uncertainty, is taken as the total uncertainty on the measurement. The final measured temperature dependence of the ionization electron drift velocity is

$$\frac{1}{v}\frac{\Delta v}{\Delta T} = -(2.32 \pm 0.16)\ \%/\mathrm{K}, \tag{7.8}$$

which is similar to the value reported by Walkowiak, though consistency between the two measurements is impossible to assess without an uncertainty on Walkowiak's measurement.

7.4 Electron Lifetime Measurements

Electron lifetime measurements using ambient radioactive sources, specifically ^{208}Tl γ decays, were discussed in Chapter 5. The same technique is used in this section to determine the electron lifetime before, during, and after the ^{220}Rn injection period in the 2x2 Demonstrator. This provides a means to assess the suitability of ^{220}Rn as a calibration source while also quantifying any potential impact of the injection on LAr purity. In particular, it is necessary to verify that the introduction of ^{220}Rn does not introduce a measurable degradation in electron lifetime, especially if it is to be used as a calibration handle.

The electron lifetime measurements based on ^{208}Tl γ decays have already been validated in Chapter 5 through comparison with dedicated purity monitor data. This validation enables electron lifetime measurements made using ^{208}Tl γ decays during the ^{220}Rn injection period to serve as a benchmark for evaluating the accuracy of BiPo-based electron lifetime estimates. It has already been shown that BiPo candidates, selected with two-pulse coincidence using associated scintillation light waveforms, are consistent with ^{212}Bi β decays, as confirmed through comparisons between data and MC simulation (Figure 7.21). However, this interpretation is robust only above approximately 46 ke$^-$, likely due to limitations in β light yield explored by Figure 7.19 (right). This motivated the chosen charge lower limit of 46 ke$^-$ to minimize potential effects that may influence the charge spectrum at lower reconstructed charges as a function of drift coordinate (which may bias the measurement). In addition, using only high reconstructed charges ensures that the charge-light matching purity (Figure 7.19, left) is close to 100%.

Figure 7.38 presents electron lifetime measurements obtained before, during, and after the injection period using both prospective ^{208}Tl γ decays (through associated electrons produced in Compton scattering) and BiPo candidates. The measurements using ^{208}Tl γ decays use the same charge and drift-time selections in Chapter 5, reconstructed charge greater than 53 ke$^-$ and reconstructed drift time between 60 µs and 160 µs. In contrast, the BiPo measurements use a wider drift-time range, 20 µs to 180 µs, to improve statistical precision, and applies a lower charge threshold of 46 ke$^-$ ensuring agreement between data and MC simulation for the chosen range.

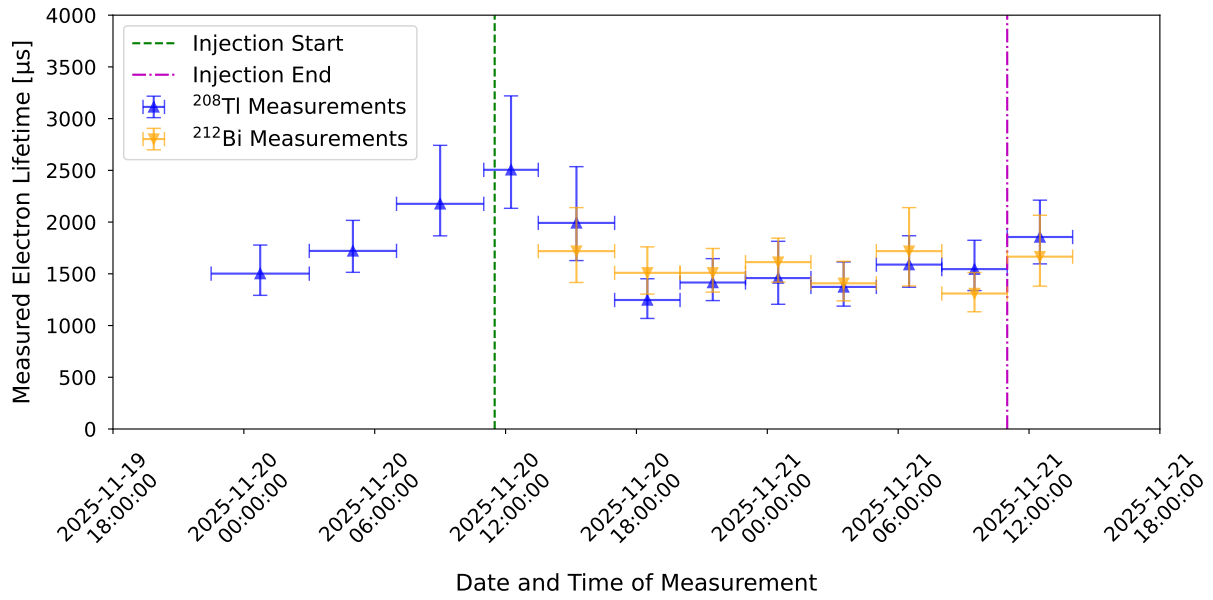

Figure 7.38: Electron lifetime measurements before, during, and after the ^{220}Rn injection period using two different sources of radioactivity. Both measurements utilize the diffusion correction determined using simulated ^{208}Tl γ decays.

The results indicate good agreement between the two measurement approaches throughout the injection period. An apparent reduction in the extracted electron lifetime is observed following the onset of ^{220}Rn injection. One possible explanation is the continuous introduction of impurities during injection, partially offset by removal through the LAr recirculation system, resulting in a lower equilibrium purity level. However, intermittent cryogenic instabilities occurring during the same period has not been excluded as the primary cause or at least a contributing factor. Further investigation is therefore required to determine the origin of the observed reduction in electron lifetime.

Overall, the results indicate that any impact of the ^{220}Rn injection on LAr purity is limited and does not compromise detector operation. This supports the feasibility of using ^{220}Rn-based calibration techniques in LArTPCs. Additional dedicated studies, either in future runs of the 2x2 Demonstrator or at the Colorado Argon Test Stand (CATS) at CSU, would be valuable for further quantifying the impact of ^{220}Rn injection on LAr purity.

7.5 ^{220}Rn Mixing and LAr Flow Dynamics

7.5.1 ^{220}Rn $\rightarrow$ ^{216}Po $\rightarrow$ ^{212}Pb Tagging

The decay chain involving ^{220}Rn and ^{216}Po proceeds via successive α decays with energies of approximately 6.4 MeV and 6.9 MeV, respectively. These correlated decays are commonly referred to as "RnPo" pairs due to their short temporal separation. In particular, ^{216}Po has a half-life of approximately 145 ms, enabling the possibility of identifying time-correlated decay pairs in a LArTPC provided that the rate of random coincident activity remains sufficiently low.

The identification of such RnPo pairs is of interest for characterizing LAr flow patterns. The half-life of ^{220}Rn (55.6 s) is short compared to typical mixing timescales in large detectors, implying that full volumetric mixing is unlikely before a significant fraction of decays occur. Consequently, an excess of RnPo activity is expected in regions proximal to the injection point, which may also coincide with LAr inlet locations. This spatial dependence provides a potential probe of LAr flow dynamics within the detector.

Furthermore, the approximately 145 ms half-life of ^{216}Po allows for the possibility that the two α decays occur at spatially separated locations due to motion in the liquid associated with convective currents. Given the sub-second separation between decays, the expected displacement is typically small (at most at the centimeter scale), but nevertheless potentially measurable. The observed spatial separation in the two decays can be used to estimate local LAr flow velocities. This technique has been demonstrated in smaller-scale liquid xenon experiments such as XENON100 [113].

In practice, however, the detection of RnPo decays in the 2x2 Demonstrator is challenging. The α decays produce relatively small ionization signals (roughly 5–6 ke$^-$). As shown in Figure 7.39, the majority of pixel channels have thresholds that are too high to reliably register such signals. Fewer than 5% of channels satisfy the required sensitivity, with the exception of Module-2, where this fraction increases to approximately 8%. This limited coverage leads to a correspondingly low overall detection efficiency for RnPo charge signatures.

Figure 7.40 shows the LRS trigger rate during the initial phase of the ^{220}Rn injection period at the 2x2 Demonstrator. Multiple short-lived spikes in LRS trigger rate are observed within the first

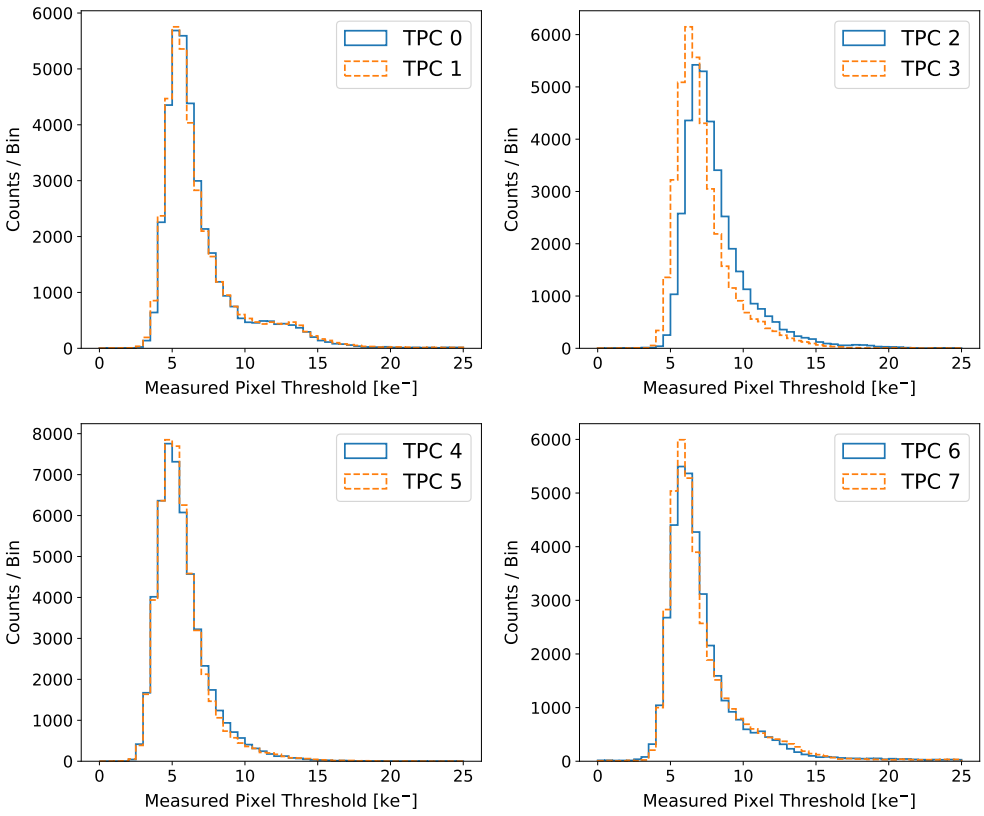

Figure 7.39: Measured pixel-by-pixel thresholds in the 2x2 Demonstrator. The thresholds are obtained using specialized data-taking runs in which the LArPix periodic reset is disabled and the CRS is triggered by accumulated "leakage" charge. (Top left) Module-0, (Top Right) Module-1, (Bottom Left) Module-2, (Bottom Right) Module-3.)

few hours of injection. These features decay within approximately 10–15 min, a timescale that is more consistent with the decay of RnPo pairs than with the longer lived ^{212}Pb or ^{212}Bi components of the ^{220}Rn decay chain. Also illustrated in Figure 7.40 are the corresponding optical hit rates in the bottom ACL detectors, which show a uniquely prominent set of rate spikes. In particular, specific TPC regions in Module 1 (TPCs 2 and 3) and Module 3 (TPC 7) show the largest spikes, which is likely due to their proximity to the injection site. The fact that the signals appear to be predominantly associated with the lower regions of the detector, as indicated by the response of the bottom ACLs, further suggests that ^{220}Rn activity may be entering the active volume from below. In such a scenario, limited vertical mixing would result in a substantial fraction of decays occurring before full homogenization of the injected activity is achieved.

