

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Measurement of Muon Capture on Argon with a Pixelated Liquid Argon Time Projection Chamber

Permalink

<https://escholarship.org/uc/item/7k13v4b0>

Author

Madigan, Peter

Publication Date

2023

Peer reviewed|Thesis/dissertation

Measurement of Muon Capture on Argon with a Pixelated Liquid Argon Time Projection Chamber

by

Peter Sean Madigan

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Physics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Kam-Biu Luk, Chair

Dr. Daniel Dwyer

Professor Daniel McKinsey

Professor Kai Vetter

Fall 2023

Measurement of Muon Capture on Argon with a Pixelated Liquid Argon Time Projection Chamber

Copyright 2023
by
Peter Sean Madigan

Abstract

Measurement of Muon Capture on Argon with a Pixelated Liquid Argon Time Projection Chamber

by

Peter Sean Madigan

Doctor of Philosophy in Physics

University of California, Berkeley

Professor Kam-Biu Luk, Chair

Studies of fundamental interactions have played a principle role in the development of our understanding of particle physics, but the Standard Model (SM) remains incomplete. The phenomenon of neutrino oscillation lies outside of the prescriptions of the SM, driving current experimental efforts. The Deep Underground Neutrino Experiment (DUNE) is being built to provide a comprehensive study of neutrino oscillation and to search for beyond-the-Standard-Model (BSM) physics. To accomplish this, DUNE utilizes a massive-scale far detector, a sophisticated near detector, and a high-intensity, broad-spectrum neutrino beam. The near detector adopts a new design of liquid argon time projection chamber (LArTPC) with pixelated charge readout and a modular structure to resolve events in the high-intensity neutrino beam. A prototype of this detector design was successfully operated at the University of Bern, Switzerland in 2021. From the cosmic-ray data collected with this prototype, a study of muon capture on argon has been performed, resulting in the first measurement of the Huff factor for argon ($R_H = 1.29 \pm 0.15$) and measurements of the muon capture rate ($\lambda_c = 1.53 \pm 0.19 \mu\text{s}^{-1}$) and muon disappearance rate ($\lambda_d = 2.11 \pm 0.24 \mu\text{s}^{-1}$) in argon. These results can help to reduce the uncertainty in neutrino-nucleus cross-section modelling at low Q^2 .

Contents

Contents	i
List of Figures	iv
List of Tables	xvi
1 Introduction	1
1.1 A brief history of weak interactions	2
1.1.1 Fermi’s theory	2
1.1.2 Muon decay	3
1.1.3 V-A interaction	5
1.2 Muon capture and muon decay	7
1.2.1 Electroweak Interaction	8
1.2.2 Muon properties	12
1.2.3 Muon decay	13
1.2.4 Muon capture	15
1.2.5 Muon capture on nuclei	17
1.2.6 Muon decay in orbit	17
1.3 Relationship to neutrino experiments	20
1.3.1 Deep Underground Neutrino Experiment (DUNE)	20
1.3.2 Neutrino cross-sections	24
1.4 Existing measurements	26
2 Detector technology	28
2.1 Tracking detectors	28
2.2 Liquid argon time projection chambers (LArTPCs)	31
2.2.1 Charged-particle signals	32
2.2.2 LArTPC design	34
2.3 ArgonCube concept	35
2.3.1 DUNE LAr-ND	35
2.3.2 ArgonCube	36
2.3.3 LArPix	37

2.3.4	ArcLight/LCM	40
2.3.5	Field cage	41
3	Simulation	44
3.1	Cosmic-ray event generation	44
3.2	Particle propagation and interaction	47
3.3	larnd-sim	49
3.3.1	Scintillation and ionization model	51
3.3.2	LRS response model	53
3.3.3	CRS response model	55
4	Data collection	61
4.1	Module 0 prototypes	61
4.1.1	The SingleCube demonstrator	62
4.1.2	Module 0 high-voltage test	63
4.1.3	Module-0 detector	65
4.1.4	Trigger and timing configuration	67
4.1.5	Run 1	68
4.1.6	Run 2	75
4.2	Data samples	76
4.3	Module 0 simulation	78
4.3.1	CORSIKA parameters	79
4.3.2	EdepSim parameters	79
4.4	Simulation tuning	81
4.4.1	Noise overlay	82
4.4.2	Threshold tuning	84
4.4.3	Light impulse model	87
4.4.4	Scintillation model	87
4.4.5	Light LUT and PDE	88
4.4.6	Light noise overlay	90
4.5	Simulation sample	91
5	Data processing and analysis	93
5.1	Event definitions	93
5.1.1	Charge-event definition	94
5.1.2	Light-event definition	97
5.1.3	Light- and charge-event matching	98
5.2	Signal processing	101
5.2.1	Charge readout	101
5.2.2	Light readout	112
5.3	Muon reconstruction	118
5.3.1	Stopping-muon trajectory reconstruction	119

5.3.2	Stopping-muon profiled dQ/dx reconstruction	123
5.4	Muon decay reconstruction	129
5.4.1	Muon decay classification	129
5.4.2	Decay time reconstruction	131
6	Capture rate measurement	135
6.1	Event selection	136
6.1.1	Stopping muons	136
6.1.2	Muon decay	140
6.1.3	Event selection summary	141
6.2	Fit	146
6.3	Systematic effects	152
6.4	Results	158
7	Conclusions	164
7.1	Summary of results	164
7.1.1	Global comparison	164
7.2	Future work	168
	Bibliography	170
A	Systematics summary	187
B	Event gallery	190
B.1	Stopping muon sample	191
B.2	Muon-decay sample	197
B.3	CORSIKA muon sample	200
B.4	CORSIKA muon-decay sample	206

List of Figures

1.1	Beta decay spectrum of ^{226}Ra , as measured by Wooster and Ellis [1].	2
1.2	First observational evidence of the mesotron decaying to an electron [6]. A mesotron is seen entering the cloud chamber at location A and decays to an electron at location F, which exits the chamber at location G.	4
1.3	First observation of a pion decaying into a muon [16]. An incoming short pion track (m_1) is observed along with the subsequent muon track (m_2) exiting the emulsion. Arrows indicate track deflections of $> 2^\circ$	5
1.4	Shape of electron energy spectrum from μ -decay under different assumptions for the underlying interaction, as calculated by Michel [17]. Curve B most closely resembles modern measurements.	6
1.5	Feynman diagram of first-order contribution to muon decay in the SM. An electron is emitted from a decaying negative muon, along with an anti-electron neutrino and a muon neutrino.	13
1.6	Comparison of the Goulard-Primakoff model with experimental results on muon capture. This figure was taken from [56] and the parameters used for the comparison are listed in the text.	18
1.7	Modified decay electron spectrum for muon-Fe decay-in-orbit (DIO) under different assumptions for the muon decay coupling (S = scalar, V = vector, T = tensor). The vacuum decay spectrum (VF) is shown for reference. Units of energy are $(m_\mu + m_e)/m_\mu$	19
1.8	Summary of neutrino cross-section measurements and predictions for the three interaction channels described in the text, provided by the NUANCE neutrino interaction generator. Figure is taken from Ref. [88].	25
2.1	Example neutrino interaction as imaged by the ICARUS LArTPC [98].	28
2.2	Characteristic energy loss per unit length for different particle species, taken from Ref. [99].	29
2.3	Comparison of liquid (black, thick) and gaseous (red, thin) argon scintillation spectra. The dominant dimer decay occurs at ~ 128 nm. Taken from Ref. [112]	33
2.4	Diagram of LArPix front-end and self-triggering logic [144].	38

2.5	Minature LArPix-instrumented anode showing a LArPix-v2 application-specific integrated circuit (ASIC) (left) mounted onto a multi-layer printed circuit board (PCB) with charge sensitive pixels on the opposite side (right).	38
2.6	Top view of a Pixel Array Control, Message, and Network (PACMAN) card. The high-density connector on the left of the board interfaces directly to the PCB vacuum-feedthrough. The ethernet, direct current (DC) power supply, pulse-per-second (PPS), and trigger signal interface via the connectors on the right of the board.	40
2.7	(Left) Diagram of the ArcLight operating principle: 1) a UV photon is down-shifted to ≈ 400 nm, 2) at this wavelength, the photon passes through the dichroic filter, 3) the 400 nm photon is further down-shifted to ≈ 500 nm, 4) and 5) the green photon is trapped within the wavelength shifting (WLS) medium to be read out by a silicon photo-multiplier (SiPM) at 6). (Right) Diagram of the light collection module (LCM) operating principle: 1) a UV photon is down-shifted to ≈ 400 nm, 2) the photon is further down-shifted to ≈ 500 nm, 3) the photon is trapped within the WLS fiber by total internal reflection, and 4) photon is readout by a SiPM at one end of the fiber. Taken from [149].	41
2.8	Partially-instrumented anode-plane assembly with 1 ArcLight panel (left) beside 3 LCMs (right) used in the Module-0 high-voltage (HV) test, described in more detail in Sec. 4.1	42
2.9	Resistive field cage used in the Module 0 tests. The HV cathode voltage is applied within the circular opening at the top of the Module (shown to the left in this photo).	43
2.10	Resistive field cage prototype produced by the the Laboratory for High-Energy Physics (LHEP) group at the University of Bern, taken from Ref. [157].	43
3.1	Total muon flux at the surface of the Earth, taken from [165].	45
3.2	Muon charge ratio at the surface of the Earth, taken from [165].	46
3.3	East-west asymmetry in the muon charge ratio as observed by the WILLI detector [169].	47
3.4	(Left) Comparison of total calculation time for the recombination factor as a function of the number of track segments using graphics processing unit (GPU) and central processing unit (CPU) hardware. (Right) Comparison of induced current calculation time for GPU and CPU hardware. Both images are taken from Ref. [182].	51
3.5	The local-field inductive current model is calculated with a finite-element simulation, described in Ref. [200]. This figure shows the resulting drift paths (a) and the induced current (b) for various starting points of the test charge. Taken from Ref. [183].	56

3.6	Dipole approximation of the z -component of the weighting field described by Eq. 3.21, calculated with terms out to $n = 5$ and a cathode positioned at $z = 300$ mm. Arrows indicate the negative gradient direction. Under an assumption of uniform drift in $-\hat{z}$, in regions where $\partial W/\partial z > 0$ an electron will induce a negative current on the pixel, where in regions with $\partial W/\partial z < 0$ induce a positive current.	57
3.7	(left) Example fit to analog monitor response to impulse-like charge injection. (right) Fit results for three channels at room temperature and liquid nitrogen temperature. Saturation for large amplitude pulses modifies the response shape resulting in poor fit χ^2 assuming an exponential-only response. Impacted points (>500 mV ΔV_{out}) are excluded from the quoted averages.	59
3.8	Change in hit-charge distribution as a function of the self-trigger threshold (1 DAC ≈ 2 ke $^-$ eq.). Baseline thresholds are ≈ 4.5 ke $^-$ eq.	60
4.1	Rendering of the Module-0 detector, as constructed. Taken from Ref. [149]. . . .	62
4.2	(Left) View of the SingleCube drift volume and cathode prior to assembly, each dimension is 30 cm. (Right) View of the pixel tile, left, and ArcLight pane, right, used for the SingleCube test.	63
2.8	Partially-instrumented anode-plane assembly with 1 ArcLight panel (left) beside 3 LCMs (right) used in the Module-0 HV test.	64
4.3	(Left) Fully assembled Module 0. (Right) Module 0 being lowered into the cryostat.	65
4.4	Bottom view of Module 0 active volume, prior to complete assembly. Labeled regions are: A) pixel anode tile, B) ArcLight module and C) LCMs protruding into the drift region, and D) resistive cathode and field cage.	66
4.5	Diagram of timing synchronization system used for Module 0, taken from Ref. [207]. A PPS signal is generated by a global positioning system (GPS) unit (lower right) and distributed to both the light analog-to-digital converter (ADC)s and the PACMAN cards (not shown). If either ADC triggers, a BUSY signal is generated and passed to the other ADC and the PACMAN cards.	68
4.6	Accumulated events during Run 1 along with average trigger rates. Taken from Ref. [207].	69
4.7	Example of the results on one tile (time projection chamber (TPC) 1, tile 2) from a pedestal run with all pixels enabled (1 ADC ≈ 3.9 mV). Here x and y are defined relative to the anode tile, with the readout connectors falling along $x \approx -150$ mm and the sensitive area in the direction of $-(\hat{x} \times \hat{y})$	71
4.8	Example of the results on one tile (TPC 1, tile 2) from a pedestal run with only "good" pixels enabled (1 ADC ≈ 3.9 mV). Here x and y are defined relative to the anode tile, with the readout connectors falling along $x \approx -150$ mm and the sensitive area in the direction of $-(\hat{x} \times \hat{y})$	72
4.9	Disabled charge readout system (CRS) channels for Run 1 (top) and Run 2 (bottom).	73

4.10	Measured self-trigger thresholds in the two operating configurations of Module 0 (1 ke ⁻ ≈ 4 mV).	74
4.11	Light readout system gain including SiPM, cryo-preamp, variable gain amplifier (VGA), and ADC – (left) ArcLight SiPMs and (right) LCM SiPMs.	75
4.12	Cross-section of the GDML description of the Module 0 detector in the experimental hall at the University of Bern [196]. The central gray volume is the outside of the cylindrical cryostat that houses the Module 0 detector and includes a detailed model of the detector construction. This detail is not readily shown in this figure, instead refer to Fig. 4.1 for a close-up view. The blue volume surrounding the cryostat is a pit cut into the concrete on-site building and is filled with air. The brown rectangular volumes are simple representations of the building concrete.	78
4.13	(Left) Single-pixel reconstructed dQ/dx for minimum ionizing particle (MIP)-like tracks averaged over all pixels during a medium-threshold run, compared pre- and post-overlay. (Right) Fractional residual of single-pixel reconstructed dQ/dx pre- and post-overlay. Best fit values are quoted in the text.	83
4.14	Reconstructed track overlap with individual pixels normalized to the total analyzed track length. The distribution is approximately uniform below 4.4 mm.	84
4.15	2D illustration of the mechanism that causes the dx distribution shown in Fig. 4.14. On the left, straight-line track segments are drawn with uniform offsets in x (blue) and a square region is highlighted, indicating the overlapping segment of the track with a pixel (black). On the right, the resulting dx probability distribution function (PDF) is shown for tracks with the corresponding orientation. As the track crosses the pixel at a larger angle, clipping occurs in which the track intersects with the corner of the pixel. Because the relative position between the track and the pixel is random and uniform, the clipped-segment length is also random and uniform.	85
4.16	Self-trigger thresholds used in simulation, extracted from MIP-like tracks: (top) medium-threshold configuration, and (bottom) high-threshold configuration.	86
4.17	SiPM impulse response model extracted from an average of waveforms collected during an light emitting diode (LED) calibration run.	88
4.18	1D averaged profile of reweighted light look-up table (LUT) used in simulation for a SiPM in each light detector type. Each LUT bin is 4.4 cm × 4.7 cm × 3.8 cm. Drift axis is z and vertical axis is y	89
4.19	Relative light yield for MIP-like tracks observed in Module 0 compared to the Geant4-based LUT described in Ref. [196].	89
4.20	Example of a light waveform with and without noise overlay: (top) simulated waveform with a light signal present, (bottom) similar amplitude waveform from Module 0 data.	90
4.21	Muon kinetic energy and angular distribution created by the sampling approach defined in the text.	92

5.1	Self-trigger rate as a function of PPS timestamp along with the cut applied to remove SYNC-induced self-triggers.	97
5.2	Clock correction applied to light readout system (LRS) PPS timestamps. Figure (a) and (b) show the PPS timestamp of each digitizer compared to the event mean, along with the best fit used by the clock drift correction. Figure (c) shows the residual timestamp after the clock drift correction has been applied. Three bands are present due to the global trigger timing scheme.	99
5.3	Self-reported, uncalibrated SYNC timestamps from each PACMAN card during a typical run. The peaks are not aligned with the nominal 10^7 ticks/s due to different clock frequencies of the PACMAN cards.	100
5.4	(Left) Distribution of pedestal calibration values used for Module 0 data processing. (Right) Width of V_{ADC} distribution on each channel, extracted from pedestal data.	102
5.5	Comparison of injected vs. output pulse amplitudes using a dedicated test pulse circuit at room temperature and while submerged in liquid nitrogen. The reported ratio of -6.77 ± 0.04 corresponds to a gain of $4.52 \pm 0.03 \mu\text{V}/\text{e}$	104
5.6	Example of track reconstruction and $dx_{2\text{D}}$ (solid blue) calculation for the validation of the front-end gain calibration.	105
5.7	Median and width of the dQ/dx distribution for each pixel using reconstructed MIP-like tracks. An increased width and an additional 5% variation in the median is present in the data due to digital-analog crosstalk and channel-to-channel gain variations that are not simulated.	106
5.8	Electron lifetime during operations of Module 0. Only standard runs at 500 V/cm field are included, so gaps are present. Taken from Ref. [207].	107
5.9	(Left) Trigger timing of the integrated signal on a pixel for a track parallel to the anode ($\theta_{\text{drift}} = 0$). (Right) Bias in the reconstructed t_0 for tracks parallel to the anode plane using a weighted sum of the trigger time.	108
5.10	Two candidate muon-decay events demonstrating triggering on the inductive signal prior to the arrival of charge. Along the drift direction (z), hits of a smaller amplitude (purple) can be seen preceding the track. The increase in the dQ/dx near the end-point of a stopping muon enhances this effect.	108
5.11	Calculated pixel response to a track parallel to the anode plane, using the model described in Section 3.3.3. The base model includes no far-field effect. The curves of the far-field model compare the impact of increasing the proportion of the track that is included in the pixel-response calculation (1 pixel = 4.4 mm). Moving from left-to-right, the total drift time is decreased, demonstrating how the far-field effects change the signal shape and reduce the total integrated charge as the charge is deposited closer to the anode.	109

5.12	Improvement in track reconstruction using likelihood-based track merging. Efficiency is defined as the total fraction of hits produced by the primary particle that are associated to the same reconstructed track. Purity is defined as the fraction of hits in the reconstructed track that are produced by particles other than the primary particle.	111
5.13	Application of the dedicated 10-MHz pickup filter to a heavily impacted waveform. (a) shows the original waveform and the 100 ns periods taken from the pretrigger window, (b) shows the individual periods and their average, (c) demonstrates the extrapolation of the noise template – resulting in the noise-subtracted waveform in (d).	113
5.14	Comparison of the standard deviation of pretrigger samples before and after applying 10-MHz pickup filter.	113
5.15	Impulse response function extracted from a typical ArcLight (left) and LCM (right) SiPM.	115
5.16	Global trigger alignment of a single event in which ADC1 self-triggers, and ADC0 is triggered by the subsequent global trigger signal.	116
5.17	Low-amplitude LED signal distribution with short (left) and long (right) integration time used to extract the individual SiPM gains.	117
5.18	2D diagram illustrating the greedy tracking algorithm used to reconstruct muon decays. From the current trajectory step j , a cylindrical region (yellow) defined by the two reconstruction parameters Δ and α is projected forward along the current trajectory direction \vec{d}_j . Hits that fall within this region (red) are used to calculate the next trajectory step.	119
5.19	Stopping-muon candidate, reconstructed using the tracking algorithm described in the text. The lower images show a zoomed perspective of the same event near the stopping point. A significant gap in the data (≈ 5 cm) is successfully traversed by the tracking algorithm. The reconstructed stopping point using the dQ/dx profile (Sec. 5.3.2) is highlighted with a star.	121
5.20	View of the same event as in Fig. 5.19, zoomed into the region near the stopping point.	122
5.21	Diagram of dx used for the stopping-muon dQ/dx calculation at each trajectory point. The outer-most box (dashed grey) indicates the boundary for the trajectory point that is set by the hits associated with the step (light blue). The region highlighted in red indicates a portion of the track that crossed a disabled pixel and thus was not included in the dx . The interpolated length of the trajectory that is used for the dx is highlighted in dark blue.	123
5.22	Comparison of dQ/dx as a function of the reconstructed range for muon decay events. For reference, the naive Birks' continuous slowing down approximation (CSDA) model for dQ/dx is also shown, which was used for separation of through-going and stopping muons.	124
5.23	MIP-like classifier score using post-fit background contributions extrapolated from the stopping muon selection to pre-cut selection (described in Sec. 6.1.3).	126

5.24	Proton-like classifier score using post-fit background contributions extrapolated from the stopping-muon selection to the pre-cut selection (described in Sec. 6.1.3).	127
5.25	Residual vertex reconstruction error from stopping point reconstruction in each dimension (a-c) and the 3D residual (d). The residual in z is defined such that $\delta z < 0$ is biased closer to the anode and $\delta z > 0$ is biased closer to the cathode.	128
5.26	Michel likelihood score as a function of the three input parameters: x -axis is $\cos \theta$ with respect to the electron axis; y -axis is $\cos \theta$ with respect to the muon axis; and the plot grid shows the distance from the reconstructed stopping point, increasing from left-to-right, top-to-bottom.	130
5.27	Michel likelihood score as a function of the parallel and transverse directions to the muon axis. Plots from left to right show the change in the discriminator behavior as a function of the reconstructed electron direction.	131
5.28	Background likelihood score as a function of the two input parameters (left) and the parallel and transverse directions to the muon axis (right). Note that the background likelihood is defined such that a decreasing score indicates an increasing likelihood of being a background-induced hit.	131
5.29	principle component analysis (PCA)-derived templates used to subtract the prompt-light contribution from the waveform. An oscillation is present immediately after the prompt signal in the 0th PCA component due to an imperfect impulse model used for the Wiener deconvolution. The shape of the 0th component of the PCA decomposition is very close to the mean signal shape, the shape of the 1st component is dominated by leading-edge differences between waveforms that result in a horizontal shift with respect to the trigger edge, and the 2nd component starts to account for differences in the singlet-to-triplet ratio.	132
5.30	Example of PCA prompt-signal subtraction on a candidate muon-decay event. The left-most frame shows the summed waveform of the SiPMs on each detector (top) and the overall signal sum (bottom). The middle frame shows the PCA fit to the prompt signal. And the right-most frame shows the residual waveform after the prompt signal has been subtracted.	133
6.1	Pre-fit, pre-selection proton-likelihood score, with the simulation exhibiting a prominent peak of stopping-proton events not present in the Module 0 data.	138
6.2	Change in best-fit value for R_H as a function of fiducial-volume cut. The overall statistical uncertainty on R_H is shown with dashed lines. The relative change in the statistical error after modifying the fiducial cut is shown with the error bars ($\sigma_{\text{rel}}^2 = \sigma_{\text{nominal}}^2 - \sigma_{\text{change}}^2 $) [224].	153
6.3	Change in the shape of the decay-like nhit distributions ($(n_{\text{nom.}} - n_{\text{slow}})/n_{\text{nom.}}$) between the nominal muon decay sample and the slow-decay samples, where $n_{\text{norm.}}$ and n_{slow} are normalized to the number of events in each sample, respectively. The corresponding χ^2/ndf are 49.76/40, 35.45/37, and 27.65/38 for the Run 1 medium threshold, Run 1 high threshold, and Run 2 high threshold samples, respectively.	155

6.4	Post-fit muon-decay time distributions for each data sample. The overall χ^2/ndf for the fit, including the muon-decay nhit distributions, is 229.8/237.	159
6.5	Post-fit muon-decay nhit distributions for each data sample. The overall χ^2/ndf of the fit, including the muon-decay time distributions, is 229.8/237.	160
7.1	Summary of published experimental measurements [165, 226, 227] of the cosmic-ray muon charge ratio plotted against muon kinetic energy and compared against this work.	165
7.2	Comparison of the measured Huff factor in different nuclei, Refs. [64, 71], and theoretical calculations, Refs. [60, 63]. The results of this work using different disappearance rate constraints have been offset from $Z = 18$ for clarity.	166
7.3	Global summary of measured μ^- disappearance rate for argon. The results from this work are compared to the experimental measurements described in Refs. [93, 94, 96, 97].	166
7.4	Global summary of the measured μ^- capture rate on argon. The results from this work are compared to the experimental measurements described in Refs. [93, 94, 96, 97] and the Goulard-Primakoff formula [56]. A Huff factor of 1.04 ± 0.05 was assumed in interpreting the previous measurements of the μ^- disappearance rate as a μ^- capture rate. The error bars were increased accordingly, as indicated.	167
B.1	Stopping-muon candidate events from the medium-threshold Module 0 run 1 data sample (batch 1). Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM. .	191
B.2	Stopping-muon candidate events from the medium-threshold Module 0 run 1 data sample (batch 2). Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM. .	192

- B.3 Stopping-muon candidate events from the high-threshold Module 0 run 1 data sample (batch 1). Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM. . . 193
- B.4 Stopping-muon candidate events from the high-threshold Module 0 run 1 data sample (batch 2). Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM. . . 194
- B.5 Stopping-muon candidate events from the high-threshold Module 0 run 2 data sample (batch 1). Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM. . . 195
- B.6 Stopping-muon candidate events from the high-threshold Module 0 run 2 data sample (batch 2). Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM. . . 196
- B.7 Muon-decay candidate events from the medium-threshold Module 0 run 1 data sample. Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM. . . . 197

- B.8 Muon-decay candidate events from the high-threshold Module 0 run 1 data sample. Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM. . . . 198
- B.9 Muon-decay candidate events from the high-threshold Module 0 run 2 data sample. Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM. . . . 199
- B.10 Stopping-muon candidate events from the CORSIKA-generated medium-threshold Module 0 run 1 sample (batch 1). Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM. All four of these events are stopping μ^+ with a subsequent decay. . . . 200
- B.11 Stopping-muon candidate events from the CORSIKA-generated medium-threshold Module 0 run 1 sample (batch 2). Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM. From left-to-right, top-to-bottom, these events are a μ -capture event, a μ^+ -decay event, a μ -capture event, and an electromagnetic (EM)-shower background event. . . . 201

- B.12 Stopping-muon candidate events from the CORSIKA-generated high-threshold Module 0 run 1 sample (batch 1). Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM. From left-to-right, top-to-bottom, these events are a through-going μ background event, a μ^+ -decay event, a μ^+ -decay event, and a μ^- -decay event. 202
- B.13 Stopping-muon candidate events from the CORSIKA-generated high-threshold Module 0 run 1 sample (batch 2). Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM. From left-to-right, top-to-bottom, these events are a μ^+ -decay event, a μ^- -capture event, a μ^+ -decay event, and a μ^+ -decay event. 203
- B.14 Stopping-muon candidate events from the CORSIKA-generated high-threshold Module 0 run 2 sample (batch 1). Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM. From left-to-right, top-to-bottom, these events are a μ^- -decay event, μ^+ -decay event, a μ^+ -decay event, and a μ^+ -decay event. 204
- B.15 Stopping-muon candidate events from the CORSIKA-generated high-threshold Module 0 run 2 sample (batch 2). Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM. From left-to-right, top-to-bottom, these events are a μ^- -decay event, an EM-shower background event (with a pile-up μ), an EM-shower background event, and a μ^+ -decay event. 205

- B.16 Muon-decay candidate events from the CORSIKA-generated medium-threshold Module 0 run 1 sample. Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM. From left-to-right, top-to-bottom, these events are a μ^+ -decay event, a μ^+ -decay event, a μ^- -decay event, and a μ^- -decay event. 206
- B.17 Muon-decay candidate events from the CORSIKA-generated high-threshold Module 0 run 1 sample. Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM. From left-to-right, top-to-bottom, these events are a μ^- -decay event, a μ^+ -decay event, a μ^- -decay event, and a μ^+ -decay event. 207
- B.18 Muon-decay candidate events from the CORSIKA-generated high-threshold Module 0 run 2 sample. Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM. From left-to-right, top-to-bottom, these events are a μ^+ -decay event, μ^+ -decay event, a μ^- -decay event, and a μ^+ -decay event. 208

List of Tables

1.1	Weak isospin and weak hypercharge of fermions in the SM.	11
1.2	Summary of experimental measurements of the unpolarized electron spectrum decay parameters ρ , η , ξ , and δ , taken from the review [24]. *Because ξ enters into the decay spectrum alongside the muon polarization, there is an inherent degeneracy between $(\xi > 0), (P_\mu > 0)$ and $(\xi < 0), (P_\mu < 0)$. Experimental results constrain only the product of these.	14
1.3	Mixing angles as reported by NuFit 5.2 (2022) [83], where $\Delta m_{3l}^2 = \Delta m_{31}^2 (\Delta m_{32}^2)$ for normal (inverted) ordering.	22
1.4	Summary of existing muon capture measurements for argon.	26
2.1	Summary of liquid-argon properties. Where applicable, values are reported for liquid phase at 87.3 K, 1 atm, and 500 V/cm.	31
2.2	Summary of operational differences between LArPix-v1 and -v2 ASICs.	39
3.1	Table of measured Birks' model parameters, A and k , for liquid argon.	52
4.1	Summary of Module 0 datasets marked as "good-for-analysis".	77
4.2	Parameters for the physics model of CORSIKA v7.7400 used to generate the cosmic-ray flux at the Bern site.	80
4.3	Geant4 and EdepSim configuration parameters used to simulate cosmic rays in Module 0.	81
4.4	Summary of simulation datasets used in this analysis.	91
5.1	PACMAN clock corrections used to synchronize charge data.	101
5.2	Gain calibration windows used for Module 0 LRS single photo-electron (SPE) calibration.	118
6.1	Total number of events after each selection with the relative contributions by particle species and muon disappearance channel predicted by simulation, post-fit. Predictions are extrapolated from the background vetoed sample using the efficiency estimated with the simulation and a reweight factor determined by the fit. This extrapolation can result in an efficiency estimate > 1.00 for a particle species.	142
6.2	Post-fit efficiency of each cut predicted with simulation.	143

6.3	Post-fit cumulative efficiency of each cut predicted with simulation.	144
6.4	Post-fit relative fraction of the sample labelled 'total' in Table 6.1 after each selection predicted with simulation	145
6.5	Summary of components that make up the model prediction in the likelihood function. The free parameters f , μ , σ , and λ represent relative normalization, scale, smearing, and exponential-decay constant, respectively. This table does not fully represent the free parameters used in the fit, which include additional constraints between the parameters. These are described in more detail at the end of Sec. 6.2.	147
6.6	Fit results of the toy Monte Carlo (MC) fit closure test.	152
6.7	Summary of measurement uncertainties for the fit with no external constraints.	156
6.8	Summary of measurement uncertainties for the fit with a λ_d constraint from Ref. [95].	157
6.9	Summary of measurement uncertainties for the fit with a λ_d constraint from Ref. [96].	157
6.10	Summary of measurement uncertainties for the fit with a $R_H = 0.99$ [63] constraint. The constraint assumes no error on R_H	158
6.11	Fractional covariance matrix C_{rel} of parameters of interest.	161
6.12	Fractional covariance matrix C_{rel} of parameters of interest with a λ_d constraint from Ref. [95].	161
6.13	Fractional covariance matrix C_{rel} of parameters of interest with a λ_d constraint from Ref. [96].	161
6.14	Fractional covariance matrix C_{rel} of parameters of interest with a $R_H = 0.99$ [63] constraint.	162
6.15	Summary of the best-fit μ^+/μ^- ratio under different constraints.	162
6.16	Summary of the best-fit Huff factor, R_H , under different constraints.	163
6.17	Summary of the best-fit disappearance rate, λ_d , under different constraints.	163
6.18	Summary of the best-fit capture rate, λ_c , under different constraints.	163
A.1	Summary of systematics considered for this measurement.	187

List of Acronyms

- ADC** analog-to-digital converter
- ARAPUCA** Argon R&D Advanced Program at UniCAmp
- ASIC** application-specific integrated circuit
- BESS** Balloon-borne Experiment with a Superconducting Spectrometer
- BNB** Booster Neutrino Beam
- BSM** beyond-the-Standard-Model
- CAPRICE** Cosmic AntiParticle Ring Imaging Cherenkov Experiment
- CC** charged-current
- CDF** cumulative probability distribution function
- CERN** European Council for Nuclear Research
- CMOS** complementary metal-oxide semiconductor
- CMS** Compact Muon Solenoid
- CORSIKA** COsmic Ray SIMulations for KAscade
- CPU** central processing unit
- CRS** charge readout system
- CSA** charge-sensitive amplifier
- CSDA** continuous slowing down approximation
- CUDA** Compute Unified Device Architecture
- CVC** conserved vector current
- DAC** digital-to-analog converter
- DAQ** data acquisition system
- DAR** decay-at-rest
- DBSCAN** density-based spatial clustering of applications with noise
- DC** direct current

DDR double data-rate

DFT discrete Fourier transform

DIO decay-in-orbit

DIS deep inelastic scattering

DM dark matter

DUNE Deep Underground Neutrino Experiment

EM electromagnetic

EMI electromagnetic interference

FIFO first-in first-out

FFT Fast Fourier transform

FNAL Fermi National Accelerator Laboratory

FPGA field-programmable gate array

FWHM full-width at half-maximum

GDML Geometry Description Markup Language

GPS global positioning system

GPU graphics processing unit

HV high-voltage

IO input/output

IP internet protocol

JINR Joint Institute for Nuclear Research, Dubna

LAr liquid argon

LArTPC liquid argon time projection chamber

LED light emitting diode

LBNF Long-Baseline Neutrino Facility

LBNL Lawrence Berkeley National Laboratory

LCM light collection module

LHEP the Laboratory for High-Energy Physics

LRS light readout system

LSB least-significant bit

LUT look-up table

LY light yield

MIP minimum ionizing particle

MC Monte Carlo

NC neutral-current

NTP network time protocol

ppb parts-per-billion

PACMAN Pixel Array Control, Message, and Network

PCA principle component analysis

PCB printed circuit board

PDE photon detection efficiency

PDF probability distribution function

PE photo-electrons

PL programmable logic

PMNS Pontecorvo-Maki-Nakagawa-Sakata

PMT photomultiplier tube

ppm parts-per-million

PPS pulse-per-second

PSD pulse shape discrimination

PSI Paul Scherrer Institute

PVC polyvinyl chloride

QC quality control
QCD quantum chromodynamics
QE quasi-elastic
QED quantum electrodynamics
QMC quantum Monte Carlo
QY charge yield
RANSAC random sample consensus
RES resonant
ROI region-of-interest
RPA random-phase approximation
RTD resistive temperature detector
SCE space-charge effect
SiPM silicon photo-multiplier
SLAC SLAC Linear Accelerator Laboratory
SM Standard Model
SNR signal-to-noise ratio
SPE single photo-electron
SRAM static random-access memory
SURF Sanford Underground Research Facility
TCP transmission control protocol
TPB tetra-phenyl-butadiene
TPC time projection chamber
UART universal asynchronous receiver-transmitter
UDP User Datagram Protocol
VGA variable gain amplifier
WIMP weakly-interacting massive-particle
WLS wavelength shifting

Acknowledgments

First and foremost, I have to thank my advisor, Kam-Biu. Your relentless optimism and support kept me moving forward through many uncertain and challenging circumstances. I was worried and saddened in 2020 when you told me that you would be retiring from Berkeley and relocating to the Hong Kong University of Science and Technology for the remainder of my PhD, but since then you've managed to continue to play an active and indispensable role in my work. Thank you for always making space for me in your schedule and dutifully meeting with me to discuss my progress or just the latest events in our lives. How you are able to be successful physicist and still provide sincere 1:1 mentorship to your students is beyond me. I feel blessed and lucky to have had you as an advisor. And to the extent that this dissertation can be considered readable is due entirely to your detailed review of my (*very*) rough drafts. Thank you for correcting all of my misplaced hyphens. I never would have made it without you as an advisor.

I also have to thank my second advisor, Dan. You pushed me to do more than I thought I could. I would have never stuck around had you not whole-heartedly embraced my interest in instrumentation. Your do-it-yourself approach to physics is inspiring, though my suspicion is that it's only because you have such a boundless intellect that you make it work. You've taught me that complex problems are best solved with simple solutions, and no obstacle cannot be overcome with a clever approach. Thank you for teaching me so much, and for guiding me over the years.