A selection based on LRS trigger data is used to identify RnPo decay candidates. A dedicated

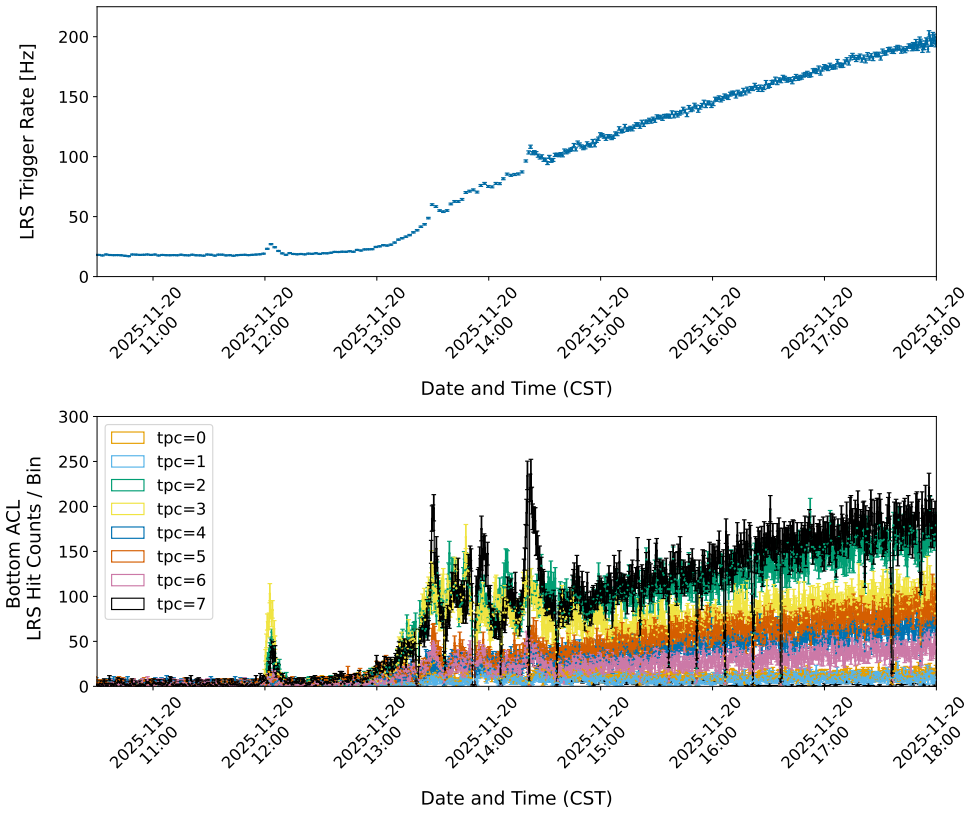

Figure 7.40: (Top) Total LRS trigger rates early in the injection period and (Bottom) hit counts overtime for lower ACL detectors broken down by TPC.

procedure is required to mitigate pile-up effects, which become more significant at higher trigger rates. The final short-lived LRS trigger rate spike occurs when the total LRS trigger rate reaches approximately 100 Hz, corresponding to an average of roughly 20 events within a 200 ms time window. The event selection criteria are defined as follows:

1. Events must be triggered by the lower ACL light detector modules.

2. Matched LRS trigger pairs must originate from the same TPC.

3. The time separation between matched LRS triggers must be less than 2 s.

4. The total light yield, calculated as the integral from 2 µs to 5.2 µs with respect to the beginning of the readout window, of the matched LRS trigger pair must agree within 100 PE.

5. No charge activity may be associated with either LRS trigger.

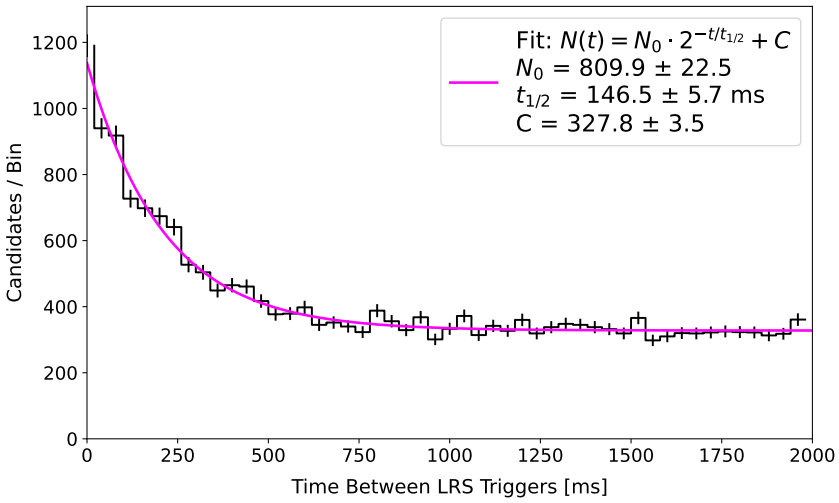

Figure 7.41: Distribution of time differences between matched LRS triggers for RnPo candidates. The extracted half-life is consistent with the expected $\sim$ 145 ms half-life of ^{216}Po.

The light-yield consistency requirement exploits the fact that the two α decays in the RnPo sequence have similar energies and are expected to occur at nearly the same spatial location. As a result, the detected photoelectron (PE) yields are expected to be comparable. The requirement of no associated charge activity is motivated by the low efficiency for detecting ionization signals from α decays given the CRS threshold distribution. A separate selection, described below, is therefore required to identify charge deposits associated with these events.

Figure 7.41 shows the distribution of RnPo candidate counts as a function of the time difference between matched LRS triggers. An exponential fit combined with a constant background term is used to extract the effective half-life, yielding 146.5 $\pm$ 5.7 ms, consistent with the expected ^{216}Po half-life of approximately 145 ms. A non-negligible contribution from accidental coincidences is observed as a flat component in the distribution; this is accounted for by the constant term in the fit. After subtraction of the estimated background contribution, approximately 4,700 RnPo candidates are identified in the early phase of the ^{220}Rn injection period.

Following identification of RnPo candidates using LRS information, an additional selection is applied to associate ionization signals with each decay. The criteria are as follows:

1. Two low-energy charge clusters are required within a 1 cm radius in three dimensions.

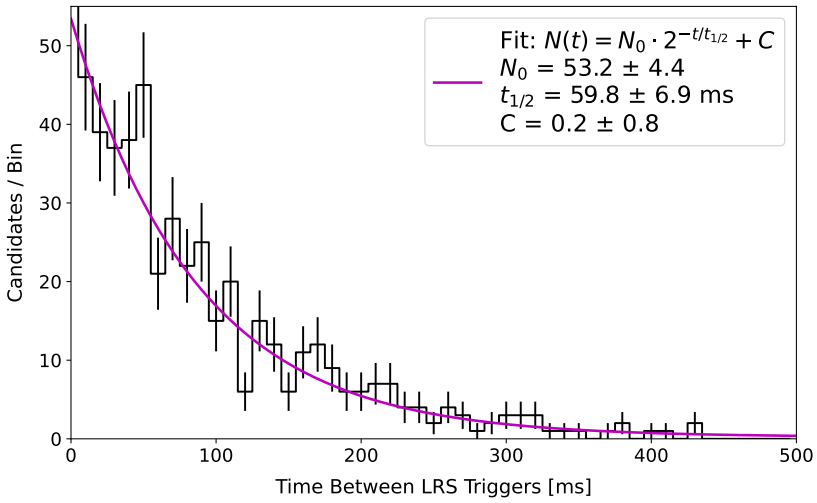

Figure 7.42: Distribution of time differences between matched LRS triggers for RnPo candidates with an additional requirement of an associated charge cluster for each trigger.

2. Each charge cluster must be matched to an LRS trigger within a 1 s time window.

3. Each charge cluster must consist of a single hit (one charge packet) with a reconstructed charge between 1 ke^- and 10 ke^-.

The distribution of time differences for charge-light-matched RnPo candidates is shown in Figure 7.42. The number of selected RnPo candidates is significantly reduced relative to the light-only selection, as expected given the high CRS thresholds. The extracted effective half-life is approximately 60 ms, substantially lower than the expected value. Even if the initial ^{220}Rn decay occurs within the sensitive region of a low-threshold pixel channel, the subsequent ^{216}Po atom may drift sufficiently far such that its decay products are detected by a different pixel channel. If that pixel has a higher self-trigger threshold, the signal from the ^{216}Po α decay may not be detected. This introduces a bias in which longer time separations are preferentially lost, leading to a systematic suppression of reconstructed RnPo pairs. As a result, the extracted effective half-life would be biased towards shorter values. The agreement between the light-only measurement and the expected half-life confirms the presence of RnPo decays, indicating that the deviation observed in the charge-based selection is most likely attributable to threshold-induced inefficiencies.

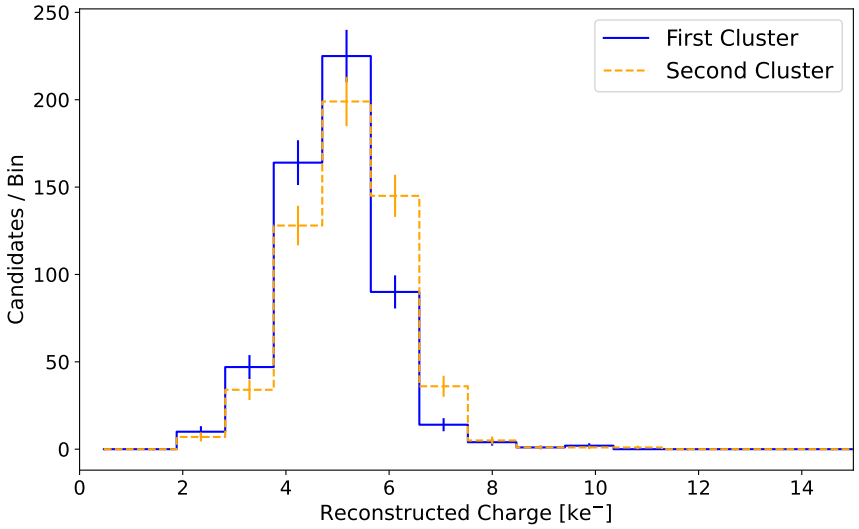

Figure 7.43: Reconstructed charge distributions for charge–light matched RnPo candidates. Each point corresponds to an individual cluster associated with an RnPo decay. It should be noted that the measured half-life is expected to be significantly biased due to CRS inefficiencies related to variations in the CRS thresholds.

Figure 7.43 shows the reconstructed charge distributions for clusters associated with RnPo candidates, separated into the earlier and later decays in each pair. Under ideal conditions, such signals would exhibit approximately Gaussian reconstructed charge distributions; however, threshold effects can introduce an asymmetry in the distributions. In addition, a systematic offset is observed between the two populations, with the later cluster exhibiting slightly higher reconstructed charge on average. This is consistent with the energy difference between the two α decays, as the ^{216}Po decay has approximately 8% higher energy (at 6.906 MeV) than the ^{220}Rn decay (at 6.405 MeV).