Next, the members (and former members) of the Lawrence Berkeley National Laboratory (LBNL) neutrino group provided a unique and welcoming environment for me to grow. During my time in the group, we saw many dramatic changes – uncertain funding situations, a pandemic and all that followed, a change in the group lead, the final days of Daya Bay. But it is a testament to the wonderful people who are part of the group that we still managed to stay nimble and cohesive throughout. Brooke, thank you for unflinchingly taking up the LArPix torch and keeping me company during the quiet hours of the lab shutdown. The Module 0 runs would never have been a success without you, and I always felt better knowing that you were on shift. Your commitment to physics is inspiring, and I look forward to seeing all that you achieve. Cheng-Ju, thank you for being such a welcoming and personable force in the group. You made me feel like I had a place, whether it was in Berkeley or at CERN during the first protoDUNE run. I believe it was you who originally mentioned muon capture as a potential topic for my thesis, which came at a pivotal time for me. And of course, also thank you for introducing me to all of the best food around Saint-Genis. Herb, thank you for your compassionate hard-ball questions during my qualifying exam prep and your continued interest in muon physics. After your questions, the real deal was a cake walk. I hope to have a quarter of your vibrancy when (*and if!*) I reach your age. Roberto, we made a great team on the larnd-sim paper, and I miss having you in the group. Thank you for your openness whenever I would come at you with my simulation woes. Also thank you for navigating the software environment at Fermilab and getting the CORSIKA simulation set up for Module 0. To the LBNL engineering team (Armin, Carl, Dario, Thorsten, and Andrew), Module 0 never

would have happened without your creative solutions and constant guidance and patience – thank you. It’s your ingenuity and passion that has inspired my shift to engineering. Chris, thank you for always having an hot topic to discuss at lunch. Your stimulating presence in our group has been missed since you left. Thank you for your guidance and support in navigating the world of neutrino cross-sections when I was first starting to look at muon capture. Sam, thank you for introducing me to proper coding techniques and being my connection to the grad student union. Keep up the good fight! Callum, Kevin, Stephen, Matt, thanks for always being up for a social event - though, I’m still disappointed I never got to hear Stephen at open-mic night at the Starry Plough.

While the LBNL group provided the charge readout for Module 0, nigh everything else was provided by the Bern group. Not only did the Bern team set us up for the first demonstration of LArPix in their cryostat, but they provided the driving force behind ArgonCube and much-needed fresh ideas in LArTPC design. Due to the constraints of the pandemic, the on-site shifts during the Module 0 runs were filled entirely by the Bern team. Thank you for providing me a home-away-from-home during my visits, and for always taking me out for a few-too-many-drinks between shifts or meetings. Especially thank you to Roman, Livio, and Saba for fielding my many questions about the detector configuration during the run periods. And many thanks to Livio and Patrick for producing and maintaining the GDML files and the LRS look up tables used for the simulation in this thesis. Michele and Igor, thank you for organizing the runs and pushing the ArgonCube effort forward. Without your extension of the Module 0 run, this thesis would not have been possible.

Within DUNE, Alysia, thank you for introducing me to neutrino physics and being the best undergraduate advisor I could have asked for. The Joint Institute for Nuclear Research, Dubna (JINR) team (Nikolay and Sasha), thank you for performing the light calibration and providing me with all the information I needed about the LCMs. The SLAC Linear Accelerator Laboratory (SLAC) team (Kazu, Yifan, Hiro, James), thank you for being gracious hosts during my brief tour through machine learning, you were always my go-to people at DUNE collaboration meetings. Pedro, thank you for your always clear-sighted advice and for lending me your old LBNL desk. Andy, thank you for developing the run scripts for the cosmic ray simulation. Mike and Lane, thank you for providing me the electron lifetime calibration for the Module 0 runs. The protoDUNE gang: Francesca, Chris, Milo, Casandra, Bryan, James, Serhan, Maura, Jake, Flor, Heng-Ye, Rob, Geoff and others, thank you for welcoming me with open arms while I was at CERN and making my time there something that I will cherish forever.

Finally, my committee members Kai Vetter, Dan McKinsey, and Wick Haxton took time to sit on my qualifying exam committee (and for Kai, after only having a couple days notice) and didn’t give me too many hardball questions – thank you! Kai and Dan were also willing to serve on my dissertation committee and read through all 170+ pages (*you did read it all, right?*). Thank you both.

To the friends that I’ve made along the way, without your support and welcome diversions I would have not made it to this point. My fellow beer-people: Brad and Marie, thank you for keeping things fresh every week and reminding me that there is life outside of Quantum

problem-sets. QinQin and Jakob, I can't imagine going through grad school without you around – thank you for the many delicious treats and always being up for a board game. And the Kurt, Maggie, Laura, Shawn crew, thank you for letting me tag along and making Berkeley feel like home.

To my family, Mom, Dad, and Nate, thank you each for shaping me and standing by me every step of the way. Mom, thank you for pushing me to do more, even though I excel at doing nothing. You've always believed in me and your unwavering conviction was the only thing that kept me going. Dad, our long conversation on the phone have been some of my most cherished moments in the past years and gave me comfort and distraction when I needed them most. And Nate, your words of encouragement carried me through some of the hardest times. You've each played a considerable role in getting me to this point.

And finally, Andrea, thank you for emailing me 9 years ago. I'm almost certain that you had no idea what you would be getting into at the time. Thank you for supporting me through it all and for filling our home with good food and good people. You have persevered as much as I have, and I am thankful that we made it to the other side. There is no challenge that I would want to face without you, and nothing that we could not overcome together. You make my life infinitely richer – *thank you*.

Chapter 1

Introduction

While not widely known, weak interactions play a central role in nature. They drive the processes that power the sun and produce the elements critical to life on Earth. We now know that seemingly disparate phenomena like inertia, electromagnetism, and nuclear reactions can all be linked via the fundamental symmetries of nature. Late in the 19th century, physicists were grappling with ideas that today are taken for granted, such as the periodic table and the charge distribution within atoms. It's because of their efforts that radioactivity was discovered, and we are now in a place to ask (and perhaps even get answers to) profoundly fundamental questions like *why does matter exist?* and *what was the universe like at the earliest times?* – both linked to the nature of the weak interaction.

To gain insight into these questions, the Deep Underground Neutrino Experiment (DUNE) is being built. As part of its research program, this experiment will study a phenomenon that arises through the weak interaction – *neutrino oscillation*. But in order to improve our understanding of neutrino oscillation, DUNE must control the systematic uncertainties of its measurements to a high degree. Systematics enter into the measurements, in part, through uncertainties in neutrino-nucleus interaction models. This thesis demonstrates the successful operation of a pixelated liquid argon time projection chamber (LArTPC), Module 0. The design of this detector is such that it is capable of resolving neutrino interactions in a high-rate environment, enabling the DUNE near detector to use a similar detector technology as the DUNE far detector. This substantially reduces the reliance on neutrino-nucleus interaction modelling for neutrino oscillation measurements. In addition, this thesis presents a study of a weak nuclear process, *muon capture*, from the data collected during the operation of Module 0 in 2021. Two observables, the *muon-decay nhit* and the *muon-decay time*, were calculated for approximately 600,000 muons stopping in Module 0. A maximum-likelihood fit to these data resulted in a measurement of the muon capture rate ($\lambda_c = 2.53 \pm 0.19 \mu\text{s}^{-1}$), the muon disappearance rate ($\lambda_d = 2.11 \pm 0.24 \mu\text{s}^{-1}$), and the ratio of the bound muon decay rate to the vacuum decay rate¹ ($\lambda_{\text{DIO}}/\lambda_{\text{DAR}} \equiv R_H = 1.29 \pm 0.15$) for argon. These results can help to reduce the uncertainty in neutrino-argon cross-section modelling at low Q^2 .

¹This is also referred to throughout the text as the *Huff factor*.

In Chapter 1, I summarize the historical context and central experimental results that have led to our understanding of weak interactions. This is followed by a description of the weak interaction as it is understood within the Standard Model (SM). From this, I present the relevant components that feature into muon capture and the relevance of muon capture for DUNE. In Chapter 2, I describe the detector technology used in the measurement – a liquid argon time projection chamber (LArTPC). A detailed detector simulation for this technology was developed and is described in Chapter 3. Chapter 4 describes the operation of the prototype detector and the data and simulation used in this study of muon capture. From there, I describe the analysis techniques used to process and analyze the data in Chapter 5. Chapter 6 then describes a selection used to identify muon-capture and muon-decay events within the data. The maximum-likelihood fit is also described, along with a discussion of systematic effects. This chapter then concludes with the results of the fit for λ_c , λ_d , and R_H , including a variety of alternative constraint scenarios. As part of the fit, a measurement of the cosmic-ray muon charge ratio was also obtained ($R = 1.22 \pm 0.06$) and is included in the results. Finally, I conclude by placing the measurements of this thesis in context with others and provide some strategies for improving the measurement in Chapter 7.

1.1 A brief history of weak interactions

1.1.1 Fermi’s theory




FIG. 1.

Figure 1.1: Beta decay spectrum of ^{226}Ra , as measured by Wooster and Ellis [1].

In 1914, James Chadwick performed the first measurement of the energy spectrum of the ^{226}Ra β -decay chain and observed a continuous spectrum (Fig. 1.1). Based on the observed

particles of the decay, ${}^A_Z X \rightarrow {}^A_{Z-1} X + e^-$, energy conservation would suggest a discrete energy equal to the difference in the nuclear mass energies of the β -decay, contrary to the observation. Additional measurements by Wooster and Ellis supported this result [1, 2], removing the possibility of secondary interactions leading to the observed continuum.

In 1930, Pauli proposed an explanation to Lise Meitner in a letter, suggesting that the decay also involved a secondary, neutral form of radiation, later called the *neutrino*. In this proposal, the decay is a 3-body decay ${}^A_Z X \rightarrow {}^A_{Z-1} X + e^- + \nu^2$, and so the released energy can be shared between the electron and the neutrino, producing a continuous spectrum. After the discovery of the neutron in 1932 by Chadwick [4], Fermi developed a formal theory surrounding an interaction involving a 4-point vertex with a $p \rightarrow n$ and an $e \rightarrow \nu$ transformation. Beta decay could then be explained as the reaction involving this vertex, $n \rightarrow p + \bar{\nu} + e^-$. With the experimental data on β -decay rates that he had available at the time, he estimated that the interaction strength would be phenomenally small [5],

$$G_F \sim 4 \times 10^{-50} \text{cm}^3 \text{erg}.$$

1.1.2 Muon decay

During this same period, the nuclear theory community was struggling to explain some of the features of the nuclear binding force. In particular, pp scattering experiments suggested a short-range attractive interaction in addition to the Coloumbic force, and the similarity of the binding energies of ${}^3\text{H}$ (pnn) and ${}^3\text{He}$ (ppn) suggested that this force was of about the same strength as the nn interaction [7]. Yukawa [8] was the first to provide a solid theoretical foundation to these features – he proposed a massive particle that drove $n \rightarrow n$, $p \rightarrow p$, and $p \leftrightarrow n$ interactions and whose mass would lead to an exponential suppression of the force with distance. He estimated the mass required to explain the experimental data would need to be $\approx 200 \times m_e$.

Soon after, Anderson and Neddermeyer [9] discovered a new kind of highly-penetrating charged particle with a mass between that of an electron and proton as part of their cosmic-ray experiments. This particle was believed to be of the same kind proposed by Yukawa. However, in the subsequent years, attempts to measure its decay and interaction rates proved confounding. The results of Ehrenfest, Freon, Johnson, and Pomerantz [10, 11], suggested that the flux of these particles varied as a function of the zenith according to a decay or interaction rate on the order of $2 \mu\text{s}$. This could only be explained by a weak interaction with the atmosphere at most, much weaker than required to be interpreted as the particle predicted by Yukawa. While at the same time Montgomery, Ramsey, and Cowie performed an experiment to search for a possible Fermi-like decay, but their non-observation [12] seemed to suggest that these particles interacted strongly when they came to rest.

²Using modern conventions, the neutrino state emitted from β -decay is actually interpreted to be an *anti*-neutrino $\bar{\nu}$, but at the time, the mathematical foundations of antimatter had only just been developed [3] and were not widely accepted.




Figure 1.2: First observational evidence of the mesotron decaying to an electron [6]. A mesotron is seen entering the cloud chamber at location A and decays to an electron at location F, which exits the chamber at location G.

During 1940s, the picture started to clear. Williams and Roberts [6] captured photographic evidence of the muon decaying into an electron, Fig. 1.2, that offered evidence in favor of a Fermi-like interaction involving an electron. Calculations made by Yukawa and Okayama pointed out that it was likely that the meson would only capture after stopping in a material and would nearly always capture on a nucleus. Further, Tomonaga and Araki [13] pointed out that the capture rates of positive and negative mesons would be dramatically different due to the Coloumb attraction between the nucleus and the stopped meson. Positive mesons would experience a repulsive force from the positively charge nucleus and thus would be unlikely to interact with a nucleon, whereas the negative meson would experience an attractive force and be pulled into the nucleus after slowing. This was demonstrated in the work by Conversi, Pancini, and Piccioni in 1944 [14], in which a magnetic field was used to select events which stopped in an iron absorber. Here they observed markedly different behavior of the positive and negative mesons³. Their follow up study in 1946 [15] showed that not only did the negative meson decay differently, there was also a strong dependence on

³Publications from this period broadly apply the name meson or mesotron to all of the newly discovered particles with a mass between that of the electron and that of the proton. However, in modern parlance, the label of *meson* is only applied to strongly-interacting bosons consisting of a quark-antiquark pair. The important distinction here is that the particle most relevant to this thesis, the *muon*, is no longer considered to be a meson. The original naming is used here for historical accuracy.

Z^- – with a cross-section that was incompatible with the predictions for the Yukawa meson. This conclusively showed that the positive and negative muon decay rates were different, but more importantly, their work demonstrated that the negative muon decay occurred when stopping in matter. Thus it occurred with a relatively high rate compared to the capture – much higher than would have been expected if it was the particle carrier of the nuclear binding force.




Figure 1.3: First observation of a pion decaying into a muon [16]. An incoming short pion track (m_1) is observed along with the subsequent muon track (m_2) exiting the emulsion. Arrows indicate track deflections of $> 2^\circ$.

Finally in 1947, emulsions from Perkins and Lattes, Occhialini, and Powell [16] determined that the confusion surrounding the mesotron arose because there were in fact two different particles - the pion and the muon - with a relatively small mass difference ($\approx 30 \text{ MeV}/c^2$).

1.1.3 V-A interaction

With the identity of the muon resolved, studies into the nature of the particle and its interactions proceeded. The observed decay to an electron was studied by Steinberger [18] and Hincks and Pontecorvo [19]. From the initial observations of cosmic rays, it was known that the decay did not produce any additional visible energy other than the decay electron. Steinberger, et al. showed by observing delayed coincidences through a series of absorber plates that the emitted electron energy is continuous, falling between 0 and $\sim m_\mu c^2/2$ and consistent with a 3-body decay to near massless particles. By adding a thin, high-Z absorber to their experiment, Hincks and Pontecorvo [20] were also able to rule out the possible presence of a decay photon. Because the interaction strength determined by the decay lifetime was of the same order as Fermi's β -decay interaction, it also suggested *universality* of the force governing both β -decay, muon decay, and perhaps pion decay. In 1949, Michel [17]




Figure 1.4: Shape of electron energy spectrum from μ -decay under different assumptions for the underlying interaction, as calculated by Michel [17]. Curve B most closely resembles modern measurements.

showed that the shape of the muon decay energy spectrum would change depending on the form that the interaction coupling takes, shown in Fig. 1.4. But it was not until much later that precise enough measurements would be possible.

Up until this point, parity was generally assumed to be conserved in the fundamental interactions. However, in 1956 while investigating the parity-conserving and parity-non-conserving decay modes of what appeared to be the same particle, Lee and Yang [21] reviewed the experimental evidence for parity conservation and found that there were no conclusive measurements that could either support or refute it. This spurred Wu, et al. in 1957 [22] to search for direct evidence of parity violation. In her experiment, she used a polarized ^{60}Co β -decay source to measure the anisotropy of the decay electron angular distribution. If parity were conserved, decay electrons would be emitted isotropically from the polarized source. But what she found was instead a strong preference for emitting electrons in the opposite direction from the polarization, providing proof that parity is violated and also the parity violating component of the interaction must have an opposite sign as the parity conserving component.

In 1958, Goldhaber, et al. [23] performed a direct measurement of the neutrino helicity via the decay of ^{152}Eu . In their experiment, they used the decay chain of $^{152}\text{Eu} \rightarrow ^{152}\text{Sm}^* + \nu_e \rightarrow ^{152}\text{Sm} + \gamma$, which based on simple arguments of angular momentum conservation produces, for back-to-back decay, a circularly polarized gamma ray with the same polarization as the neutrino helicity. By using resonant scattering of the 960-keV gamma ray on a Sm_2O_3 target, they were able to select back-to-back decays and measured a net negative polarization of these gamma rays, showing that the neutrino has a negative helicity.

Later, precise measurements of the muon decay parameters⁴ heavily constrained the

⁴As originally presented in Ref. [17], there are four dimensionless parameters $K_{1,2,3,4}$ that determine the

possible forms of the weak interaction. The high-energy portion (>20 MeV) of the spectrum first precisely measured by Dudziak, et al. [25] constrained $\rho = 0.752 \pm 0.004$. In 1969, Plano [26] exposed a magnetized hydrogen bubble chamber to a beam of μ^+ and from the anisotropy of the electron distribution, measured $|\xi| = 0.94 \pm 0.07$ and $\delta = 0.78 \pm 0.05$. Measurements of the low-energy portion of the spectrum (<10 MeV) performed by Derenzo [27], showed that η was consistent with 0. By measuring the polarization of the electron, Buhler, et al. [28] were able to access the ξ' parameter ($= 1.05 \pm 0.30$). And finally in 1967, Bogart measured the radiative muon decay ($\mu^+ \rightarrow e^+ + \gamma + \nu_e + \bar{\nu}_\mu$) and found the $\bar{\eta}$ parameter consistent with zero. Combined, these results imply a rather simple structure for the weak interaction, namely the $V-A$ interaction.

Since then, the study of neutrino interactions with the Gargamelle bubble chamber [29] showed that weak interactions also include a neutral-current interaction, e.g. $\nu_\mu + \text{hadron} \rightarrow \nu_\mu + \text{hadron}$. And later, the discovery of the Z^0 boson with a mass of $M_Z = 94 \pm 2$ GeV/ c^2 [30, 31] and the discovery of the W^\pm bosons with a mass of $M_W = 81 \pm 5$ GeV/ c^2 [32, 33] by UA1 and UA2 experiments, provided evidence towards a unified interaction of the electromagnetic and weak forces, as put forward by Weinberg [34]. This model has been confirmed with the discovery of the Higgs boson in 2012 by the CMS and ATLAS experiments [35, 36], which causes the symmetry-breaking that gives rise to the masses of the weak vector bosons (Z^0 and W^\pm).

In this unified interaction, the original coupling proposed by Fermi is determined by the underlying interaction strength g_W and is suppressed at low energies by the large masses of the weak vector bosons

$$G_F = \frac{\sqrt{2}}{8} \frac{g_W^2}{M_W^2} = \frac{1}{\sqrt{2}v} \quad (1.1)$$

which are induced by the vacuum expectation value of the Higgs field v .

1.2 Muon capture and muon decay

From the experimental and theoretical developments described, we now have a clear understanding of weak interactions. One such interaction is *muon capture* whereby a muon is converted to a neutrino with a corresponding conversion of a nucleon. This is the main topic of this section and the interaction studied in this thesis. To understand this process, it's valuable to first consider muon decay, which is described after an overview of electroweak theory within the SM. Finally, because of nucleon-nucleon interactions, both muon decay and muon capture are modified when they take place inside of a nuclei. Some of these modifications are described at the end of this section.

shape of the non-polarized muon decay spectrum. In later formulations of the theoretical spectrum, a slightly different parameterization is used, with ρ, η, ξ , and δ being the commonly-referenced decay parameters. Later measurements could probe the polarized muon decay spectrum, which is sensitive to more exotic interactions that require additional parameters, e.g. ξ' and $\bar{\eta}$ [24].

1.2.1 Electroweak Interaction

The SM is a quantum field theory with an assumed $SU(3) \times SU(2)_L \times U(1)$ local gauge symmetry and a global Lorentz symmetry [37]. Electroweak interactions are produced by the $SU(2)_L \times U(1)$ subgroup, while the strong interactions are produced by the $SU(3)$ subgroup.

Within the electroweak sector, the spin-1/2 fermions are arranged into left-handed doublets

$$\psi_{l,L} = \begin{pmatrix} \nu_{eL} \\ e_L \end{pmatrix}, \begin{pmatrix} \nu_{\mu L} \\ \mu_L \end{pmatrix}, \begin{pmatrix} \nu_{\tau L} \\ \tau_L \end{pmatrix} \quad (1.2)$$

$$\psi_{q,L} = \begin{pmatrix} u_L \\ d_L \end{pmatrix}, \begin{pmatrix} c_L \\ s_L \end{pmatrix}, \begin{pmatrix} t_L \\ b_L \end{pmatrix} \quad (1.3)$$

and right-handed singlets

$$\psi_{l,R} = (e_R), (\mu_R), (\tau_R) \quad (1.4)$$

$$\psi_{q,R} = (u_R), (c_R), (t_R), (d_R), (s_R), (b_R), \quad (1.5)$$

where handed-ness is defined by the field's behavior under the chiral projection operators

$$P_L \psi \equiv \frac{1 - \gamma_5}{2} \psi \equiv \psi_L \quad (1.6)$$

$$P_R \psi \equiv \frac{1 + \gamma_5}{2} \psi \equiv \psi_R. \quad (1.7)$$

By introducing the covariant derivative

$$D_\mu \equiv \partial_\mu - ig_W \frac{\sigma_a}{2} W_\mu^a - i \frac{g_Y}{2} Q_Y B_\mu, \quad (1.8)$$

where σ_a are the three Pauli matrices ($a = 1, 2, 3$) and μ is the Lorentz index ($\mu = 0, 1, 2, 3$), a Lorentz-invariant, gauge-symmetric Lagrangian can be produced

$$\mathcal{L}_{EW} = \sum_\psi \bar{\psi} \gamma^\mu D_\mu \psi - \frac{1}{4} W_a^{\mu\nu} W_{\mu\nu}^a - \frac{1}{4} B^{\mu\nu} B_{\mu\nu}, \quad (1.9)$$

where the sum \sum_ψ iterates over all of the fermion fields $\psi_{l,(L,R)}$ and $\psi_{q,(L,R)}$ and $W_{\mu\nu}^a$ and $B_{\mu\nu}$ are the field strength tensors of the gauge fields W_μ^a and B_μ introduced by the $SU(2)$ and $U(1)$ symmetries, respectively,

$$W_{\mu\nu}^a \equiv \partial_\mu W_\nu^a - \partial_\nu W_\mu^a + 2ig_W \epsilon^{abc} W_\mu^b W_\nu^c \quad (1.10)$$

$$B_{\mu\nu} \equiv \partial_\mu B_\nu - \partial_\nu B_\mu, \quad (1.11)$$

where ϵ^{abc} is the anti-symmetric Levi-Civita tensor

$$\epsilon^{abc} = \begin{cases} +1 & \text{even permutation of } (a, b, c) \\ -1 & \text{odd permutation of } (a, b, c) \\ 0 & a = b, b = c, \text{ or } a = c . \end{cases} \quad (1.12)$$

The coupling constants g_W and g_Y correspond to the $SU(2)_L$ weak isospin coupling constant and $U(1)$ weak hypercharge coupling constant that determine the strength of the interactions. These are also often expressed in terms of the weak mixing angle $\theta_W \equiv \tan^{-1}(g_Y/g_W)$ [24].

To explain the suppression of the weak interaction at low energies and the mass of the gauge fields, a complex scalar-field doublet ϕ is introduced with a weak hypercharge $Q_Y = 1$

$$\phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix}. \quad (1.13)$$

A gauge-invariant Lagrangian that this field allows is

$$\mathcal{L}_H = (D_\mu \phi)(D^\mu \phi) - \lambda(|\phi|^2 - v^2/2)^2, \quad (1.14)$$

where v and λ are constant free parameters. The potential term $\lambda(|\phi|^2 - v^2/2)^2$ has a minimum at a non-zero value for $|\phi|$, and thus the nominal $U(1)$ symmetry of the vacuum state is *spontaneously broken*. Because the Lagrangian is invariant under $SU(2)$, a transformation can be applied to enforce that

$$U\phi = \begin{pmatrix} 0 \\ \frac{v}{\sqrt{2}} + \frac{H}{\sqrt{2}} \end{pmatrix}, \quad (1.15)$$

such that H is a real scalar field and $H = 0$ minimizes the potential. With this choice, the potential term introduces a mass term for H , plus two self-interaction terms

$$\mathcal{L}_H = (D_\mu \phi)(D^\mu \phi) - \frac{1}{2} \left(\sqrt{2\lambda}v \right)^2 H^2 + \lambda v H^3 + \frac{\lambda}{2} H^4. \quad (1.16)$$

Taking $H = 0$, the kinetic term can be expanded to show that this also introduces quadratic terms for the gauge fields

$$\mathcal{L}_H = \frac{v^2 g_W^2}{2} \frac{1}{4} [(W_\mu^1)^2 + (W_\mu^2)^2] + \frac{v^2}{2} \left(\frac{g_Y}{2} B_\mu - \frac{g_W}{2} W_\mu^3 \right)^2 + \dots \quad (1.17)$$

which can be re-parameterized to reveal a mass for three of the four gauge fields

$$\mathcal{L}_H = -\frac{1}{2} \left(\frac{v g_W}{2} \right)^2 W_\mu^+ W^{\mu-} - \frac{1}{2} \left(\frac{v \sqrt{g_W^2 + g_Y^2}}{2} \right)^2 Z_\mu Z^\mu + \dots \quad (1.18)$$

where the new fields are defined as

$$W_\mu^+ \equiv \frac{W_\mu^1 - iW_\mu^2}{\sqrt{2}} \quad (1.19)$$

$$W_\mu^- \equiv \frac{W_\mu^1 + iW_\mu^2}{\sqrt{2}} \quad (1.20)$$

$$Z_\mu \equiv \cos \theta_W W_\mu^3 - \sin \theta_W B_\mu \quad (1.21)$$

$$A_\mu \equiv \sin \theta_W W_\mu^3 + \cos \theta_W B_\mu \quad (1.22)$$

These fields are identified as the force carrying fields for the weak interactions (W_μ^\pm , Z_μ) and for electromagnetism (A_μ).

The interactions that are produced in this theory correspond to the three observed electroweak interactions: charged-current weak interactions, neutral-current weak interactions, and quantum electrodynamics (QED). Charge-current weak interactions involve a change in the weak isospin and are mediated by the W_μ^\pm fields

$$g_W \bar{\psi}_L \gamma^\mu \frac{\sigma^1 \pm i\sigma^2}{\sqrt{2}} W_\mu^\pm \psi_L, \quad (1.23)$$

while the neutral-current weak interactions and QED preserve weak isospin. Neutral-current weak interactions and QED appear with a similar form

$$\frac{g_W}{2 \cos \theta_W} \bar{\psi} \gamma^\mu (g_V - g_A \gamma^5) Z_\mu \psi \quad (\text{weak neutral-current}) \quad (1.24)$$

$$g \sin \theta_W Q \bar{\psi} \gamma^\mu A_\mu \psi \quad (\text{QED}), \quad (1.25)$$

where $g_V \equiv I_W - 2Q \sin^2 \theta_W$ and $g_A \equiv I_W$ with I_W and Q being the weak isospin and electric charge of each fermion, respectively, as listed in Table 1.1. But importantly, this implies that QED couples the left- and right-handed chiral components of the fields together and has a universal interaction strength ($e \equiv g \sin \theta_W$), while weak neutral-current interactions do not. In weak neutral-current interactions, only the vector component is involved for right-handed particles.

The allowed interactions with the scalar field ϕ under the assumed gauge symmetries also produces the mass terms for the fermions

$$\mathcal{L}_Y = - \sum_\psi y_\psi \psi_L^\dagger \phi \psi_R + \text{h.c.} \quad (1.26)$$

$$= - \sum_\psi \frac{y_\psi v}{\sqrt{2}} \psi_L^\dagger \psi_R + \text{h.c.} + \dots \quad (1.27)$$

where the Yukawa coupling constants y_ψ are free parameters that determine the mass of each of the fermions. As there are no right-handed neutrinos within the SM, there is no corresponding mass term that is allowed.

Table 1.1: Weak isospin and weak hypercharge of fermions in the SM.

Fermion	Weak isospin (I_W)	Weak hypercharge (Q_Y)	Electric charge ($Q \equiv I_W + \frac{1}{2}Q_Y$)
$\nu_{eL}, \nu_{\mu L}, \nu_{\tau L}$	+1/2	-1	0
e_L, μ_L, τ_L	-1/2	-1	-1
u_L, c_L, t_L	+1/2	+1/3	+2/3
d_L, s_L, b_L	-1/2	+1/3	-1/3
e_R, μ_R, τ_R	0	-2	-1
u_R, c_R, t_R	0	+4/3	+2/3
d_R, s_R, b_R	0	-2/3	-1/3

Examining Table 1.1, there are natural groupings of the fermions into three *generations* which have the same charge in weak interactions: (u, c, t) , (d, s, b) , (e, μ, τ) , and $(\nu_e, \nu_\mu, \nu_\tau)$. Grouping by generation admits a natural eigenbasis defined by the coupling to the vacuum expectation value of the Higgs field

$$\mathcal{L}_{Y,\text{quark}} = -\frac{v}{\sqrt{2}} (u_L^\dagger \quad c_L^\dagger \quad t_L^\dagger) \begin{pmatrix} y_u & 0 & 0 \\ 0 & y_c & 0 \\ 0 & 0 & y_t \end{pmatrix} \begin{pmatrix} u_R \\ c_R \\ t_R \end{pmatrix} + \text{h.c.} + \dots \quad (1.28)$$

However, as the fermions share the same weak charges, weak interactions do not intrinsically prefer this eigenbasis, and thus a unitary transformation

$$\begin{pmatrix} u' \\ c' \\ t' \end{pmatrix} = \mathbb{U}_u \begin{pmatrix} u \\ c \\ t \end{pmatrix} \quad (1.29)$$

is allowed. The resulting interaction

$$\frac{g_W}{\sqrt{2}} (u_L^\dagger \quad c_L^\dagger \quad t_L^\dagger) \mathbb{U}_u^\dagger \bar{\sigma}_a^\mu W_\mu^{a+} \mathbb{U}_d \begin{pmatrix} d_L \\ s_L \\ b_L \end{pmatrix} + \text{h.c.}, \quad (1.30)$$

with $\bar{\sigma}_a^\mu = (\mathbb{1}, -\sigma_a)$, defines a flavor basis according to the matrix $\mathbb{U}_u^\dagger \mathbb{U}_d \equiv V_{\text{CKM}}$. If one family of fermions is massless, then the vacuum eigenstates are degenerate and one can freely chose a basis such that $\mathbb{U}_{+1/2}^\dagger \mathbb{U}_{-1/2} = \mathbb{1}$. Such a scenario is predicted by the SM for the neutrino states. However, if both families of fermions are massive, this freedom is not available, and so the weak charged-current interactions will introduce coupling across the particle generations, e.g. $u\bar{s} \rightarrow u\bar{u} + \nu_e + e^+$.

Overall, we can summarize the features of the electroweak components of the SM as follows:

1. There are two groups of spin-1/2 fermions, *leptons* and *quarks*.
2. An assumed $SU(3) \times SU(2)_L \times U(1)$ local gauge symmetry gives rise to interactions with spin-1 bosons – one per generator of each group.
3. Quarks are charged under the full $SU(3) \times SU(2)_L \times U(1)$, while leptons are charged only under $SU(2)_L \times U(1)$.
4. An additional $SU(2) \times U(1)$ -charged complex scalar field with a non-zero vacuum expectation value breaks the $U(1)$ symmetry, giving rise to masses for three of the spin-1 bosons and the fermions with right-handed partners.
5. Diagonalization of the three flavor states of each group of fermions in the mass basis results in off-diagonal mixing of the flavor states involved in weak interactions, leading to flavor-changing interactions.

1.2.2 Muon properties

We now focus on one particular particle within the SM, the muon (μ). As described earlier, the muon was discovered in 1937. It shares many of the same properties as the electron, namely, its spin (1/2), its charge ($\pm e$), and its weak interactions $I_{W,(L,R)} = (\pm 1/2, 0)$. However, it differs substantially in its mass

$$m_\mu = 0.1134289257(25) \text{ u [38]} \approx 105.7 \text{ MeV}/c^2 \quad (1.31)$$

$$\approx 206 \times m_e. \quad (1.32)$$

Because of this, it can decay weakly, predominately through the charged-current reaction $\mu^- \rightarrow \nu_\mu + \nu_e + e^-$, and with a mean lifetime of

$$\tau_\mu = 2.1969811(22) \text{ } \mu\text{s [24]}. \quad (1.33)$$

Additional, radiative and multi-electron decays have also been observed

$$\begin{aligned} \mu^- &\rightarrow e^- + \bar{\nu}_e + \nu_\mu + \gamma && (6.0 \pm 0.5) \times 10^{-6} \% \text{ [39]} \\ \mu^- &\rightarrow 2e^- + \bar{\nu}_e + \nu_\mu + e^+ && (3.4 \pm 0.4) \times 10^{-3} \% \text{ [40]}. \end{aligned}$$

And it has a magnetic moment of

$$\vec{\mu}_\mu = -\frac{g}{2} \frac{e\vec{S}}{m_\mu} \quad (1.34)$$

where $g = 2 + 0.002331841080(108)$ [41].

Due to their short lifetime, muons are only found naturally as a component of secondary particles produced in high-energy collisions with the Earth by stable particles that have been accelerated by violent astrophysical phenomena, discussed more in Sec. 3.1. In laboratories,

they can be produced in large numbers via the dominant pion decay mode $\pi^- \rightarrow \mu^- + \bar{\nu}_\mu$ at particle accelerators. For this, a high-energy proton beam is typically accelerated to energies above the threshold kinetic energy of $\approx 289 \text{ MeV}/c^2$ and impinged on a target material. Since the pion lifetime is short, they decay before traveling a substantial distance and the resulting muons can be selectively focused by a magnetic field. This makes separation of μ^+ and μ^- relatively easy with the reversal of the current that generates the magnetic field. One consequence of the 2-body pion decay and the left-handed coupling of the weak charged-current (CC) interaction is that muon beams are highly polarized, being produced in a nearly pure right-handed helicity state.

1.2.3 Muon decay




Figure 1.5: Feynman diagram of first-order contribution to muon decay in the SM. An electron is emitted from a decaying negative muon, along with an anti-electron neutrino and a muon neutrino.

In the SM, only weak charged-current interactions allow a change in a charged particle's flavor. For a muon, the only fermions with a smaller mass are the electron and the neutrinos, so its decay must proceed to some combination of these states. The lowest-order decay involves the Feynman diagram shown in Fig. 1.5. Ignoring the SM-allowed interactions that were outlined in the earlier sections, the three-body decay amplitude can be represented in general as a sum over all possible forms of the interaction

$$i\mathcal{M} = \frac{4G_F}{\sqrt{2}} \sum_i g_i (\bar{u}_{\nu_\mu} O_i^\gamma u_\mu) (\bar{u}_e O_{i\gamma} v_{\nu_e}) \quad (1.35)$$

where i iterates over the Lorentz symmetry properties of O^γ ($S = \text{scalar}$, $P = \text{pseudoscalar}$, $V = \text{vector}$, $A = \text{pseudovector}$, $T = \text{tensor}$). Michel [17] first showed that this implies, for the electron emitted from muon decay and with a polarization measured along the arbitrarily-chosen electron polarization axis \hat{n} ,

$$\left(\frac{d^2\Gamma}{dx d\cos\theta} \right)_{\mu^\pm} = \frac{m_\mu}{4\pi^3} W^4 G_F^2 \sqrt{x^2 - x_0^2} [F(x) \pm P_\mu G(x) \cos\theta] [1 - \hat{n} \cdot \vec{P}_e(x, \theta)], \quad (1.36)$$

where the constants $W \equiv (m_\mu^2 + m_e^2)/2m_\mu$, $x \equiv E_e/W$, and $x_0 \equiv m_e/W$ represent the maximum electron energy, the relative electron energy, and the minimum electron energy, respectively. P_μ represents the magnitude of muon polarization and $\vec{P}_e(x, \theta)$ is the polarization of the outgoing electron state. The functions $F(x)$ and $G(x)$ are defined as [24]

$$F(x) \equiv x(1-x) + \frac{2}{9}\rho(4x^2 - 3x - x_0^2) + \eta x_0(1-x) \quad (1.37)$$

$$G(x) \equiv \frac{1}{3}\xi\sqrt{x^2 - x_0^2} \left[1 - x + \frac{2}{3}\delta \left(4x - 4 + \sqrt{1 - x_0^2} \right) \right] \quad (1.38)$$

and are parameterized by dimensionless quantities ρ , η , ξ , and δ , briefly mentioned in Sec. 1.1.2. These quantities can be non-trivially related back to the original couplings g_i involved in the interaction. The electron polarization function $\vec{P}_e(x, \theta)$ can also be non-trivially related back to the form of the interaction through the dimensionless constants.⁵

In Eq. 1.35, there are 10 complex free parameters g_i . Since an overall normalization can be absorbed into G_F , there are 19 real, independent values that can be measured by experiments. The SM predicts $g_V = -g_A = 1$ with all other parameters = 0, thus the muon decay spectrum is used as a stringent test on the SM.