Figure 7.44 shows the a scatter plot of reconstructed RnPo candidates within the 2x2 Demonstrator. The events are strongly concentrated near the bottom regions of the detector, despite no explicit spatial selection criteria being applied. This reinforces the conclusion drawn from the light-only analysis that the dominant production region of RnPo decays lies near the bottom of the TPCs, consistent with ^{220}Rn entering the detector preferentially from below and undergoing limited vertical mixing prior to decaying to ^{216}Po.

A higher fraction of reconstructed candidates is observed in Module-2 relative to the light-only

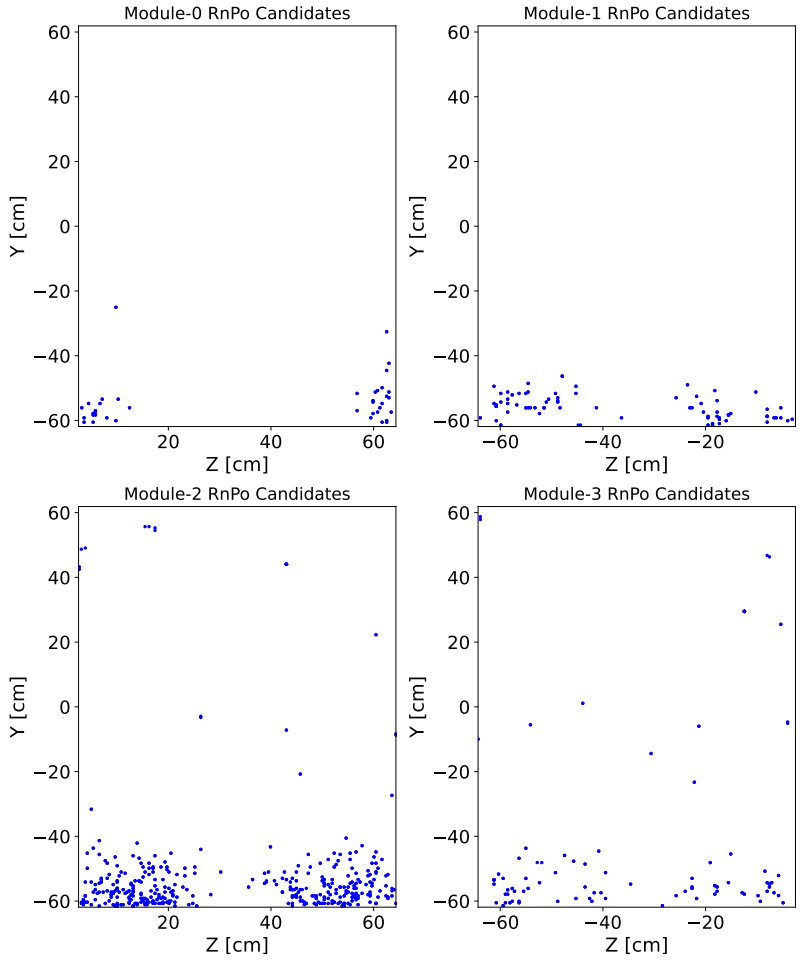

Figure 7.44: Spatial distribution of charge-light-matched RnPo candidates within the detector volume.

selection. This is consistent with the lower CRS thresholds in that module, which increase the efficiency for detecting low-energy ionization signals. A reduced reconstruction efficiency is also evident in regions near the center of each TPC, likely caused by light collection efficiency effects that suppress RnPo candidate identification. Module-0 in particular exhibits reduced efficiency, consistent with lower light collection efficiency compared to the other modules. These effects collectively reflect the strong dependence of the reconstruction efficiency on light collection efficiency and CRS thresholds.

7.5.2 Analysis of Charge Data

As discussed in Chapter 5, the continuous readout provided by the LArPix system enables access to information that is not readily obtainable using LRS triggering alone. In particular, low-energy electron activity (from β decays or γ decay Compton scatters) distributed across the full detector volume is difficult to reconstruct with high efficiency using the LRS. Section 5.7 demonstrated that charge-only reconstruction, without requiring a coincident light signal, can reveal spatial features that are otherwise obscured in light-triggered selections.

In the context of ^{220}Rn injection, this distinction is particularly relevant. The LRS-triggered dataset is intrinsically biased towards high light yield processes, most notably α decays. This is further accentuated in the 2x2 Demonstrator, where the LRS was typically operated with relatively high thresholds, resulting in selected low-energy electron candidates that are strongly concentrated near the light detection system. Consequently, charge-light-matched samples are dominated by α-induced signals, primarily from ^{212}Bi and ^{212}Po decays. In contrast, β and γ interactions, such as those from ^{212}Pb and ^{208}Tl, are selected with relatively low efficiency. Access to these decay channels is nevertheless useful in probing LAr flow dynamics in the cryostat.

To isolate spatial features arising specifically from the ^{220}Rn injection, a data-driven background subtraction procedure is employed. For each TPC, the charge cluster spatial distributions in y and z are constructed using a 90-minute period prior to injection, as described in Section 5.7. These distributions serve as the background reference. The same quantities are then constructed for a 90-minute interval near the end of the injection period. A bin-by-bin subtraction of the background from the injection-period distributions is performed to isolate injection-induced contributions. The resulting distributions are shown in Figures 7.45–7.48.

Several consistent features are observed across all TPCs. Most notably, a pronounced excess is present at the lower boundary of each TPC. This is consistent with the interpretation that activity enters the detector preferentially from below. This observation is further supported by the RnPo analysis in Section 7.5.1, which identified a concentration of short-timescale correlated decays near the bottom of the TPCs, consistent with limited vertical mixing from the ^{220}Rn and ^{216}Po

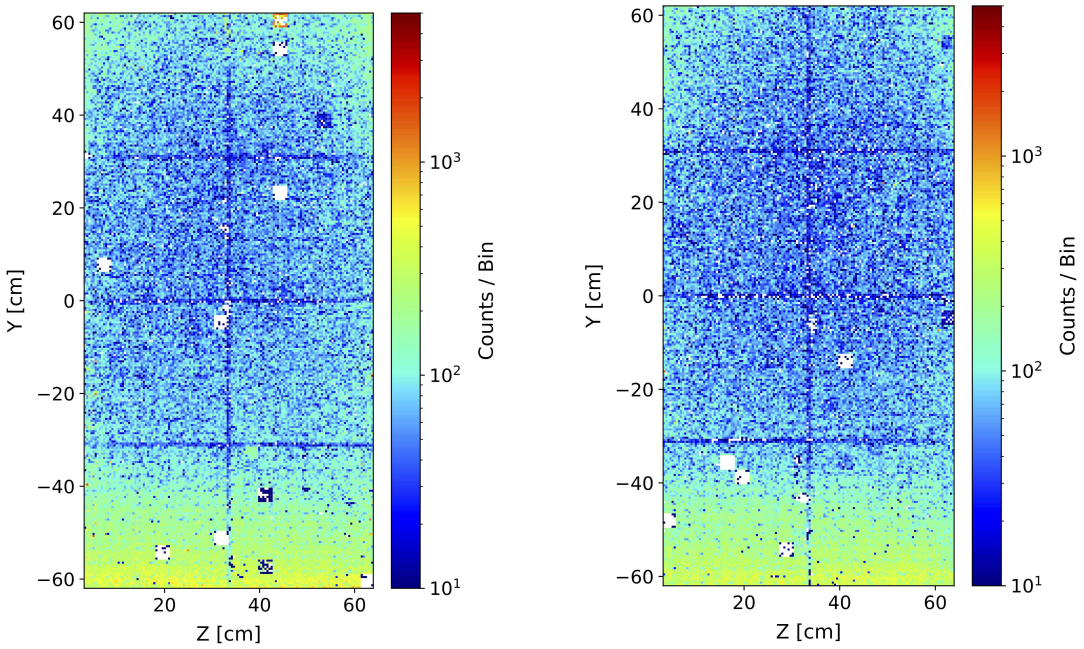

Figure 7.45: Background-subtracted spatial distribution of charge-light-matched low-energy charge clusters during ^{220}Rn injection in Module-0 of the 2x2 Demonstrator using a charge-only selection. (Left) TPC-0, (Right) TPC-1.

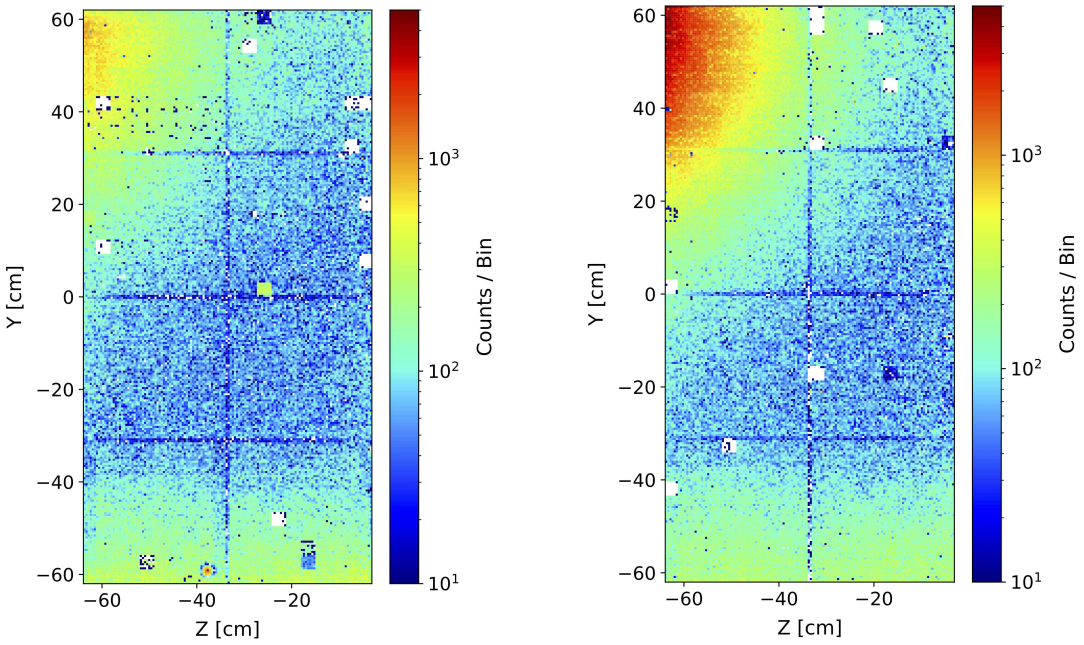

Figure 7.46: Background-subtracted spatial distribution of charge-light-matched low-energy charge clusters during ^{220}Rn injection in Module-1 of the 2x2 Demonstrator using a charge-only selection. (Left) TPC-2, (Right) TPC-3.