Table 1.2: Summary of experimental measurements of the unpolarized electron spectrum decay parameters ρ , η , ξ , and δ , taken from the review [24]. *Because ξ enters into the decay spectrum alongside the muon polarization, there is an inherent degeneracy between $(\xi > 0), (P_\mu > 0)$ and $(\xi < 0), (P_\mu < 0)$. Experimental results constrain only the product of these.

Parameter	SM prediction	Leading measurement	Reference
ρ	3/4	0.74977 ± 0.00026	[42]
η	0	0.071 ± 0.037	[43]
δ	3/4	0.75049 ± 0.00034	[42]
ξ^*	1	$1.00084 \pm \begin{smallmatrix} 0.00168 \\ 0.00069 \end{smallmatrix}$	[42]

Table 1.2 summarizes the current best measurements of the primary Michel parameters. Overall, there is good agreement between the SM and the experimental results.

⁵The full expressions for \vec{P}_e and g_i in terms of the Michel decay parameters are not reproduced here because they are quite lengthy and add little to the general understanding of muon decay. For the sake of brevity, I will instead refer you to Ref. [24] for the detailed expressions.

1.2.4 Muon capture

If a muon comes to rest in a material, there is the added possibility of an interaction between the muon and the medium. Owing to the repulsive Coulomb force between the nucleus and positive muons, the only modification to the μ^+ -decay arises from the formation of an exotic atom called *muonium*. This state is the result of the muon stripping an electron from one of the weakly bound orbitals of the material, and temporarily forming a μ^+e^- bound state. Muonium is an interesting system for testing predictions of QED as its energy levels are precisely calculable and are not subject to complicating effects that arise from the strong interaction [44]. However, because the binding energy of muonium is small compared to the Q -value of the decay, the μ^+ -decay proceeds largely unperturbed by this initial state.

For negative muons, there is instead an attractive Coulomb force between the nucleus and the muon. Because the bound-state energy of a negatively-charged particle in a Coulomb potential is proportional to the reduced mass of the system

$$\mu = \frac{m_\mu m_A}{m_\mu + m_A} \quad (1.39)$$

where m_μ is the muon mass and m_A is the nuclear mass, the binding energy of the muon is roughly $200\times$ that of an electron. Thus, the atom preferentially will bind to the muon, displacing one of its electrons. This process happens quickly $< O(\text{ns})$ once the muon kinetic energy reaches 10s of eV [45]. The capture occurs into one of the high-angular momentum orbitals and decays to the ground state via the emission of Auger electrons and X-rays [46]. Once in the ground state, the muon wavefunction has a width of $O(\hbar/Z\alpha m_\mu c) \sim 250 \text{ fm}/Z$ and so interactions with the nucleus are particularly relevant.

Muon capture is the dominant process whereby the muon interacts via a charged-current interaction converting a proton to a neutron and emitting a neutrino, $\mu^- + p \rightarrow \nu_\mu + n$. The kinematics of the reaction follow basic 2-body scattering with energy of the emitted neutrino given by

$$E_\nu = \frac{1}{2} \frac{(m_p + m_\mu)^2 - m_n^2}{m_p + m_\mu} \approx 99 \text{ MeV} \quad (1.40)$$

assuming the neutrino is massless. This puts the $q^2 (\equiv (k_\mu - k_\nu)^2)$ of the reaction at

$$q^2 = m_\mu^2 - 2m_\mu E_\nu \approx -0.88m_\mu^2. \quad (1.41)$$

The status of muon capture has been summarized most recently in the exhaustive review by Measday [47], and I will highlight some of the key points from this here. In the simplest case of muon capture on hydrogen, the capture rate can be calculated via the matrix element

$$\mathcal{M} = \frac{2G_F V_{ud}^{\text{CKM}}}{\sqrt{2}} (V^\mu - A^\mu)^\dagger L_\mu \quad (1.42)$$

where V_{ud}^{CKM} is the diagonal element of the CKM mixing matrix, mentioned earlier, and the currents are given by

$$L_\mu \equiv \bar{\psi}_{\nu_\mu} \gamma_\mu P_L \psi_\mu \quad (1.43)$$

$$V_\mu^\dagger \equiv \bar{\psi}_n \left(g_V \gamma_\mu + \frac{ig_M}{2m_p} \sigma_{\mu\nu} q^\nu + \frac{g_S}{m_\mu} q_\mu \right) \psi_p \quad (1.44)$$

$$A_\mu^\dagger \equiv \bar{\psi}_n \left(g_A \gamma_\mu + \frac{ig_T}{2m_p} \sigma_{\mu\nu} q^\nu + \frac{g_P}{m_\mu} q_\mu \right) \gamma^5 \psi_p \quad (1.45)$$

where the form factors g_V , g_A , etc. arise due to the inner structure of the proton and neutron and are functions of the energy scale of the interaction q^2 . This can be related back to the low-energy extension of the electroweak theory described in Sec. 1.2.1 by examining the relevant matrix element that arises in the SM

$$\mathcal{M} = \frac{g^2}{2} Q^\mu \frac{1}{q^2 - m_W^2} L_\mu \quad (1.46)$$

with $Q_\mu \equiv \bar{\psi}_d \gamma_\mu V_{\text{CKM}}^\dagger P_L \psi_u$ and taking the limit $q^2 \rightarrow 0$. With this perspective, the proton and neutron states can be thought of as a single bound, interacting valence quark (u or d , respectively) that is *dressed* by strong interactions. This is the modern understanding of *isospin*, where the weak Hamiltonian is expressed in terms of a vector and axial-vector current constructed from isospin doublets

$$V_\mu^j = \bar{\Psi} \gamma_\mu \frac{\vec{\tau}_j}{2} \Psi \quad (1.47)$$

$$A_\mu^j = \bar{\Psi} \gamma_\mu \gamma^5 \frac{\vec{\tau}_j}{2} \Psi \quad (1.48)$$

$$\Psi \equiv \begin{pmatrix} \psi_p \\ \psi_n \end{pmatrix}, \quad (1.49)$$

with $\vec{\tau}_j \equiv \vec{\sigma}_j$. From this, the charge-changing operators are defined $V_\mu^\pm \equiv V_\mu^1 \pm iV_\mu^2$, $A_\mu^\pm \equiv A_\mu^1 \pm iA_\mu^2$, and the corresponding conservation law analogous to charge conservation ($\partial_\mu (\bar{\psi} \gamma_\mu \vec{\tau}_j \psi) = 0$) can be deduced. This is known as the conserved vector current (CVC) hypothesis [48]. Violation of CVC is expected to be $O((m_p - m_n)/2m_p)$ where it is the mass difference between the proton and neutron that breaks the isospin symmetry [49].

From this approximate symmetry, a variety of approximate rules can be derived for low-energy weak interactions. ⁶ Relevant to muon capture, this implies $g_S, g_T \rightarrow 0$ as $q^2 \rightarrow 0$, and the vector and magnetic form factors, g_V and g_M , can be related to the electromagnetic form factors F_1^p , F_1^n , F_2^p and F_2^n which have been measured precisely in electron scattering experiments. The axial form factor at $q^2 = 0$ is precisely constrained by neutron beta decay, $g_A = 1.2755 \pm 11$ [51], and its shape with q^2 has been determined by neutrino scattering on hydrogen [52]. The pseudoscalar form factor has only been measured via muon capture and is determined to be $g_P(-0.88m_\mu^2) = 8.06 \pm 0.55$ [53].

⁶See Ref. [50] for a good review.

1.2.5 Muon capture on nuclei

The previous section described muon capture in the case of hydrogen, i.e. capture with a pure proton state. In moving to muon capture on nuclei, one would naively expect that the muon capture rate would be proportional to Z^4 due to the reduced volume of the muon wavefunction and the increased number of protons. This is generally true, although typically an effective Z (Z_{eff}) is introduced to account for the charge distribution of the nucleus – an effect more pronounced for larger Z . The formalism described in the previous section can be applied to nuclear capture. However, modifications of the weak form factors due to the many-body physics of the nucleus are not well understood [47]. In a nuclear environment, there exist *exchange currents* whereby a virtual particle (e.g. pion, kaon, etc.) can be exchanged between the interacting nucleon and the nuclear medium during the interaction. Calculations of these effects are challenging due to the fact that they arise within the non-perturbative regime of quantum chromodynamics (QCD).

Primakoff [54] made an early, non-relativistic attempt to calculate the Z -dependence of muon capture. Because muon capture has a large Q -value compared to typical nuclear energies, he assumed that all nuclear transitions that involved the appropriate isospin change ($\Delta I = -1$) contribute to the matrix element and that terms involving the final-state energy could be replaced with averaged values over the final nuclear states. This led to the Primakoff formula, which most prominently shows the effect of the Pauli exclusion principle on the final-state neutron

$$\lambda_c = Z_{\text{eff}}^4 X_1 \left[1 - X_2 \left(\frac{A - Z}{2A} \right) \right] \quad (1.50)$$

where the constants $X_1 \approx 170 \text{ s}^{-1}$ and $X_2 \approx 3.125$ [47] are generally determined by experiment. This was later extended to a four parameter model [55]

$$\lambda_c = Z_{\text{eff}}^4 K \left[1 + \beta_1 \frac{A}{2Z} - \beta_2 \left(\frac{A - 2Z}{2Z} \right) - \beta_3 \left(\frac{A - Z}{2A} + \frac{A - 2Z}{8Az} \right) \right] \quad (1.51)$$

where a global fit provides $K = 261 \pm 9 \text{ s}^{-1}$, $\beta_1 = -0.040 \pm 0.002$, $\beta_2 = -0.26 \pm 0.02$, and $\beta_3 = 3.24 \pm 0.01$ [56].⁷ This generally reproduces measurements of even- Z nuclei to $\pm 10\%$ and odd- Z nuclei to $\pm 20\%$, see Fig. 1.6. There are some modern evaluations of the muon capture rate on nuclei using quantum Monte Carlo (QMC) techniques [57] and the random-phase approximation (RPA) [58, 59] which offer some improvement over the Goulard-Primakoff equation, but still only achieve agreement to within $\sim 10\%$.

1.2.6 Muon decay in orbit

Because negative muons are bound to the nucleus, the decay is modified by three main features. First, because the muon is bound, there is less energy available for the final-state

⁷The uncertainties quoted here are the difference in the best fit values with and without the TRIUMF data listed in Table V of Ref. [56].




Figure 1.6: Comparison of the Goulard-Primakoff model with experimental results on muon capture. This figure was taken from [56] and the parameters used for the comparison are listed in the text.

particles, reducing the total phase space. The phase space of the free decay is proportional to m_μ^5 , whereas the phase space of the bound decay is proportional to $(m_\mu - E_b)^5$. Assuming a simple Coulomb potential, $E_b = \frac{1}{2}m_\mu(Z\alpha)^2$, so the ratio of the bound decay to the free decay considering only the available phase space follows

$$R_H \equiv \frac{\lambda_{\text{DIO}}}{\lambda_{\text{DAR}}} \approx 1 - \frac{5}{2}(Z\alpha)^2. \quad (1.52)$$

Thus the not-so-insignificant binding energy tends to reduce the decay rate. Second, because the muon has a non-zero velocity and momentum distribution with respect to the nucleus, it will exhibit relativistic time dilation from the lab frame and cause a center-of-mass boost to the decay products. This also has the effect of reducing the decay rate and Doppler-broadening energy distributions of the decay products. Finally, the Coulomb attraction of the nucleus modifies the wavefunction of the outgoing electron from the naive plane-wave approximation, and the finite size of the nucleus modifies the initial muon state away from a pure Coulomb potential state. As the outgoing electron is attracted to the nucleus, it should be expected that this enhances the overlap of the initial muon and final electron wavefunctions, thus increasing the decay rate. While on the other hand, the finite nuclear size broadens the muon wavefunction, slightly reducing the decay rate.

Overall, estimating the balance of these effects is not trivial and typically relies on experimental measurements of nuclear charge distributions or ad-hoc approximations to the nuclear potential or wavefunctions. But it has been done by a number of authors [60, 61, 62, 63, 64, 65], with the most widely known being Huff's calculation in 1961 [63], and giving the namesake to the Huff factor R_H in Eq. 1.52.

The calculation of the decay rate comes down to the calculation of the following matrix

element

$$\mathcal{M} \propto \int d^3\vec{r} e^{-i(\vec{k}+\vec{k}')\cdot\vec{r}} \left(\bar{u}_r^{\nu\mu}(\vec{k}) \gamma^\lambda P_L u_s^\mu(\vec{r}) \right) \left(\bar{u}_{E,j,s'}^e(\vec{r}) \gamma_\lambda P_L v_{r'}^{\nu e}(\vec{k}') \right), \quad (1.53)$$

which differs from the free muon case (Eq. 1.35) only in that the muon and electron wavefunctions have been expressed into a spatial basis and the V-A interaction is now assumed. The neutrino-half of this matrix element is well understood and does not pose challenges for the calculation as the neutrino masses are negligible and the neutrino is unaffected by the nucleus. However, various assumptions must go into determining the muon and electron wavefunctions. The results of this are quite varied, but generally present a series expansion for low- Z nuclei in the form of $R_H = 1 - \kappa(Z\alpha)^2$, where κ ranges from 1/2 to 11/2 [65]. The work of Chatterjee, et al. [60] differs from the others in indicating an enhancement in the decay rate near $Z \sim 20$ arising from the electron wavefunction. Huff's calculation [63] is perhaps the most reliable, using nuclear charge distributions derived from electron scattering data [66] to determine the electron and muon wavefunctions. His results suggests that the effects mentioned before largely cancel, producing only a slight deviation from unity out to fairly large Z , e.g. $R_H = 0.975$ for $Z = 23$.

The energy spectrum of the decay electron is also modified by these same effects from that of free muon decay, shown in Fig. 1.7. Porter and Primakoff provided the first estimates of this in their paper [67], but there has been recent interest in this problem due to its relevance for calculating the background for search for lepton number violating $\mu^- \rightarrow e^-$ conversion. In general, the attractive force of the nuclear charge shifts the spectrum to lower energies, while the ground-state momentum of the muon produces a tail beyond that of the free-muon decay end-point.

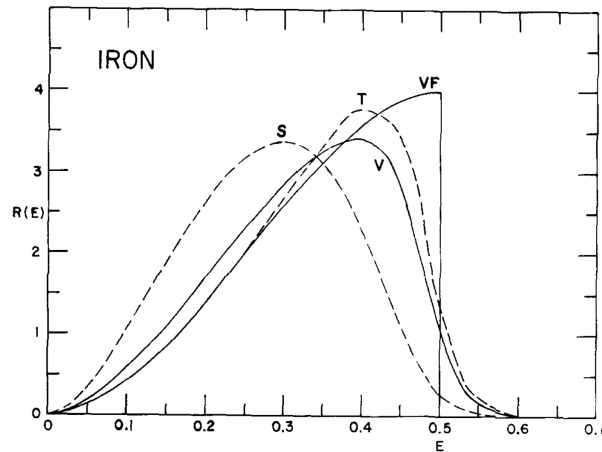


Figure 1.7: Modified decay electron spectrum for muon-Fe decay-in-orbit (DIO) under different assumptions for the muon decay coupling (S = scalar, V = vector, T = tensor). The vacuum decay spectrum (VF) is shown for reference. Units of energy are $(m_\mu + m_e)/m_\mu$.

Relatively few experimental measurements on the Huff factor and the muon DIO spectrum exist. On the Huff factor, there was some excitement and controversy related to the results of [68] and [69] due to an apparent sharp peak in the DIO rate in iron. These results disagreed with both [70] and the prediction from Huff for iron. And the results of Blair, et al. [71] also corroborated the Huff calculation for high- Z and did not exhibit evidence of a peak. At this point, it is generally assumed that the peak was caused by backgrounds from nuclear de-excitation photons that were included when counting the number of DIO events, thus inflating the apparent decay rate. For the high- Z discrepancy, the earlier results did not correct for the fraction of decay electrons that fall below the energy threshold required to avoid de-excitation photons from capture events. This is particularly important at high- Z , where the electron loses a substantial fraction of energy when escaping from the deep nuclear potential. As an example, the fraction of decay electrons below 10 MeV is $\approx 30\%$ of the total electron spectrum for $Z = 82$ [71].

For the decay spectrum, precise measurements do exist on an Al target produced by the TWIST collaboration [72]. They were able to measure down to electron energies of ≈ 20 MeV and were sensitive enough to probe second-order radiative corrections. Overall, they observe good agreement with theoretical expectations once including these corrections. Other measurements exist for heavier nuclei within the context of $\mu^- \rightarrow e^-$ experiments, but they are limited to the very high-energy tail of the electron spectrum [73, 74, 75, 76].

1.3 Relationship to neutrino experiments

At this point, measurements of muon capture on hydrogen are extremely precise [53] and lead the constraint on the pseudoscalar form factor g_P . Many measurements exist for nuclear muon capture across Z and include measurements of both the capture rate on nuclear isotopes as well as differential measurements of final states. As a higher-energy process than β -decay, it provides a useful test of high-energy physics models at low energies and in nuclear environments, while also providing information about high-energy matrix elements used for β -decay and other nuclear-effect calculations.

This section introduces DUNE and discusses some of the systematic uncertainties that are involved in neutrino oscillation experiments. In Sec. 1.3.2, these uncertainties are linked to muon capture, highlighting how measurements of muon capture can provide input to neutrino cross-section models. Finally in Sec. 1.4, the existing measurements of muon capture are summarized for the target material in DUNE, argon, and the disagreement in the most recent results is briefly discussed.

1.3.1 Deep Underground Neutrino Experiment (DUNE)

DUNE is an upcoming neutrino experiment that aims to measure the neutrino mass hierarchy, the CP -violating phase δ_{CP} , supernova neutrinos, and search for beyond-the-Standard-Model (BSM) physics.

In recent decades, it has been concluded that contrary to the prediction of the SM, neutrinos have a small, non-zero mass [77, 78, 79, 80]. Because of this, a flavor mixing matrix analogous to the V_{CKM} matrix mentioned earlier can be introduced into the weak leptonic interactions. Thus there are two natural bases to describe neutrino states: the mass eigenstates $|\nu_1\rangle, |\nu_2\rangle, |\nu_3\rangle$, generated by an as-of-yet unknown mechanism; and the flavor eigenstates $|\nu_e\rangle, |\nu_\mu\rangle, |\nu_\tau\rangle$, which are involved in weak interactions. Generally, this phenomenon is represented by a 3×3 unitary matrix U_{PMNS} called the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) mixing matrix [81, 82]. This matrix has 4 or more degrees of freedom, including 3 real mixing angles θ_{12}, θ_{23} , and θ_{13} and at least one complex phase δ_{CP} . The mixing angles represent the relative fractions of the neutrino mass eigenstates that are produced in weak interactions. And based on an argument of the number of inconsequential phase rotations that can be applied to the states involved in the $\bar{l}\sigma_a^\mu W_\mu^{a-} U_{\text{PMNS}}\nu$ interaction term, it can be shown that at least one complex phase is allowed within the PMNS matrix, δ_{CP} . Additional CP -violating phases are allowed if the neutrino is a Majorana particle, but no evidence has yet to be found that suggests the neutrino is not a Dirac particle.

Our understanding of this mixing matrix is largely complete from measurements of neutrino oscillation. Within these experiments, a neutrino source is produced in a pure $|\nu_{e,\mu,\tau}(E)\rangle$ state and propagates according to the mass eigenstates $|\nu_{1,2,3}(E)\rangle$. For oscillations in vacuum, this results in each mass eigenstate accumulating a different phase

$$|\nu_\alpha(L, E)\rangle \approx \sum_j e^{-im_j^2 L/2E} U_{\alpha j}^* |\nu_j(E)\rangle \quad (1.54)$$

where $U \equiv U_{\text{PMNS}}$, m_j are the neutrino masses, L is the distance from the production point, and α is the flavor of the lepton involved in the interaction that produced the neutrino. The resulting probability of the neutrino interacting and producing a lepton flavor β is then

$$|\langle \nu_\beta(E) | \nu_\alpha(L, E) \rangle|^2 = \left| \sum_j U_{\alpha j}^* U_{\beta j} e^{-im_j^2 L/2E} \right|^2 \quad (1.55)$$

which for 2-neutrino oscillation simplifies to be

$$P_{\alpha \rightarrow \beta}(L) \approx \sin^2 2\theta_{ij} \sin^2 \left(\frac{\Delta m_{ij}^2 L}{4E} \right) \quad (1.56)$$

with $\Delta m_{ij}^2 \equiv m_i^2 - m_j^2$ being the mass-squared difference of the two primary neutrino states involved in the oscillation and θ_{ij} the mixing angles in the U_{PMNS} matrix parameterization, defined by $\sin^2 \theta_{ij} \equiv |U_{ij}|^2 / [1 - (1 - \delta_{ei}\delta_{j3})|U_{e3}|^2]$. Thus, there is a non-zero probability of observing neutrinos of a different flavor at a distance from the source that depends on the flavor mixing matrix parameters, the neutrino mass-squared differences, and the factor L/E . This phenomenon is well established and the mixing angles have been measured and are summarized in Table 1.3.

Table 1.3: Mixing angles as reported by NuFit 5.2 (2022) [83], where $\Delta m_{3l}^2 = \Delta m_{31}^2 (\Delta m_{32}^2)$ for normal (inverted) ordering.

Parameter	Best fit		Units
	normal	inverted	
θ_{12}	$33.41^{+0.75}_{-0.72}$	$33.41^{+0.75}_{-0.72}$	$^\circ$
θ_{23}	$42.2^{+1.1}_{-0.9}$	$49.0^{+1.0}_{-1.2}$	$^\circ$
θ_{13}	$8.58^{+0.11}_{-0.11}$	$8.57^{+0.11}_{-0.11}$	$^\circ$
δ_{CP}	232^{+36}_{-26}	276^{+22}_{-29}	$^\circ$
Δm_{21}^2	$7.41^{+0.21}_{-0.20}$	$7.41^{+0.21}_{-0.20}$	10^{-5} eV^2
Δm_{3l}^2	$+2.507^{+0.026}_{-0.027}$	$-2.498^{+0.032}_{-0.025}$	10^{-3} eV^2

However, of particular interest is the complex phase δ_{CP} which violates CP-symmetry. This can be observed experimentally by examining the difference in oscillations of neutrino and anti-neutrino states which contains a term depending on δ_{CP}

$$\Delta P \equiv P_{\mu \rightarrow e}(L) - P_{\bar{\mu} \rightarrow \bar{e}}(L) = 4 \sum_{i < j} \Im [U_{\mu i} U_{ei}^* U_{\mu j}^* U_{ej}] \sin \left(\frac{\Delta m_{ij}^2 L}{2E} \right) \quad (1.57)$$

$$\supset -16 J_{CP}^{\max} \left(\frac{\Delta m_{21}^2 L}{4E} \right) \sin(\delta_{CP}) \sin^2 \left(\frac{\Delta m_{31}^2 L}{4E} \right) \quad (1.58)$$

with $J_{CP}^{\max} \equiv c_{12}s_{12}c_{23}s_{23}c_{13}^2s_{13}$, $c_{ij} \equiv \cos \theta_{ij}$, $s_{ij} \equiv \sin \theta_{ij}$, and $\delta_{CP} \equiv -\arg U_{e3}$ [24]. The value of δ_{CP} has been measured recently by the T2K and Nova experiments but with large uncertainties.

Also of interest is the neutrino *mass hierarchy*, which is the ordering of the neutrino mass eigenstates. As measurements of neutrino mass through oscillation are only sensitive to the differences in the neutrino masses, only the relative masses are known and the lightest neutrino could be massless. The experimental constraints are listed in Table 1.3 and allow the possibility of either two scenarios: $m_1 < m_2 < m_3$ or $m_3 < m_1 < m_2$ also called the *normal* and *inverted* hierarchies. Here the mass states 1 and 2 are defined by the smallest mass-squared difference $\Delta m_{21}^2 > 0$ and the third mass state is defined such that $|\Delta m_{31}^2| \approx |\Delta m_{32}^2|$. Because existing constraints on the mass hierarchy come from oscillation measurements of the form in Eq. 1.56, there is an inherent degeneracy in the sign of Δm_{ij}^2 , limiting experiments' ability to determine the mass ordering.

To overcome this, the *matter effect* [84, 85] can be exploited. When a neutrino mass eigenstate propagates in matter, the electron flavor component evolves differently than the

others due to the coherent charged-current interaction $\nu_e + e \rightarrow e + \nu_e$. This has the effect of introducing a potential term into the effective Hamiltonian only for the electron flavor component of $V = \sqrt{2}G_F n_e$ where n_e is the number density of electrons in the medium. This modifies the flavor makeup of the propagating eigenstates, resulting in a CP -violating oscillation probability term of

$$\begin{aligned} \Delta P = & -16 \operatorname{sign}(\Delta m_{31}^2) J_{\text{CP}}^{\max} \sin(\delta_{\text{CP}}) \\ & \times \frac{|\Delta m_{21}^2|}{|\Delta m_{31}^2|} \sin(|\Delta m_{31}|^2 L/4E) \\ & \times \frac{\sin(|A||\Delta m_{31}|^2 L/4E)}{|A|} \times \frac{\sin((1-A)|\Delta m_{31}|^2 L/4E)}{1-A} \end{aligned} \quad (1.59)$$

where $A \equiv 2EV/\Delta m_{31}^2$. The form used in Eq. 1.59 deliberately shows that with the inclusion of matter effects the oscillation probability difference is then sensitive to not only δ_{CP} through its magnitude, but also to the mass hierarchy determined by $\operatorname{sign}(\Delta m_{31}^2)$.

DUNE will measure these parameters by using a broad-band ν ($\bar{\nu}$) beam and a long baseline such that the CP -violating and matter effect are prominent. As a neutrino source, DUNE will use the Long-Baseline Neutrino Facility (LBNF) beam based at Fermi National Accelerator Laboratory (FNAL), produced by a 1.2-MW 120-GeV proton beam. The proton beam will be focused onto an actively-cooled 2-m graphite target to produce a secondary hadron beam, predominately consisting of pions and kaons. A series of three magnetic-field structures called *horns* surrounding the target will then selectively focus positive or negative secondary particles. These then decay within a 220-m helium-filled *decay pipe* to produce a beam of predominantly muon neutrinos or anti-neutrinos, depending on the direction of the current in the magnetic horns [86].

A near-detector complex will be placed 574 m from the target and 60 m underground on the beam axis, and will host a suite of detectors used to measure the neutrino beam characteristics and constrain uncertainties in the beam flux and neutrino interaction models. The primary component of the near detector is the ND-LAr detector consisting of a movable, 5m×7m×3m modular LArTPC, scaling up the technologies described in Ch. 2. Additional detectors will be employed to constrain nuclear-physics models, the on-axis flux, improve the coverage of the ND-LAr detector, and provide more detailed measurements of ν -argon interactions [87].

A far detector will be situated 1500 m underground at the Sanford Underground Research Facility (SURF) in South Dakota, providing a oscillation baseline of 1300 km. The far detector will use four 10-kt fiducial-mass detector modules, with at least three being LArTPCs. The large-scale and low-background nature of the far-detector modules also enables a broader physics program than just δ_{CP} and the mass hierarchy, including the potential detection of supernova neutrinos and searches for a variety of BSM physics [86].

1.3.2 Neutrino cross-sections

The observed neutrino spectrum of flavor α and energy E at the DUNE far-detector is given by

$$n_{\alpha}^{\text{far}}(E) = \int \Phi_{\alpha}^{\text{far}}(E') \sigma_{\alpha}^{\text{far}}(E') D_{\text{far}}(E', E) dE' \quad (1.60)$$

where $\Phi_{\alpha}^{\text{far}}$ is the neutrino flux of flavor α , $\sigma_{\alpha}^{\text{far}}$ is the ν -argon cross-section, and D_{far} is a smearing matrix that accounts for detector and reconstruction effects. The far-detector flux is represented by a fictitious un-oscillated flux and the neutrino oscillation probability, $\Phi_{\beta}^{\text{far}}(E) = \sum_{\alpha} \Phi_{\alpha}^{\text{far,unosc}}(E) P_{\alpha \rightarrow \beta}(E)$, providing access to the oscillation parameters.

Many of the uncertainties can be controlled by measuring the correlated near-detector neutrino spectrum

$$n_{\alpha}^{\text{near}}(E) = \int \Phi_{\alpha}^{\text{near}}(E') \sigma_{\alpha}^{\text{near}}(E') D_{\text{near}}(E', E) dE' \quad (1.61)$$

where the oscillation probabilities can be assumed to be small. Within DUNE, many aspects of the near-detector design aim to reduce the impact of uncertainties by making the near detector as similar as possible to the far detector. In this scenario and at a fixed neutrino energy, $\Phi_{\alpha}^{\text{near}} \approx \Phi_{\alpha}^{\text{far,unosc}}$, $\sigma_{\alpha}^{\text{near}} \approx \sigma_{\alpha}^{\text{far}}$, and $D_{\text{near}} \approx D_{\text{far}}$, and the measurement of the oscillation probability can be made directly using the far-to-near ratio. However, even in the scenario that these approximations could be made exact, the fact that the neutrino oscillation probability is a function of energy means that the far detector ultimately observes a different flux than the near detector, and cancellation will never be perfect. Because of this, changes in the neutrino cross-section as a function of energy can introduce non-trivial changes in the number of observed events at the far detector, distorting the extracted neutrino oscillation probability. In addition, the detector model is in actuality better represented by a sum over neutrino interaction types

$$\sum_{\text{int}} \sigma_{\alpha}^{\text{int}}(E) D_{\text{det}}^{\text{int}}(E, E') \approx \sigma_{\alpha}(E) D_{\text{det}}(E', E), \quad (1.62)$$

which is not cleanly separable and introduces additional cross-section correlated systematic uncertainties into the detector model. Near-detector measurements of the flux also rely on knowledge of the neutrino cross-section and thus a bias in the cross-section model can introduce a change in the near- and far-detector flux models. Finally, cross-sections differ quite significantly between ν and $\bar{\nu}$ leading to effects that can mimic δ_{CP} -associated oscillations. Therefore, strong understanding of neutrino cross-sections will be critical in the ultimate success of DUNE.

Neutrino interactions within the $O(\text{GeV})$ regime are highly complex, depending strongly on the 4-momentum transferred to the struck particle. They are typically classified into *quasi-elastic (QE)*, *resonant (RES)*, and *deep inelastic scattering (DIS)*, shown in Fig. 1.8




Figure 1.8: Summary of neutrino cross-section measurements and predictions for the three interaction channels described in the text, provided by the NUANCE neutrino interaction generator. Figure is taken from Ref. [88].

At the lowest energies, below about 2 GeV, QE scattering dominates. In these events, the neutrino scatters off of a quasi-free nucleon. In CC interactions, the nucleon undergoes a change in isospin ($\nu_l + n \rightarrow l^- + p$ and $\bar{\nu}_l + p \rightarrow l^+ + n$). And in neutral-current (NC) interactions, the nucleon b is unchanged ($\nu_l + b \rightarrow \nu_l + b$ and $\bar{\nu}_l + b \rightarrow \bar{\nu}_l + b$). The cross-sections for these interactions are expressed following a similar formalism as was introduced in Sec. 1.2.4 in Eq. 1.42. Namely, the nucleon current is represented using electroweak form factors that depend on the inner structure of the nucleon and are modified in the nuclear environment.

Above about 300 MeV, resonant processes dominate, whereby the struck nucleon is excited into a resonant state and quickly decays, typically to a pion and a baryon. Particularly challenging for the reconstruction of these interactions is the possibility of *final-state interactions*, which involve the re-interaction of the secondary particles with the nuclear medium. When this occurs, many processes can occur, e.g. the pion can be absorbed, mimicking a QE-like interaction. This leads to difficulties in making clear measurements of a given interaction cross-section.

Finally, at high energies (above \sim GeV), neutrinos interact by scattering from an individual parton contained within the nucleon. This is called DIS. Within this regime, form factors can be represented using sums over parton distribution functions which are probability distribution functions (PDFs) over the fraction of the nucleon momentum $x \equiv Q^2/2(p_l \cdot q)$, e.g.,

$$F_2(x, Q^2) = 2 \sum_q xq(x, Q^2) + x\bar{q}(x, Q^2), \quad (1.63)$$

where $Q^2 \equiv -q^2 \equiv -(p_\nu - p_l)^2$, q (\bar{q}) is the (anti-)quark parton distribution function, and

F_2 is the magnetic form factor mentioned in Sec. 1.2.4.

As outlined in Sec. 1.2, nuclear effects modify the nucleon form factors in weak interactions. The same modifications that are found in muon capture can be also applied to neutrino scattering, and thus muon capture can provide a direct, independent benchmark for neutrino QE cross-sections at $Q^2 \approx 0.01 \text{ GeV}^2$. While this is much lower than the typical neutrino energy in DUNE, forward QE neutrino scattering produces a non-negligible number of events with small Q^2 even at these higher energies. As observed by a number of neutrino experiments [89, 90], and most recently in ν -argon by the MicroBooNE experiment [91], neutrino scattering data suggests significant disagreements between model predictions and the observed cross-section, particularly within this forward region. This effect is likely related to the quenching of the axial form factor, as predicted in RPA calculations [92], and has been introduced as a systematic into neutrino oscillation results by T2K and MINER ν A. The scale of this systematic is set by comparisons of the theoretical predictions and measurements of muon capture on the target material. Thus improving the experimental measurements of muon capture can provide an independent means to more reliable neutrino cross-section models.

1.4 Existing measurements

Muon capture on argon has been measured a number of times since the early 1970s. These experimental results are summarized in Table 1.4.

The first measurement [93] used a muon beam produced by the European Council for Nuclear Research (CERN) 600-MeV synchrotron and stopped in gaseous deuterium target, doped with trace argon gas. Because the muon binds more strongly to the argon nucleus, a

Table 1.4: Summary of existing muon capture measurements for argon.

Ref.	Muon source	Target material	τ_d (ns)	λ_c (μs^{-1})	R_H used
[93]	μ^- -beam (CERN)	$\text{D}_{2(g)} + \text{Ar}_{(\text{trace})}$	-	1.20 ± 0.08	1
[94]	μ^- -beam (CERN)	$0.95\text{Ar}_{(g)} + 0.05\text{C}_3\text{H}_{8(g)}$	537 ± 32	1.41 ± 0.11	0.99
[95]	μ^- -beam (JINR)	$\text{Ar}_{(s)}$	568 ± 6	-	-
[96]	μ^- -beam (PSI)	$\text{Ar}_{(g)}$	616.9 ± 6.7	-	-
[97]	cosmic-ray μ^\pm	$\text{Ar}_{(l)}$	625 ± 48	$1/(0.871 \pm 0.093)$	0.988

well-established transfer reaction $\mu^-(\text{low-}Z) + (\text{high-}Z) \rightarrow (\text{low-}Z) + \mu^-(\text{high-}Z)$ can occur. They reported a capture rate of $\lambda_c = 1.20 \pm 0.08 \mu\text{s}^{-1}$ from the time spectrum of the decay electrons, but did not include the small effects of the modified muon DIO process in their calculation. It was measured again a few years later by Carboni, et al. [94] via muons stopping in a low-pressure argon gas with a 5% mix of propane. They included measurements of the prompt muonic X-ray spectrum from the direct atomic muon capture process as well as a transfer reaction attributed to $\mu^-H + \text{Ar} \rightarrow H + \mu^-\text{Ar}$. They found the total capture rate to be $\lambda_c = 1.41 \pm 0.11 \mu\text{s}^{-1}$ ($\tau_d = 537 \pm 32$ ns), in reasonable agreement with the previous result.