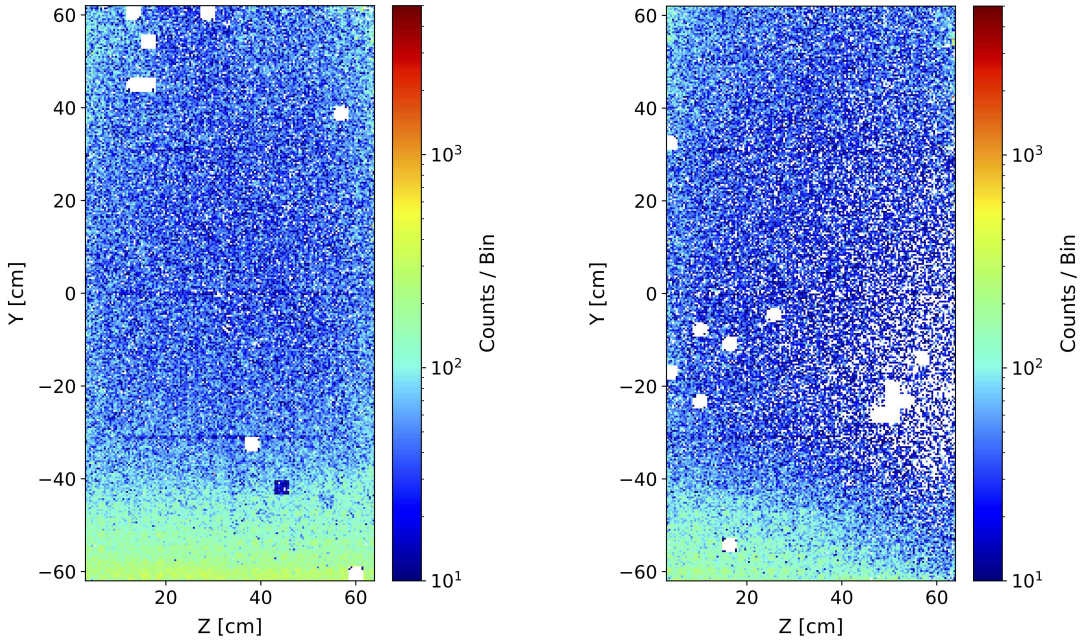

Figure 7.47: Background-subtracted spatial distribution of charge-light-matched low-energy charge clusters during ^{220}Rn injection in Module-2 of the 2x2 Demonstrator using a charge-only selection. (Left) TPC-4, (Right) TPC-5.

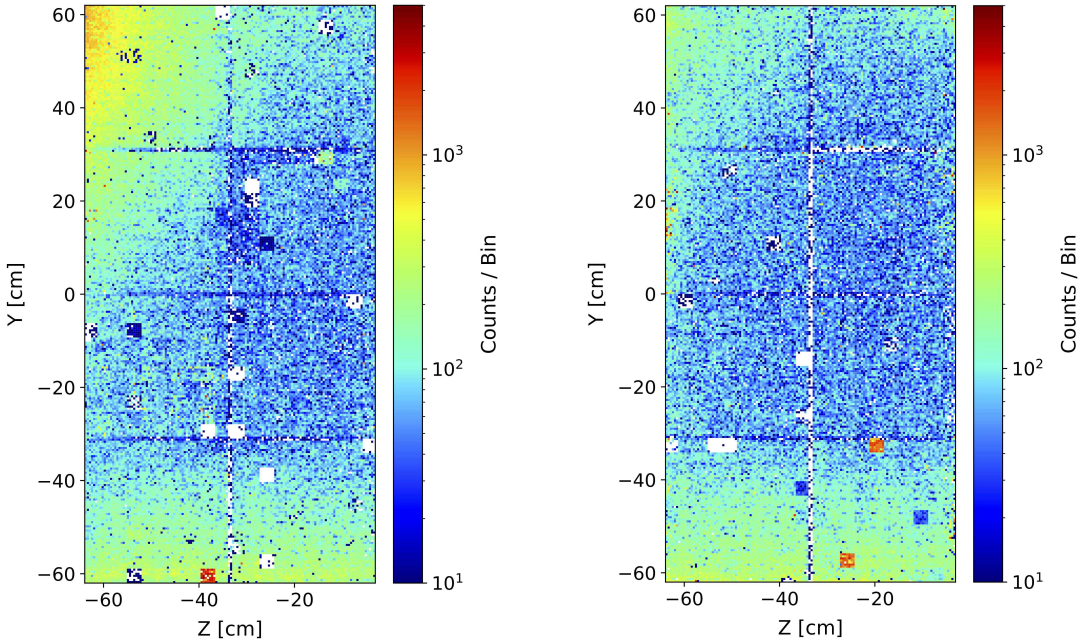

Figure 7.48: Background-subtracted spatial distribution of charge-light-matched low-energy charge clusters during ^{220}Rn injection in Module-3 of the 2x2 Demonstrator using a charge-only selection. (Left) TPC-6, (Right) TPC-7.

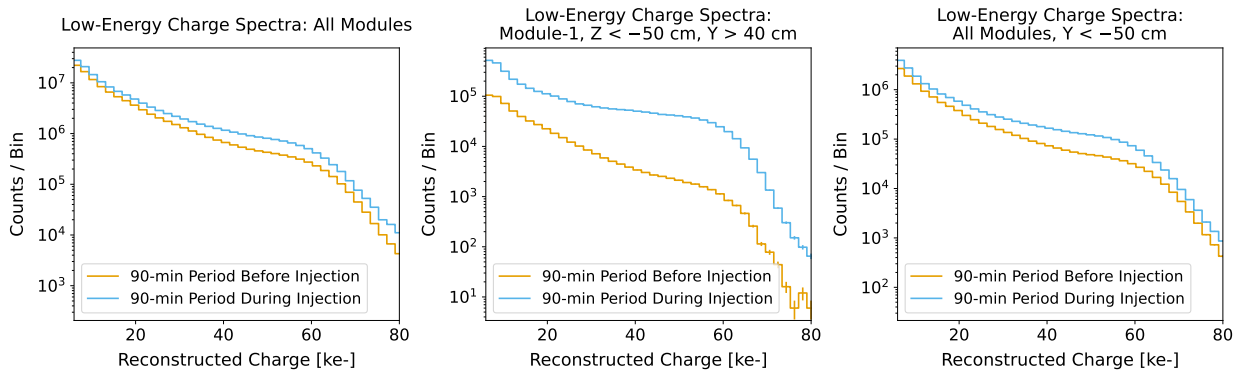

Figure 7.49: Comparison of charge spectra during ^{220}Rn injection versus pre-injection background for three areas inside the 2x2 Demonstrator.

isotopes. In addition, earlier observations of enhanced cathode-side activity near the lower field cage further support the mixing of decay products from the bottom-end of the detector. Taken together, these results provide coherent evidence that ^{220}Rn decay chain activity is introduced predominantly through the lower regions of the 2x2 Demonstrator where it can mix throughout the entire volume of each TPC.

There is also a distinct feature observed near the top of the TPCs that is not present in the charge-light-matched selection. In particular, a pronounced excess of activity appears to originate from the top-left region of several TPCs. This effect is most prominent in Module-1 and Module-3, which are closest to the ^{220}Rn injection point. The geometry of the detector relative to the injection site is shown in Figure 7.2. Module-1 TPC-3 is the closest to the injection location and correspondingly exhibits the strongest enhancement in the top-left region.

The strength of this feature is observed to increase over the course of the injection period. In addition, related but less pronounced excesses are present in both the top-left and top-right corners of several TPCs, though these are not uniformly visible across all detector regions. These structures are indicative of non-trivial LAr flow patterns originating from the upper regions of the TPCs.

Figure 7.49 shows charge spectra for selected regions of the detector, comparing a period of time during the injection to an equally long period of time prior to the injection. When integrated over the full detector volume, the spectrum exhibits an overall enhancement of approximately 32% in comparison to the pre-injection background distribution. The most pronounced local excess

is observed in the top-left region of TPC-3 in Module-1, where the rate reaches approximately 5.7 times the background level. In contrast, regions near the bottom of the detector show a more modest increase, roughly 1.6 times the background rate.

A common feature across all examined regions is an enhancement of the Compton edge associated with ^{208}Tl γ decays in the high-charge portion of the spectrum. Because ^{208}Tl is produced in the ^{220}Rn decay chain, its increased presence during injection is expected. The charge spectra further indicate that the excess in Module-1 is dominated by Compton scatters associated with ^{208}Tl γ decays, consistent with this interpretation. One possibility is that a large fraction of isotopes of the ^{220}Rn decay chain linger in the gas-phase and re-introduce themselves later on, preferentially near the top of the liquid, which would be consistent with the observations of the excesses near the top. A complete understanding of these effects will likely require additional dedicated studies of fluid transport and decay chain evolution in the detector. Nevertheless, the observed features provide valuable constraints on both the ^{220}Rn injection behavior and the LAr flow dynamics in the 2x2 Demonstrator.

7.5.3 LAr Flow Dynamics During LAr Top-Off

Following the ^{220}Rn injection period, the cryostat underwent an unrelated LAr top-off procedure. While this process temporarily reduces the overall LAr purity due to the introduction of less purified argon, the effect is reversible through subsequent recirculation. It should be noted that top-off operations are not part of standard running conditions, as the 2x2 Demonstrator cryogenic system is designed to re-condense argon boil-off gas to maintain a stable liquid level. However, during the second run of the 2x2 Demonstrator, intermittent cooling instabilities necessitated occasional LAr top-offs to maintain an adequate liquid level and ensure stable high-voltage operation. These commissioning activities formed part of the broader goals of this run, which serve as a system validation phase ahead of the anticipated 2026 beam run. Identified issues are therefore under active mitigation for future operation at the time of writing this thesis.

The introduction of additional LAr with different initial purity provides an opportunity to study

flow and mixing dynamics using CRS data. In particular, the continuous readout capability enables observation of spatial and temporal variations in ionization activity on timescales of several minutes. Figures 7.50 and 7.51 show the evolution of low-energy activity spatial locations during the LAr top-off period for two TPCs. To highlight changes over time, each distribution is compared to a reference distribution taken approximately five hours after the start of the top-off, with the reference-subtracted distributions shown in the two figures.

The observed distributions evolve significantly throughout the LAr top-off process, with noticeable changes occurring on timescales of approximately 15 minutes. The top-off procedure itself involves the injection of filtered LAr through four sprayers located at the top of each module. The top field shell panel contains perforations that allow the incoming LAr to diffuse into the TPC active volume. Each TPC exhibits broadly similar temporal evolution. Although the LAr top-off operation was halted at approximately 15:38, the spatial distributions do not return to a stable configuration for several hours thereafter.

The results show clear patterns consistent with LAr flow patterns originating at the top and bottom of the TPC active volume. These are locations where known perforations exist in the field structures that allow LAr to flow into or out of the TPC active volume. Overall, these results demonstrate that the continuous nature of the CRS readout potentially allows for measuring precise LAr flow rates in the y–z plane. Additional studies will be required to further explore this possibility in the future.

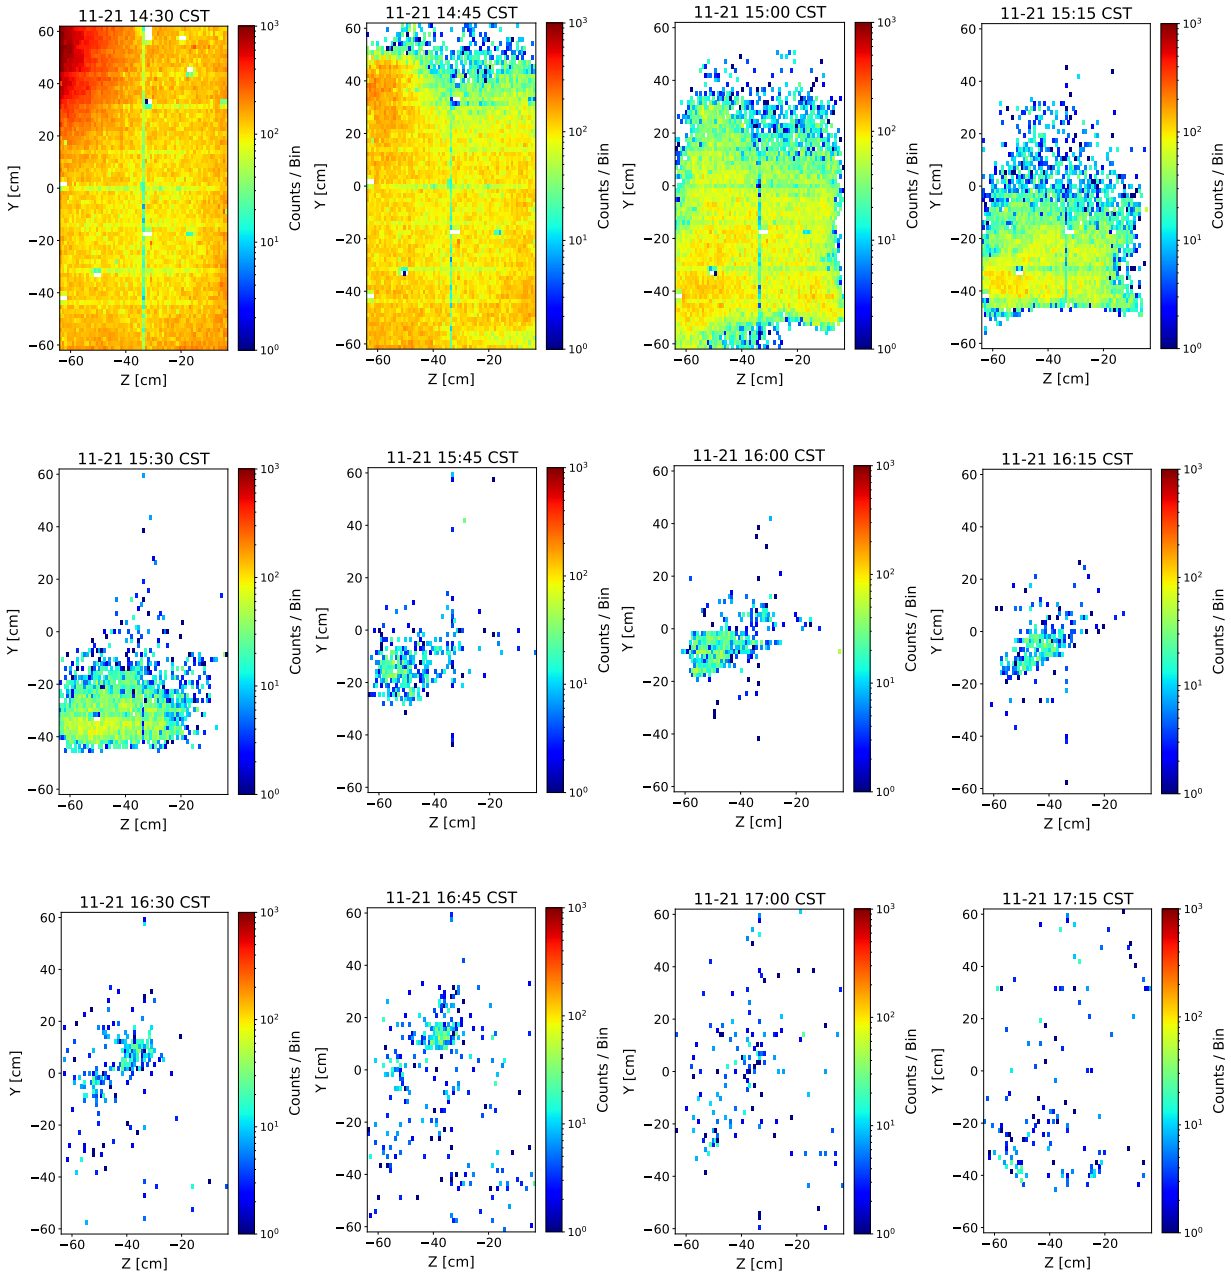

Figure 7.50: Spatial distributions of low-energy charge activity over time during a LAr top-off for Module-1 TPC-3 in the 2x2 Demonstrator. A reference distribution is subtracted from each of the above spatial distributions to accentuate LAr flow effects. Each plot contains 15 min of data after the time shown in the title.

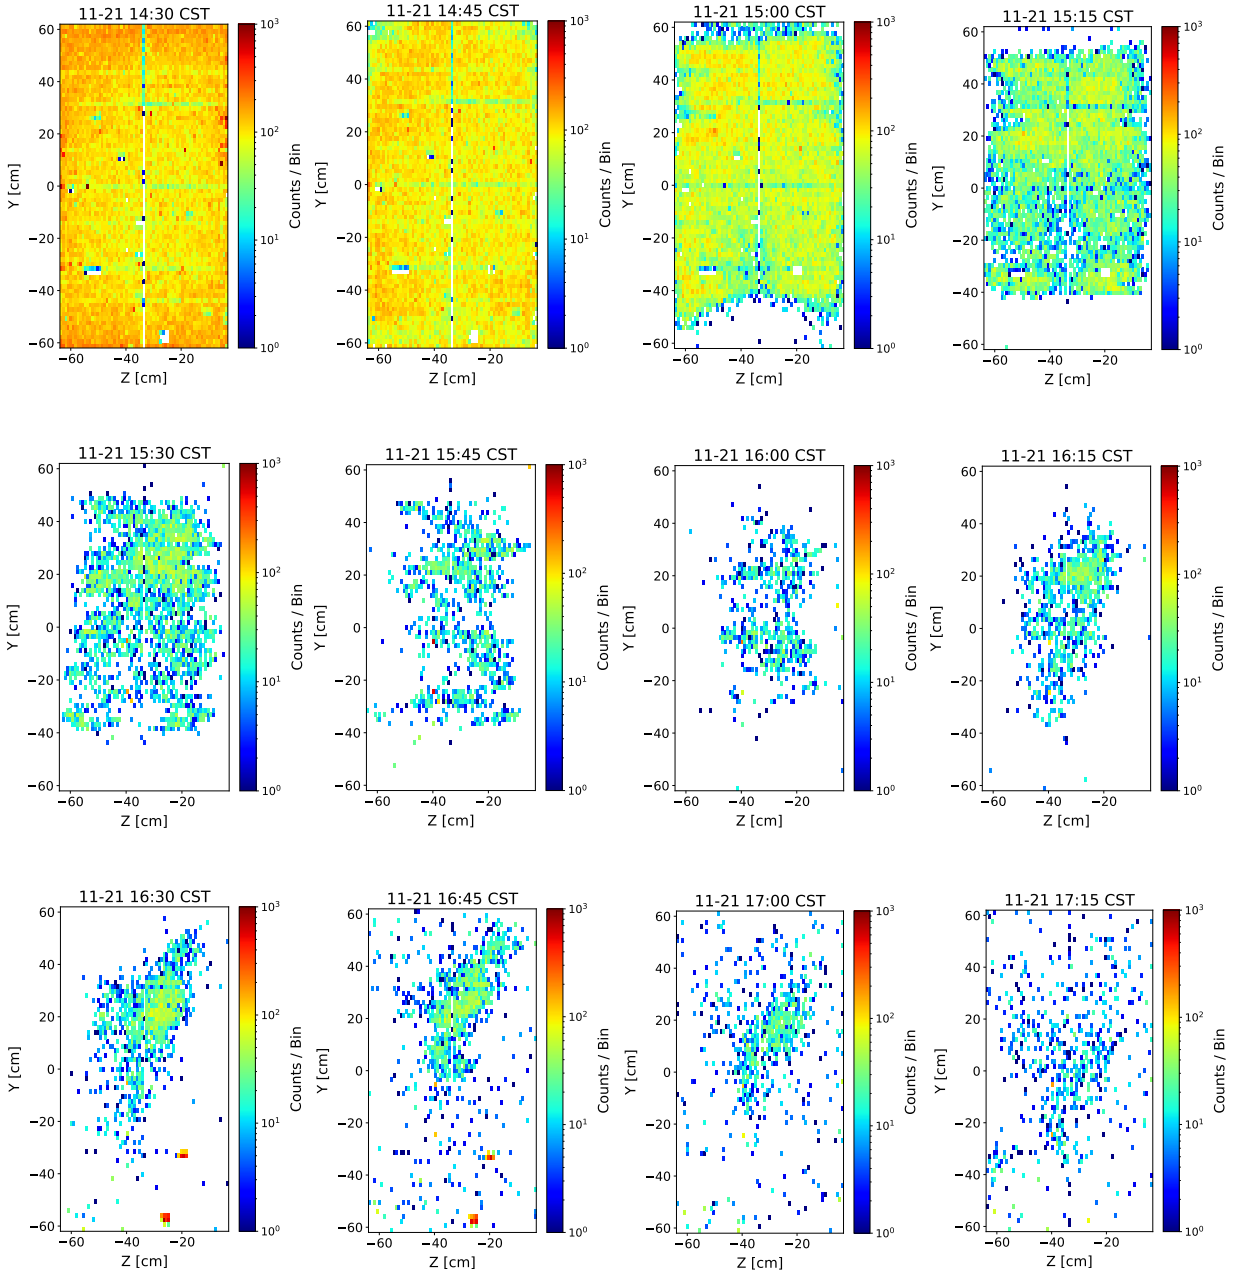

Figure 7.51: Spatial distributions of low-energy charge activity over time during a LAr top-off for Module-3 TPC-7 in the 2x2 Demonstrator. A reference distribution is subtracted from each of the above spatial distributions to accentuate LAr flow effects. Each plot contains 15 min of data after the time shown in the title.

Chapter 8

Summary and Conclusions

Detector calibration is a cornerstone activity required for the successful operation and physics output of LArTPCs. The primary objective of this thesis was to investigate and validate calibration techniques that leverage low-energy radioactivity. Developing these alternative methods is crucial for next-generation experiments such as DUNE, which will operate in a deep underground environment with a severely suppressed cosmic-ray muon flux.

The 2x2 Demonstrator operated at Fermilab offered a unique, low-background environment to pioneer novel calibrations utilizing sources of low-energy radioactivity. This work successfully demonstrated that ambient radiological backgrounds, as well as intentionally introduced isotopes from the ^{220}Rn decay chain, can be reliably reconstructed and utilized to characterize key detector performance metrics.

Ambient Low-Energy Radioactivity Calibrations

Using data collected by the 2x2 Demonstrator, this work established that low-energy ambient radioactivity can be detected and correlated across both the charge and light collection systems. This joint detection enables full 3D event reconstruction, which in turn facilitates several distinct detector calibrations:

- **Electric-Field Distortion Monitoring:** The 3D reconstruction along the drift coordinate allowed for monitoring of average electric-field distortions near TPC boundaries.
- **Charge-Scale Calibration:** Exploiting the prominent energy signature of ^{208}Tl γ-decays, a charge-scale calibration was developed to improve data-to-simulation agreement. This calibration resulted in an approximately 4.3$\%$ reduction in the simulated charge scale to match the experimental data.

- **Electron-Lifetime Measurements:** The distinct Compton edge of the ^{208}Tl spectrum provided a standalone metric for monitoring electron lifetime. The resulting purity measurements demonstrated excellent agreement with independent data from a dedicated purity monitor.

While differences in ultimate detector geometries mean that the exact implementation of these ambient calibrations may be different for the DUNE Far Detector, these results conclusively demonstrate that ambient low-energy backgrounds are a highly valuable tool for evaluating detector performance in low-background environments.

Small-Scale Validation: The Colorado Argon Test Stand

To explore novel calibration techniques under controlled conditions, a small-scale LArTPC test stand, the Colorado Argon Test Stand (CATS), was developed at CSU. CATS integrates a cubic-foot LArTPC utilizing pixel-based charge readout (LArPix technology) operated in tandem with a custom X-ARAPUCA-inspired photon detection system. Significant progress was achieved in establishing charge-light event coincidence, a capability essential for investigating the ^{220}Rn injection calibration technique.