More recently, Mamedov, et al. [95] and Klinskikh, et al. [96] achieved precision measurements using improved muon beam facilities. Mamedov, et al. used a standard approach of measuring the decay-electron disappearance rate of muons stopping in a frozen argon target placed in the 125-MeV synchrocyclotron at the Joint Institute for Nuclear Research, Dubna (JINR). A magnetic field of ≈ 0.18 T was used to measure the residual muon polarization. Their target chamber was made from 99.9% pure copper ($\tau_d(\text{Cu}) = 163.5 \pm 1.0$ ns [56]) and was included as a component in their fit. They measured $\tau_d = 568 \pm 6$ ns and no measurable residual muon polarization $< 1.5\%$. Klinskikh, et al. used a 25-45 MeV muon beam generated by the Paul Scherrer Institute (PSI) meson factory, stopped in a gaseous argon target. Different from previous measurements, their result used the decay spectrum of the de-excitation photons produced by the final-state excited nuclei. As such, they were also able to provide a measurement of the relative yield final-state nuclei in μ -capture on argon. Their result was $\tau_d = 616.9 \pm 6.7$ ns, substantially higher than the previous measurements. Thus we are left in an unsatisfying position, with two precision measurements that disagree at 5 standard deviations. This certainly motivates further experimental exploration.

It is also worth mentioning that recently, the LArIAT experiment published a measurement of the muon capture rate in argon using cosmic-ray muons stopped in a liquid argon detector [97] similar to the technique used in this thesis. They observed a disappearance rate of $\tau_d = 625 \pm 48$ ns, also higher than previous measurements, albeit with a significantly larger error.

Chapter 2

Detector technology

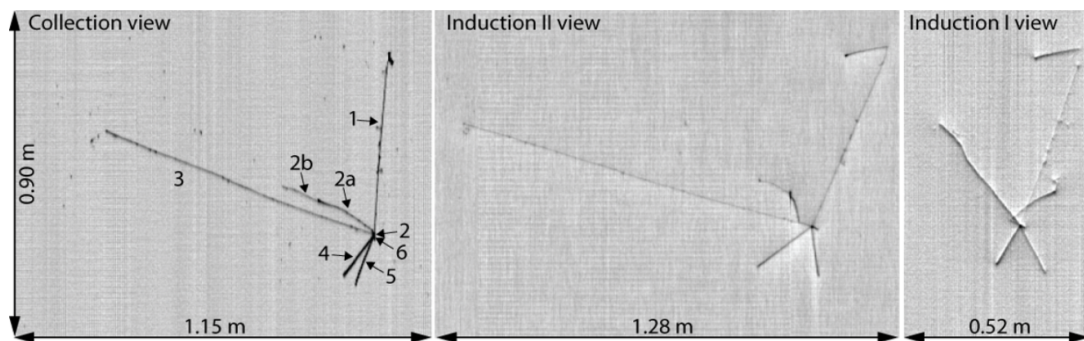


Figure 2.1: Example neutrino interaction as imaged by the ICARUS LArTPC [98].

Inseparable from the development of our theoretical understanding of physics has been the development of advanced detector technologies that enable us to observe high-energy particle interactions. This chapter describes one such technology – the liquid argon time projection chamber (LArTPC). The general characteristics and operating principles of LArTPCs are provided in Sections 2.1 and 2.2. Because of these characteristics, LArTPCs lend themselves particularly well to the study of neutrino interactions and is the reason that they will be the primary detector technology used for DUNE. Specific to the DUNE near detector, high interaction rates require improvements over the traditional design. These have been realized by the ArgonCube concept, discussed in Sec. 2.3.

2.1 Tracking detectors

From very early on, it has been understood that observing a particle’s trajectory is particularly valuable when reconstructing high-energy interactions. A wealth of information is available from the path of the particle and the amount of ionization that the particle




Figure 2.2: Characteristic energy loss per unit length for different particle species, taken from Ref. [99].

produces. Magnetic fields can be used to perform charge identification and momentum measurement via the particle's cyclotron radius

$$rB = \frac{p_{\perp}}{q} \quad (2.1)$$

where r is the radius of curvature, p_{\perp} is the component of the particle's momentum perpendicular to the magnetic field B , and q is the charge of the particle. The identification of different particle species and calorimetry can be achieved by observing the energy loss along the trajectory, which depends on the particle's mass, charge, and velocity, shown in Fig. 2.2. Typically with tracking detectors, a low-density detection material is used so as not to significantly perturb the particle during the measurement. However, high-density materials have also been employed as a target medium, to induce electromagnetic (EM) or hadronic showers, or measure the particle energy.

The earliest forms of tracking detectors were Geiger-like detectors in which secondary ionization produced by a large electric field in a gaseous medium to induce an electron avalanche within the detector when a particle traverses it. An array of wires was used

to provide positional information, as only the wires nearest to the avalanche produce an appreciable signal. This provides positioning, but is limited to saturated signals (i.e. no information about the energy loss along the trajectory) and has fundamental limitations in their spatial resolution due to the construction. Mult-wire proportional chambers [100] improved on the first point by operating below the Gieger regime, but above the secondary ionization threshold, producing electron gain. Thus the signal produced is proportional to the ionization deposited in the detector, enabling particle discrimination by the dE/dx . As it takes some time for the avalanche to collect on the wire ($O(\mu s)$), the spatial resolution was improved with the introduction of the drift chamber, which considers time it takes for the ionization to drift to the wire.

The time projection chamber (TPC) took the principle of the drift chamber to the extreme. It was introduced by Nygren [99] as a tracking detector consisting of:

1. a low-mass tracking material capable of supporting large electric fields without breakdown and with low electron affinity
2. an $O(100 \text{ V/cm})$ electric field applied uniformly across a large volume of the target material
3. and a sensitive, multi-electrode anode capable of detecting signals at the scale of tens to tens of thousands of electrons.

When a charged particle passes through the gas mixture, it liberates electrons from the parent atoms through ionization. In the presence of an electric field, the free electrons drift towards the anode. If the material has a low electron affinity, the electrons can travel largely unimpeded to the anode over a timescale of $O(> \mu s)$. At the anode, the electrons can be sensed either by direct collection, or via the inductive current caused by the changing electric field as the electrons drift. With a highly segmented anode, each segment of the anode will see a different signal caused mostly by the electrons close to or collected on the sensing element. Because of this, one can reconstruct the drifting electrons' position at the time they reach the anode by using the relative amplitude of the signal on each anode segment. Combined with knowledge of the initial event time t_0 , one can then calculate the original ionization position z_0 within the detector volume with the simple formula

$$z_0 = (t_{\text{drift}} - t_0)v_{\text{drift}} + z_{\text{anode}} \quad (2.2)$$

where z_{anode} is the position of the anode, v_{drift} is the drift velocity of the charge cloud, and t_{drift} the arrival time of the charge at the anode. The ability to reconstruct the three-dimensional position is limited practically only by the degree of anode segmentation, the uniformity of the electric field, and the diffusion of the electron cloud as it travels through the medium.

2.2 Liquid argon time projection chambers (LArTPCs)

Table 2.1: Summary of liquid-argon properties. Where applicable, values are reported for liquid phase at 87.3 K, 1 atm, and 500 V/cm.

Property	Value	Unit	Ref.
Atomic number	18		
Molar mass	39.94793 ± 0.00003	g/mol	[101]
Boiling point	87.303 ± 0.002	K	[102]
Freezing point	83.8 ± 0.3	K	[102]
Density	1.396 ± 0.001	g/mL	[102]
Dielectric constant	1.504 ± 0.001		[103]
Dielectric strength	>10	kV/cm	[104]
Ionization W-value	23.6 ± 0.3	eV/e ⁻	[105]
Scintillation W-value	19.5 ± 0.2	eV/photon	[106]
Fano factor	0.107 ± 0.005		[107]
Radiation length	14.0	cm	[24]
Moliere radius	10.0	cm	[24]
Nuclear interaction length	85.8	cm	[24]
Critical energy	32	MeV	[24]
e ⁻ mobility	318.6	cm ² /V s	[108]
e ⁻ drift velocity	0.1593	cm/ μ s	[108]
Longitudinal e ⁻ diffusion coefficient	5 ± 1	cm ² /s	[108, 109]
Transverse e ⁻ diffusion coefficient	12	cm ² /s	[110]
Ion mobility	1×10^{-3}	cm ² /V s	[111]
Ion drift velocity	5×10^{-7}	cm/ μ s	[111]
Peak scintillation wavelength	128	nm	[112]
Index of refraction (at 128 nm)	1.358 ± 0.003		[113]
Rayleigh scattering length (at 128 nm)	99.1 ± 2.3	cm	[113]


Proposed by Rubbia in 1977 [114], the LArTPC extends the TPC by using the uniquely beneficial properties of liquid argon as a target medium within a TPC. Foremost of all, liquid argon is dense, enabling a large target mass. It is also relatively cheap and easy to handle, having a boiling point just above that of readily-available liquid nitrogen. This makes it possible to acquire and store the large quantities needed for rare-event detection. And most importantly, it has the characteristics needed to operate as a TPC. Namely, it has a low electron affinity and high electron mobility. There are also very few environmental risks,

being a common component of the atmosphere, or exposure risks, besides the immediate burn and oxygen deficiency hazards of a cryogenic liquid, so control and disposal is a manageable task. An additional feature that is appealing for dark matter and neutrino experiments is that not only does argon produce ionization electrons, it is an excellent scintillator and is transparent to its own scintillation light. A significant fraction of the scintillation light is emitted within a short time after the initial energy deposition, enabling LArTPCs to be used as self-triggered detectors. These features facilitate the construction a large-mass, self-triggered detector with excellent energy, time, and spatial resolution at a relatively low cost.

2.2.1 Charged-particle signals

When a high-energy charged particle passes through liquid argon, it scatters off of the bound electrons of the argon atom. Typically, the energy transferred to the electron is small, of the order of a few ten to hundreds of eV [24]. If the energy is smaller than the atomic binding energy, the argon atom is left in an excited atomic state. In pure argon, the excited argon atoms form dimer molecules wherein an excited electron is shared between two argon atoms [115]. Dimer formation occurs within a short timescale \sim ps [116], and the subsequent dimer decay leads to the production of scintillation light. However, if the energy is larger than the atomic binding energy, the electron is freed from the parent argon ion. The electron can subsequently recombine with another argon ion, producing an excited argon atom that leads to further scintillation light. Or the electron can escape, and will drift under the influence of an electric field. Occasionally, scattered electrons are given enough energy to cause secondary ionization.

After the initial production of excited argon dimers and ionization electrons, depending on the spin state of the excited electron, these dimer states are either a singlet ($^1\Sigma_u^+$) or triplet ($^3\Sigma_u^+$) state. The radiative transitions of these states then occur, with decay times of ~ 7 ns and ~ 1600 ns respectively [117],



The emitted light spectrum for both gaseous and liquid argon scintillation light is shown in Fig. 2.3, which both scintillate via similar dimer decay mechanisms. The resulting UV scintillation photons are emitted isotropically with a spectrum peaked at ~ 128 nm [118]. Overall, liquid argon generates approximately 10^4 isotropic primary photons per MeV of deposited energy. The relative population of the singlet and triplet states depends on the ionization density, resulting in a time profile that varies with the dE/dx , and thus the particle species, with a typical ratio of ~ 0.3 . Of this, about 30% is emitted within the first 15 ns after the deposition time, with the remainder emitted slowly across the next few μs , due to the slow triplet decay. This is the basis for pulse shape discrimination




Figure 2.3: Comparison of liquid (black, thick) and gaseous (red, thin) argon scintillation spectra. The dominant dimer decay occurs at ~ 128 nm. Taken from Ref. [112]

(PSD) techniques that have been used for background rejection in dark-matter experiments. Additionally, the triplet scintillation light can be quenched in the presence of contaminants due to additional competing decay modes with shorter time scales [119, 120]. Liquid argon is largely transparent to the scintillation light with an attenuation length of ≈ 1 m [121, 113] and with a refractive index of 1.4 at 128 nm [113]. Thus the scintillation light propagates freely through the liquid argon, impacted only by Rayleigh scattering¹.

At the same time, between collisions with argon atoms, the free electrons are accelerated under the influence of the applied electric field. At macroscopic scales, the effect is a net collective motion of the electron cloud that can be described by a constant group velocity $v_d \sim \text{mm}/\mu\text{s}$, antiparallel with the electric field, and two terms σ_{\parallel} and σ_{\perp} corresponding to diffusion parallel and transverse to the applied field [123].

The charge and light production and detection characteristics of LArTPCs are thus dependent on a number of effects:

- **Primary particle:** the velocity and mass of the primary particle determines the energy transferred to the atomic electrons, which modifies the charge to light yield ratios, the number of excitations per track length, as well as the initial population of singlet and triplet states. This enables analysis techniques that separate different particle species using their characteristic energy loss [124, 125], their light signal timing [126], and by combining the light and charge signals [127].

¹This dispersive effect is small at short length scales ($\lambda_R \approx 95$ cm) [122, 113].

- **Applied electric field:** both the direction and amplitude of the electric field determines the density and motion of electrons after the initial energy deposit, as well as the total amount of scintillation light. This can distort and confuse the strategies mentioned above if it is not uniform and its magnitude is not well known.
- **Impurities:** contaminants in the argon open other decay modes for excited argon atoms and potential binding sites for freed electrons, impacting both charge and light yields and signal characteristics. This typically results in signal loss, but can also be exploited to modify the signal characteristics for a specific purpose [128, 129]. The build-up of charged argon ions can also lead to distortions in the electric field [130].
- **Medium properties:** the density and temperature of the argon medium modifies the excitations per track length and the propagation of light and electrons within the bulk volume. This can generally be controlled for by in-situ monitoring of the argon, but non-trivial fluid dynamics of the argon can introduce complexities.

2.2.2 LArTPC design

The development of LArTPCs relies on a number of liquid-argon properties, summarized in Table 2.1. To maximize their utility, LArTPC designs generally aim to use liquid argon with sub-parts-per-billion (ppb) contamination and a maximally large and uniform electric field. Recent experiments have demonstrated that liquid argon can be purified to extreme degrees, permitting electron lifetimes > 10 ms [131]. The space-charge effect (SCE) from long-lived charged ions produced by the ionizing particles is particularly challenging for long-drift detectors on the surface [91], but is not expected to be a challenge for detectors with even a moderate overburden [87]. Electric fields are typically the limiting factor in the size and operational stability of LArTPCs. Long drifts greater than 3 m have been achieved [131, 132], but not without difficulty.

Besides in dark matter (DM) experiments, the scintillation light has been somewhat neglected in designing large-scale detectors. This is primarily due to the cost and challenge of instrumenting the high-field regions of the detector, which is necessary to collect the isotropic scintillation light. The focus has instead been to collect sufficient light to generate a trigger for the more sensitive charge readout. There are some new efforts to improve the light readout capabilities for large-scale detectors [133, 134].

2.2.2.1 LArTPC examples

The original concept has been demonstrated with a number of functional LArTPCs, targeting neutrino detection as an application. The pioneering ICARUS [98] detector was the largest LArTPC ever built (760 tons) only until very recently when it was surpassed by the 1:20th scale prototype DUNE detector, protoDUNE [131]. The LArIAT [135] and ArgoNueT detectors were small LArTPCs that made important contributions to argon cross-section measurements [136, 137] and made critical measurements of energy loss for a variety of

particle species [124, 97]. The miniCAPTAIN [138] detector was another small LArTPC that made the first measurements of neutron-argon cross-sections at neutron energies of 100-800 MeV [139]. The WArP and DarkSide experiments [140, 141] demonstrated that LArTPCs could be used to search for dark matter and have placed competitive limits on weakly-interacting massive-particle (WIMP) models. More recently, the μ BooNE experiment has published the first single- and double-differential measurements of neutrino-argon cross-sections in the Booster Neutrino Beam (BNB) beam [91].

2.3 ArgonCube concept

2.3.1 DUNE LAr-ND

The LArTPC has been selected as the primary detector technology for DUNE, described previously in Sec. 1.3.1. DUNE pushes the boundaries for LArTPC design – at the far detector site, it will use at least three LArTPC detector modules, each $20\times$ larger than protoDUNE; and at the near detector site, the event rate will be larger than any LArTPC has seen.

At the far detector, the design can follow the same recipe developed by the ICARUS detector:

1. Large monolithic liquid argon volume
2. A series of sense-wire arrays at the anode, oriented at angles with respect to each other – typically two induction-sensing and one collection-sensing are used
3. A field cage structure of connected conducting panels forming a resistive chain
4. Light collection from sensitive detectors placed behind the sense-wire planes

This approach works well for large-scale LArTPCs with low event rates. Due to the reasons outlined in Sec. 1.3.2, it is highly advantageous to use a near-detector technology as similar as possible to the far-detector. Thus, DUNE requires a LArTPC component at the near-detector site [87].

However, for the near-detector site, a number of practical challenges arise when operating a monolithic, wire-based LArTPC due to the high-rate environment. Intrinsically, the charge signals in LArTPCs have high spatial resolution (\sim mm), but poor time resolution ($d/(\sqrt{12}v_d) \sim 100\mu$ s). But on the other hand, the light signals typically have poor spatial resolution (limited by the light collection area), but fast time resolution (\sim ns). In a low-rate environment, light signals can be matched unambiguously to charge signals, enabling interactions to be reconstructed with the spatial resolution of the charge signals and time resolutions of the light signals. But in a high-rate environment, the light signals can no longer be matched unambiguously to the charge signals, degrading the detector performance. Further, wire-based readout encounters intrinsic ambiguities when reconstructing three-dimensional

positions from charge that arrives simultaneously at the anode plane. This ambiguity can often be resolved in low-rate environments by using multiple wire planes and the assumption of sparsity of the energy deposition [142], however the performance degrades as the density increases (such as within high-energy showers or high-rate environments). Thus the liquid argon component of the DUNE near detector must utilize a different LArTPC technology than the far detector.

2.3.2 ArgonCube

The ArgonCube concept overcomes the challenging environment at the DUNE near-detector by adopting a modular approach. Instead of a single- or dual-TPC volume supported by a single cathode, the liquid argon active volume is broken into independent modules, with each volume containing its own cathode, dual TPCs, and anodes. Within a single module, the drift length is relatively short – reducing the requirements on the cathode voltage, the electron lifetime, and space charge. But more importantly, each module is individually instrumented with optical detectors and optically separated from the other modules. Charge-light ambiguities scale with the number of events per TPC and are therefore improved. To realize this concept, the ArgonCube collaboration has developed a number of improvements over the traditional LArTPC design.

From the electrostatics prospective, large electric fields are needed to provide adequate drift velocities and charge yields $O(500V/cm)$. But with large electric fields come risks of breakdown, in which the stored energy in the TPC is released and can cause catastrophic damage as well as a loss in livetime if not mitigated. The standard approach to mitigating these risks while achieving a uniform electric field involves the use of bulky field-shaping elements and a resistive chain between them. This is not possible within a modular design as these field-shaping elements introduce significant dead regions between modules. Instead, a resistive field shell material is used, minimizing the dead volume.

For the same reasons as the field shell, the light detection system must also have a compact form factor. Silicon photo-multipliers (SiPMs) are compact, single-photon sensitive devices that are becoming competitive with photomultiplier tubes (PMTs) in performance and that function well at cryogenic temperatures. But, the sensitive area of a single SiPM tends to be prohibitively small $O(mm)$ for direct light collection across a large area. SiPMs also normally need to be operated in low-field regions, limiting the reasonable locations for their placement to the low-voltage anode plane. In this arrangement, the SiPMs compete for real-estate with the charge readout. Instead, to increase the effective solid angle of the SiPM, light collection modules are used to trap and focus light onto the sensitive area of a small number of SiPMs. Being constructed from a dielectric material, the light collection modules can then extend into the liquid argon volume along the field cage walls, providing an effective photodetector coverage of almost 30%.

Finally, traditional LArTPC anode planes are constructed from sets of overlapping wire planes. This creates a delicate structure that must overcome forces of the wire tension and thermal stress. Long wires also introduce significant noise through their intrinsic capaci-

tance and low-frequency vibrations. These issues have posed significant problems in other LArTPCs [143]. To avoid these issues and achieve ambiguity-free 3D imaging, ArgonCube adopts the newly-developed LArPix charge readout technology, which provides native 3D charge readout using a printed circuit board (PCB)-based anode.

The ArgonCube design for the DUNE near detector (LArND) will consist of (5×7) $1 \text{ m} \times 1 \text{ m} \times 3 \text{ m}$ ArgonCube modules on a movable platform [87]. Combined, the ArgonCube design mitigates the primary challenges associated with operating a LArTPC in a high-rate environment, while also adopting technologies that reduce operational risks and improve the overall quality of the data.

The newly-developed LArPix charge readout technology is described in Sec. 2.3.3. Two light collection module designs are being pursued for ArgonCube and are described further in Sec. 2.3.4. The novel field cage material and structure is described in Sec. 2.3.5.

2.3.3 LArPix

Charge signals in LArTPCs vary between low-energy depositions of radioactive Compton scatters at the scale of 100-1000s of electrons to high-energy depositions from proton and α particles depositing 100,000s of electrons within a mm track length. Diffusion limits the intrinsic resolution of LArTPCs to $O(\text{mm})$ over meter-scale drifts, thus improving the charge readout segmentation beyond this provides little benefit. Drift signals are slow, but the time resolution of the charge readout needs to be comparable to the spatial resolution $O(\text{mm})/O(1.6 \text{ mm}/\mu\text{s}) \sim O(500 \text{ ns})$, so as not to limit the spatial resolution in any particular direction. In addition, the charge readout must have a power density low enough to avoid local boiling of the argon $O(< 100 \mu\text{W}/\text{ch})$.

LArPix meets these requirements by utilizing a custom cryogenic-compatible, mixed-signal complementary metal-oxide semiconductor (CMOS) application-specific integrated circuit (ASIC), mounted on a PCB-based anode tile with charge collection pads. Charge sensitivity is achieved within the chip using an array of integrating charge-sensitive amplifiers (CSAs), each with a dedicated, adjustable discriminator and an 8-bit analog-to-digital converter (ADC). A diagram of the front-end circuit and ADC is shown in Fig. 2.4. Each CSA input is routed via short traces and blind vias to the other side of the PCB anode tile, where a pixel pad faces the active volume of the TPC. Upon the arrival of charge at the anode, a current is driven at the input of the CSA. The CSA gain is $\approx 1\text{mV}/250 e^-$. The CSA has a quiescent voltage of $\approx 580 \text{ mV}$ and a saturation point of $\approx 1800 \text{ mV}$ giving a linear dynamic range of $\approx 300 ke^-$. Without resistive feedback, as charge collects on the pixel the output of the CSA grows until the discriminator threshold is reached. After a configurable integration delay period, internal logic triggers a track-and-hold circuit at the input of the ADC to transition into a HOLD state, latching the output voltage of the CSA. The CSA is held in reset for a short duration, and then resumes charge collection. Simultaneously the internal logic initiates the ADC cycle, producing a digitized value 9 clock cycles later. At the nominal 10-MHz clock frequency, the minimum time between two successive digitizations is $1.1 \mu\text{s}$, but is typically operated with a larger integration time $+1.5 \mu\text{s}$ to reduce the impact

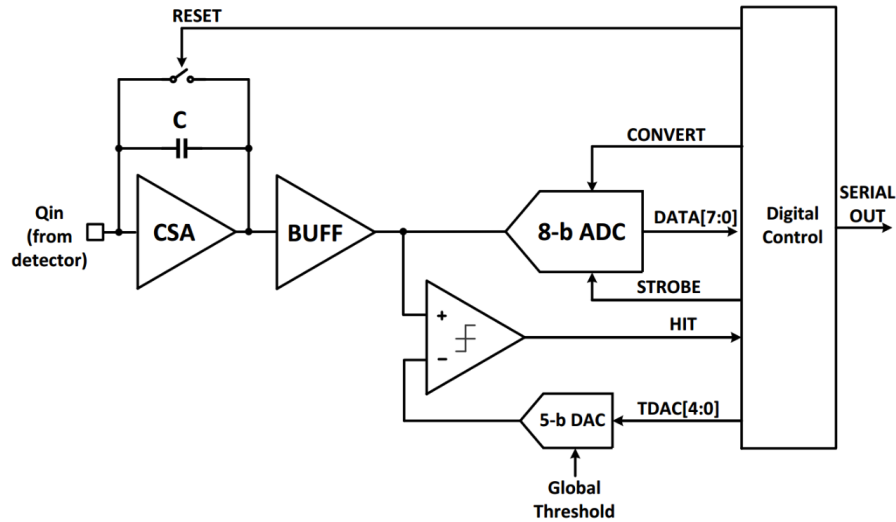


Figure 2.4: Diagram of LArPix front-end and self-triggering logic [144].




Figure 2.5: Miniature LArPix-instrumented anode showing a LArPix-v2 ASIC (left) mounted onto a multi-layer PCB with charge sensitive pixels on the opposite side (right).

of the 100-ns deadtime during the CSA reset. An internal count-since-last-reset of the 10-MHz clock is latched at the completion of the digitization. The timestamp and ADC value are packaged with a unique identifier for the chip and channel into a digital packet and put into an internal first-in first-out (FIFO) buffer. The buffer is drained by custom chip-to-chip universal asynchronous receiver-transmitter (UART) communication protocol to enable digital multiplexing of signals across many chips. Because LArTPC signals have a sparse, line-like signal, configurable low-bandwidth (typically, 2.5 Mbaud) readout is employed, reducing power consumption. This communication protocol is bi-directional, enabling the setting of local configuration registers within each ASIC.

The first pixelated PCB anode was demonstrated in Ref. [145], while true 3D readout was first demonstrated by our Lawrence Berkeley National Laboratory (LBNL) group using the LArPix-v1 ASIC [146]. Subsequent improvements have been made to the performance and resiliency of the ASIC based on the experience that was gained from this prototype. Specifically, the v1 ASIC used a daisy-chained chip-to-chip communication network and hard-coded chip identifiers set by soldered pin connections. This introduced significant single-point-failure risks, where a single chip-to-chip link could be damaged or malfunctional, making the entire system inoperable. The v2 ASIC used a novel networking scheme, Hydra-IO [147], in which each ASIC also acts as a switch, routing data to/from neighbors depending on the configuration of the chip. The v2 ASIC also introduced pseudo-differential signaling for digital signals to reduce digital-to-analog cross-talk, described in more detail in Sec. 4.4.1. Table 2.2 lists some of the primary operational differences of the LArPix-v1 and -v2 ASICs.

Table 2.2: Summary of operational differences between LArPix-v1 and -v2 ASICs.

	LArPix-v1	LArPix-v2
Channel count	32/ASIC	64/ASIC
Power consumption (nominal supply voltage)	294 $\mu\text{W}/\text{channel}$	234 $\mu\text{W}/\text{channel}$
CSA noise with 2.2 pF total input capacitance	675 e^- eq.	750 e^- eq.
Interchip IO	single-ended CMOS	pseudo-differential CMOS
Interchip network	Daisy chain	Hydra ²
Chip identifier	Hard-wired	Configurable
Internal bias currents	No	Yes
Internal ADC reference voltages	No	Yes
FIFO type	flip-flop array	SRAM
FIFO depth	2048	2048

To instrument the TPC module, an array of anode tiles are mounted to a G10 frame with 4 stainless steel stand-offs epoxied into blind-tapped holes on the tile. The mixed-signal anode tile are 8-layers with the bottom metal etched with filleted-square pixels, running from edge to edge with a 4.4-mm pitch. Each tile is instrumented with an array of LArPix chips (density $\approx 0.11 \text{ cm}^{-2}$) with up to 64 active channels per chip, achieving pixel densities of $O(0.07 \text{ mm}^{-2})$. Each anode tile has 4 redundant data and power connections via a twisted-pair ribbon cable, reducing the risk of single-point failure.

Ribbon cables are connected to a custom PCB vacuum-feedthrough assembly and a PACMAN card, shown in Fig. 2.6. Each PACMAN card supplies power and $32 \times$ UART

²Each ASIC is connected to its neighbors, and the ASIC configuration determines how data is routed.




Figure 2.6: Top view of a Pixel Array Control, Message, and Network (PACMAN) card. The high-density connector on the left of the board interfaces directly to the PCB vacuum-feedthrough. The ethernet, direct current (DC) power supply, pulse-per-second (PPS), and trigger signal interface via the connectors on the right of the board.

interfaces for up to 8 LArPix tiles. Auxiliary timing and logic signals generated by the PACMAN card allow for synchronization of the LArPix chips to within 1 clock cycle (100 ns). Within the PACMAN, round-robin FIFO buffers handle merging the datastreams received on the UART interfaces and timing and monitoring data generated internally by the PACMAN card. Data packets are buffered within a 900-MB circular buffer within the PACMAN double data-rate (DDR) memory, enabling buffering for up to ≈ 1 s at the maximum data rate. Data packaging is handled by the PACMAN card. Data messages are broadcasted to the data acquisition system (DAQ) using the ZeroMQ messaging library [148] and standard 1-Gb/s ethernet and transmission control protocol (TCP)/internet protocol (IP) layers.

2.3.4 ArcLight/LCM

To collect the 128-nm scintillation photons, Argoncube uses two light detector designs: ArcLight and LCM, modeled after the Argon R&D Advanced Program at UniCamp (ARAPUCA) light-trap design [150]. A schematic of these designs are shown in Fig. 2.7. Between the two designs, the light collection principle is functionally the same. A dielectric light-trap panel is attached to the LArPix anode tile at a right angle relative to the anode, extending into the active volume. The ArcLight modules use a uniform panel, while the LCM uses a set of 3 smaller polyvinyl chloride (PVC) panels mounted with U-shaped fiber bundles. This provides almost 30% collection-area coverage for light emitted within the TPC. With the light-trap design, light incident upon the dielectrics are first converted via wavelength-shifting coating, absorbing and re-emitting the 128-nm VUV scintillation photons to ≈ 400 nm [151]. To




Figure 2.7: (Left) Diagram of the ArcLight operating principle: 1) a UV photon is down-shifted to ≈ 400 nm, 2) at this wavelength, the photon passes through the dichroic filter, 3) the 400 nm photon is further down-shifted to ≈ 500 nm, 4) and 5) the green photon is trapped within the wavelength shifting (WLS) medium to be read out by a SiPM at 6). (Right) Diagram of the light collection module (LCM) operating principle: 1) a UV photon is down-shifted to ≈ 400 nm, 2) the photon is further down-shifted to ≈ 500 nm, 3) the photon is trapped within the WLS fiber by total internal reflection, and 4) photon is readout by a SiPM at one end of the fiber. Taken from [149].

trap the re-emitted photons, a 3M DF-PA Chill dichroic filter [152] (ArcLight) and total internal reflection (LCM) is used to trap the re-emitted photons within the dielectric medium. The ArcLight panels are built from a EJ280 [153] wavelength shifter and so a secondary absorption and re-emission occurs, further shifting the light from ≈ 400 nm to ≈ 500 nm. Figure 2.8 shows the two light detector designs mounted on the anode used in the Module-0 high-voltage (HV) test (described in Sec. 4.1).

Readout electronics are identical between the two light detector designs. Photons are read out using SiPMs soldered on PCBs mounted on the non-active side of the LArPix tiles. Each SiPM instrumentation board (E-board) provides ≈ 50 -V bias voltages and a cryo-compatible preamp to drive the SiPM signals to the readout boards, which are mounted outside of the cryostat.

The total photon detection efficiency of the readout has been measured to be $\approx 0.2\%$ (2%) for the ArcLight (LCM) readout modules [154, 155]. Thus a light signal of 10 MeV provides a fast (<20 ns) signal with an amplitude of ≈ 20 photo-electrons (PE) (200 PE) for a module instrumented with ArcLights (LCMs).

2.3.5 Field cage

The field cage is formed by four drift panels and a cathode panel. The panels use a carbon-loaded polyimide film with a sheet resistance of ≈ 1 G Ω /cm². The resistive sheet is mounted to the FR4 drift panels with an etched-metal pattern to provide structural support and




Figure 2.8: Partially-instrumented anode-plane assembly with 1 ArcLight panel (left) beside 3 LCMs (right) used in the Module-0 HV test, described in more detail in Sec. 4.1

reduce electrical non-uniformities. The temperature and field dependence of the resistive sheet material has been measured by SLAC Linear Accelerator Laboratory (SLAC) [156] and are considered in the overall system design. The cathode panel is constructed of the same resistive sheet material with an etched metal pattern. Figure 2.9 shows the field cage design that was used in the prototype testing described in Chapter 4. A single cathode HV connection is made at the top of the module with a custom HV cable. The HV cable is designed to pass through a vacuum port at the top of the cryostat, minimizing electrical interfaces. A prototype of the field cage with a resistive sheet was first built and tested by the University of Bern, demonstrating feasibility at small-scale [157]. This prototype TPC is shown in Fig. 2.10.




Figure 2.9: Resistive field cage used in the Module 0 tests. The HV cathode voltage is applied within the circular opening at the top of the Module (shown to the left in this photo).




Figure 2.10: Resistive field cage prototype produced by the the Laboratory for High-Energy Physics (LHEP) group at the University of Bern, taken from Ref. [157].

Chapter 3

Simulation

The measurement described in Chapter 6 relies on a detailed understanding of the relevant primary-particle flux, particle interactions, and detector physics. To this end, a simulation chain was assembled to utilize pre-existing and pre-vetted models, where possible. The simulation chain consisted of three stages:

1. primary particle generation
2. particle-interaction modelling
3. and the detector response.

The primary particles of interest to the measurement are cosmic-rays, thus a dedicated cosmic-ray primary particle generator, COsmic Ray SIMulations for KAScade (CORSIKA) [158], was used. A description of this generator is provided in Sec. 3.1. To model the particle propagation and interactions within the detector, the widely-used Geant4 package [159] was used with a wrapper, EdepSim [160], that reduces the overhead in developing Geant4-based simulations. The relevant features of Geant4 and EdepSim are described in Sec. 3.2, including the models used to simulate muon capture and decay processes. Finally, the large number of channels present in a pixelated LArTPC makes detector response simulation challenging. So a dedicated detector simulation making use of graphics processing unit (GPU)-acceleration was built to simulate the response of LArTPCs to particle interactions. This is described in Sec. 3.3.

3.1 Cosmic-ray event generation

The Earth is under constant bombardment from charged particles that originate from high-energy extra-solar sources. Predominantly consisting of H^+ and He^{2+} ions, their spectra follow a characteristic power law of -2.75 above ≈ 10 GeV/nucleon and level off to a roughly flat distribution below [161, 162, 163, 164]. Once reaching the upper atmosphere, the high-energy particles interact with the nuclei of the atmosphere, causing a cascade of secondary

particles. Across the primary spectrum, pions are the most numerous secondary particle, with kaons playing a more prominent role above ≈ 100 GeV/n. Because of their relatively short lifetime, most of the secondary mesons decay before reaching the surface. This produces a tertiary flux of muons and neutrinos via the dominant pion decay mode $\pi^\pm \rightarrow \nu_\mu^{(-)} + \mu^\pm$.




Figure 3.1: Total muon flux at the surface of the Earth, taken from [165].

For the muon component, the energy spectrum follows that of the primary spectrum, approximately flat below 1 GeV/c and with a characteristic power law between 10-100 GeV/c. At higher energies, the secondary pions become more likely to interact than decay, reducing the number of muons relative to other particle species and steepening the energy spectrum. At the surface, an asymmetry in the ratio of μ^+/μ^- is present due to the fact that the primary particles are positively charged. The resulting charge ratio is mostly flat within a muon momentum range of 1 GeV/c – 100 GeV/c with a value of 1.268 ± 0.013 [166]. Below 1 GeV/c, geomagnetic effects play a significant role and the charge ratio depends strongly on the latitude at which the experiment is carried out [167, 165, 168], shown in Fig. 3.2. For muons with a zenith angle approaching 90° , this effect is enhanced, producing a prominent east-west asymmetry, shown in Fig. 3.3.




Figure 3.2: Muon charge ratio at the surface of the Earth, taken from [165].