Initial validation tests performed in CATS confirmed that substantial ^{220}Rn activity could be successfully introduced and detected in LAr via the light readout system. Comparative testing of two distinct ^{220}Rn sources revealed that a commercial source from Pylon Electronics yielded higher detection rates in LAr. This crucial validation paved the way for the subsequent scale-up and deployment of the Pylon source at the 2x2 Demonstrator in late 2025.

^{220}Rn Injection in the 2x2 Demonstrator

The deployment of the ^{220}Rn injection system at the 2x2 Demonstrator represented a critical milestone, scaling the technique into a large-volume DUNE prototype. Following a 24 hour injection window, data were collected over five days to track the spatial distribution and temporal decay of the daughter isotopes.

The test yielded several key insights and quantitative conclusions:

- **Clear Indications of the Presence of ^{220}Rn:** Light readout system trigger rates spiked during injection and returned to background levels over approximately five days. A large population of events featuring a high prompt light fraction (f_{prompt}) confirmed the presence of the expected α decays.

- **BiPo Event Selection:** Double-peak scintillation signals characteristic of the rapid BiPo sequence were clearly identified. The effective half-life extracted from the peak-to-peak time differences demonstrates agreement with multiple independent measurements, proving that the sample is a high-purity sample of BiPo decays. In addition, isolating these distinct temporal signatures significantly enhanced the selection purity of charge-light-matched events.

- **LAr Purity During Injection:** Utilizing a template-based pulse-shape discrimination (PSD) method, independent electron lifetime measurements were performed using both electron-like and BiPo-like activity. The data revealed a slight degradation in LAr purity coinciding with the start of the injection. This effect was attributed to the high gaseous argon carrier flow rates required to prevent isotope plateout along the lengthy injection lines. The purity impact is expected to be significantly lessened by reducing flow rates, shortening piping length, and reducing pipe diameter. Additional studies will be required to fully understand the impact using different injection configurations.

- **Drift Speed and Electric-Field Calibration:** A significant portion of the ^{220}Rn daughter isotopes plated out onto the TPC cathodes. This localized, near-cathode event sample provided an ideal calibration candle to measure average electron drift speeds across all TPCs. While the TPCs showed broad agreement, notable local variations in drift speed and E-field spatial uniformity were observed near TPC and pixel-tile boundaries.

- **Temperature Dependence of Drift Speed:** Exploiting minor pressure-induced temperature fluctuations during the injection period, the temperature dependence of the electron drift

speed in liquid argon was measured to be $-(2.32 \pm 0.16)\,\%/\mathrm{K}$. This is an important quantity to measure for DUNE LArTPCs at the nominal E-field of 500 V/cm, which helps with understanding the importance of stable cryogenic conditions during detector operations.

- **Mapping Liquid Argon Flow Patterns:** Tracking the spatial distribution of ^{220}Rn $\rightarrow$ ^{216}Po $\rightarrow$ ^{212}Pb candidate events, consisting of two successive α decays separated by a fraction of a second, revealed a strong spatial clustering near the bottom of the TPCs. Due to the short half-lives of ^{220}Rn and ^{216}Po, these isotopes decay before they can mix uniformly, acting as localized flow tracers. This clustering, along with localized charge activity excesses and enhanced cathode plateout, proves that the injected isotopes primarily enter the active volume through known perforations in the bottom field cage panel. Conversely, an independent charge excess observed near the top corner closest to the injection site highlighted a secondary entry pathway through the top field cage panel.

Finally, a LAr top-off procedure performed after the injection inadvertently demonstrated a powerful, unanticipated capability: the continuous, high-granularity readout of the LArPix system is uniquely suited for mapping highly transient, short-timescale LAr fluid dynamics and mixing patterns.

Summary of Impact and Future Outlook

This dissertation has comprehensively evaluated and demonstrated the efficacy of low-energy radioactive calibrations in LArTPCs using both natural ambient backgrounds and artificially introduced isotopes. The successful deployment of the ^{220}Rn injection technique, moving systematically from small-scale validation in CATS to a large-scale demonstration in the 2x2 Demonstrator, establishes a mature foundation for this calibration pathway.

To optimize this system for deployment in DUNE ND-LAr and the DUNE Far Detector, ongoing and future studies at CSU will continue to use the CATS facility to evaluate alternative injection geometries and lower carrier gas flow rates aimed at mitigating LAr purity impacts. Ultimately, the methodologies and results presented in this work provide a viable, robust framework for calibrat-

ing next-generation neutrino detectors deep underground, successfully overcoming the limitations imposed by the relative scarcity of cosmic-ray muons.

Bibliography

[1] Allan D Franklin and Alysia D Marino. *Are there really neutrinos?: an evidential history*. crc Press, 2020.

[2] Laurie M Brown. The Idea of the Neutrino. *Physics Today*, 31(9):23–28, 1978.

[3] Frederick Reines and CL Cowan Jr. Detection of the Free Neutrino. *Physical Review*, 92(3):830, 1953.

[4] Mary K Gaillard, Paul D Grannis, and Frank J Sciulli. The Standard Model of Particle Physics. *Reviews of Modern Physics*, 71(2):S96, 1999.

[5] Yoshiyuki Fukuda, T Hayakawa, E Ichihara, K Inoue, K Ishihara, Hirokazu Ishino, Y Itow, T Kajita, J Kameda, S Kasuga, et al. Evidence for oscillation of atmospheric neutrinos. *Physical Review Letters*, 81(8):1562, 1998.

[6] S.M. Bilenky and B. Pontecorvo. Lepton mixing and neutrino oscillations. *Physics Reports*, 41(4):225–261, 1978.

[7] Agnese Giaz. Status and perspectives of the JUNO experiment, 2018.

[8] Ke Abe, Ke Abe, H Aihara, A Aimi, R Akutsu, C Andreopoulos, I Anghel, LHV Anthony, M Antonova, Y Ashida, et al. Hyper-Kamiokande Design Report. *arXiv preprint arXiv:1805.04163*, 2018.

[9] Babak Abi, Roberto Acciarri, Mario A Acero, Giorge Adamov, David Adams, Marco Adinolfi, Zubayer Ahmad, Jhanzeb Ahmed, Tyler Alion, S Alonso Monsalve, et al. Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume II: DUNE Physics. *arXiv preprint arXiv:2002.03005*, 2020.

[10] Francesco Capozzi, Shirley Weishi Li, Guanying Zhu, and John F. Beacom. DUNE as the Next-Generation Solar Neutrino Experiment. *Phys. Rev. Lett.*, 123:131803, Sep 2019.

[11] John D Vergados, Hiroyasu Ejiri, and Fedor Šimkovic. Theory of neutrinoless double-beta decay. *Reports on Progress in Physics*, 75(10):106301, 2012.

[12] Michelle J Dolinski, Alan WP Poon, and Werner Rodejohann. Neutrinoless double-beta decay: Status and prospects. *Annual Review of Nuclear and Particle Science*, 69(1):219–251, 2019.

[13] Paul Langacker. Grand unified theories and proton decay. *Physics Reports*, 72(4):185–385, 1981.

[14] M Miyajima, T Takahashi, S Konno, T Hamada, S Kubota, H Shibamura, and T Doke. Average energy expended per ion pair in liquid argon. *Physical Review A*, 9(3):1438, 1974.

[15] Ettore Segreto. Properties of liquid argon scintillation light emission. *Physical Review D*, 103(4):043001, 2021.

[16] R. L. Workman and Others. Review of Particle Physics. *PTEP*, 2022:083C01, 2022.

[17] NEST Collaboration. Noble Element Simulation Technique (NEST). https://nest.physics.ucdavis.edu/, n.d.

[18] Babak Abi, Roberto Acciarri, Mario A Acero, Giorge Adamov, David Adams, Marco Adinolfi, Zubayer Ahmad, Jhanzeb Ahmed, Tyler Alion, S Alonso Monsalve, et al. Volume I. introduction to DUNE. *Journal of Instrumentation*, 15(08):T08008–T08008, 2020.

[19] DUNE collaboration, A Abed Abud, B Abi, R Acciarri, MA Acero, MR Adames, G Adamov, D Adams, M Adinolfi, A Aduszkiewicz, et al. Design, construction and operation of the ProtoDUNE-SP Liquid Argon TPC. *Journal of Instrumentation*, 17(01):P01005, 2022.

[20] B Abi, A Abed Abud, R Acciarri, MA Acero, G Adamov, M Adamowski, D Adams, P Adrien, M Adinolfi, Z Ahmad, et al. First results on ProtoDUNE-SP liquid argon time

projection chamber performance from a beam test at the CERN Neutrino Platform. *Journal of instrumentation*, 15(12):P12004–P12004, 2020.

[21] DUNE collaboration et al. Scintillation light detection in the 6-m drift-length ProtoDUNE Dual Phase liquid argon TPC. 2022.

[22] Dune Collaboration et al. The DUNE far detector vertical drift technology Technical design report. *Journal of Instrumentation*, 19(8):T08004, 2024.

[23] A. Abed Abud et al. Performance of a modular ton-scale pixel-readout liquid argon time projection chamber, 2024.

[24] David Griffiths. *Introduction to Elementary Particles*. John Wiley & Sons, 2020.

[25] Clyde L Cowan Jr, Frederick Reines, FB Harrison, HW Kruse, and AD McGuire. Detection of the free neutrino: a confirmation. *Science*, 124(3212):103–104, 1956.

[26] Gianfranco Bertone and Dan Hooper. History of dark matter. *Rev. Mod. Phys.*, 90:045002, Oct 2018.

[27] Laurent Canetti, Marco Drewes, and Mikhail Shaposhnikov. Matter and Antimatter in the Universe. *New Journal of Physics*, 14(9):095012, 2012.

[28] Michael Dine and Alexander Kusenko. Origin of the matter-antimatter asymmetry. *Rev. Mod. Phys.*, 76:1–30, Dec 2003.

[29] Raymond Davis. Solar Neutrinos. II. Experimental. *Phys. Rev. Lett.*, 12:303–305, Mar 1964.

[30] J.N. Abdurashitov et al. Results from sage (the russian-american gallium solar neutrino experiment). *Physics Letters B*, 328(1):234–248, 1994.

[31] Gallex Collaboration et al. Solar neutrinos observed by GALLEX at Gran Sasso. *Phys. Lett. B*, 285:376, 1992.

[32] A. Bellerive, J.R. Klein, A.B. McDonald, A.J. Noble, and A.W.P. Poon. The sudbury neutrino observatory. *Nuclear Physics B*, 908:30–51, 2016. Neutrino Oscillations: Celebrating the Nobel Prize in Physics 2015.

[33] Jun Cao and Kam-Biu Luk. An overview of the Daya Bay reactor neutrino experiment. *Nuclear Physics B*, 908:62–73, 2016.

[34] Andreas Piepke. KamLAND: A reactor neutrino experiment testing the solar neutrino anomaly. *Nuclear Physics B-Proceedings Supplements*, 91(1-3):99–104, 2001.

[35] JK Ahn, Reno Collaboration, et al. RENO: An Experiment for Neutrino Oscillation Parameter theta_13 Using Reactor Neutrinos at Yonggwang. *arXiv preprint arXiv:1003.1391*, 2010.

[36] Y Abe, JC Dos Anjos, JC Barriere, E Baussan, I Bekman, M Bergevin, TJC Bezerra, L Bezrukov, E Blucher, C Buck, et al. Improved measurements of the neutrino mixing angle θ 13 with the Double Chooz detector. *Journal of High Energy Physics*, 2014(10):1–44, 2014.