For this analysis, it is also necessary to predict the hadronic and electromagnetic components, which are non-negligible in a surface-level detector. Thus to simulate the muon flux and the background contributions, CORSIKA v7.7400 was used [158]. CORSIKA was originally developed to simulate high-energy cosmic-ray showers of $O(10^{15} \text{ GeV})$ but has been extended to lower energies. It includes models for the primary-particle flux, primary interactions, and the development of the shower within the atmosphere and Earth's magnetic field.

Generally, the most relevant experimental comparisons of CORSIKA come from the measurements of the cosmic-ray muon flux at low energies, performed by surface level operations of the CAPRICE, BESS, and HEAT experiments [170]. CORSIKA is able to reproduce the trend of the μ^+/μ^- charge ratio towards unity for low-energy muons, but over-predicts the flux of muons at low energies [170]. CORSIKA is also able to reproduce the east-west effect of the geomagnetic rigidity cutoff as studied by the WILLI detector [171].

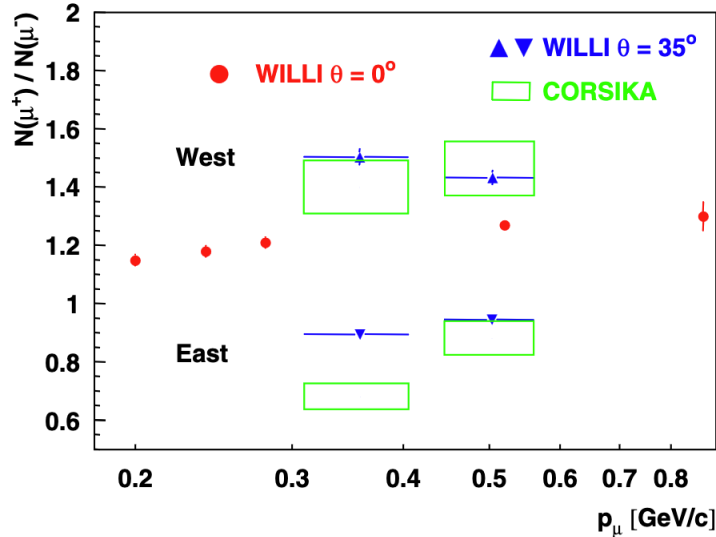


Figure 3.3: East-west asymmetry in the muon charge ratio as observed by the WILLI detector [169].

3.2 Particle propagation and interaction

The Geant4 simulation toolkit [159] is a widely-used and well-established particle simulation suite that provides tools for simulating particle propagation and interactions across a wide range of energies. As input, Geant4 requires information about the materials, geometry, particle properties, and interactions. A large number of default physics models have been implemented, covering particle decay processes, electromagnetic physics, and hadronic interactions. Of particular interest to this analysis, Geant4 includes comprehensive models of continuous energy deposition processes, described in detail within Sec. 1, 3, 4, 8, and 9 of Chapter 4 in Ref. [172]. I will highlight here a few of the most relevant features.

Generally within Geant4, energy deposition and particle interactions are handled in two ways, as a continuous processes along each trajectory simulation step (e.g. δ -ray production) or as a discrete processes at the end of each trajectory step (e.g. Compton scattering). Step lengths are determined as the lesser of the total cross-section of a particle times the number density of the medium or from a step-size limit where a maximum step length is determined by

$$\Delta S_{\text{limit}} = \alpha R + \rho(1 - \alpha)(2 - \rho/R) \quad (3.1)$$

where α is a fixed, unitless parameter, ρ is the minimum step length, and R is the current continuous slowing down approximation (CSDA) range of the particle. As $R \rightarrow \infty$, the αR term dominates and the step size is a fixed fraction α of the particle range. Conversely as $R \rightarrow 0$, the ρ term dominates, with the step limit surpassing the particle range at $R \sim \rho$.

Continuous energy loss for heavy particles is produced following a modified Bethe-Bloch model with corrections

$$\frac{dE}{dx} = 2\pi r_e^2 m_e n_e \frac{z^2}{\beta^2} \left[\ln \left(\frac{2m_e \beta^2 \gamma^2 T_{\text{cut}}}{I^2} \right) - \beta^2 \left(1 + \frac{T_{\text{cut}}}{T_{\text{max}}} \right) - \delta - \frac{2C_e}{Z} + F \right], \quad (3.2)$$

where r_e is the classical electron radius, m_e is the electron mass, n_e is the electron density of the material, z is the incident particle charge ($z_{\mu^\pm} = \pm 1$), $\gamma = \left(\sqrt{1 - \beta^2} \right)^{-1}$ and $\beta = v/c$ are the Lorentz factors of the incident particle, T_{cut} is a configurable threshold for delta-ray production, T_{max} is the maximum energy transfer to the scattered electron, I is the mean excitation energy, δ is the density effect correction, C_e is the shell correction function, Z is the nuclear charge of the medium, and F contains higher-order corrections. Energy loss fluctuations along the step are simulated using a simplified 2-state model of atomic excitation and ionization, described in Sec. 4.3.2 of [172]. Delta-ray production is generated by pulling random samples from high-energy tail of the Bethe-Bloch cross-section (above T_{cut}) and using energy conservation in 2-body scattering to fix the angle of the δ -ray and the parent particle.

Electrons and positrons use the Berger-Seltzer model for continuous energy loss

$$\frac{dE}{dx} = 2\pi r_e^2 m_e n_e \frac{1}{\beta^2} \left[\ln \left(\frac{2(\gamma + 1)}{(I/m_e)^2} \right) + F^\pm(T_{\text{cut}}/m_e, T_{\text{max}}/m_e) - \delta \right] \quad (3.3)$$

where $F^\pm(\tau_{\text{cut}}, \tau_{\text{max}})$ encapsulates the different energy losses due to Møller scattering and Bhabha scattering for electrons and positrons, respectively. In addition to electron scattering, bremsstrahlung radiation is particularly important for electrons and positrons. Within Geant4, bremsstrahlung for electrons and positrons is generated according to the calculation in [173].

High-energy photons do not lose energy continuously over their trajectory, but instead they lose energy via discrete processes: photoelectric absorption, Compton scattering, and pair production. Geant4 includes models for each of these processes. Photoelectric absorption dominates at low energies (< 1 MeV) and involves the absorption of the photon by a bound electron, resulting in its ionization and a small nuclear recoil. The cross-section is simulated using the expansion in Ref. [174]. The emitted-electron angle is sampled from the Sauter-Gavrila distribution [175] and atomic de-excitation can be simulated using low-energy EM interaction models. Compton scattering and pair production are particularly relevant to the simulation of electrons emitted from muon decay, which often produce Bremsstrahlung photons with energies ≈ 10 MeV. For a 10-MeV photon in argon, the pair production and Compton scattering cross sections are roughly equal, and are both much larger than the photoelectric cross-section. The Compton cross-section is simulated using an empirical fit to the data provided in Ref. [176]. The energy and angle of the scattered photon are sampled using the Klein-Nishina formula

$$\frac{d\sigma}{d\epsilon} = \pi r_e^2 \frac{m_e}{E_{\text{inc}}} Z (\epsilon^{-1} + \epsilon) \left(1 - \frac{\epsilon \sin^2 \theta}{1 + \epsilon^2} \right), \quad (3.4)$$

where $\epsilon \equiv E_{\text{scat}}/E_{\text{inc}}$ and E_{inc} and E_{scat} are the incident and scattered photon energies [177], and the scattering angle θ is determined by energy conservation in the 2-body interaction and by assuming the initial-state electron is at rest. Above $E_\gamma \approx 2m_e$, pair production can occur, where the photon is converted into a e^+e^- pair within the presence of a Coulomb field. Geant4 uses the model described in Ref. [178] to simulate pair production and includes corrections for the electron potential amongst other things.

Muon decay is modeled using the standard V-A theory, described in Sec. 1.2.3. Namely, the electron energy is generated in the rest frame of the muon according to

$$d\Gamma = \frac{G_F^2 m_\mu^5}{192\pi^3} 2\epsilon^2 (3 - 2\epsilon) \quad (3.5)$$

where $\epsilon = 2E/m_\mu$ is the maximum of the electron energy E relative to the muon mass m_μ . Muon polarization is neglected, and its impact on the angular distribution of the electron is neglected. For muon DIO and muon capture, the electromagnetic cascade from the initial capture to the ground state is simulated according to the models in Refs. [179, 180]. Once in the ground state, the total disappearance rate is calculated using the vacuum decay rate with a correction factor of

$$Q_{\text{DIO}} = 1 - (Z\alpha)^2 \times [A - B(Z\alpha)] \quad (3.6)$$

with $A = 0.868699$ and $B = 0.708985$ for $Z \geq 14$ and a capture rate of

$$\lambda_c(^{40}\text{Ar}) = 1.2727 \pm 0.0650 \mu\text{s}^{-1}, \quad (3.7)$$

both extracted from the table in [56]. For nuclei which are not listed in the table, the Goulard and Primakoff formula, described in Sec. 1.2.5, is used [55].

For DIO, the decay electron energy is generated according to Eq. 3.5, but includes the effect of the binding energy of the hydrogen-like S1 muon-argon state, modifying the endpoint and the decay rest frame. A binding energy of $E_b = (13.6 \times 10^{-6} \text{ MeV}) Z_{\text{eff}}^2 \mu$ is used with the reduced mass of the muon-argon system $\mu \equiv m_{\text{Ar}} m_\mu / (m_{\text{Ar}} + m_\mu)$ and $Z_{\text{eff}} = 0.923 \times 18$, taking into account a finite-nuclear size correction for argon [56].

For nuclear capture, the muon is assumed to be at rest, and a proton chosen from a Fermi density distribution¹ for the nuclei and is converted to a neutron. Energy conservation is used to fix the neutrino energy in the muon-proton rest frame. The final-state neutrino and neutron are then boosted to the lab frame. The nuclear relaxation of the neutron-³⁹Cl* system is performed by the native Geant4 PRECO model, which is valid for the energies relevant to muon capture ($< 100 \text{ MeV}$).

3.3 larnd-sim

Energy deposition information alone is insufficient to fully model the detector response. Some examples include:

¹See Ref. [172] Section 27.1.2 for a detailed description of the Fermi density model used for modelling muon capture in Geant4.

- liquid-argon scintillation introduces a characteristic delay to each photon emitted from the original energy deposition
- photon detection efficiency (PDE) is small enough that shot noise is non-negligible
- electron attachment across the drift volume modifies the charge yield as a function of position
- readout electronics have transfer functions that introduce finite delays and signal modifications

These effects were modelled and applied on an event-by-event basis within a Monte Carlo (MC) simulation to include their impact on the final result. Because these effects are independent from the details of the particle source and particle interaction processes, this component of the simulation was separated from the other parts of the simulation chain and treated in a stand-alone fashion.

There have been previous efforts to generalize many components of the detector simulation stage for LArTPCs under the LArSoft framework [181]. However, this work has largely been focused on wire-based readout. With use of a pixelated detector, the required number of charge readout channels increases dramatically. With the increase in the number of channels, the simulation burden of the detector response increases proportionally. Module 0 has 78,400 individual readout channels, which is $5\times$ more than that of the largest LArTPC operated to date². Many of the calculations involved in the detector simulation, however, can be decoupled either across the individual channels or across small segments of the particle trajectories. Thus, parallel computing strategies can be applied and expected to offer meaningful improvements.

With this in mind, we developed `larnd-sim`, a general simulation of LArPix-based LArTPCs [183] that uses GPUs for the bulk of the computation [184]. It is written in python [185] and built using the Numba [186] and CuPy [187] packages. CuPy uses Compute Unified Device Architecture (CUDA) Toolkit [188] libraries with a NumPy-compatible [189] interface and provides general algorithms and array-based data structures. For the algorithms within `larnd-sim` that are more complex, Numba enables just-in-time compilation of dedicated CUDA kernels. By developing custom algorithms for GPUs, we were able to achieve multiple orders of magnitude improvement over equivalent algorithms compiled for a CPU, highlighted within the two plots of Fig. 3.4. This work was published in Ref. [183], but this section will summarize some of the key components of the ionization, scintillation, and detector modelling that is included within `larndsim`.

There are a number of smearing effects that come into play after the energy deposition information is known. After the `edep-sim` simulation stage, Landau fluctuations in the dE/dx have already been included, along with the production of macroscopic delta rays and other interaction processes. So within `larnd-sim`, there are generally four layers that must be simulated:

²ProtoDUNE at a size of $6\text{ m} \times 7\text{ m} \times 7\text{ m}$ [131].




Figure 3.4: (Left) Comparison of total calculation time for the recombination factor as a function of the number of track segments using GPU and central processing unit (CPU) hardware. (Right) Comparison of induced current calculation time for GPU and CPU hardware. Both images are taken from Ref. [182].

1. Recombination effects that modify the conversion of deposited energy into the two signal channels – charge and light.
2. Transport effects that govern the propagation of the signal to the sensing detectors – for charge, the electron drift and attachment; and for light, the propagation, absorption, and re-emission along the photon path.
3. Electrical response effects that determine the electronic signal that is observed based on the input charge or light signal.
4. Trigger effects that determine when, how, and with what resolution the electrical signal is read out.

These are explained in the following subsections.

3.3.1 Scintillation and ionization model

As mentioned previously, energy is deposited by exciting or ionizing argon atoms. After the initial ionization, the charge cloud undergoes recombination, in which some electrons recombine with argon ions. The recombination process forms excited states that decay, producing additional scintillation light. The density of the ionization is approximately proportional to the deposited energy, but also depends on the energy loss of the parent particle. This is represented by a recombination factor \mathcal{R} , related to the dE/dx via

$$\frac{dQ}{dx} = \frac{\mathcal{R}}{W_i} \frac{dE}{dx} \quad (3.8)$$

Table 3.1: Table of measured Birks' model parameters, A and k , for liquid argon.

Reference	Source	Electric field	A	k
Scalettar [192]	β^- (364 keV)	0.075–1.5 kV/cm	0.83 ± 0.01	0.179 ± 0.003
ICARUS(3 ton) [191]	μ, p (cosmic)	0.25–0.5 kV/cm	0.800 ± 0.003	0.0486 ± 0.0006
ICARUS(T600) [191]	μ (cosmic)	0.5 kV/cm	0.81 ± 0.05	0.055 ± 0.005
ArgoNeuT [124]	p, d (NuMI)	0.5 kV/cm	0.806 ± 0.010	0.052 ± 0.001

where W_i is the average energy required to produce one ionization electron. For liquid argon, $W_i = 23.6 \text{ eV/e}^-$.

The recombination factor is modelled in LArTPCs typically using either the Birks' model [190] or the modified Box model [124]. In the Birks' model, it is

$$\mathcal{R}_{\text{Birks}} = \frac{A}{1 + (k/\epsilon)dE/dx} \quad (3.9)$$

where ϵ is the electric field times the liquid argon density. The parameters A and k have been measured by other liquid argon experiments and are summarized in Table 3.1. The value measured by ICARUS of $A = 0.800$ and $k = 0.0486 \text{ kV/MeV g/cm}^3$ was used for the baseline simulation [191].

In the modified Box model, the recombination factor is

$$\mathcal{R}_{\text{Box}} = \frac{\ln \alpha + (\beta/\epsilon)dE/dx}{(\beta/\epsilon)dE/dx} \quad (3.10)$$

where the parameters β and α have been measured by ArgoNeuT to be 0.93 and 0.207 kV/MeV g/cm³, respectively [124]. Both models are included within `larnd-sim` and used to explore systematic effects related to the recombination model.

For scintillation light, the absolute light yield L is benchmarked at zero field where it can be assumed that all of the ionization electrons recombine and thus the light yield is directly proportional to the energy loss

$$\frac{dL}{dx}(\mathcal{E} = 0) = \frac{1}{W_l} \frac{dE}{dx} \quad (3.11)$$

where $W_l = 19.5 \pm 0.2 \text{ eV/photon}$ [106] is the average deposited energy per photon and \mathcal{E} is the electric field. The evolution of the light and charge yield as a function of the applied field has been studied by other experiments [193], as well as in Module 0 [149]. It is well described by

$$\frac{dL}{dx}(\mathcal{E}) = \left(\frac{1}{W_l} - \frac{\mathcal{R}(\mathcal{E})}{W_i} \right) \frac{dE}{dx}. \quad (3.12)$$

where $\mathcal{R}(\mathcal{E})$ is the recombination factor as a function of the field and W_i is the average deposited energy per ionization electron. Because the recombination factor is a function of the dE/dx , the independent charge and light energy resolution suffer from fluctuations in the total yield due to natural fluctuations in dE/dx . The ultimate energy resolution that can be achieved from a theoretical standpoint thus requires combining the charge and light signals or operating in one of the extremes ($\mathcal{E} = 0$ or $\mathcal{E} \rightarrow \infty$). Obviously, there are other practical considerations that limit the use of detectors at these extrema. The combined energy resolution can be described by introducing a Fano factor

$$\sigma_{Q+L}^2(N_e) = F(N_i + N_e), \quad (3.13)$$

where F is 0.10 ± 0.01 [107], and N_e and N_i is the total number of excited argon atoms and ionization electrons, respectively. The energy resolution is thus nominally limited only by the energy lost to thermal excitations that are not observed. Naively, the correlated fluctuations between the charge, light, and thermal components should behave according to multinomial statistics,

$$\sigma_i^2 = p_i(1 - p_i)(N_i + N_e) \quad (3.14)$$

with a constant probability for each component p_i derived from the recombination model in Eq. 3.12. However, measurements performed by dark-matter experiments in liquid xenon suggest that there may be additional effects (such as the initial-energy distribution of the ionization electrons) that cause a further enhancement in the recombination fluctuations [194]. There is currently no published measurement of the recombination fluctuations yet for liquid argon. In the context of LArTPCs used for high-energy neutrino imaging, readout noise of the electronics and the fluctuations that arise from the low photodetector acceptance thus far dominate over these intrinsic fluctuations. Thus within `larnd-sim`, these effects were not included.

3.3.2 LRS response model

Within the light readout system (LRS), the UV scintillation photons undergo two conversion processes to enhance the SiPM PDE. First via tetra-phenyl-butadiene (TPB) and then via a WLS plastic to a wavelength of ≈ 400 nm that has a higher quantum-conversion efficiency. The primary component of the TPB conversion happens rapidly between about 1 to 10 ns, however there is some evidence of small contributions from longer time components [195]. Within `larnd-sim`, the scintillation photons are not individually tracked as this would introduce unnecessary computational burden and, with typical PDEs being $O(0.1-1\%)$, would be highly inefficient. Instead, the propagation time and the relative acceptance (including reflections and re-emission times) is off-loaded to a look-up table (LUT) which encodes the average acceptance and average time profile for photons arriving from a particular location in the TPC volume. A dedicated Geant4 simulation was created to generate a LUT for

this measurement and is described in Ref. [196]. A large number of isotropic, uniformly-distributed photons are generated within $4.4 \text{ cm} \times 4.7 \text{ cm} \times 3.8 \text{ cm}$ voxels of the active volume, and the average acceptance is calculated via

$$\langle \text{vis} \rangle_{ij} = \frac{N_{\text{det},ij}}{N_i} \quad (3.15)$$

where $N_{\text{det},ij}$ is the number of simulated photons that reach the photosensor j from voxel i and N_i is the total number of photons generated in voxel i .

The LUT does not contain the effect of the primary liquid-argon scintillation time profile, which depends on the liquid-argon purity. This is instead modelled within `larnd-sim` as a two-component exponential and convolved with the time profile generated by combining the true particle deposition time and the propagation and reemission time profile of the LUT.

Because the PDE is small, Poisson fluctuations in the number of primary electron-hole pairs produced within the SiPM are relevant. These are directly simulated within each simulation time-tick interval. After-pulsing and cross-talk are also included using a branching Poisson model [197]. In this model, each primary avalanche has a fixed probability of producing a secondary avalanche some time later. Each secondary avalanche can also cause tertiary avalanches with the same probability, and so on, until no more avalanches are produced. Within `larnd-sim`, the time spectrum of these secondaries are exponentially distributed with respect to the primary avalanche. In Module 0, the observed correlations in the light yield fluctuations within 50-ns intervals at $> 400 \text{ ns}$ compared to the yield within $< 50 \text{ ns}$ showed excess variations in the light yield. The excess was proportional to $\sqrt{N_{<50\text{ns}}}$ and consistent with an after-pulsing probability of 30%. This is larger than Hamamastu specification of 10% [198] by about a factor of 3. It is possible that this excess noise fraction is caused by something other than after-pulsing or cross-talk, but in the absence of any other obvious explanation, the after-pulsing probability within the simulation was increased and assumed to appropriately model the excess noise. The development of the PE-induced avalanche and response of the front-end amplifier is simulated by convolving the effective photo-current ($\text{PE}/\mu\text{s}$) with a template derived by averaging the waveforms obtained during light emitting diode (LED) calibration runs.

Noise is simulated on each channel by interpolating a Fast Fourier transform (FFT) extracted from calibration data. For each FFT bin i of frequency f_i , a sinusoid with a random phase ϕ_i

$$n_i(t) = A_i \sin(2\pi f_i t - \phi_i) \quad (3.16)$$

is generated, with an amplitude A_i such that the noise power per frequency is equal to the calibration FFT. This generates noise on channels that is uncorrelated and stationary, but reflective of the observed noise spectrum. Triggering and digitization is then performed on the resulting waveform, producing a trigger across the detector module if any LCM sum crosses the set threshold. Delays in the trigger logic between the two ADC units are not included in the simulation, nor is deadtime between subsequent triggers.

3.3.3 CRS response model

For the ionization electrons that escape recombination with the help of the applied field, three primary effects must be simulated: the electron drift, electron attenuation from impurities, and the pixels' response to the electron drift. The influence of the external field causes the electrons to drift with a velocity of $O(1)$ mm/ μ s in the direction of the anode, introducing a readout delay proportional to the initial position. Diffusion of the charge cloud occurs over the drift, with different characteristics parallel and transverse to the applied field. The electrons follow the electric field lines, which can be distorted due to irregularities in the field shaping components of the detector as well as due to the SCE. The SCE also plays a role in the recombination, but this is typically quite small. Along the drift path, electronegative contaminants, such as O_2 , can temporarily capture drift electrons resulting in an attenuation factor

$$\frac{dN_e}{dt_{\text{drift}}} = -\frac{N_e}{\tau_e} \quad (3.17)$$

$$\tau_d = \frac{1}{\sum_X k_X[X]} \quad (3.18)$$

where $k_X[X]$ is the attachment coefficient times the molar concentration of a contaminant X . Also during the drift, the electron cloud modifies the local electric field and induces a current on the pixels given by the Shockley-Ramo theorem [199]. This can be related to the voltage observed on a CSA via the induced current

$$I(t) = \int \left[\rho(\vec{x}, t) \vec{v}_d(\vec{x}, t) \cdot \vec{\nabla} W(\vec{x}) \right] d^3\vec{x} \quad (3.19)$$

where $\rho(\vec{x}, t)$ is the electron cloud density, $W(\vec{x})$ is the so-called *weighting potential*, and $\vec{v}_d(t)$ is the local drift velocity that depends on the local electric field. The weighting potential $W(\vec{x})$ is a specific electrostatic potential that is produced by fixing the voltage on the conductive sensing element to a unit potential and forcing all other conductors to ground. Within the Shockley-Ramo theorem, this field determines the relative response of the sensing element as a function of position and the drift velocity.

With `larnd-sim`, the drift field is treated as uniform, except for a rectangular region within 0.5 cm of the anode and within 0.66 cm laterally from the pixel. For relatively small drift lengths and the uniform field produced by the resistive sheet used in Module 0, the SCE is relatively small – limited to ~ 1 cm distortions near the corners of the cathode. Thus, no SCE or field non-uniformities are considered. Electron attachment is applied using the drift time derived from the uniform drift velocity, and electron re-emission and ion drift is ignored. For the induced current, a 3D response model

$$\mathcal{I}_{\text{near}}(\vec{x}_\perp, t) \equiv \vec{v}_d[\vec{x}_\perp](t) \cdot \vec{\nabla} W_{\text{near}}(\vec{x}_d[\vec{x}_\perp](t)) \quad (3.20)$$

is estimated using a finite-element simulation of a 3D geometrical likeness of a pixel embedded within a pixel array. First, the weighting potential $W_{\text{near}}(\vec{x})$ within the drift volume was




Figure 3.5: The local-field inductive current model is calculated with a finite-element simulation, described in Ref. [200]. This figure shows the resulting drift paths (a) and the induced current (b) for various starting points of the test charge. Taken from Ref. [183].

solved at all points \vec{x} within the 2 cm volume extending from the pixel anode using the Poisson equation. Standard boundary conditions of $V_{\text{pixel}} = 1$ V and all other conductors at 0 V were used. The full electrostatic field was then calculated using an applied uniform electric field. A unit test charge was then propagated from a series of source points \vec{x}_{\perp} at a fixed drift distance from the pixel and according to the local field, deriving the drift path $\vec{x}_d[\vec{x}_{\perp}](t)$ and the drift velocity $\vec{v}_d[\vec{x}_{\perp}](t)$. The gradient of the weighting potential along the drift path was then evaluated and combined with the drift velocity to produce the look-up table $\mathcal{I}_{\text{near}}(\vec{x}_{\perp}, t)$ used by the simulation. Figure 3.5 illustrates this calculation for a few test charge starting positions. More details about the finite-element simulation and near-field model can be found in Ref. [200].

In addition to the local-field response, the lack of a shielding grid in front of the pixels introduces a far-field response, in which a current is induced by drifting charge at long-ranges. These can modify the integrated current signal by up to $\sim 10\%$ and lead to a phenomenon in which pixels self-trigger before the charge arrives at the anode. The finite-element approach used for the near-field encounters computational difficulties at the length scales required to calculate long-range currents. Instead, these were added to the current model using a dipole approximation and the method-of-image technique. The Shockley-Ramo weighting potential requires the cathode and anode potentials to be fixed at 0 everywhere except for the pixel-of-interest. At distances from the anode at which the pixel can be approximated by a point (when the pixel pitch is much smaller than the distance from the pixel), these boundary conditions lead to a dipole-like solution, with modifications near the cathode. This field can




Figure 3.6: Dipole approximation of the z -component of the weighting field described by Eq. 3.21, calculated with terms out to $n = 5$ and a cathode positioned at $z = 300$ mm. Arrows indicate the negative gradient direction. Under an assumption of uniform drift in $-\hat{z}$, in regions where $\partial W/\partial z > 0$ an electron will induce a negative current on the pixel, where in regions with $\partial W/\partial z < 0$ induce a positive current.

be approximated by using a series of dipole fields, reflected about the anode and cathodes ($z = 0$ and $z = l$ respectively)

$$\frac{\vec{\nabla} W_{\text{far}}(\vec{x})}{C} = \frac{3(\hat{z} \cdot \hat{x})\hat{x} - \hat{z}}{x^3} + \sum_{n=1}^{\infty} \sum_{\pm} \frac{3(\hat{z} \cdot \hat{x}_n^{\pm})\hat{x}_n^{\pm} - \hat{z}}{(x_n^{\pm})^3} \quad (3.21)$$

$$\vec{x}_n^{\pm} = \vec{x} \pm 2^n l \hat{z} \quad (3.22)$$

where \vec{x} is the test point far from the pixel at $\vec{x} = 0$, C is an arbitrary normalization constant, and \vec{x}_n^{\pm} is the position relative to the image dipoles placed at $\vec{x} = \pm 2^n l \hat{z}$. The residual transverse field at the anode is precisely zero for each term in the n summation, meeting the boundary requirement at the anode, with the second boundary condition at the cathode violated only by the \pm summation term that does not cancel. This term converges to 0 as $n \rightarrow \infty$. Figure 3.6 shows the z -component of this field for a 30-cm drift.

A resulting current response model was then calculated using the dipole field, under the assumption of uniform drift in \hat{z}

$$\mathcal{I}_{\text{far}}(\vec{x}_{\perp}, t) = v_d \left. \frac{\partial W}{\partial z} \right|_{\vec{r}=(\vec{x}_{\perp}, l - v_d t \hat{z})}. \quad (3.23)$$

The near- and far-field response models were combined by fixing the normalization constant C at a semi-arbitrary reference drift time $t_{\text{ref}} = 6.7 \mu\text{s}$

$$\mathcal{I}_{\text{near}}(\vec{x}_{\perp}, t_{\text{ref}}) = \mathcal{I}_{\text{far}}(\vec{x}_{\perp}, t_{\text{ref}}). \quad (3.24)$$

To calculate the induced current by a simulated charge deposition, the combined response model is integrated with the simulated charge density using a MC sampling technique

$$I(t) \approx \frac{Q}{N} \sum_i \mathcal{I}(\vec{x}_{\perp,i}, t_i) \quad (3.25)$$

where Q is equal to the total deposited charge along the track segment, N is the number of MC samples, and \vec{x}_i and t_i are chosen to lie along the track segment. The perpendicular distance \vec{x}_{\perp} is the transverse distance of the sample point from the pixel-area normal axis. The expected diffused charge distribution for a point-like charge, neglecting electron attachment, is

$$\rho(\vec{x}, t) = \frac{q_e n_e}{4\pi t \sqrt{D_{\perp} D_{\parallel}}} e^{-\frac{(z-v_d t)^2}{4D_{\parallel} t}} e^{-\frac{(x_{\perp})^2}{4D_{\perp} t}}, \quad (3.26)$$

with the two terms D_{\perp} and D_{\parallel} representing the transverse and longitudinal diffusion coefficients, is simulated by separating the transverse and longitudinal components

$$\rho(\vec{x}, t) = q_e n_e p_{\text{norm}}(z; \mu = v_d t, \sigma^2 = 2D_{\parallel} t) p_{\text{norm}}(x_{\perp}; \mu = 0, \sigma^2 = 2D_{\perp} t) \quad (3.27)$$

where $p_{\text{norm}}(x; \mu, \sigma^2)$ is the PDF for a normal distribution. The smearing produced by diffusion is then simulated by applying random Gaussian translations in the transverse and longitudinal directions of \vec{x}_i , corresponding to the influence of diffusion at each time tick and drift time. Diffusion constants of $D_{\perp} = 8.8 \text{ cm}^2/\text{s}$ and $D_{\parallel} = 4.0 \text{ cm}^2/\text{s}$, were used for the baseline simulation. For the relatively short drift length of the ArgonCube design, diffusion is much smaller than the pixel pitch and is not particularly impactful.

A detailed model of the pixel response and the LArPix self-trigger logic was implemented within `larnd-sim`. Because LArPix employs an integrating front end and does not have a shielding plane, the signal $V_{\text{sig}}[i]$ that is read out is a function of the entire drift

$$V_{\text{sig}}[i] = G \int_{t_0}^{t_i} I_{\text{tot}}(t) dt + \text{noise} \quad (3.28)$$

$$= G \int_{t_0}^{t_i} \left(\int \rho(\vec{x}, t) \vec{v}_d(t) \cdot \vec{\nabla} W(\vec{x}) d^3 \vec{x} + I_{\text{leak}} \right) dt + \text{noise} \quad (3.29)$$

where I_{leak} is the leakage current through the transistor, G is the CSA gain, t_0 is the time of last reset, and t_i is the digitization time. The leakage current has been measured on the bench and depends strongly on the temperature [146], becoming negligibly small ($\sim 10^{-4} \text{ e}^-/\mu\text{s}$) at liquid-argon temperatures. The noise term can be broken into two components, depending on how the noise impacts samples taken without resetting the CSA between samples. The *correlated* noise component is fully correlated for samples that are taken without a CSA reset between samples, while the *uncorrelated* noise component is not. When operating a LArPix

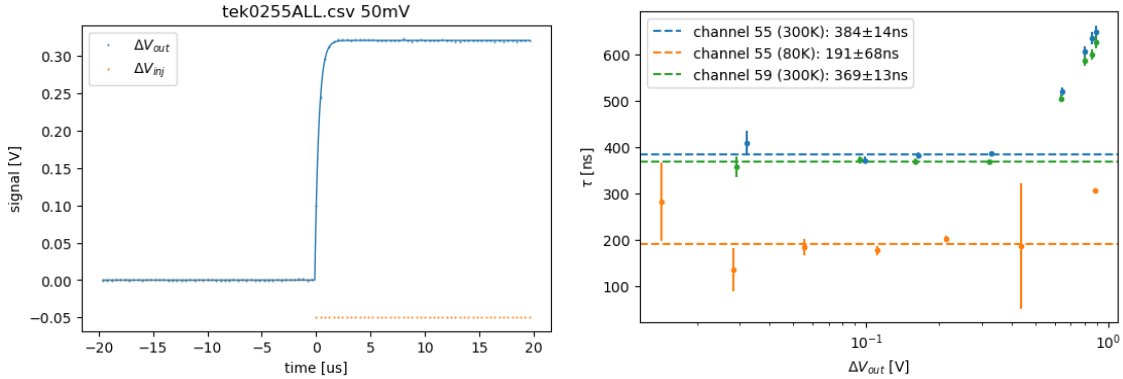


Figure 3.7: (left) Example fit to analog monitor response to impulse-like charge injection. (right) Fit results for three channels at room temperature and liquid nitrogen temperature. Saturation for large amplitude pulses modifies the response shape resulting in poor fit χ^2 assuming an exponential-only response. Impacted points (>500 mV ΔV_{out}) are excluded from the quoted averages.

channel in the standard self-trigger mode, the total charge noise in a sample is given by the quadrature sum of these components

$$\sigma^2 = \sigma_{\text{correlated}}^2 + \sigma_{\text{uncorrelated}}^2. \quad (3.30)$$

Benchtop measurements and in-situ measurements using pedestal calibration runs show variations in individual channels but with a typical total noise of $\sigma(\text{LAR}) \approx 950 e^- \text{ eq.}$ and a correlated component that dominates with $\sigma_{\text{correlated}}^2 / \sigma_{\text{uncorrelated}}^2 \approx 4$.

The CSA has a finite bandwidth that reduces the response to high-frequency signals. This can be constrained by examining the output of the analog monitor to impulse-like charge injection signals, shown in Fig. 3.7. The resulting waveform from the analog monitor is then a step function convolved with the response of the CSA. The capacitive loading of the analog monitor is non-negligible and contributes to the resulting signal, adding an additional convoluting factor. Thus the constraint from this measurement is an upper bound on the true CSA bandwidth. The bandwidth was observed to change as a function of temperature, increasing to 5 MHz at liquid-argon temperatures. Within `larnd-sim`, a convolution of an exponential function with a time constant of 100 ns is applied to simulate this effect, with an assumed error of ± 100 ns.

Simulation of the self-trigger logic is performed on every pixel within a 5-pixel radius from the projection of a charge deposition onto the anode. Self-trigger thresholds modify the output of a pixel significantly – lower thresholds cause pixels to trigger earlier on a signal, and when low enough, they can trigger multiple times. Figure 3.8 shows the hit-charge distribution for minimum ionizing particle (MIP)-like tracks as a function of the self-trigger threshold. As the threshold is lowered, the distribution shifts to lower charge in part because smaller energy deposits can be read out, but predominantly due to the shift

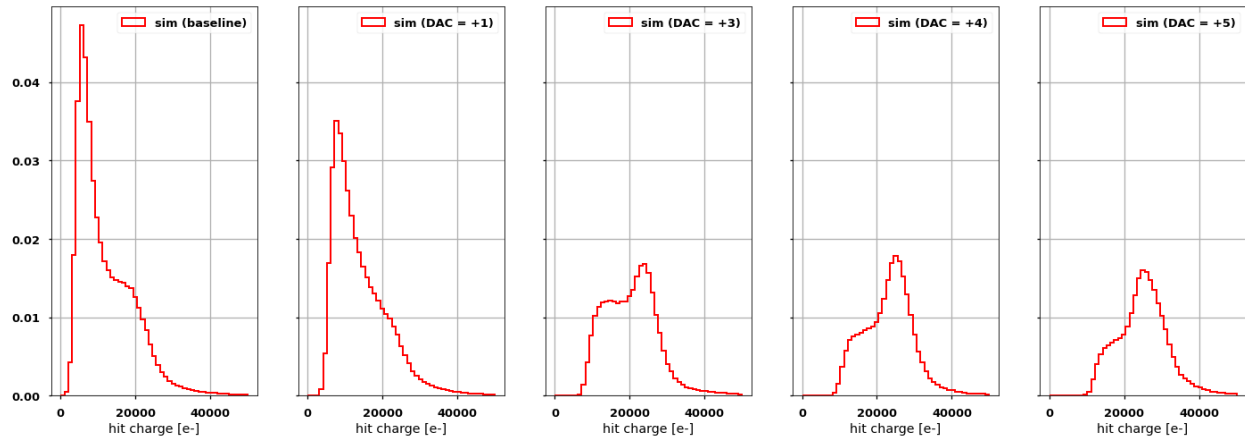


Figure 3.8: Change in hit-charge distribution as a function of the self-trigger threshold (1 DAC $\approx 2 \text{ ke}^- \text{ eq.}$). Baseline thresholds are $\approx 4.5 \text{ ke}^- \text{ eq.}$

towards multiple triggers per energy deposit. Additionally, $\approx 100 \text{ ns}$ of dead time results after each self-trigger. Depending on how this dead time overlaps with the signal, more or less of the total charge can be lost. `larnd-sim` takes into account the relative self-trigger time and the dead time between self-triggers by generating a waveform for each pixel and overlaying noise. The waveform is then updated upon each self-trigger, mimicking the logic of the trigger-reset-digitize cycle.