[37] Mark G Aartsen, M Ackermann, J Adams, JA Aguilar, M Ahlers, M Ahrens, D Altmann, K Andeen, T Anderson, I Ansseau, et al. The IceCube Neutrino Observatory: instrumentation and online systems. *Journal of Instrumentation*, 12(03):P03012–P03012, 2017.

[38] K Abe, Y Hayato, T Iida, M Ikeda, C Ishihara, K Iyogi, J Kameda, K Kobayashi, Y Koshio, Y Kozuma, et al. Solar neutrino results in Super-Kamiokande-III. *Physical Review D—Particles, Fields, Gravitation, and Cosmology*, 83(5):052010, 2011.

[39] S Fukuda, Y Fukuda, Takashi Hayakawa, E Ichihara, Masaki Ishitsuka, Y Itow, T Kajita, J Kameda, K Kaneyuki, Shunsuke Kasuga, et al. The super-kamiokande detector. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 501(2-3):418–462, 2003.

[40] JJ Evans. The MINOS experiment: results and prospects. *Advances in High Energy Physics*, 2013(1):182537, 2013.

[41] RB Patterson. The NOvA experiment: status and outlook. *Nuclear Physics B-Proceedings Supplements*, 235:151–157, 2013.

[42] F J Hasert et al. Observation of neutrino-like interactions without muon or electron in the GARGAMELLE neutrino experiment. *Physics Letters B*, 46(1):138–144, 1973.

[43] Q R Ahmad et al. Direct evidence for neutrino flavor transformation from neutral-current interactions in the Sudbury Neutrino Observatory. *Physical Review Letters*, 89(1):011301, 2002.

[44] J. Chadwich. A Preliminary Investigation of the Intensity Distribution in the Beta-Ray Spectra of Radium B and C. *Proceedings of the Cambridge Philosophical Society*, 21:274–280, 1922.

[45] Fred L Wilson. Fermi's Theory of Beta Decay. *American Journal of Physics*, 36(12):1150–1160, 1968.

[46] B Pontecorvo. Inverse Beta Process. *Chalk River Laboratory Report PD-205*, 1946.

[47] B. Pontecorvo. Electron and Muon Neutrinos. *Zh. Eksp. Teor. Fiz.*, 37:1751–1757, 1959.

[48] M. Schwartz. Feasibility of Using High-Energy Neutrinos to Study the Weak Interactions. *Phys. Rev. Lett.*, 4:306–307, Mar 1960.

[49] Gaillard Danby, Jean Maurice Gaillard, Konstantin Goulianos, Leon M Lederman, Nariman Mistry, Melvin Schwartz, and Jack Steinberger. Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos. *Physical Review Letters*, 9(1):36, 1962.

[50] J.G. Rutherglen. 1 - SPARK CHAMBERS. In O.R. FRISCH, editor, *Progress in Nuclear Physics*, pages 1–26. Pergamon, 1964.

[51] Yung-Su Tsai. Decay correlations of heavy leptons in e++ e-→ l++ l-. *Physical Review D*, 4(9):2821, 1971.

[52] Martin L Perl. The discovery of the tau lepton and the changes in elementary-particle physics in forty years. *Physics in Perspective*, 6(4):401–427, 2004.

[53] K Kodama, N Ushida, C Andreopoulos, N Saoulidou, G Tzanakos, P Yager, B Baller, D Boehnlein, W Freeman, B Lundberg, et al. Final tau-neutrino results from the donut experiment. *Physical Review D—Particles, Fields, Gravitation, and Cosmology*, 78(5):052002, 2008.

[54] ATLAS Collaboration. Observation of a New Particle in the Search for the Standard Model Higgs Boson with the ATLAS Detector at the LHC. *Phys. Lett. B*, 716:1–29, 2012.

[55] CMS Collaboration. Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC. *Phys. Lett. B*, 716:30–61, 2012.

[56] Makoto Kobayashi and Toshihide Maskawa. CP-violation in the renormalizable theory of weak interaction. *Progress of theoretical physics*, 49(2):652–657, 1973.

[57] Raymond Davis, Don S. Harmer, and Kenneth C. Hoffman. Search for Neutrinos from the Sun. *Phys. Rev. Lett.*, 20:1205–1209, May 1968.

[58] Yoshiyuki Fukuda, T Hayakawa, E Ichihara, K Inoue, K Ishihara, Hirokazu Ishino, Y Itow, T Kajita, J Kameda, S Kasuga, et al. Evidence for oscillation of atmospheric neutrinos. *Physical review letters*, 81(8):1562, 1998.

[59] John N Bahcall, Raymond Davis, Peter Parker, Alexei Smirnov, and Roger Ulrich. Measurement of the rate of ve+ d→ p+ p+ e-interactions produced by 8b solar neutrinos at the sudbury neutrino observatory. In *Solar Neutrinos*, pages 451–456. CRC Press, 2018.

[60] Raymond Davis, Don S. Harmer, and Kenneth C. Hoffman. Search for Neutrinos from the Sun. *Phys. Rev. Lett.*, 20:1205–1209, May 1968.

[61] Nobel Prize Organization. Raymond Davis Jr. – Photo gallery. NobelPrize.org. Nobel Prize Outreach 2026., 2026.

[62] V. A. Kuz'min. Detection of Solar Neutrinos by Means of the Ga71(u, e-Ge71 Reaction. *Zh. Eksp. Teor. Fiz.*, 49:1532–1534, 1965.

[63] Sage Collaboration et al. Measurement of the solar neutrino capture rate by sage and implications for neutrino oscillations in vacuum. *arXiv preprint astro-ph/9907131*, 1999.

[64] M. Cribier et al. Results of the whole gallex experiment. *Nuclear Physics B - Proceedings Supplements*, 70(1):284–291, 1999. Proceedings of the Fifth International Workshop on topics in Astroparticle and Underground Physics.

[65] A. Abud et al. Deep Underground Neutrino Experiment (DUNE) Near Detector Conceptual Design Report. *Instruments*, 5(4), 2021.

[66] Fulvio Mauri. Icarus: A second-generation proton decay experiment and neutrino observatory at the gran sasso laboratory. In Hans Volker Klapdor-Kleingrothaus, editor, *Beyond the Desert 2003*, pages 941–948, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg.

[67] Roberto Acciarri, C Adams, R An, A Aparicio, S Aponte, J Asaadi, M Auger, N Ayoub, L Bagby, B Baller, et al. Design and construction of the microboone detector. *Journal of Instrumentation*, 12(02):P02017, 2017.

[68] Valentina De Romeri, Enrique Fernandez-Martinez, and Michel Sorel. Neutrino oscillations at dune with improved energy reconstruction. *Journal of High Energy Physics*, 2016(9):1–25, 2016.

[69] Masataka Fukugita and Tsutomu Yanagida. Barygenesis without grand unification. *Physics Letters B*, 174(1):45–47, 1986.

[70] A. D. Sakharov. Violation of CP Invariance, C Asymmetry, and Baryon Asymmetry of the Universe. *Pisma Zh. Eksp. Teor. Fiz.*, 5:32–35, 1967. [JETP Lett. 5, 24 (1967)].

[71] Pran Nath and Pavel Fileviez Perez. Proton stability in grand unified theories, in strings and in branes. *Physics reports*, 441(5-6):191–317, 2007.

[72] KS Babu, E Kearns, U Al-Binni, S Banerjee, DV Baxter, Z Berezhiani, M Bergevin, S Bhattacharya, S Brice, and R Brock. Working group report: baryon number violation. Technical report, 2013.

[73] K. Abe, Y. Haga, Y. Hayato, M. Ikeda, K. Iyogi, J. Kameda, Y. Kishimoto, M. Miura, S. Moriyama, et al. Search for proton decay via $p \rightarrow e^+\pi^0$ and $p \rightarrow \mu^+\pi^0$ in 0.31 megaton $\cdot$ years exposure of the Super-Kamiokande water Cherenkov detector. *Phys. Rev. D*, 95:012004, Jan 2017.

[74] KS Babu, E Kearns, U Al-Binni, S Banerjee, DV Baxter, Z Berezhiani, M Bergevin, S Bhattacharya, S Brice, R Brock, et al. Baryon Number Violation. *arXiv preprint arXiv:1311.5285*, 2013.

[75] B. Abi et al. Supernova neutrino burst detection with the Deep Underground Neutrino Experiment. *The European Physical Journal C*, 81(5):423, 2021.

[76] Adam Burrows and James M Lattimer. Neutrinos from sn 1987a. *Astrophysical Journal, Part 2-Letters to the Editor (ISSN 0004-637X), vol. 318, July 15, 1987, p. L63-L68.*, 318:L63–L68, 1987.

[77] James M Lattimer and A Yahil. Analysis of the neutrino events from supernova 1987a. *Astrophysical Journal, Part 1 (ISSN 0004-637X), vol. 340, May 1, 1989, p. 426-434.*, 340:426–434, 1989.

[78] A. Abed Abud et al. Supernova pointing capabilities of DUNE. *Physical Review D*, 111:092006, 2025.

[79] E Morikawa, R Reininger, P Gürtler, V Saile, and P Laporte. Argon, krypton, and xenon excimer luminescence: From the dilute gas to the condensed phase. *The Journal of chemical physics*, 91(3):1469–1477, 1989.

[80] Shinzou Kubota, A Nakamoto, Tan Takahashi, T Hamada, Eido Shibamura, Mitsuhiro Miyajima, Kimiaki Masuda, and Tadayoshi Doke. Recombination luminescence in liquid argon and in liquid xenon. *Physical Review B*, 17(6):2762, 1978.

[81] Gabriela Araujo. Wavelength shifting and photon detection of scintillation light from liquid argon. 03 2019.

[82] R Acciarri, M Antonello, B Baibussinov, M Baldo-Ceolin, P Benetti, F Calaprice, E Calligarich, M Cambiaghi, N Canci, F Carbonara, et al. Oxygen contamination in liquid argon: combined effects on ionization electron charge and scintillation light. *Journal of Instrumentation*, 5(05):P05003, 2010.

[83] T Doke. Fundamental properties of liquid argon, krypton and xenon as radiation detector media. In *Experimental Techniques in High-Energy Nuclear and Particle Physics*, pages 537–577. World Scientific, 1991.

[84] Carlo Rubbia. The liquid-argon time projection chamber: a new concept for neutrino detectors. Technical report, 1977.

[85] Pedro AN Machado, Ornella Palamara, and David W Schmitz. The short-baseline neutrino program at fermilab. *Annual Review of Nuclear and Particle Science*, 69(1):363–387, 2019.

[86] B Baller, C Bromberg, N Buchanan, F Cavanna, H Chen, E Church, V Gehman, H Greenlee, E Guardincerri, B Jones, et al. Liquid argon time projection chamber research and development in the united states. *Journal of Instrumentation*, 9(05):T05005, 2014.

[87] Arkadiy Klebaner, Cristian Boffo, Saravan Kumar Chandrasekaran, Donato Passarelli, and Genfa Wu. Proton Improvement Plan – II: Overview of Progress in the Construction. *JACoW*, SRF2021, 5 2022.

[88] C Amsler, M Weber, M Luthi, F Noto, A Ereditato, Y Arbelo, M Bishai, F Barbato, D Smargianaki, F Lanni, et al. ArgonCube: a novel, fully-modular approach for the realization of large-mass liquid argon TPC neutrino detectors. Technical report, 2015.