Chapter 4

Data collection

Initial prototypes were built to demonstrate independent functionality of each of the ArgonCube technologies: charge readout [146], light readout [154, 155], and the field cage [157]. However, to reduce the risk in building a full-size detector, such as the DUNE ND-LAr, a series of prototype modules were built and operated. The first integration prototype (SingleCube) included both the charge readout system (CRS) and LRS readout and is described in Sec. 4.1.1. Following its success, a partially-instrumented HV test was performed using a minimally-instrumented ($1.2 \times 0.6 \times 0.6$) m² detector module. This is described in more detail in Sec 4.1.2.

The first fully-instrumented prototype was assembled and tested at the University of Bern in the spring of 2021. Data were collected during this period using cosmic rays and was used as the basis for the analysis described in this thesis. An overview of the data collected and some low-level quality metrics are described in Sec. 4.1. After the data collection and initial comparison with the simulation, some additional modifications and tuning of the detector simulation were required. These are described in Sec. 4.4. Finally, Sec. 4.2 and 4.5 summarize the datasets and simulation samples used in this analysis.

4.1 Module 0 prototypes

The Module 0 detector, shown in Fig. 4.1, was built as a full-size ($0.7 \times 0.7 \times 1.4$ m³) dual-drift LArTPC, demonstrating the full-integration of ArgonCube technologies. Prior to the operation of Module 0, there were significant unknowns that could only be determined with an integrated prototype detector.

- Could a LArTPC constructed from G10, which is known to outgas contaminants, achieve adequate purity?
- Would the resistive field-cage material introduce non-uniformities in the electric field when used over large areas?
- Would the dielectric light collectors distort the drift field in the LArTPC?




Figure 4.1: Rendering of the Module-0 detector, as constructed. Taken from Ref. [149].

- Would the perforations in the field-cage allow sufficient flow such that 80,000 readout channels could be operated without inducing local liquid argon (LAr)-boiling?
- Would electromagnetic interference (EMI) impact the capabilities of the CRS or LRS?
- Could data received from the CRS and LRS be properly associated (and achieve ns-scale timing)?

The Module 0 detector was built to investigate these design challenges and demonstrate that they could be overcome.

4.1.1 The SingleCube demonstrator

In 2020, the first integration test of the charge and light readout was performed with the SingleCube demonstrator [201]. The SingleCube demonstrator was a $(30\text{ cm})^3$ single-drift TPC read out with a single LArPix anode tile and a single ArcLight panel. The field cage was constructed using PCB-based design and instrumented with a single 900-cm^2 pixel tile




Figure 4.2: (Left) View of the SingleCube drift volume and cathode prior to assembly, each dimension is 30 cm. (Right) View of the pixel tile, left, and ArcLight pane, right, used for the SingleCube test.

and ArcLight panel, shown in Fig. 4.2. The detector was assembled and installed within a cryostat at LHEP of the University of Bern.

The demonstrator was operated continuously over a 5-day period, collecting cosmic-ray data at fields between 100 V/cm and 1000 V/cm. During this test, an electron lifetime between 500 μs and 1.2 ms was achieved. Light and charge signals were time-correlated using the trigger logic outlined in Sec. 4.1.4. Overall, the CRS operated with 95% of pixels active, with the remaining fraction manually disabled due to elevated self-trigger rates. An analysis of the cosmic-ray muon dQ/dx was performed using a principle component analysis (PCA)-based track reconstruction, described in Sec. 5.2.1.6. This analysis demonstrated $< 8\%$ deviation in the reconstructed dQ/dx as a function of the track orientation, indicating high-fidelity 3D imaging from the CRS.

4.1.2 Module 0 high-voltage test

After the successful operation of the SingleCube demonstrator, a dedicated test of the HV field cage was performed. This test used the same field cage as would be operated in Module 0 but a subset of the readout electronics. This reduced the risk of HV instabilities damaging the electronics in the readout systems, had they occurred. The overall dimensions were $(0.7 \times 0.7 \times 1.4) \text{ m}^3$ with each TPC containing an volume of $0.30 \times 0.62 \times 1.24 \text{ m}^3$. SLAC designed and produced the TPC field cage and HV system for Module 0. The cathode was constructed from a 25- μm Kapton XC sheet [202], which was connected to the HV supply via a custom cable feedthrough and spring-loaded connector. The field cage was produced from 6-mm FR-4 sheets with a 100- μm -thick Kapton DR8 [203] laminate. The top




Figure 2.8: Partially-instrumented anode-plane assembly with 1 ArcLight panel (left) beside 3 LCMs (right) used in the Module-0 HV test.

and bottom panels of the field cage were perforated with 4-mm holes to provide adequate liquid circulation within the active volume. Dedicated, separate HV ground connections were maintained to separate the current of the HV supply from the sensitive detectors and provide protection from overvoltage. For this test, the module was instrumented with 4 LRS modules (1 ArcLight panel and 3 LCMs) and a single CRS pixel tile, shown in Fig. 2.8. The remainder of the anode was populated with grounded dummy tiles, in order to maintain a uniform electric field.

In September of 2021, the HV test proceeded, evacuating and then filling the Module 0 cryostat with liquid argon. A HV scan was successfully performed between 250 V/cm and 1 kV/cm without catastrophic breakdown. At 1 kV/cm, some instabilities in the cathode voltage occurred, but the detector recovered without intervention. Charge and light data was collected and successfully associated between the subsystems. The pixels along the outer edge of the CRS tile encountered early challenges due to pickup from the neighboring dummy tiles. This was remedied by modifications to the grounding of the HV and dummy tiles. The electron lifetime was measured by the CRS and found to be ≈ 2 ms throughout the run, adequate for LArTPC operation. This test retired the primary risks of the field cage design, namely, adequate fluid recirculation and HV breakdown.



Figure 4.3: (Left) Fully assembled Module 0. (Right) Module 0 being lowered into the cryostat.

4.1.3 Module-0 detector

The fully-instrumented Module 0 detector, shown in Fig. 4.3, was operated in two run periods at the end of March and at the end of June in 2021. During the first run period, more than 25 million cosmic-ray triggers were collected during the 8-day run period with different field configurations and charge readout thresholds. The second run period included an additional sample of more than 14 million cosmic-ray triggers, with data taken under a single threshold and field configuration but different fluid recirculation rates.

The charge readout for each of the modules' TPCs consisted of (2×4) $31 \text{ cm} \times 32 \text{ cm}$ LArPix anodes. Each anode tile was instrumented with 4,900 (4.43×4.43) mm pixels, for




Figure 4.4: Bottom view of Module 0 active volume, prior to complete assembly. Labeled regions are: A) pixel anode tile, B) ArcLight module and C) LCMs protruding into the drift region, and D) resistive cathode and field cage.

a total active channel count of 78,400. Each TPC was read out by a single PACMAN card mounted to a dedicated cryostat flange. Each PACMAN card provided a continuous data stream on a ZMQ PUB socket [148], accessed by the DAQ using an ethernet connection and TCP/IP address set up on a local network. Each PACMAN was configured to receive a timing signals to enable offline timing alignment, described later in Sec. 4.1.4.

The light readout for each TPC used 4 ArcLight panels and 12 LCM panels, mounted in the configuration shown in Fig. 4.4. Each ArcLight (LCM) panel extended $30\text{ cm} \times 1\text{ cm}$ into the active volume and was read out using 6 (2) low-crosstalk Hamamastu S13360-6050 SiPMs [198]. The SiPM waveforms were driven out of the cryostat with cryo-compatible preamps over shielded minicoaxial cables. The signals were passed through a variable gain amplifier (VGA) prior to the 100-MHz digitization to balance the relative signal amplitudes of the ArcLights and LCMs, which have a significantly different PDE. A 10-bit 100-MHz 64-channel differential ADC was developed by JINR to digitize and transmit the waveform data to the DAQ over a 10-Gb/s optical link, providing $1/f_{\text{Nyquist}} = 20\text{ ns}$ sampling resolution. Due to the cabling layout, each digitizer read out the light modules on adjacent field cage

walls in each TPC, consisting of 4×6 ArcLight SiPMs, $4 \times 3 \times 2$ LCM SiPMs. Auxillary, summed-waveform channels combined the 6 SiPM signals on each ArcLight and the 6 SiPM signals on each group of 3 LCMs. These channels were used for triggering, which is described in the next section. The global trigger signal and the global PPS SYNC signal were also digitized to enable offline timing alignment. Thus a total of 58 out of 64 ADC channels were utilized per ADC.

Additional monitoring infrastructure was used to insure stable operation and provide safety interlocks in the event of a malfunction. Four-wire resistive temperature detectors (RTDs) were installed at a variety of positions within the detector and cryostat to monitor the cooling rate during the initial fill and provide a redundant liquid level measurement. Pressure monitors were installed to monitor the ullage volume, the fill and vent lines, and the vacuum insulation of the cryostat. A dedicated capacitive level sensor was also used for continuous liquid level measurement near the operational fill level. The HV subsystem was monitored using the integrated monitoring of the HV power supply and ammeter measurements of the HV return current. A gas chromatograph monitored the presence of trace gases within the ullage volume down to ≈ 10 parts-per-million (ppm). Purification of the liquid argon was performed throughout the run with a molecular sieve material (RCI-DRI 4A Mol-Sieve [204]) and an activated copper oxide material (Q-5 copper catalyst [205]) according to the purification strategy laid out in Ref. [206]. A $5\text{-}\mu\text{m}$ particulate filter at the output of the filtration system was used to remove fine particulates of the purifier materials from the detector volume.

4.1.4 Trigger and timing configuration

To match the fast light signals and the slow charge signals, a means of timestamping and matching the charge and light readout data was required. A block diagram of the timing system used in the Module 0 test is shown in Fig. 4.5. The time synchronization was performed using a global PPS signal, generated by a GPS receiver. The PPS signal was then distributed to both the LRS and CRS.

Within the LRS, each ADC was triggered by both the global trigger and the PPS signal. The PPS signal was digitized on one of the ADC channels, enabling offline identification of each PPS interval. Within each ADC, a local 100-MHz timestamp is read out with each trigger, counting the number of clock cycles since the ADC was powered on. Between the 100-MHz timestamp and the PPS signal, an offline analysis can be used to reconstruct the trigger time within the PPS interval, correcting for the ADC clock frequency and drift.

For the CRS, the PPS signal was received by each PACMAN card. From this, a SYNC pulse was triggered and passed to the timing system within the PACMAN controller field-programmable gate array (FPGA) and distributed to the LArPix chips. The SYNC pulse reset the PACMAN controller timestamp and LArPix timestamps in tandem, maintaining the same counter value across all subcomponents. On each synchronization event, a dedicated data word was inserted into the PACMAN datastream by the PACMAN card containing the timestamp just prior to reset. This timestamp was monitored for missed PPS signals,

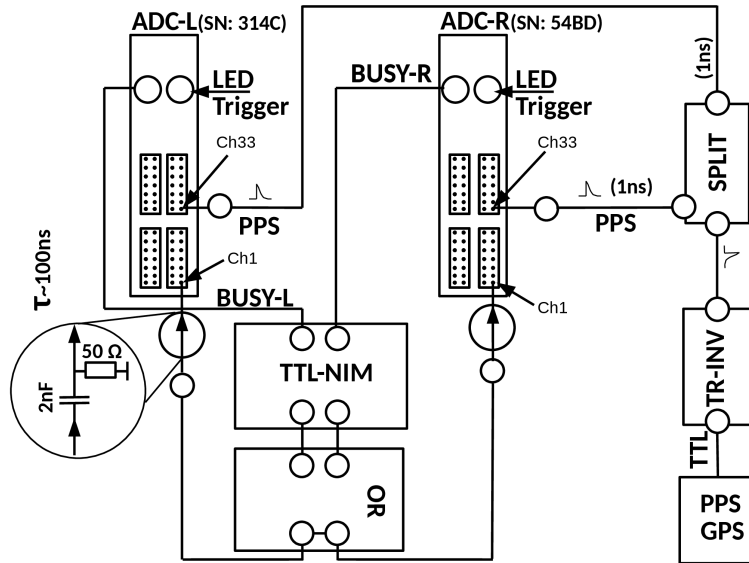


Figure 4.5: Diagram of timing synchronization system used for Module 0, taken from Ref. [207]. A PPS signal is generated by a global positioning system (GPS) unit (lower right) and distributed to both the light ADCs and the PACMAN cards (not shown). If either ADC triggers, a BUSY signal is generated and passed to the other ADC and the PACMAN cards.

but none were found during operation. When operated in self-trigger mode, LArPix data packets include the timestamp of digitization, which was used to determine the self-trigger time within the PPS interval.

To identify time windows with activity, a global trigger signal was generated by the LRS. Each group of 6 SiPMs across a group of 3 LCM detectors were summed. A threshold of 1500 ADC counts was set on the summed LCM waveforms, initiating a global trigger if any group of 6 SiPMs crossed this threshold. Only the LCM detectors were used to trigger Module 0 due to the presence of pickup noise in the ArcLight detectors. The global trigger is then provided to the other LRS ADC and to both PACMAN controllers. At the other LRS ADC, the global trigger forced a digitized waveform to be read out, with a delay of ≈ 300 ns, varying slightly due to the rising edge of the waveform. At the PACMAN card, the local timestamp was latched and inserted into the datastream as a flag for offline analysis.

4.1.5 Run 1

The first Module 0 run took place between March 27th and April 13th of 2021, with the cumulative data shown in Fig. 4.6. For the first two days of the run, the detector was evacuated to reduce outgassing of the detector materials during operations. The evacuation




Figure 4.6: Accumulated events during Run 1 along with average trigger rates. Taken from Ref. [207].

was followed by the cool-down and filling with liquid argon, which took 2 days. After the initial fill, the DAQ was commissioned and the system was brought to a stable operating configuration by April 1st. During the initial commissioning, one E-board was found to be disfunctional, along with two additional SiPM channels, resulting in a total active SiPM count of 88/96 (91.7%). The cathode HV ramped to the nominal 500 V/cm field in a two-stage process, temporarily halting at 250 V/cm to ensure HV stability. Data were collected for 2 days at the nominal field and with the charge system in a high-threshold configuration. The CRS self-trigger threshold was then lowered to improve performance at the expense of higher trigger rates. Data were collected in this configuration for 2 days. A HV scan with steps of 50 V/cm and 40-min exposures was performed with an additional 6-hour exposure taken at the maximum field of 1 kV/cm. Data were collected for the remainder of the run in a mix of the high- and medium-threshold configurations of the CRS.

4.1.5.1 Pedestal calibrations

To calibrate the CRS, *pedestal* runs were taken periodically at an interval of about 1 hour. During these calibration runs, LArPix was operated in forced-digitization mode, in which each ASIC would cycle through triggering one of its 64 channels at a fixed interval, uncorrelated with any other activity. These calibration runs provide the input data for the pedestal calibration described in Sec. 5.2.

A few features in the pedestals were found, as shown in Fig. 4.7 which shows the average ADC value, the variation in ADC value, and the trigger rate obtained during one of these calibration runs. Small patterns of channels were found to be missing during some runs, but the exact channels that were missing varied run-to-run. This is attributed to the bit errors described in Sec. 5.1.1.2. Localized regions were also impacted by increased leakage current, evidenced by an enhanced mean ADC in the pedestal run, likely introduced by damage to the ASIC during assembly. The channels most significantly impacted also often appeared to introduce significant leakage current on neighboring channels, resulting in slightly larger impacted regions. These channels were flagged during the commissioning phase and disabled. A shift and enhanced noise in the pedestal was also found for pixels near the edge of each pixel tile and was attributed to coupling between the pixel and the neighboring tile ground. Later studies were able to link the enhanced pedestal noise to excessive ground impedance between neighboring pixel tiles. These channels were also disabled for the run to improve data quality. Figure 4.9 shows all of the disabled channels for both Run 1 and Run 2. In total, the number of disabled channels during Run 1 was 6,959 of 78,400, providing an active area of 91.2% of the total anode area.

Finally, digital-analog crosstalk was enhanced for pixels near the input/output (IO) pins, resulting in the repetitive pattern shown in the standard deviation of the ADC values in Fig 4.7. Figure 4.8 shows the pedestal across channels on an anode tile, including only the channels that were enabled for self-triggering data. Overall, pedestal values were extremely stable over the course of the run with the typical channel pedestal varying between calibration runs by less than 1 mV.

4.1.5.2 Charge thresholds

Thresholds were tuned with a dedicated algorithm such that each pixel produced a similar, < 2 Hz self-trigger rate with null drift field. However, self-trigger rates were found to differ between the tuning runs and self-triggering runs due to a feedback loop caused by digital-analog cross talk. The increased self-trigger rate caused by charge deposition signals present with a non-null field produced a high data rate, making self-triggering runs more susceptible to cross-talk instabilities. To avoid this, the initial data was collected in the high-threshold configuration in which the global threshold on all pixels was raised by a fixed amount relative to the tuned value. Subsequent optimization by disabling particularly noisy channels allowed this offset to be reduced, but not totally eliminated, resulting in higher thresholds than nominally could be expected based on the intrinsic LArPix front-end noise. An analysis of the hit charge distribution on individual pixels determined that the median self-trigger threshold during these data were ≈ 24 mV (≈ 48 mV) for the runs in the medium (high) threshold configuration, shown in Fig. 4.10. A bi-modal distribution is present in the self-trigger threshold due to digital-analog crosstalk, which impacts channels nearest the IO pins of the ASIC most substantially.




Figure 4.7: Example of the results on one tile (TPC 1, tile 2) from a pedestal run with all pixels enabled ($1 \text{ ADC} \approx 3.9 \text{ mV}$). Here x and y are defined relative to the anode tile, with the readout connectors falling along $x \approx -150 \text{ mm}$ and the sensitive area in the direction of $-(\hat{x} \times \hat{y})$.

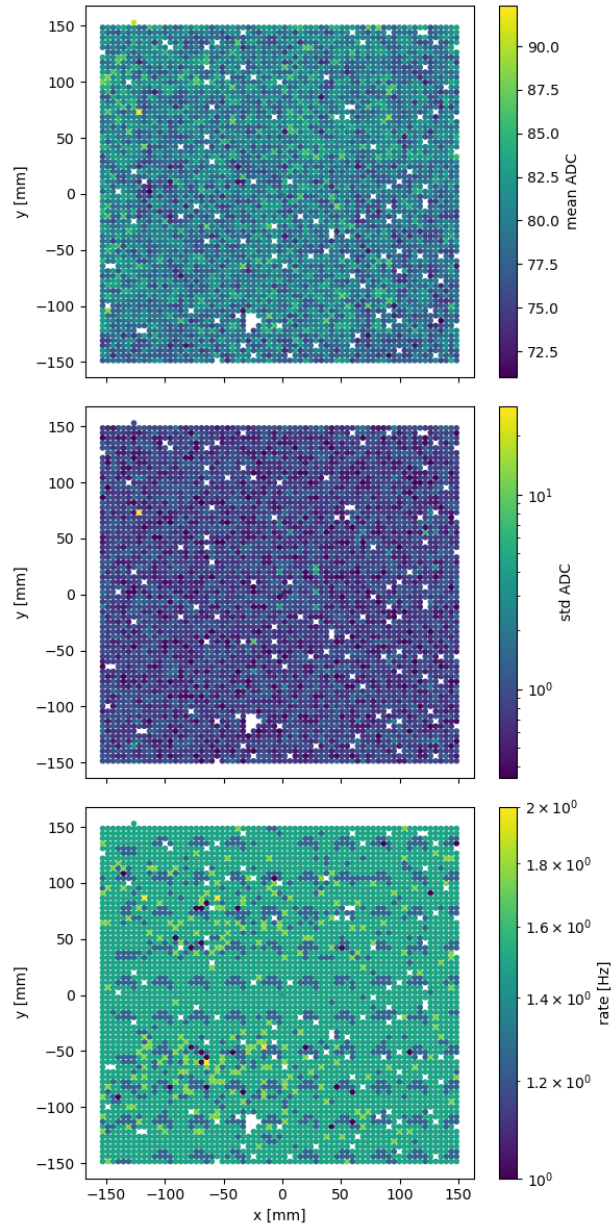


Figure 4.8: Example of the results on one tile (TPC 1, tile 2) from a pedestal run with only "good" pixels enabled ($1 \text{ ADC} \approx 3.9 \text{ mV}$). Here x and y are defined relative to the anode tile, with the readout connectors falling along $x \approx -150 \text{ mm}$ and the sensitive area in the direction of $-(\hat{x} \times \hat{y})$.

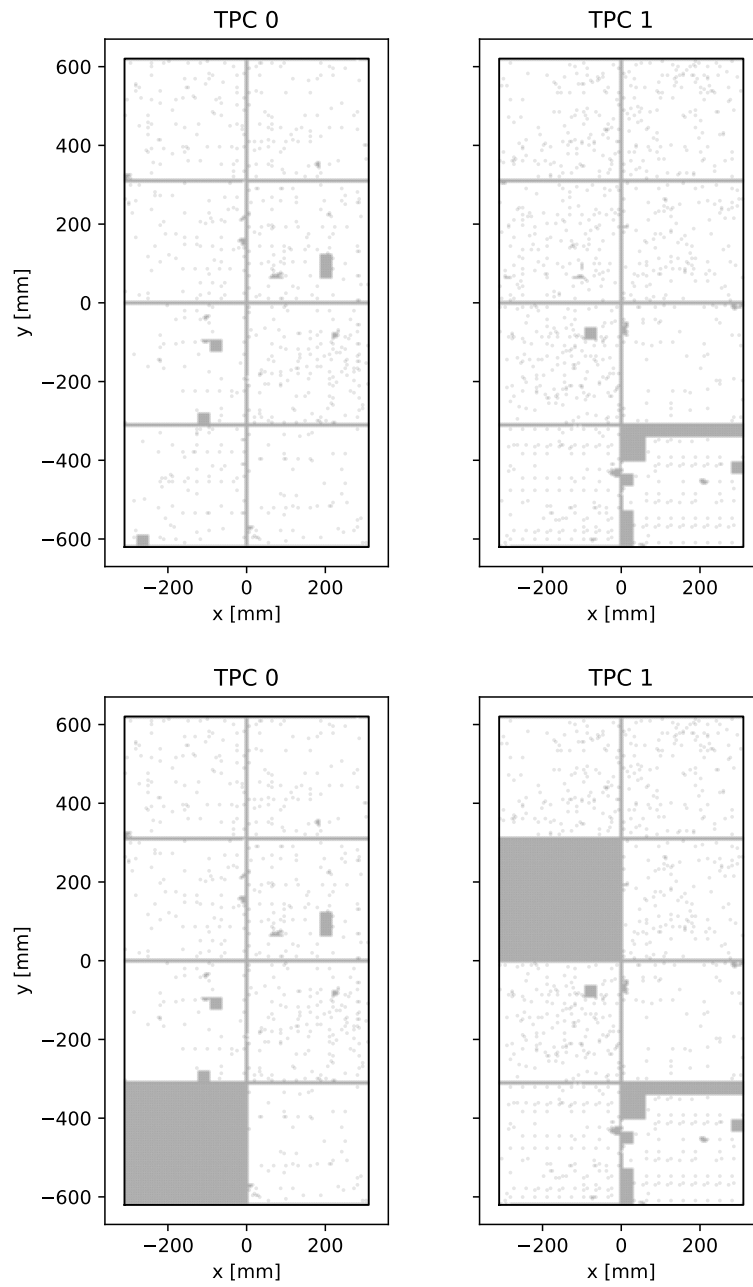


Figure 4.9: Disabled CRS channels for Run 1 (top) and Run 2 (bottom).




Figure 4.10: Measured self-trigger thresholds in the two operating configurations of Module 0 ($1 \text{ ke}^- \approx 4 \text{ mV}$).

4.1.5.3 Light calibration runs




Figure 4.11: Light readout system gain including SiPM, cryo-preamp, VGA, and ADC – (left) ArcLight SiPMs and (right) LCM SiPMs.

Prior to data collection, SiPM bias voltages were tuned using an LED installed on each anode. During these runs, the LED pulse amplitudes were adjusted to produce single photoelectron (SPE) signals on the SiPM-under-test. The bias voltage of the SiPM-under-test was set such that each SiPM provided approximately the same gain, as determined by the analysis described in Sec. 5.2.2. The VGAs amplified the signals from the cryo-preamps and enabled balancing the relative signal amplitude of the ArcLight and LCM detectors. A setting of 31 dB was used for the ArcLight and 21 dB was used for the LCM detectors to avoid saturation of the ADCs. At these settings, the ArcLight detector channels saturated at a signal amplitude of ≈ 6 PE/tick and the LCMs at an amplitude of ≈ 60 PE/tick.

4.1.6 Run 2

The second operation of Module 0 took place between June 14th and 26th of 2021. This data collection period used the *piston-purge* cryostat preparation strategy [208]. During this procedure, gaseous argon is used to purge atmospheric impurities by using a natural gravitational concentration gradient. This enabled high liquid-argon purity without exposing the cryostat to negative gauge pressures – a necessity for large-volume membrane cryostats [209]. High-purity gaseous argon was flowed between the 14th and 20th to achieve sufficient volume exchanges prior to cooldown. The detector was cooled and filled from the 21st to the 22nd and re-commissioning began on the 23rd. Data was taken using the high-charge threshold configurations determined in the first run period, and the nominal drift field of 500 V/cm. The LRS DAQ was configured with a longer readout window of 1024 samples to reduce deadtime that negatively impacts muon decay events. Triggering was performed with slightly larger LRS trigger thresholds of 2500 ADC counts (≈ 4.2 PE) to reduce overall

data volume. During re-commissioning, high self-trigger rates were encountered on two of the anode tiles. Due to the primary goals of this run – namely, testing the viability of the piston-purge preparation and stability of the charge readout under different rates of fluid recirculation – it was decided that operating the detector without these tiles enabled would be sufficient to achieve these goals. Thus the effective active area for this run was significantly lower than the previous run with 24,271/78400 (30.9%) of pixels disabled (see Fig. 4.9). Based on the stability of the pedestal calibration results from Run 1, continuous pedestal calibration subruns were not performed during this run. Data was collected in this configuration between the 23rd and the 25th, accumulating approximately 37 hours of livetime. An additional 12 hours of auxiliary data was collected including fully digitized waveforms of the analog signals of 4 pixels, read out by an oscilloscope. In this dataset, the external trigger markers in the PACMAN datastream indicated the presence of an analog signal on one of the pixels connected to the oscilloscope and a light signal within $\pm 150 \mu\text{s}$. This dataset was used as an independent validation of the charge simulation described in Sec. 3.3.3. However, precise analysis was challenged by an unknown transfer function introduced by the analog signal driver and long cable. Finally, data was taken under low-flow and no-flow configurations of the recirculation system to test charge readout stability with reduced cooling. The detector was then emptied and warmed on June 26th.

4.2 Data samples

Throughout data collection, subruns were marked by a human shifter as "good-for-analysis" based on the status of the light and charge DAQ monitoring, as well as by instruction from the on-site and off-site subsystem experts. Generally, runs were marked as "good-for-analysis" if the detector was operating stably in one of the two charge threshold configurations, and if subsystem experts were not performing debugging or calibration runs. Two kinds of detector instabilities were identified during operation by online data quality monitoring. On the CRS, occasional interruptions were encountered when a single channel would begin producing a continuous stream of self-triggers. These events would require manual intervention to stop the DAQ, re-configure the LArPix ASICs, and then resume operation. In a few instances, the pixel required disabling prior to resuming operation, but this was limited to a rate of ≈ 0.1 pixels/hour and so introduced a negligible loss in active area. The cause of these "hot" pixels is still unknown and is under investigation, but all subruns that were impacted by the presence of a hot pixel have been removed for this analysis. On the LRS, a bug in the light DAQ would occasionally fail to save any data to the output file from the beginning of the subrun. Active monitoring by the shifters was sufficient to identify these events and restart the DAQ before significant down time had elapsed. Any runs without light data were not included in this analysis.

Table 4.1 summarizes the datasets used for calibration and analysis described in this thesis.

Table 4.1: Summary of Module 0 datasets marked as "good-for-analysis".

Sample nickname	Run	CRS threshold (ke ⁻ eq.)	LRS threshold (ADC)	LRS window (samples)	Drift field (V/cm)	Subrun count	Events collected
High threshold	1	≈ 12	1500	256	500	81	11,124,904
Medium threshold	1	≈ 6	1500	256	500	104	24,495,996
Run 2	2	≈ 12	2500	1024	500	53	8,502,698
Pedestal	1	-	-	-	0	77	≈ 200/ch., ea.
LED calibration	1	-	-	256	0	15	≈ 10,200 ea.

4.3 Module 0 simulation




Figure 4.12: Cross-section of the GDML description of the Module 0 detector in the experimental hall at the University of Bern [196]. The central gray volume is the outside of the cylindrical cryostat that houses the Module 0 detector and includes a detailed model of the detector construction. This detail is not readily shown in this figure, instead refer to Fig. 4.1 for a close-up view. The blue volume surrounding the cryostat is a pit cut into the concrete on-site building and is filled with air. The brown rectangular volumes are simple representations of the building concrete.

To simulate Module 0, the simulation chain described in Ch. 3 was used. A simplified experimental hall geometry was implemented in Geometry Description Markup Language (GDML), including 50 cm of concrete overburden, the concrete pit containing the cryostat, and the cylindrical cryostat with a 2-cm thick steel lid. Within the liquid argon, a detailed model of the G10 sleeve and field cage structure were included along with brackets that support the field cage and light detectors. A cross-sectional rendering of this geometry is shown in Fig. 4.12. The base coordinate system used, placed the origin at the center of the

two active volumes within the cathode. The x -coordinate was defined to lie parallel to the cathode and parallel to the horizon. The y -coordinate was defined to point in the direction of the zenith, and the z -coordinate is parallel to the electric field, with TPC-0 at positive z and TPC-1 at negative z . The specific configuration of Geant4 and CORSIKA are described in Sec. 4.3.1 and 4.3.2, respectively, along with modifications to the event generators that were used to increase sample statistics.

4.3.1 CORSIKA parameters

Table 4.2 contains the specific CORSIKA parameters that were used for the Module 0 cosmic-ray flux simulation. No other customization of the CORSIKA physics models or primary-flux spectra were used to produce the reference simulation samples.

The output of CORSIKA is a collection of particles passing through a horizontal plane at the Earth’s surface. However, relatively few of these simulated particles actually enter the active volume of Module 0, resulting in a significant simulation inefficiency. Preliminary tests showed that CORSIKA was substantially slower than the later stages of the simulation chain, so some improvement in efficiency was warranted. Nominally, the horizontal plane could be reduced to a region closer to the detector volume, but this under-represents the number of high-angle particles, and so was not pursued. Instead, a re-sampling procedure was developed. In particular, a $10\text{ m} \times 10\text{ m}$ area 10 m above the detector was used to select particles from the CORSIKA-generated showers. For each shower, a random particle in the shower and a random point within the active volume of Module 0 were selected. All of the particles in the event were then translated horizontally such that the vertical projection of the selected particle passes through the selected point. This produces a sample with a substantially improved efficiency. Specifically, each translated CORSIKA shower was distributed into 11×11 separate events by translating the origin by a fixed horizontal translation of $\pm(i\Delta x, j\Delta z)$, where i and j run from 0 to 5, Δx is the length of the active volume in x , and Δz is the width of the active volume in z . This moderately improves the simulation efficiency for events where multiple scattering causes the particle to deviate from the vertical projection. However, it should be noted that within this re-sampling scheme, events are not fully statistically independent, particularly for the angular distributions.

4.3.2 EdepSim parameters

A Geant4 wrapper, EdepSim [160], was used to reduce the overhead in building the Geant4 simulation for Module 0. The EdepSim wrapper includes basic utilities for event generation, data persistence, and backtracking to truth information. Edep-sim version 3.2.0 was used with Geant4 version 10.6.1 with very little customization of the physics model. Table 4.3 contains the non-default parameters that were used in the Module 0 simulation.

To reduce the step size to the length scale of interest for this analysis, values for the Genat4 step size of $\alpha = 0.001$ and $\rho = 1\text{ mm}$ were used (see Sec. 3.2 for details). This results in relatively short and uniform step lengths. For example, a 300-MeV muon in the liquid

Table 4.2: Parameters for the physics model of CORSIKA v7.7400 used to generate the cosmic-ray flux at the Bern site.

Parameter name	Value	Description
PRMPAR	14	Primary particle type [PDG code]
ESLOPE	-2.7	Slope of primary particle energy spectrum
ERANGE	$1.3 - 10^5$ GeV	Range of primary particle energies
THETAP	$0 - 84.9^\circ$	Range of primary zenith angle
PHIP	$-180 - 180^\circ$	Range of primary azimuth angle
QGSJET	true	Enable QGSJET II-04 model for high energy interactions
QGSSIG	true	Enabled QGSJET II-04 model for high energy cross-sections
OBSLEV	550 m	Altitude of Bern, Switzerland
CURVOUT	false	Disable simulation of Earth's curvature
MAGNET	$(21.793, 42.701)$ μT	Earth's magnetic field direction (true North, vertical)
HADFLG	$(0, 0, 0, 0, 0, 2)$	Parameters for the hadronic interaction and fragmentation model
ECUTS	0.05 GeV	Energy cutoff for low-energy particles
MUMULT	true	Use Molière's theory for muon multiple scattering
ELMFLG	$(\text{false}, \text{true})$	Model parameters for electromagnetic component (use EGS4)
STEPFC	1.0	Multiple scattering step length factor for e^\pm
ARRANG	0°	Coordinate orientation
ATMOD	1	Use US standard atmosphere density model [210]

argon of Module 0 was simulated with step lengths varying from 1 mm to 2.4 mm. This step length was chosen because energy deposition fluctuations are handled at the Geant4 stage rather than the detector simulation stage, thus the length scale of energy deposition fluctuations must be similar to the position resolution of the detector. If larger or smaller step-sizes are used, the recombination models of Sec. 3.3.1 are not valid. Future modifications to the detector-simulation model that include energy deposition fluctuations could allow for a more flexible step size at the Geant4 stage.

Table 4.3: Geant4 and EdepSim configuration parameters used to simulate cosmic rays in Module 0.

Name	Value	Unit	Description
Physics list	QGSP_BERT		Sets the physics model used to calculate hadronic, EM, and decay processes (EdepSim default)
e^\pm, γ tracking cut	1	mm	Range cut for tracking γ, e^\pm (EdepSim default)
/edep/db/set/gammaThreshold	10	keV	Threshold for saving γ -ray truth information
/edep/db/set/neutronThreshold	1	MeV	Threshold for saving neutron truth information
/edep/db/set/trajectoryDeposit	10	keV	Threshold for saving truth charged-particle truth information
/edep/hitLength/volTPCActive	1	mm	Threshold for merging Geant4 trajectories
/process/eLoss/StepFunction	0.001		Energy loss step length α and ρ ,
	1	mm	described in text

The γ -ray truth tracking threshold was lowered to preserve annihilation photon truth information ($E_\gamma \approx 511$ keV), which was investigated as a potential avenue for muon charge separation. The neutron tracking threshold was lowered to track secondary neutrons produced by muon capture. And the general particle trajectory tracking was lowered to separate low-energy δ -rays from parent particles. The energy loss step function parameters were modified to produce ≈ 1 -mm steps, as described earlier.