[89] D.A. Dwyer, M. Garcia-Sciveres, D. Gnani, C. Grace, S. Kohn, M. Kramer, A. Krieger, C.J. Lin, K.B. Luk, P. Madigan, C. Marshall, H. Steiner, and T. Stezelberger. LArPix: Demonstration of low-power 3D pixelated charge readout for liquid argon time projection chambers. *Journal of Instrumentation*, 13(10):P10007, oct 2018.

[90] S. Bertolucci. KLOE at DANE. *Nuclear Physics B - Proceedings Supplements*, 37(1):43–50, 1994. Proceedings of the Third KEK Topical Conference on CP Violation, its Implications to Particle Physics and Cosmology.

[91] L Aliaga, L Bagby, B Baldin, A Baumbaugh, A Bodek, R Bradford, WK Brooks, D Boehnlein, S Boyd, H Budd, et al. Design, calibration, and performance of the MINERvA detector. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 743:130–159, 2014.

[92] Roman Berner, Yifan Chen, Antonio Ereditato, Patrick P. Koller, Igor Kreslo, David Lorca, Thomas Mettler, Ting Miao, Francesco Piastra, James R. Sinclair, and Michael S. Weber. First operation of a resistive shell liquid argon time projection chamber: A new approach to electric-field shaping. *Instruments*, 3(2), 2019.

[93] Martin Auger, Yifan Chen, Antonio Ereditato, Damian Goeldi, Igor Kreslo, David Lorca, Matthias Luethi, Thomas Mettler, James Sinclair, and Michele Weber. ArCLight—A compact dielectric large-area photon detector. *Instruments*, 2(1):3, 2018.

[94] N Anfimov, R Berner, I Butorov, A Chetverikov, D Fedoseev, B Gromov, D Korablev, I Kreslo, K Kuznetsova, A Olshevskiy, et al. Development of the light collection module for the liquid argon time projection chamber (LArTPC). *Journal of Instrumentation*, 15(07):C07022, 2020.

[95] AA Machado, E Segreto, D Warner, A Fauth, B Gelli, R Maximo, A Pissolatti, L Paulucci, and F Marinho. The x-arapuca: An improvement of the arapuca device. *Journal of instrumentation*, 13(04):C04026, 2018.

[96] G. Botogoske and on behalf of the DUNE collaboration. Dune photon detection system. *Journal of Instrumentation*, 20(06):C06034, jun 2025.

[97] A.A. Machado and E. Segreto. Arapuca a new device for liquid argon scintillation light detection. *Journal of Instrumentation*, 11(02):C02004, feb 2016.

[98] M.A. Arroyave, B. Behera, F. Cavanna, A. Feld, F. Guo, A. Heindel, C.K. Jung, K. Koch, D. Leon Silverio, D.A. Martinez Caicedo, C. McGrew, A. Paudel, W. Pellico, R. Rivera, J. Rodríguez Rondon, S. Sacerdoti, P. Shanahan, W. Shi, D. Torres Muñoz, D. Totani, C. Uy, C. Vermeulen, and H. Vieira de Souza. Characterization and novel application of power over fiber for electronics in a harsh environment. *Journal of Instrumentation*, 19(10):P10019, oct 2024.

[99] H.V. Souza and on behalf of the DUNE collaboration. Light detection with power and signal transmission over fiber. *Journal of Instrumentation*, 18(02):C02029, feb 2023.

[100] Salvatore Amoruso, M Antonello, P Aprili, F Arneodo, A Badertscher, B Baiboussinov, M Baldo Ceolin, G Battistoni, B Bekman, P Benetti, et al. Study of electron recombination in liquid argon with the ICARUS TPC. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 523(3):275–286, 2004.

[101] R Acciarri, C Adams, J Asaadi, B Baller, T Bolton, C Bromberg, F Cavanna, E Church, D Edmunds, A Ereditato, et al. A study of electron recombination using highly ionizing particles in the ArgoNeuT Liquid Argon TPC. *Journal of Instrumentation*, 8(08):P08005–P08005, 2013.

[102] M Szydagis, N Barry, K Kazkaz, J Mock, D Stolp, M Sweany, M Tripathi, S Uvarov, N Walsh, and M Woods. NEST: a comprehensive model for scintillation yield in liquid xenon. *Journal of Instrumentation*, 6(10):P10002–P10002, 2011.

[103] Matthew Szydagis, Grant A Block, Collin Farquhar, Alexander J Flesher, Ekaterina S Kozlova, Cecilia Levy, Emily A Mangus, Michael Mooney, Justin Mueller, Gregory RC Rischbieter, et al. A review of basic energy reconstruction techniques in liquid xenon and argon detectors for dark matter and neutrino physics using NEST. *Instruments*, 5(1):13, 2021.

[104] Steven D Biller, Daniel Kabat, Richard C Allen, Gerhard Bühler, and Peter J Doe. Effects of oxygen and nitrogen on drifting electrons in a liquid argon tpc. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 276(1-2):144–150, 1989.

[105] R Andrews, W Jaskierny, H Jöstlein, C Kendziora, S Pordes, and T Tope. A system to test the effect of materials on electron drift lifetime in liquid argon and the effect of water. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 608(2):251–258, 2009.

[106] George Bakale, Ulrich Sowada, and Werner F Schmidt. Effect of an electric field on electron attachment to sulfur hexafluoride, nitrous oxide, and molecular oxygen in liquid argon and xenon. *The Journal of Physical Chemistry*, 80(23):2556–2559, 1976.

[107] Yichen Li, C Bromberg, M Diwan, S Kettell, S Martynenko, X Qian, V Paolone, J Stewart, C Thorn, and C Zhang. Parameterization of electron attachment rate constants for impurities in LArTPC detectors. *Journal of Instrumentation*, 17(11):T11007, 2022.

[108] Mitchell M. Mote. Measurement of space charge effects and energy calibration in protodune-sp. volume NuFact2021. Fermi National Accelerator Laboratory (FNAL), Batavia, IL (United States), 03 2022.

[109] P Abratenko, M Alrashed, R An, J Anthony, J Asaadi, A Ashkenazi, S Balasubramanian, B Baller, C Barnes, G Barr, et al. Measurement of space charge effects in the MicroBooNE LArTPC using cosmic muons. *Journal of Instrumentation*, 15(12):P12037–P12037, 2020.

[110] P Benetti, F Calaprice, E Calligarich, Mario Cambiaghi, F Carbonara, F Cavanna, AG Cocco, F Di Pompeo, N Ferrari, Giuliana Fiorillo, et al. Measurement of the specific activity of 39Ar in natural argon. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 574(1):83–88, 2007.

[111] P Abratenko, O Alterkait, D Andrade Aldana, L Arellano, J Asaadi, A Ashkenazi, S Balasubramanian, B Baller, A Barnard, G Barr, et al. Demonstration of new mev-scale capabilities in large neutrino lartpcs using ambient radiogenic and cosmogenic activity in microboone. *Physical Review D*, 111(3):032005, 2025.

[112] Florian Jörg, Shengchao Li, Jochen Schreiner, Hardy Simgen, and Rafael F Lang. Characterization of a 220Rn source for low-energy electronic recoil calibration of the XENONnT detector. *Journal of Instrumentation*, 18(11):P11009, 2023.

[113] E. Aprile et al. Results from a calibration of XENON100 using a ^{220}Rn source. *Phys. Rev. D*, 95(7):072008, 2017.

[114] P Abratenko, O Alterkait, D Andrade Aldana, L Arellano, J Asaadi, A Ashkenazi, S Balasubramanian, B Baller, G Barr, D Barrow, et al. Measurement of ambient radon progeny decay rates and energy spectra in liquid argon using the MicroBooNE detector. *Physical Review D*, 109(5):052007, 2024.

[115] DUNE Collaboration. Operation of a Modular 3D-Pixelated Liquid Argon Time-Projection Chamber in a Neutrino Beam. *Instruments*, 10(1), 2026.

[116] Martin Ester, Hans-Peter Kriegel, Jorg Sander, Xiaowei Xu, et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *kdd*, volume 96, pages 226–231, 1996.

[117] Glenn F. Knoll. *Radiation Detection and Measurement*. John Wiley & Sons, Hoboken, NJ, 4 edition, 2010.

[118] Adam Abed Abud, Babak Abi, Roberto Acciarri, Mario A Acero, Márcio Rostirolla Adames, George Adamov, Mark Adamowski, David Adams, Marco Adinolfi, Cris Adriano, et al. Highly-parallelized simulation of a pixelated LArTPC on a GPU. *Journal of instrumentation*, 18(04):P04034, 2023.

[119] Stanisław Węglarczyk. Kernel density estimation and its application. In *ITM web of conferences*, volume 23, page 00037. EDP Sciences, 2018.

[120] P Abratenko, M Alrashed, R An, J Anthony, J Asaadi, A Ashkenazi, S Balasubramanian, B Baller, C Barnes, G Barr, et al. Measurement of space charge effects in the microboone lartpc using cosmic muons. *Journal of Instrumentation*, 15(12):P12037, 2020.

[121] S Amoruso et al. Analysis of the liquid argon purity in the ICARUS T600 TPC. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 516(1):68–79, 2004.

[122] Nippon Sanso Matheson, Inc. Nippon sanso matheson — the gas professionals, 2026. Accessed: 2026-07-04.

[123] D. M. Webber et al. Measurement of the positive muon lifetime and determination of the fermi constant to part-per-million precision. *Physical Review Letters*, 106:041803, 2011.

[124] Laura Manenti, Linda Cremonesi, Francesco Arneodo, Anastasia Basharina-Freshville, Mario Campanelli, Anna Holin, Ryan Nichol, and Ruben Saakyan. Performance of different photocathode materials in a liquid argon purity monitor. *Journal of Instrumentation*, 15(09):P09003, 2020.

[125] Samuel J Fogarty. Development of a liquid argon purity monitoring system. Master's thesis, Colorado State University, 2023.

[126] O Bunemann, TE Cranshaw, and JA Harvey. Design of grid ionization chambers. *Canadian Journal of Research*, 27(5):191–206, 1949.

[127] W. Maneschg, L. Baudis, R. Dressler, K. Eberhardt, R. Eichler, H. Keller, R. Lackner, B. Praast, R. Santorelli, J. Schreiner, M. Tarka, B. Wiegel, and A. Zimbal. Production and characterization of a custom-made 228th source with reduced neutron source strength for the borexino experiment. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 680:161–167, 2012.

[128] Eckert & Ziegler Isotope Products. *Eckert & Ziegler AF-228-A2 Thorium-228 (228Th) Alpha Calibration Source*. Eckert & Ziegler AG, Berlin, Germany, 2026.

[129] Pylon Electronics Inc. *Pylon Electronics TH-1025 Thoron (228Th) Flow-Through Radioactive Source*. Pylon Electronics, Ottawa, Ontario, Canada, 2016.

[130] Harry Bateman. The solution of a system of differential equations occurring in the theory of radio-active transformations. *Proceedings of the Cambridge Philosophical Society*, 15:423–427, 1910.

[131] P Abratenko, J Anthony, L Arellano, J Asaadi, A Ashkenazi, S Balasubramanian, B Baller, C Barnes, G Barr, J Barrow, et al. Observation of radon mitigation in microboone by a liquid argon filtration system. *Journal of instrumentation*, 17(11):P11022, 2022.

[132] W. Walkowiak. Drift velocity of free electrons in liquid argon. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 449(1):288–294, 2000.

[133] National Institute of Standards and Technology. NIST Chemistry WebBook, Saturation Properties for Argon — Temperature Increments, 2026. Accessed: 2026-05-12.