4.4 Simulation tuning

After the initial comparison of the Module 0 data and the simulation was performed, some additional features needed to be incorporated into `larnd-sim`. First, digital-analog cross-talk introduces a non-Gaussian noise contribution, which is observed to enhance the tails of the dQ/dx distribution. A noise overlay mimicking the expected random statistics of digital-analog cross-talk was fit and applied as a post-process overlay to the simulation. Second, the charge thresholds in the data varied by a factor of $\approx 2\times$ channel-to-channel, as shown in Fig. 4.10. Because the channel response depends strongly on the charge threshold, it was necessary to include channel-to-channel variations in the simulation. Third, the avalanche characteristics of the SiPMs and scintillation characteristics were unknown prior to the Module 0 run, so a model was extracted directly from the LED calibration data. Fourth, the

PDE in the light LUT was poorly constrained and variations between light collection modules were observed, so the light LUT was tuned to account for these differences. And finally, a significant 10-MHz pickup signal from the LArPix clock was observed on a number of light SiPM channels, so a dedicated overlay using a pickup-signal template derived from the data was applied. These are described in detail in the following sections.

4.4.1 Noise overlay

Because the charge readout system uses a mixed-signal PCB to handle both the charge collection and the digital multiplexing, digital-analog cross-talk is a concern. Of particular relevance for the CRS are signals with frequencies less than the LArPix front-end bandwidth (≈ 5 MHz) and greater than the periodic reset frequency (2.5 kHz). One such signal is the digital chip-to-chip communication which used a signaling frequency of 2.5 MHz during Module 0 operations. Capacitive coupling between the pixel pad, trace, and chip pins and the digital signal lines can induce non-negligible cross-talk, e.g., a capacitance of only 100 aF to a signal line with a voltage swing of 1.8 V will induce a cross-talk signal of about 1000 e⁻, larger than the intrinsic noise of the readout. A variety of mitigation efforts at the hardware-level were included: pseudo-differential signalling that balances the positive swing cross-talk with a complementary negative swing, reduced digital-communication voltage, careful layout and PCB stack-up to reduce pixel capacitance, and grounding the CSA pins nearest to the IO pins, but digital-analog cross-talk was still found to meaningfully contribute to the overall charge readout noise. As future ASIC designs target eliminating this noise contribution, digital-analog cross-talk was not included in the dedicated simulation and instead was applied as part of a post-process overlay. The post-process overlay also incorporated channel-to-channel gain variations, which were not included in the simulation.

The excess noise model that was applied to each channel i and hit j is separated into two parts. First, a channel-to-channel gain factor

$$S_{\text{gain}}[i] \sim \text{norm}(\mu = \mu_{\text{gain}}, \sigma = \sigma_{\text{gain}}) \quad (4.1)$$

where i is the pixel index, μ_{gain} represents a potential overall scale shift, and σ_{gain} represents the global variation in the gain between channels. And second, a cross-talk noise term constructed from three independent random variates

$$S_{\text{XT}}[i] \sim \text{norm}(\mu = \mu_{\text{XT}}, \sigma = \sigma_{\text{XT}}) \quad (4.2)$$

$$N[ij] \sim \text{Bernoulli}(p = p_{\text{XT}}) \quad (4.3)$$

$$A[ij] \sim \text{Bernoulli}(p = p_{\text{asymm}}) \quad (4.4)$$

where j is the pixel trigger index, μ_{XT} represents the average scale of the cross-talk signal, σ_{XT} is the channel-to-channel variation in the cross-talk scale, p_{XT} is the probability that a given hit is impacted by cross-talk, and p_{asymm} is the asymmetry between positive and negative cross-talk. The new hit charge after the overlay was then updated to be

$$Q_{\text{overlay}}[ij] = S_{\text{gain}}[i]Q_{\text{sim}}[ij] + S_{\text{XT}}[i]N[ij](2A[ij] - 1). \quad (4.5)$$

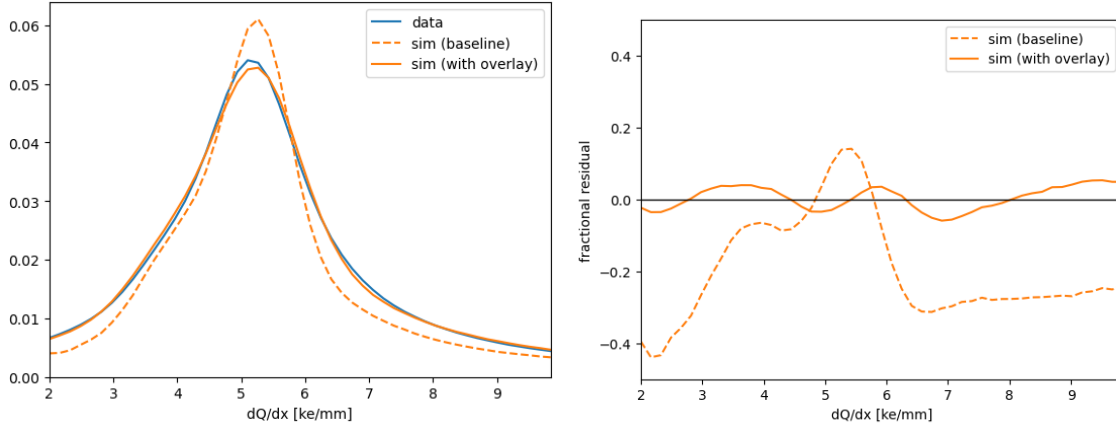


Figure 4.13: (Left) Single-pixel reconstructed dQ/dx for MIP-like tracks averaged over all pixels during a medium-threshold run, compared pre- and post-overlay. (Right) Fractional residual of single-pixel reconstructed dQ/dx pre- and post-overlay. Best fit values are quoted in the text.

The parameters of this overlay were determined by fitting the dQ/dx distribution for long, MIP-like tracks generated by CORSIKA to the Module 0 data. Figure 4.13 shows a comparison of the dQ/dx distribution before and after applying the noise overlay. The best-fit parameters for the medium-threshold sample were

$$\begin{aligned}
 \mu_{\text{gain}} &= 1.0107 \pm 0.0005 \\
 \sigma_{\text{gain}} &= 6.85 \pm 0.17\% \\
 \mu_{\text{XT}} &= 6.7 \pm 0.5 \text{ ke}^- \\
 \sigma_{\text{XT}} &= 2.5 \pm 1.1 \text{ ke}^- \\
 p_{\text{XT}} &= 24.5 \pm 0.8\% \\
 p_{\text{asymm}} &= 0.489 \pm 0.011,
 \end{aligned} \tag{4.6}$$

in reasonable agreement with naive expectations. This model was compared with other, simpler noise models:

- $S[i]Q_{\text{sim}}[ij]$ only
- $S[i]Q_{\text{sim}}[ij] + \text{norm}(\mu, \sigma)[ij]$
- $Q_{\text{sim}}[ij] + S_{\text{XT}}[i]N[ij](2A[ij] - 1)$ only,

but the combined model produced significantly better results. The best-fit parameters from the high-threshold data and the medium-threshold sample were also commensurate, suggesting that this model is a good reflection of the underlying excess noise processes. Due to practical considerations when the noise overlay was applied, channel correlations were only

preserved across a grouping of 32 events. Beyond this, a new scale factor was produced for each channel.

4.4.2 Threshold tuning






Figure 4.14: Reconstructed track overlap with individual pixels normalized to the total analyzed track length. The distribution is approximately uniform below 4.4 mm.

Charge thresholds modify the digitization time, hit charge, and the number of self-triggers substantially. Coupled with noise variations, this introduces a spatially-dependent variation in the CRS characteristics. To reliably mimic this effect in the simulation, the self-trigger thresholds were measured on individual pixels and included in the simulation as a per-channel configuration parameter.


To measure the self-trigger thresholds for each pixel, the hit-charge distribution was used. The hit-charge distribution is made up of the overlap of the track with the pixel (dx) convolved with charge- and energy-smearing effects (e.g. dE/dx fluctuations, recombination fluctuations, channel noise, etc.). The dx distribution for reconstructed MIP-like tracks from the CORSIKA medium-threshold sample is shown in Fig. 4.14. There is a peak near the pixel pitch (4.4 mm) as most tracks have a small zenith angle and thus overlap with the pixel parallel to one of the pixel pitch axes, as shown in Fig. 4.15a. However, for tracks not perfectly aligned with the vertical or horizontal axes, a portion of the track crosses the corner of the pixels, as shown in Fig. 4.15a and 4.15b. For these regions, an approximately uniform continuum between $0-\sqrt{2}d_{\text{pixel}}$ is formed.



(a) Tracks parallel to one of the pixel pitch directions.



(b) Tracks slightly off-parallel to the pixel pitch direction.



(c) Tracks diagonal to the pixel pitch directions.

Figure 4.15: 2D illustration of the mechanism that causes the dx distribution shown in Fig. 4.14. On the left, straight-line track segments are drawn with uniform offsets in x (blue) and a square region is highlighted, indicating the overlapping segment of the track with a pixel (black). On the right, the resulting dx PDF is shown for tracks with the corresponding orientation. As the track crosses the pixel at a larger angle, clipping occurs in which the track intersects with the corner of the pixel. Because the relative position between the track and the pixel is random and uniform, the clipped-segment length is also random and uniform.

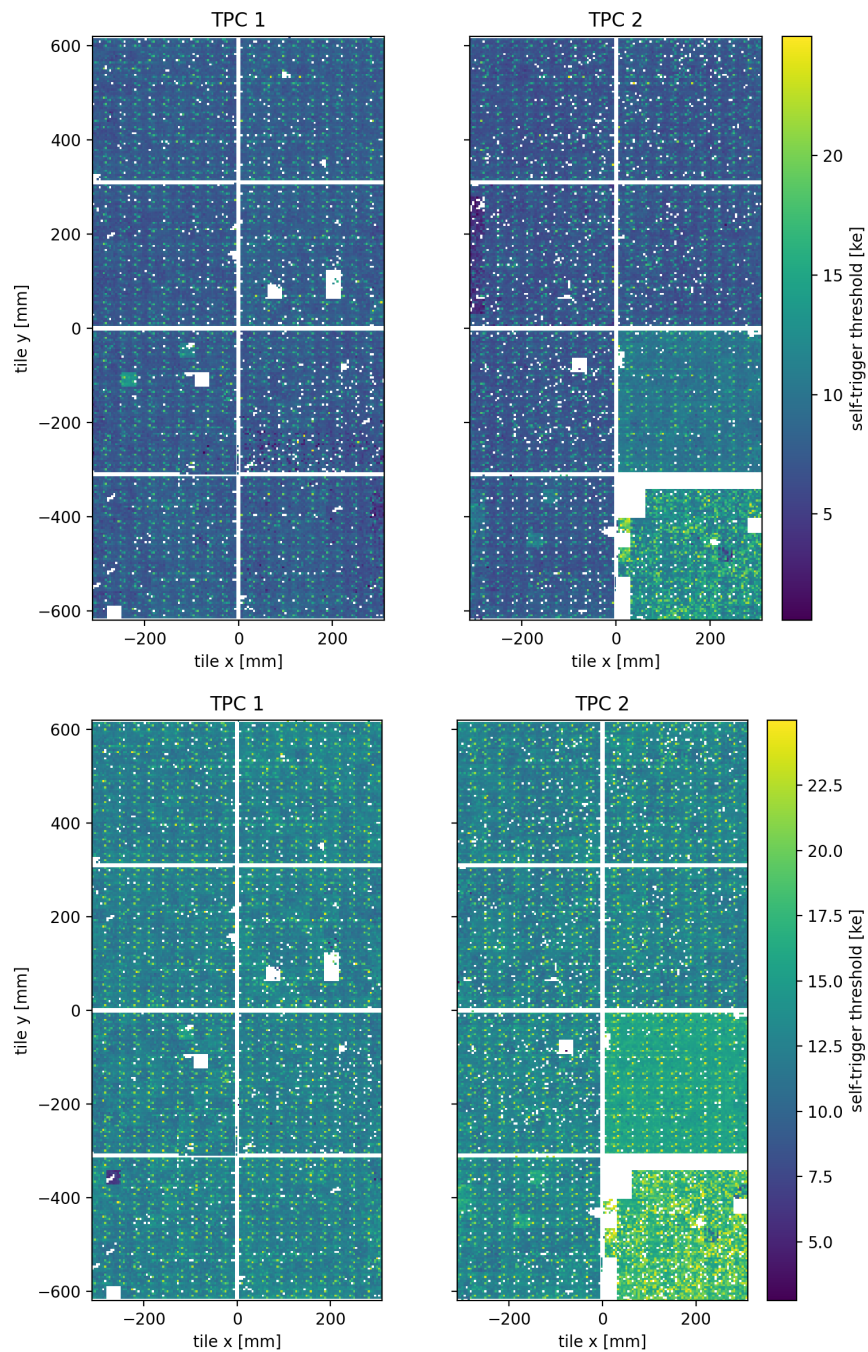


Figure 4.16: Self-trigger thresholds used in simulation, extracted from MIP-like tracks: (top) medium-threshold configuration, and (bottom) high-threshold configuration.

Shown previously in Fig. 3.8, the hit-charge distribution follows a similar pattern as the dx distribution except for a broadening of the peak and the addition of a sharp rising-edge that corresponds to the self-trigger threshold. By examining this rising edge, the underlying uniformity of the dx distribution can be exploited to measure the pixel threshold. This method is complicated somewhat by the presence of induced signals, which may trigger the pixel when the track does not overlap with the pixel at all, violating the assumption of uniformity. However, these signals are slowly varying compared to the collection signals and so only contribute by sharpening the rising edge at the charge threshold. Using the 50% point of the rising edge of this distribution on each individual pixel, the extracted self-trigger thresholds were validated with `larnd-sim` and found to have $<5\%$ systematic bias.

Figure 4.16 shows a 2D map of the pixel thresholds measured in each of the Module 0 datasets. A couple features of the noise environment are reflected in this map. First, pixels near the corners of each chip see enhanced digital-analog cross-talk. Because the thresholds were set to produce approximately equal self-trigger rates, the pixel thresholds are higher on these pixels (> 10 ke). And second, the two tiles in the lower left quadrant of TPC 2 also show globally higher self-trigger thresholds, likely due to pickup injected via the PACMAN power supply or the cable and grounding.

These thresholds were used as input for the `larnd-sim` simulation.

4.4.3 Light impulse model

To generate realistic waveforms, the calibration LED data was used to generate an average SPE waveform. To create the average waveform, each trigger in the calibration run was aligned by extrapolating the rising edge of the signal to the zero-crossing. The average waveform was then area normalized. As no significant variation was found between the SiPMs, the average waveform was used for the impulse response of the LRS simulation, as described earlier. Figure 4.17 shows this impulse response model.

It should be noted that this calibration approach does not necessarily reproduce other features of the signal timing. In particular, the propagation time of the photons in the LED calibration data is fixed, so the LUT is used to simulate this effect. In addition, the TPB and WLS re-emission probability, timing, and spectra may differ somewhat if the LED wavelength is not matched to the scintillation light spectrum. For this analysis, these complicating factors are largely ignored, as they have little impact on the final result.

4.4.4 Scintillation model

The nominal scintillation time constants of 7 ns and 1600 ns did not produce LRS waveforms that reflected the data. In particular, at times greater than $\sim 1 \mu\text{s}$ the observed light signals died off much faster than the naive expectation. The most likely cause of this is nitrogen contamination of the liquid argon. Nitrogen has a relatively low electron affinity and thus permits electron drift at much larger concentrations than other contaminants like oxygen or water. However, the presence of nitrogen does impact the scintillation light by reducing




Figure 4.17: SiPM impulse response model extracted from an average of waveforms collected during an LED calibration run.

the light yield of the triplet state. A fit to the average waveform with the impulse response deconvolved with a two-component exponential function, provides good agreement for a triplet decay time constant of 750 ± 50 ns. The substantially suppressed triplet component is consistent with a nitrogen contamination of ≈ 6 ppm, assuming a quenching factor of $0.11 \text{ us}^{-1}/\text{ppm}$ [119]. The deconvolved waveforms did not provide sufficient resolution to determine the singlet decay time constant.

4.4.5 Light LUT and PDE

The overall acceptance of the active volume of the detector was compared to the expected acceptance from the dedicated Geant4 simulation used to generate the light LUT. To compare the acceptance, a 4D histogram was filled with the total light yield on each light detector from a given 3D voxel within the detector using the reconstructed track segment hits. A 3D normalization histogram was also filled using the total charge of the track segment hits within each voxel. The acceptance was then calculated for each voxel by taking the ratio of the two histograms (equivalent to the PE/e^- yield ratio) for each voxel of the active volume. Overall, the acceptance was commensurate, although differences were apparent near the far-side of the detector with respect to the light detector, as shown in Fig. 4.18. A new light LUT after tuning was generated by reweighting each voxel by the observed difference in acceptance.

After performing the gain calibration described in Sec. 5.2.2, the overall PDE was then




Figure 4.18: 1D averaged profile of reweighted light LUT used in simulation for a SiPM in each light detector type. Each LUT bin is $4.4 \text{ cm} \times 4.7 \text{ cm} \times 3.8 \text{ cm}$. Drift axis is z and vertical axis is y .




Figure 4.19: Relative light yield for MIP-like tracks observed in Module 0 compared to the Geant4-based LUT described in Ref. [196].

tuned using the overall light yield on the 2-SiPM sum (for the LCMs) and the 6-SiPM sum (for the ArcLights). Generally, the LUT acceptance was reduced by about a factor of $3\times$ relative to the Geant4 expectation. Figure 4.19 shows the average photon detection efficiencies relative to the baseline light LUT.

4.4.6 Light noise overlay

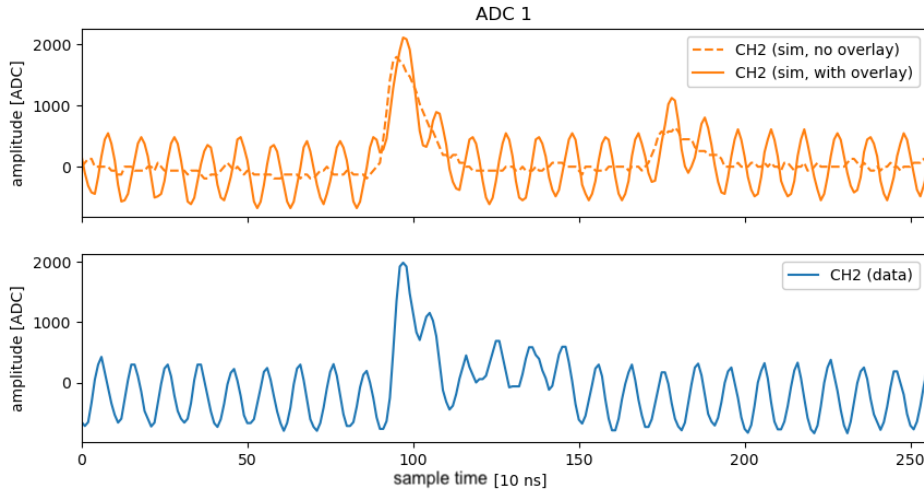


Figure 4.20: Example of a light waveform with and without noise overlay: (top) simulated waveform with a light signal present, (bottom) similar amplitude waveform from Module 0 data.

The 10-MHz LArPix clock induced significant pickup on the LRS waveforms, most dramatically on the ArcLight detectors. Offline, a dedicated filter was used to remove this pickup signal and is described in Sec. 5.2.2. However, in the interest of accurately reproducing the leakage of this signal through the filter, an overlay was applied to the simulated light waveforms to mimic this cross-talk. The noise model extracted with the dedicated filter was phase aligned by multiplying the template FFTs by the phase of the 10-MHz term of the FFT. They were then averaged for each SiPM across a typical run and amplitude normalized. The noise overlay was then applied to the simulated waveforms as

$$W_{\text{overlay}}[i, j] = W_{\text{sim}}[i, j] + A_i T_{\text{XT}}[i, j_{\text{shift}}] \quad (4.7)$$

where i is the SiPM index, j is the sample index, $W_{\text{sim}}[i, j]$ is the pre-overlay waveform, and $T_{\text{XT}}[i, j]$ is the amplitude-normalized cross-talk model. The template amplitude A_i is sampled from a normal distribution with a mean and standard deviation set to those of the filter-extracted noise model. A uniform random phase shift was applied by shifting the template by a random, discrete number of samples

$$j_{\text{shift}} \sim j + \text{floor}[n * \text{uniform}()] \pmod{n} \quad (4.8)$$

where n is the number of samples in the waveform. Figure 4.20 shows an example of a simulated waveform before and after including the noise overlay, along with a typical waveform on the corresponding SiPM from the Module 0 data.

Table 4.4: Summary of simulation datasets used in this analysis.

Sample	Disabled CRS channel list	CRS threshold (ke ⁻ eq.)	LRS threshold (ADC)	LRS window (samples)	Drift field (V/cm)	Event count
CORSIKA samples						
High threshold	Run 1	≈ 12	1500	256	500	34,716,210
Medium threshold	Run 1	≈ 6	1500	256	500	37,995,388
Run 2	Run 2	≈ 12	2500	1024	500	17,867,421
Stopping muon samples						
High threshold	Run 1	≈ 12	1500	256	500	4,490,000
Medium threshold	Run 1	≈ 6	1500	256	500	6,420,000
Run 2	Run 2	≈ 12	2500	1024	500	5,750,000

4.5 Simulation sample

A variety of simulated datasets were used for the analysis. With the tuning described in Sec. 4.4, the simulation was made to closely match the real configuration of the detector in both the high-threshold and medium-threshold operational configurations. A simulated dataset of cosmic-ray events was created for each of the detector configurations to provide approximately $5\times$ more stopping muons post-selection than the collected data. The differences in the disabled channels for Run 1 and Run 2 were included to account for the different acceptance of the detector during each period. The LRS trigger thresholds and readout windows were set according to the ADC readout settings of each run. Table 4.4 summarizes the data samples used to develop the reconstruction and perform the fit, to be described in Chapter 6.

An additional dataset of $\approx 100\%$ stopping muons was created to tune reconstruction algorithms. To generate these events, CORSIKA was not used, but rather the parameterized muon flux as a function of muon energy, described in Ref. [211]. To enhance the purity of the sample, muons were sampled from the 2D energy-zenith angle distribution and then rejected based on the following criteria. After selecting a point from the 2D distribution, a 3D point was randomly selected in a box with the same dimensions as the active volume of Module 0. The range of the muon in liquid argon using the CSDA was estimated for the sampled energy based on the range table of Ref. [212]. If the range of the muon was such that the

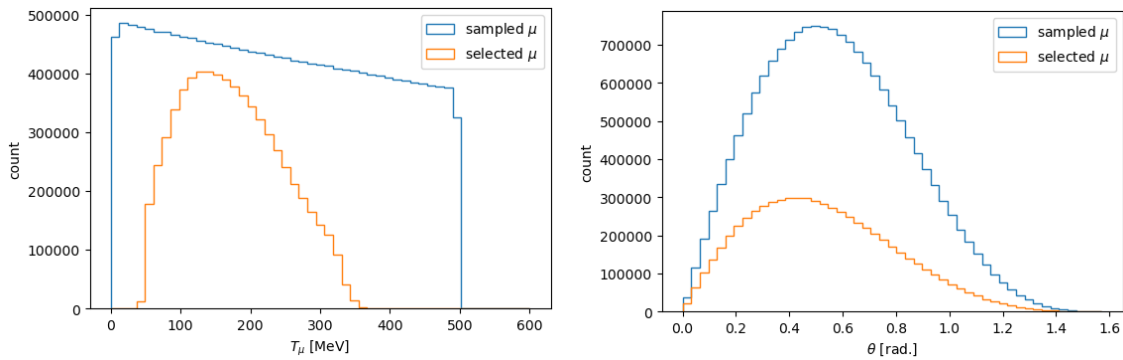


Figure 4.21: Muon kinetic energy and angular distribution created by the sampling approach defined in the text.

muon was expected to stop within the detector, the muon was selected, otherwise it was not. Figure 4.21 shows the muon kinetic energy spectrum and zenith angle distributions of the muons generated in this fashion. The charge of the muon was assigned to each randomly with an equal probability. A list of $\sim 5 \times 10^6$ of these muons were passed to the Geant4 stage and the different detector configurations were simulated. Due to multiple scattering, muons generated in this fashion could scatter out of the active volume and contributed a small systematic background ($< 0.5\%$).

Chapter 5

Data processing and analysis

There are primarily two quantities of interest that are extracted from the data. First is a metric that can separate muon decay from muon capture using the ~ 50 -MeV final-state electron. In this analysis, a generic metric called the *muon-decay nhit* is used, which is defined in Sec. 5.4.1. And second is the time delay between when the muon enters the detector and its subsequent decay. In this analysis, this is referred to as the *delayed time* or *decay time*. Obviously, these high-level quantities are not directly available in the data, and so a number of manipulations are required to extract them. This chapter describes the algorithms and calibrations used to create these quantities, as well as identify and select events that contain muon decays and captures.

Section 5.1 describes the process to sort through the self-triggered CRS and LRS data and identify activity that occurs closely in time. This forms the basis for what is defined as an "event" within this analysis. Additionally, this section describes the calibrations applied to align and synchronize the timestamps of both systems, such that activity from one subsystem can be associated with the other. Then Section 5.2 describes the manipulations performed on the digitized data to extract quantities related to the energy deposition in the detector and improve the data quality. This section contains subsections describing different algorithms applied to the CRS and LRS data. Section 5.3 describes the tracking and classification algorithms used to identify and reconstruct the muon trajectory as it comes to rest in the detector volume and to reject non-muon events. And finally, Section 5.4 describes the algorithms for separating muon decays and captures based on the event topology, along with the algorithm that determines the decay time using the scintillation light signal.

5.1 Event definitions

To begin with, the data that was recorded from the LRS and CRS systems were saved in a streaming fashion. Data was recorded to data files as it arrived at the DAQ, with no explicit integration between the different subsystems. For the CRS, messages were streamed from each of the two PACMAN cards. Each message contained a variable number of words.

Words were designated with a type, indicating the data contained within it. Generally, data arriving from LArPix was wrapped with meta-data allowing for multiplexing across multiple active UART connections with the LArPix tiles. Additional data words, originating from the PACMAN, provided information about clock synchronization and light-trigger timing. Messages were then packed into a binary dataset within an HDF5-formatted file [213]. Meta-data associated with each message, namely the unique identifier for each of the PACMAN, was saved alongside the primary dataset with a 1:1 correspondence. No effort to time order the messages was maintained – messages were appended to the primary dataset as the DAQ received and handled them. For the LRS, User Datagram Protocol (UDP) packets containing the waveforms of all enabled channels on a given ADC and metadata containing timing information were written directly to the end of a pair of binary data files, one for each ADC. Thus, in order to reconstruct and classify natural cosmic-ray events in the detector, an algorithm needed to be employed to time order and synchronize the data within the 3 output data files.

The event reconstruction proceeded by first treating the light and charge datastreams independently. As there is no trigger available to indicate the production of the particle, i.e. a beam trigger, events must be identified by selecting time slices with activity within each subsystem. Once events had been identified independently within each subsystem, they were matched to form a global event that contains both light and charge information. Because the timescale relevant for the light events ($O(1\mu s)$) is much shorter than that for the charge events ($O(100\mu s)$), global events consisted of exclusively one charge event and any number of light events that occurred within the charge-event time window.

5.1.1 Charge-event definition

At the lowest level, charge events were defined based on the number of self-triggers that occurred in the charge datastream within a period of time. Charge events were created by binning the individual LArPix data packets into a histogram based on their digitization timestamp $t_{PPS, \text{charge}}$. A 5-packet/10- μs threshold was applied to the histogram to filter out noise events and very low-energy activity, such as that from radioactive ^{39}Ar decays. The immediately neighboring time bins of each above-threshold bin were combined and declared as a "charge event". This means that the event time duration is variable and contains all charge activity in a pre-window of 0-100 μs as well as a post-window of 0-100 μs , along with all activity within at least one full cathode-to-anode drift time. There were a few complicating factors involved in the construction of the charge events which will be discussed in the following subsections.

5.1.1.1 Time ordering

A given packet from the charge system accumulates a variable delay as it moves through the FIFO buffers within the Hydra network, the FIFO buffers within the PACMAN FPGA, the circular buffer between the PACMAN data server and programmable logic (PL), and

the local network from the PACMAN to the CRS DAQ. Within the Hydra network, each UART link is limited by the UART bandwidth to a maximum of $26.4 \mu\text{s}/\text{packet}$, or about $37.8 \times 10^3 \text{ packets/s}$. This means that packets originating deeper in the network accumulate a larger delay than packets closer to the PACMAN. Additionally, when a packet is transferred between two chips, it is placed into a local FIFO queue. Thus an additional delay is incurred that depends on the current queue depth of each chip that relays the packet. The total delay that a packet accumulates through the Hydra network is then highly variable, but with a typical total delay on the scale of $O(\text{ms})$. Additionally, a bug was identified in the LArPix v2a ASIC that causes every 512th packet to be held in the FIFO buffer until 1536 additional packets pass through the buffer. This "512 bug" caused an additional, significant delay to a small number of packets, impacting $1/512 \times N_{\text{nodes}}$ of all packets ($\approx 2\%$) leaving the LArPix tiles.

To ensure good timestamp synchronicity, the PPS timestamp was reset every second during data collection. Thus, resolving the variable delay by a simple time-ordering by the PPS timestamp would introduce ambiguities in events that occur within different PPS intervals. Instead, the following algorithm was used. Packets were first placed into a running buffer. PACMAN-generated data words, produced upon each SYNC event, were flagged within the buffer and used to define the PPS intervals. Because these markers are produced at the PACMAN card and are unaffected by the Hydra-network delay, they are guaranteed to arrive in the data stream before packets that originate after a SYNC event. LArPix-generated packets are then assigned to the correct PPS interval based on the difference between the PPS timestamp and the receipt timestamp, latched when the packet arrives at the PACMAN card. For packets with a PPS timestamp greater than the receipt timestamp, a SYNC event occurred between when the packet was produced and when it arrived at the PACMAN. These packets are then assigned to the PPS interval prior to the last SYNC event. Otherwise the packet is assigned to the PPS interval after the last SYNC event. Once the PPS interval had been determined, the packets were then sorted and binned by lexicographical order on the PPS-interval number first, followed by the packet timestamp. This was practically achieved by adding $N \times 10^7$ ticks to the timestamps within each PPS interval, where N is a counter for the current PPS period.

5.1.1.2 Bit errors

It was later discovered that some of the CRS data was periodically corrupted due to bit errors in data transferred between LArPix chips. Depending on position of the bit that was affected, the impact of the bit error was highly variable. As an example, if the least-significant bit (LSB) of the ADC dataword is modified in the packet, the error would go effectively unnoticed except for a small increase in the apparent charge noise. However, if the chip id of the packet is modified, the error would cause the wrong position to be associated with the packet and would cause an error in the subsequent reconstruction of the particle position. Identifying impacted packets was possible, depending on the failure mode, but not general, as some errors resulted in a true loss of information. For instance, in the

previous example, if the bit error falls in the bits of the chip id in such a way that the chip id becomes another valid chip id within the same Hydra network, there is no way to distinguish a modified packet from an un-modified packet originating from the other chip id. That said, to identify a subset of bit errors, the receipt timestamp of the packet could be compared to its PPS timestamp. If the data rate is low and the chip FIFO buffers are empty, a packet picks up a fixed delay. This delay depends only on the network depth D_i , or number of chips within the network configuration that lie between the originating chip i and the PACMAN card. This delay was empirically determined to be

$$\begin{aligned} t_{\text{delay}} &\equiv t_{\text{receipt}} - t_{\text{PPS, charge}} \\ &\approx (28.4D_i + 0.4) \mu\text{s} \end{aligned}$$

where the 28.4 μs is slightly longer than the time duration needed to transfer one packet between two chips (26.4 μs). The expected delay was then compared to the observed delay for each packet originating from a given chip. This approach identified two significant bit errors that could be corrected.

First, the 23rd, 24th, and 25th bits of the stored timestamp within the packets arriving on a single PACMAN UART would periodically be set to a fixed, incorrect value, resulting in a large negative offset in t_{delay} . All chips upstream of Chip 80 on IO Channel 23 of PACMAN 1 (31 chips, 1519 channels) were impacted, suggesting a failure isolated to a single chip or chip-to-chip link. Bit errors of this form were present in the quality control (QC) data for the tile, but were not discovered until after the Module 0 operation. This suggests that the issue could have been caught and potentially remedied at an early stage with the appropriate QC test.

Second, packets arriving from a number of chip ids periodically had a single bit of the chip id modified to an incorrect value. This resulted in the presence of packets arriving from chip ids that were not included in the network configuration, as well as, chip ids that had a bi-modal t_{delay} distribution. All chips upstream of Chip 100 on IO Channel 20 of PACMAN 1 (14 chips, 686 channels) and Chip 96 on IO Channel 16 of PACMAN 2 (4 chips, 196 channels) were impacted. Again, this suggests a failure of a single chip or chip-chip link. Where possible, these bit errors were corrected either by using content of offending bits on other unaffected packets nearby in the datastream, or by using a lookup table of known bit errors. However, additional evidence of more widespread errors in the chip id and the timestamp were present at low rate. The total impact of these errors was roughly estimated to be $\approx 1.5\%$ of all packets by using the rate of packets with negative t_{delay} (indicative of a bit error in the timestamp bits) and by extrapolating to the full packet assuming that the impacted bits are randomly distributed throughout the packet.

5.1.1.3 SYNC noise

Significant digital-analog cross-talk was observed early in the run between the SYNC signal and the analog front-end of the LArPix chips. This coupling caused a large number of




Figure 5.1: Self-trigger rate as a function of PPS timestamp along with the cut applied to remove SYNC-induced self-triggers.

channels to trigger within a short period of time upon each synchronization of the PPS timestamp. When a large number of pixels became active, a voltage drop on the LArPix tile supply voltages caused a secondary burst of triggers across almost every channel on the tile. This secondary burst could persist for a significant duration (≈ 0.1 seconds) and had the effect of a substantially elevated trigger rate near each SYNC event. This is reflected in Fig. 5.1, which shows the total self-trigger rate of the two TPCs in Module 0 for a typical 20-minute-run period. A large peak due to this effect is seen at $t_{\text{PPS, charge}} = 0$. This excessive trigger rate posed practical challenges for event building and reconstruction during the impacted time duration. A hard cut at $t_{\text{PPS, charge}} > 0.1$ s was applied to remove packets associated with these events and reduce the reconstruction burden.

5.1.2 Light-event definition

For the light data, both ADCs were triggered based on the summed waveform of the 6 SiPMs in an LCM group. Each trigger provided a digitized time window of the current on each SiPM; a unix timestamp $t_{\text{DAQ, light}}$, defined by the network time protocol (NTP)-time data that was sent by the readout board; and a 62.5-MHz clock counter $t_{62.5 \text{ MHz}}$, initialized when the readout board was brought online. To timestamp each trigger into the same timebase as the charge readout, the PPS signal was digitized on one channel of the ADC and used as a secondary trigger for the LRS. These events were easily separable from the data. Thus, the 62.5-MHz clock timestamp of PPS events were used to create a PPS timestamp for each light trigger

$$t_{\text{PPS, light}} = (t_{62.5 \text{ MHz}} - t_{62.5 \text{ MHz, last SYNC}}) \frac{10 \text{ MHz}}{62.5 \text{ MHz}}. \quad (5.1)$$

While assembling the light events, individual triggers were read from the current position of both LRS data files and their timestamps compared. The triggers from each light ADC were matched if $\Delta t_{\text{PPS, light}} < 1 \mu\text{s}$ and the $\Delta t_{\text{DAQ, light}} < 1 \text{ s}$, forming a light event. Otherwise, partially-filled events were produced from the earlier of the two ADCs and the next trigger is read from the corresponding file. These partially-filled events were caused by two aspects of the LRS. First, the UDP transport protocol used by the light DAQ does not guarantee packet receipt and so some triggers were lost. And second, for the first second of each run period no PPS trigger been recorded and so time synchronization could not be achieved. These events were reconstructed, but not included in the analysis. In total, 0.18% and 0.06% of light triggers were impacted by packet loss and missing time synchronization, respectively.

5.1.3 Light- and charge-event matching

To facilitate the matching between the CRS and LRS, the global trigger used to trigger the LRS system was also provided to both PACMAN cards. Upon receiving a trigger pulse at the PACMAN card, a marker was embedded into the CRS datastream. This marker contained the current PACMAN PPS timestamp that corresponding to that event. The NTP-synchronized unix timestamp recorded by the DAQ ($t_{\text{DAQ, charge}}$) was used to disambiguate the PPS interval. An initial matching was performed between the PPS timestamps and the DAQ timestamps with a matching criteria of

$$|t_{\text{PPS, charge}} - t_{\text{PPS, light}}| < 10 \mu\text{s}$$

and


$$|t_{\text{DAQ, charge}} - t_{\text{DAQ, light}}| < 2 \text{ s},$$

which was sufficient to match $> 99\%$ of the embedded markers to light events. However, the uncalibrated PPS timestamps of each PACMAN card and each light ADC drift with respect to each other primarily due to slight differences in their clock frequencies. This amounts to a drift of up to a few microseconds between each PPS pulse. Thus, after the initial charge- and light-matching was performed, a timestamp correction was extracted to improve the timestamp resolution.

To calibrate the LRS timestamps, the relative timestamp between the respective ADCs was calculated

$$t_{\text{PPS,rel},i} = t_{\text{PPS},i} - \frac{1}{2} \sum_{i=1}^2 t_{\text{PPS},i}. \quad (5.2)$$

A robust random sample consensus (RANSAC) linear fit was then applied to extract the relative clock drift and reset offset of $\approx \pm 122.5 \text{ ns/s}$ and $\approx \pm 1.0 \text{ ns}$, respectively. Figures 5.2a and 5.2b show the results of this fit. Three bands are present due to a trigger delay for one




(a) Best fit ADC 0 time correction:

$$\begin{aligned}\delta t(0) &= +1.08 \pm 0.04 \text{ ns,} \\ \delta t(1\text{s}) &= -122.55 \pm 0.06 \text{ ns}\end{aligned}$$

(b) Best fit ADC 1 time correction:

$$\begin{aligned}\delta t(0) &= -1.07 \pm 0.05 \text{ ns,} \\ \delta t(1\text{s}) &= +122.48 \pm 0.08 \text{ ns}\end{aligned}$$



(c) Residual timestamp distribution for each digitizer.

Figure 5.2: Clock correction applied to LRS PPS timestamps. Figure (a) and (b) show the PPS timestamp of each digitizer compared to the event mean, along with the best fit used by the clock drift correction. Figure (c) shows the residual timestamp after the clock drift correction has been applied. Three bands are present due to the global trigger timing scheme.

of the ADCs in the events with light activity that triggers only a single LRS ADC. The drift and offset was then subtracted from the corresponding timestamp of each ADC, resulting in a trigger timestamp resolution of $\sigma = 13.53 \pm 0.04$ ns. The timestamp residuals are shown in Fig. 5.2c.




Figure 5.3: Self-reported, uncalibrated SYNC timestamps from each PACMAN card during a typical run. The peaks are not aligned with the nominal 10^7 ticks/s due to different clock frequencies of the PACMAN cards.

The timestamp correction for the CRS data was corrected using a linear interpolation

$$t_{\text{corrected}} = \frac{(t_{\text{raw}} + \langle t_{\text{offset}} \rangle)}{\langle t_{\text{SYNC}} \rangle} \quad (5.3)$$

where $\langle t_{\text{offset}} \rangle$ is an additional delay that was incurred by the reset logic in the LArPix ASIC. The fit parameters were extracted using the same RANSAC approach as with the LRS timestamp correction, and by comparing the matched CRS event timestamp to the LRS timestamp as a function of the LRS PPS timestamp. The extracted corrections that were used are quoted in Table 5.1. In addition to the global-trigger timestamp markers, in the CRS, each SYNC event produces a marker containing the timestamp just prior to the timestamp reset and was checked for the overall clock stability. The run-to-run average clock frequencies were stable to $(2.2 \pm 0.3) \times 10^{-8}$ and the second-to-second clock frequencies were stable to $(7.6 \pm 0.6) \times 10^{-8}$ on both PACMAN cards. A small number of spurious resets were observed at a rate of $(0.025 \pm 0.002)\%$ and were observed on both PACMAN cards synchronously, likely due to pickup on the inverter used to receive and distribute the GPS PPS signal.

Table 5.1: PACMAN clock corrections used to synchronize charge data.

PACMAN ID	$\langle t_{\text{offset}} \rangle$ [ticks]	$(\langle t_{\text{sync}} \rangle - 10^7)/10^7$
1	-9.597 ± 0.008	$(4.0021 \pm 0.0013) \times 10^{-6}$
2	-9.329 ± 0.007	$(1.1770 \pm 0.0013) \times 10^{-6}$

5.2 Signal processing

Prior to the reconstruction of the high-level quantities such as the trajectory and decay time, the raw data was processed into mid-level, calibrated data objects. The aim was to reduce variation between data samples and improve the signal characteristics. This section describes the manipulations to the raw data to produce inputs for the higher-level reconstruction.

5.2.1 Charge readout

As mentioned before, data from the CRS was recorded in a packetized form, where each packet corresponds to a single self-trigger of a LArPix pixel. The recorded ADC value V_{ADC} is related to the induced voltage via

$$V_{\text{ADC}} = 2^n \left(\frac{V_{\text{sig}} + V_{\text{ped}} - V_{\text{cm}}}{V_{\text{ref}} - V_{\text{cm}}} \right), \quad (5.4)$$

where n is the number of bits in the ADC value ($n = 8$), V_{sig} is the signal voltage proportional to the integrated current, V_{ped} is the quiescent pedestal voltage of the CSA, and V_{cm} and V_{ref} are reference voltages. Inverting this expression for V_{sig}

$$V_{\text{sig}} = V_{\text{ADC}} \left(\frac{V_{\text{ref}} - V_{\text{cm}}}{2^n} \right) - V_{\text{ped}} + V_{\text{cm}}. \quad (5.5)$$

The V_{ref} and V_{cm} are set by an internal 8-bit digital-to-analog converter (DAC) in the LArPix chip and are set relative to the analog voltage V_{DDA} provided by the PACMAN board

$$V_{(\text{ref,cm})} = V_{\text{DDA}} \left(\frac{V_{(\text{ref,cm}) \text{ DAC}}}{2^n} \right). \quad (5.6)$$

In bench testing, a non-negligible voltage drop was observed over the cable from the cryostat feedthrough to the LArPix anode tile. This voltage drop varied with cable length but was tuned at the PACMAN supply to provide $V_{\text{DDA}} = 1.80 \pm 0.04$ V at each anode tile.

5.2.1.1 Pedestal calibration

The pedestal voltage V_{ped} varied by ≈ 50 mV (≈ 30 mV) channel-to-channel in liquid argon (room temperature). Dedicated calibration runs were taken periodically throughout the data




Figure 5.4: (Left) Distribution of pedestal calibration values used for Module 0 data processing. (Right) Width of V_{ADC} distribution on each channel, extracted from pedestal data.

collection of the first Module 0 run. For these calibration runs, each chip was operated in a periodic-trigger mode and with the cathode HV turned off. In the periodic-trigger mode, each channel was forced to digitize at a fixed interval, irrespective of the discriminator output. Because the sampling is uncorrelated with activity on the pixel, this provides an estimate of the ADC distribution when no signal is present. A fixed, channel-dependent calibration value was extracted at the start of data collection using a truncated mean of ± 3 ADC values around the most common ADC value received on each channel during this run. The pedestal ADC value was then converted to a voltage using Eq. 5.5, setting $V_{\text{ADC}} \equiv \langle V_{\text{ADC}} \rangle$ and $V_{\text{sig}} \equiv 0$. Periodic pedestal calibration runs were taken approximately every few hours during data collection. However, the derived pedestal values were found to vary by only ≈ 1 mV run-to-run for individual channels, thus a single calibration value was deemed sufficient for each channel throughout the data period. Figure 5.4 shows the distribution of pedestal calibration values and the observed width of V_{ADC} within the pedestal calibration runs.

5.2.1.2 Charge reconstruction

The signal voltage on the pixel is related to the true charge distribution via Eq. 5.5. Unfortunately, this expression cannot be inverted trivially to determine the charge density. Instead, we make a number of approximations to obtain simple reconstruction formulae.

Since the leakage current is negligibly small at liquid-argon temperature, it is reasonable to take

$$I_{\text{leak}} \rightarrow 0. \quad (5.7)$$

Discounting hits with a total integrated charge very close to the trigger threshold, the noise can be assumed to contribute in an unbiased fashion and thus

$$\langle \text{noise} \rangle \approx 0. \quad (5.8)$$

Under the assumption of uniform drift, $\vec{v}_d(t) = v_d \hat{z}$ and thus

$$V_{\text{sig}}[i] \approx G_p v_d \int_{t_0}^{t_i} \left[\int \rho(\vec{x}', t) \frac{\partial W(\vec{x})}{\partial z} \Big|_{(\vec{x}=\vec{x}')} d\vec{x}' \right] dt, \quad (5.9)$$

where the individual pixel gain is G_p and (\vec{x}, \vec{x}') are defined relative to the pixel center. With a simple model for the electron attachment $\partial_t \rho = -\rho/\tau$ and neglecting the small impact from diffusion, the exponential time dependence can be factored out and the charge density can be assumed to evolve uniformly

$$\begin{aligned} \rho(\vec{x}', t) &\approx \rho(\vec{x}'_{\perp} + v_d t \hat{z}, t_0) e^{-t/\tau} \\ V_{\text{sig}}[i] &\approx G_p v_d \int \left[\int_{t_0}^{t_i} \rho(\vec{x}'_{\perp} + v_d t \hat{z}, t_0) e^{-t/\tau} \frac{\partial W(\vec{x})}{\partial z} \Big|_{(\vec{x}=\vec{x}')} dt \right] d\vec{x}' \end{aligned} \quad (5.10)$$

where $\vec{x}'_{\perp} \cdot \hat{z} = 0$. For a uniform track-like deposit,

$$V_{\text{sig}}[i] \approx G_p v_d \frac{dQ}{dx} \oint \left[\int_{t_0}^{t_i} e^{-t/\tau} \frac{\partial W(\vec{x})}{\partial z} \Big|_{(\vec{x}=\vec{x}'_{\perp} - v_d t \hat{z})} dt \right] d\vec{x}' \quad (5.11)$$

where the path integral is evaluated over the initial 1D track trajectory. Because $\partial_z W$ is sharply peaked at $t_0 \equiv z_0/v_d$ compared to the electron drift lifetime τ , the exponential factor can be treated as a constant term

$$V_{\text{sig}}[i] \approx G_p v_d \frac{dQ}{dx} e^{-t_0/\tau} \oint \left[\int_{t_0}^{t_i} \frac{\partial W(\vec{x})}{\partial z} \Big|_{(\vec{x}=\vec{x}'_{\perp} - v_d t \hat{z})} dt \right] d\vec{x}' \quad (5.12)$$

For multiple triggers i and multiple channels j , we can discretize the spatial integral on the pixel pitch Δ

$$V_{jk}^{\text{sig}}[i+1] \approx G_p v_d \Delta \frac{dQ}{dx} e^{-t_0/\tau} \sum_{lmn} \eta_{lmn} \left[\int_{t_i}^{t_{i+1}} \frac{\partial W(\vec{x})}{\partial z} \Big|_{(\vec{x}=\vec{x}'_{\perp} - v_d t \hat{z})} dt \right] \quad (5.13)$$

where the transverse component $\vec{x}'_{\perp}(x, y)$ can be expressed in terms of the summation index ($x'_{\perp} = (l-j)\Delta$) and ($y'_{\perp} = (m-k)\Delta$). The sparse matrix η_{lmn} is introduced to represent the discretized trajectory and evaluates to 1 in each voxel that the track passes through.

In general, one could solve

$$\frac{e^{t_0/\tau}}{G_p v_d \Delta} \sum_{ijk} \mathcal{R}_{jk}^{lmn}[i] V_{jk}^{\text{sig}}[i] \approx \frac{dQ}{dx} \eta_{lmn} \quad (5.14)$$

for an ideal response matrix. This sparse matrix form may be well-suited for more complex approaches to charge reconstruction, such as those in Ref. [214] and [215].

However, in this analysis only the collection signals were considered, $R_{jk}^{lmn} = 0$ for $l \neq j, m \neq k$, allowing us to drop the l, m, j , and k indices of the sum. Thus the unknown factors on left-hand side of Eq. 5.14 are the non-zero voxel(s) of interest n and the ionization density dQ/dx . We assume that the ionization is sparse, leading to only one non-zero term of η_n , which is chosen using a uniform weight (described in Sec. 5.2.1.5). For the track charge Q , a simple sum is used

$$Q \equiv \Delta \frac{dQ}{dx} = \frac{e^{t_0/\tau}}{G_p} \sum_i V_{\text{sig}}[i]. \quad (5.15)$$

5.2.1.3 Front-end gain calibration




Figure 5.5: Comparison of injected vs. output pulse amplitudes using a dedicated test pulse circuit at room temperature and while submerged in liquid nitrogen. The reported ratio of -6.77 ± 0.04 corresponds to a gain of $4.52 \pm 0.03 \mu\text{V}/e$.

To calibrate the front-end gain G_p , dedicated measurements were performed using a pulse-injection circuit built into a custom PCB. A voltage step was capacitively coupled directly into two inputs of a LArPix chip. This enabled injecting a known charge of $C_{\text{inj}}\Delta V$. The capacitance C_{inj} of the circuit was known to 5%. The analog monitor of the channel-under-test was then digitized across a long timescale such that the transient behavior could easily be avoided. Measurements were repeated at room temperature and while submerged in 1-atm liquid nitrogen. The gain at liquid-nitrogen temperature was measured to be

$$G_p = 4.52 \pm 0.23 \mu\text{V}/e, \quad (5.16)$$




Figure 5.6: Example of track reconstruction and dx_{2D} (solid blue) calculation for the validation of the front-end gain calibration.

shown in Fig. 5.5. The uncertainty has been inflated relative to Fig. 5.5 to account for the 5% systematic uncertainty in C_{inj} . Notably, a significant shift in the gain of $\sim 5\%$ was observed between measurements taken at room and at liquid-nitrogen temperatures.

The gain calibration was then validated using the dQ/dx of through-going muons. The CORSIKA-generated cosmic-ray sample was simulated using a gain of $4.51 \mu\text{V}/e$ and realistic channel thresholds. Track segments were reconstructed via the method described Sec. 5.2.1.6. For both the data and the simulation, a MIP-like sample was selected by requiring track segments to be greater than 10 cm in length, with an overall dQ/dx consistent with a MIP ($> 4 \text{ ke}^-/\text{mm}$, $< 8 \text{ ke}^-/\text{mm}$), and an unambiguous event t_0 .

The dQ/dx on each pixel was then calculated as

$$\frac{dQ}{dx_{\text{pixel}}} = \frac{\cos \theta_{\text{drift}} \sum_i Q_i e^{t_{\text{drift}}/\tau}}{dx_{2D}} \quad (5.17)$$

where Q_i is the charge value of each self-trigger, dx_{2D} is the overlap of the track (projected onto the anode) and the pixel, θ_{drift} is the angle between the track and the anode plane, and τ is the electron drift lifetime, described in the next section. Figure 5.6 shows a reconstructed track and dQ/dx calculation for a single pixel. Because the dQ/dx width and peak value depends on the segment length, a cut was placed on $\cos \theta_{\text{drift}}/dx_{2D}$ to require only segments that significantly overlap with a pixel ($> 4\text{mm}$, $< 5\text{mm}$). Shown in Fig. 5.7a, there is good agreement in the median dQ/dx of $(5.02 \pm 0.01) \text{ ke}^-/\text{mm}$. The contribution of channel-to-channel gain variations, which is not simulated, was determined to be

$$\sqrt{\sigma_{\text{data}}^2 - \sigma_{\text{sim}}^2} / \mu_{\text{data}} = (5.02 \pm 0.03)\%,$$

in line with the expectation of 5%, and in reasonable agreement with the extracted channel-to-channel gain variation determined in the excess-noise-model fit of $(6.9 \pm 0.2)\%$ (Sec. 4.4.1).




Figure 5.7: Median and width of the dQ/dx distribution for each pixel using reconstructed MIP-like tracks. An increased width and an additional 5% variation in the median is present in the data due to digital-analog crosstalk and channel-to-channel gain variations that are not simulated.

5.2.1.4 Electron lifetime calibration

The electron lifetime was calibrated using a dedicated sample of through-going, anode-cathode crossing muons. As described in Sec. 3.3.3, electronegative impurities can capture electrons along their drift path, resulting in an exponential attenuation factor

$$Q(t_{\text{drift}}) = Q_0 e^{-t_{\text{drift}}/\tau_e}. \quad (5.18)$$

This effect can be calibrated at the hit level by applying a correction term

$$Q_{\text{reco}} = Q_{\text{raw}} e^{t_{\text{drift}}/\tau_e}. \quad (5.19)$$

To do this, a sample of anode-cathode crossing tracks was used. Because the charge collected by a pixel is dependent on the track orientation, the dQ/dx is used instead of the hit charge and is calculated from the straight-line estimate of the anode-cathode crossing track direction. The resulting peak of the dQ/dx distribution within each drift time bin was then fit to an exponential as a function of the drift time. This calibration was performed offline for each run at 500 V/cm and the results are shown in Fig. 5.8. No significant variation was observed in either Run 1 or Run 2, with the electron lifetime rising slightly from ≈ 2 ms at the start of each run to $\approx 2.5 - 3$ ms by the end of the run. Note that there is no guarantee that Run 1 and Run 2 would have the same electron lifetime as the detector was emptied and brought to room temperature between runs. The initial preparation of the detector was also

significantly different, with Run 1 under full vacuum prior to filling the cryostat and Run 2 using a piston purge technique [208]. These results show that the piston purge technique is viable for the ArgonCube modular TPC design.




Figure 5.8: Electron lifetime during operations of Module 0. Only standard runs at 500 V/cm field are included, so gaps are present. Taken from Ref. [207].

5.2.1.5 Induced-hit merging

Due to the nature of the induced current on the pixel from a drifting charge, it is possible for a pixel to trigger, digitize, and reset multiple times on a single ionization trail. Figure 5.9 depicts an example of a signal induced on a pixel when a track-like ionization cloud reaches the anode. The shape can be considered as two components: first, a slowly rising *far-field* signal that arises due to the motion of the drifting electron cloud within the uniform field; and second, a quickly rising *near-field* signal that arises due to the motion of the electron cloud within a few pixel pitches of the anode where the field is no longer uniform. Because LArPix utilizes an integrating front-end and there is no shielding plane in front of the anode, the far-field signal can contribute a substantial fraction of the overall signal. For either relatively large signals or when operating with a low self-trigger thresholds, the discriminator can trigger on this signal, and depending on the timing, trigger multiple times for a single ionization cloud. This leads to the distinctive hit charge distribution, seen previously in Fig. 3.8. It also contributes to an apparent broadening of the particle track in the drift direction and, to a lesser extent, the transverse directions.




Figure 5.9: (Left) Trigger timing of the integrated signal on a pixel for a track parallel to the anode ($\theta_{\text{drift}} = 0$). (Right) Bias in the reconstructed t_0 for tracks parallel to the anode plane using a weighted sum of the trigger time.




Figure 5.10: Two candidate muon-decay events demonstrating triggering on the inductive signal prior to the arrival of charge. Along the drift direction (z), hits of a smaller amplitude (purple) can be seen preceding the track. The increase in the dQ/dx near the end-point of a stopping muon enhances this effect.

Figure 5.10 shows two events that demonstrate how this effect manifests in the data. Just before the drifting charge arrives at the anode, the leading edge of the integrated-charge signal causes the pixel to trigger, only partially sampling the total charge signal. A moment later, digitization occurs, capturing the remaining charge signal. The result is a track that appears broader in both the spatial dimension and the drift dimension than would be expected by diffusion alone.

This effect is also challenging to accurately simulate and model. A current is induced




Figure 5.11: Calculated pixel response to a track parallel to the anode plane, using the model described in Section 3.3.3. The base model includes no far-field effect. The curves of the far-field model compare the impact of increasing the proportion of the track that is included in the pixel-response calculation (1 pixel = 4.4 mm). Moving from left-to-right, the total drift time is decreased, demonstrating how the far-field effects change the signal shape and reduce the total integrated charge as the charge is deposited closer to the anode.

throughout the drift duration, and so, even though the induced current is small for the bulk of the drift, it integrates and contributes $O(10\%)$ to the total integrated current signal. Figure 5.11 shows the simulated effect of the far-field current varying the size and drift distance of the charge cloud. In addition to the far-field current, the shape of the induced current on the pixel depends on the assumptions of the inter-pixel electrostatics, which are difficult to ascertain. Through the self-trigger threshold, these effects manifest as differences in the trigger timing with respect to the drifting charge, as well as, in the mean number of triggers per track length. Thus, this inductive triggering introduces an threshold-dependent and dQ/dx -dependent difference in the scale and apparent location of the drift charge when reconstructing particle trajectories.

To reduce the impact of this effect, a simple hit-merging algorithm is applied to combine the near- and far-field triggers on single channels. Hits are first time-ordered and sorted by unique channel number, then any hits falling within $< 3 \mu\text{s}$ are selected, where $3 \mu\text{s}$ was chosen to be slightly greater than the minimum re-trigger time. This constrains the maximum change in the drift time for a hit to be $< 5 \text{ mm}$ for a 500 V/cm drift field. The selected hits then are merged by taking a charge-weighted average for the drift time and the charge sum for the merged charge. For sparse, track-like objects using the charge-weighted average improves the drift resolution to $\approx 300 \text{ ns}$ ($\approx 0.5 \text{ mm}$ at 500 V/cm), see Fig. 5.9.

5.2.1.6 Near-linear segment identification

A first-pass track reconstruction was used to provide an estimate of dQ/dx and track position relative to the pixel and was used as an input for the calibrations and simulation tuning. The aim of the first-pass track reconstruction was to provide track segments with very high purity, which can be used to estimate the dQ/dx and track position at a particular point. Subsequent reconstruction then used these track segments as a means to build up more complex, high-level objects.

First, hits were clustered using the *density-based spatial clustering of applications with noise* (DBSCAN) algorithm [216] with a neighbor-distance parameter of 2.5 cm and core-point-neighbor threshold of 5. The core-point-neighbor threshold of 5 hits was selected arbitrarily, but performed reasonably well. The 2.5-cm distance parameter was then determined by generating a 5th-neighbor distance distribution, which showed a prominent knee at ≈ 2.5 cm. This algorithm provides cluster labels for approximately contiguous charge depositions but is insensitive to curvature.

Each DBSCAN cluster was then fed into a RANSAC PCA estimator [217] to identify co-linear sections of the cluster. A residual radius of 8 mm was used to reject hits that showed a prominent deviation from the PCA. The RANSAC algorithm is robust, however it can be overly insensitive to gaps in a track and can erroneously include hits when a cluster contains multiple true particles. To minimize these effects, a second pass of the DBSCAN algorithm was run over the RANSAC co-linear clusters to ensure that the results from the RANSAC algorithm provide contiguous segments.

Once clusters were determined, a trajectory was estimated by an iterative approximation algorithm. First, a start and end point were determined by projecting the hits onto the cluster's primary axis and choosing the hits that lay on the extremes of this projection. Then the trajectory was linearly approximated between these points and a residual was calculated for all points. A new sample point was then found by taking the charge-weighted average of the points within a radius of 2 cm of the hit with the largest residual. The new sample point was added to the trajectory approximation at the position that is closest to its projection on the previously estimated trajectory. The process was repeated up to 16 times, or once all sample points are within 2 cm of each other. This algorithm provides a unique benefit for particle tracking which often contains elastic scatters that cause a relatively large deflection angle at a fixed point.

Due to the disabled channels during operation, reconstructed track segments were often broken at the boundaries of these insensitive regions. To mitigate this, a multi-dimensional likelihood classifier was constructed by selecting a random reconstructed track segment within the event that had a reconstructed track length of > 10 cm. A random translation was then applied to all of the hits in the event such that the selected track was contained in the detector. A mask based on the known disabled channels was then applied to the event, removing hits that fall on disabled pixels. The event was reconstructed a second time using the masked hits. The reconstructed track segments after this round of reconstruction were used to populate a probability distribution of the likelihood for a given pair of




Figure 5.12: Improvement in track reconstruction using likelihood-based track merging. Efficiency is defined as the total fraction of hits produced by the primary particle that are associated to the same reconstructed track. Purity is defined as the fraction of hits in the reconstructed track that are produced by particles other than the primary particle.

reconstructed track segments belonging to the same parent trajectory. In particular, the probability distribution for a pair of tracks is calculated as

$$p(\vec{\theta}|\text{same parent}) = \frac{N_{\text{broken pair}}(\vec{\theta})}{N_{\text{broken pair}}} \quad (5.20)$$

where a broken pair is defined as a pair of newly reconstructed track segments that share an endpoint with the selected track segment, consist of $>80\%$ of the hits belonging to the selected segment, and each segment has one endpoint reconstructed far (> 7.7 cm) from the original endpoints. The variable vector $\vec{\theta}$ represents a collection of discrimination variables that are defined on a pair of track segments, to be discussed later.

A distribution of "background" pairs was also generated using all of the other segment pairs in the event

$$p(\vec{\theta}|\text{different parent}) = \frac{N_{\text{other pair}}(\vec{\theta})}{N_{\text{other pair}}}. \quad (5.21)$$

Thus, a traditional log-likelihood discriminator could be used

$$\mathcal{L}(\vec{\theta}) = \log p(\vec{\theta}|\text{same parent}) - \log p(\vec{\theta}|\text{different parent}). \quad (5.22)$$

A cut on $\mathcal{L}(\vec{\theta})$ was then defined such that the efficiency for identifying broken track segments was $>95\%$ on the training sample.

The five variables of $\vec{\theta}$ used to discriminate broken track segments are the following: $\sin^2 \theta$, with θ being the angle between the track segment pair; d_{trans} , the maximum transverse distance of the second track from the axis of the first track; d_{missing} , the distance of the minimum displacement vector from one segment to the other that crosses enabled channels;

d_{overlap} , the quadrature sum of the 1D overlap in each of x, y, and z coordinates; and $\delta dQ/dx$, the difference in the average ionization along the segments. A 5-D histogram was populated using a typical run and smoothed using a 1-bin width Gaussian kernel.

Figure 5.12 shows the improvement in the track reconstruction efficiency using this algorithm applied to the CORSIKA-simulated sample. A majority of the tracks are reconstructed with greater than 95% of their true hits and are reconstructed with less than 10% contamination from other particles, including the contamination from δ -rays.

5.2.2 Light readout

Generally, the LRS processing consisted of first a noise subtraction algorithm to remove pick-up observed in the LRS waveforms. Then a Weiner filter was applied to remove the electronics response of the signals and produce a waveform with effectively a faster response time and better time localization. Then, the samples from the two digitizers were aligned at the O(ns) level using the global-trigger synchronization signal. And finally, a gain calibration was applied to the signal of each SiPM and the waveforms were summed to estimate the number of photons arriving at each LRS module with O(10 ns) time resolution. This section will describe these processes in more detail.

5.2.2.1 Pickup filter

Due to the location of the E-board and insufficient shielding around the LRS cold pre-amps, significant 10-MHz pickup between the LArPix clock signal and the input to the LRS pre-amp was present throughout the data collection. The ArcLight SiPM signals were more significantly impacted because they were digitized with the VGA in a higher gain setting than was used for the LCM SiPMs. A dedicated 10-MHz filter was applied offline using the pre-trigger window of each waveform. Figure 5.13 demonstrates this procedure on a waveform.

First, the pre-trigger window defined by the first 50 samples (500 ns) was interpolated and then sectioned into 100-ns periods. The periods were averaged and then extrapolated across the remainder of the readout window. This extrapolation was then subtracted removing the contribution from the periodic pickup. This filter has the benefit of phase-matching and includes higher harmonics of a potentially irregular pickup signal while preserving the frequency response within the signal frequency region. In terms of the standard deviation, Fig. 5.14 shows the reduction of noise in the pretrigger window with this filter. For the most heavily-impacted channels, the standard deviation is reduced by a factor of ~ 5 .

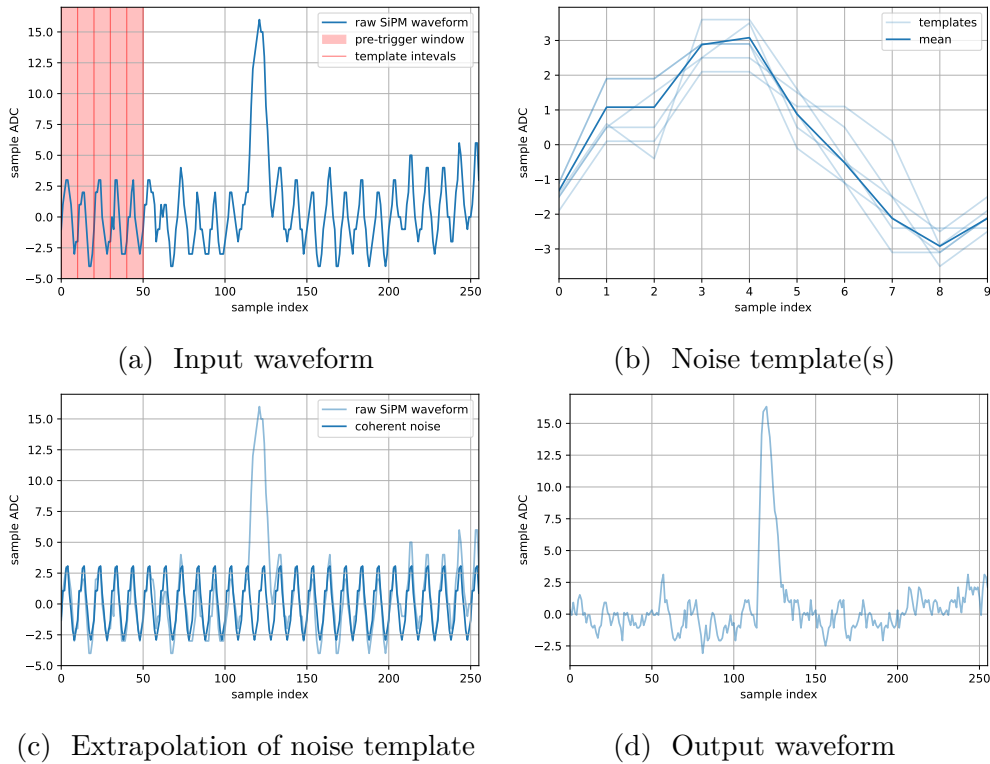


Figure 5.13: Application of the dedicated 10-MHz pickup filter to a heavily impacted waveform. (a) shows the original waveform and the 100 ns periods taken from the pretrigger window, (b) shows the individual periods and their average, (c) demonstrates the extrapolation of the noise template – resulting in the noise-subtracted waveform in (d).




Figure 5.14: Comparison of the standard deviation of pretrigger samples before and after applying 10-MHz pickup filter.

5.2.2.2 Signal deconvolution

The readout electronics introduce signal shaping through the SiPM, the amplifier, and cabling. Within the SiPM, internal resistance and capacitance reduce the signal rise time. Within the many-photon regime, after-pulsing, due to photoelectrons temporarily trapped in lattice imperfections within the SiPM [218], also introduces a slower apparent response time. At the amplifier, the RC-feedback reduces the bandwidth, shaping the signal response. Finally, imperfect impedance matching at either the front-end or along the readout cable can introduce ringing to the output signal.

In order to remove these effects and extract the relevant signal – the number of photoelectrons within a time interval – a Wiener–Kolmogorov filter was employed [219]. For the purposes of reconstructing the light signals, the filter was constructed as

$$W(f) = \frac{I(f)^* |S(f)|^2}{|I(f)|^2 |S(f)|^2 + |N(f)|^2}, \quad (5.23)$$

where $I(f)$ is the Fourier transform of the true impulse, $S(f)$ is the Fourier transform of the true signal, and $N(f)$ is the Fourier transform of the true noise. Because the observed signal is sampled, this filter is translated into a discrete Fourier transform (DFT). In this case, we can apply the same form, replacing the Fourier transforms with their discrete counterparts

$$W_k = \frac{I_k^* |S_k|^2}{|I_k|^2 |S_k|^2 + |N_k|^2}. \quad (5.24)$$

The reconstructed signal is extracted by applying this filter in frequency space

$$d_i = \text{DFT}^{-1} [\text{DFT}[s] * W]_i \quad (5.25)$$

where s_i is the input sampled waveform.

To implement this, a FFT was used. The waveform and impulse functions were padded with zeros to twice the original length in time to avoid cyclic artifacts. The noise frequency response was extracted using an average of approximately 150,000 waveforms with low amplitude (<128 ADC for LCM SiPMs and <500 ADC for ArCLight SiPMs) collected during normal data collection. The impulse response function was extracted by performing a fit of the average waveform with a moderate signal amplitude between 1000 ADC and 5000 ADC to the form of

$$f(t) = \sin(t/p_1) e^{-t/p_2} + p_0 \int_{-\infty}^{+\infty} \sin(t/p_1) e^{-t/p_2 - (\tau-t)/p_3} H(\tau) d\tau \quad (5.26)$$

where p_i are free parameters and $H(\tau)$ is the Heaviside step function. This form is motivated by the solution to an underdamped RLC-circuit response convolved with an infinitesimally short prompt signal and an exponentially decaying slow signal. Parameter p_0 is the relative contribution from the delayed component, p_1 is the oscillation period, p_2 is the RC time




Figure 5.15: Impulse response function extracted from a typical ArcLight (left) and LCM (right) SiPM.

constant, and p_3 is the slow signal time constant. This does not consider the response time of the SiPM or the propagation time of the light incident on the SiPM, which are expected to be small corrections to the signal shape. From this fit, similar results were obtained for both the LCM and ArcLight response functions and the best-fit parameters of

$$\begin{aligned}
 p_0 &= 15 \pm 10 \\
 p_1 &= 559 \pm 12 \text{ ns} \\
 p_2 &= 47.4 \pm 0.6 \text{ ns} \\
 p_3 &= 1050 \pm 70 \text{ ns}
 \end{aligned}$$

were used for all SiPM channels. Figure 5.15 shows the best fit along with the best-fit impulse for two typical channels. A periodic artifact is present in the averaged waveforms due to the 10-MHz noise mentioned in Sec. 5.2.2.

As a true signal function is not known a priori for each waveform, an assumed Gaussian with full-width at half-maximum (FWHM) of 30 ns and a total signal power equal to the average waveform power minus the average noise power was used. This provides reasonably sharp response to the singlet scintillation light while reducing the sensitivity to high frequencies where the filter is less stable.

The performance of this deconvolution was studied using the LED data collected during later operation of the Module 1 detector with a nearly identical optical readout as Module 0. In this study, one of the calibration LEDs was driven by two equal-amplitude pulses with a known time separation. The same waveform filter and deconvolution process were applied to this calibration dataset, and the relative amplitude of the two pulses was measured as a function of the time separation between them. To account for the differences in the noise and response model between Module 0 and Module 1, the double-pulse waveforms from Module 1 were deconvolved using the LED-calibrated response model of Module 1 and then convolved with the Module 0 response model. The difference in the Module 0 and Module 1 noise power spectrum could be directly added to the double-pulse data, because the LRS noise

was substantially larger in Module 0 than in Module 1. Together, this produced waveforms that were representative of Module 0 performance. No discernible change in the ratio of the pulse amplitudes was apparent after the deconvolution procedure across the range of values studied (65 ns to 500 ns). Additionally, the peak-to-valley ratio was consistent with a Gaussian pulse shape of a FWHM ≈ 35 ns, suggesting good separation of light signals with separations of < 100 ns.

5.2.2.3 Waveform alignment

Due to the trigger logic described in Sec. 4.1.4, the samples of the waveform between the two LRS ADCs in each event are not necessarily aligned in time. Further, the time resolution of the clock correction described in Sec. 5.1.2 does not provide sufficient accuracy to match sample-to-sample. Thus an additional algorithm is required to achieve sub-sample timing resolution between the two ADCs.

In addition to the SiPM waveforms, the global trigger signal was also digitized on each event. The leading edge of this pulse was then used to determine the relative alignment of the waveforms on each ADC. For each trigger, an offset was calculated by taking the slope of the rising edge on this channel, defined by the difference in the closest two samples crossing a fixed threshold of 5000 ADC, and projecting back to the 0-threshold crossing. Figure 5.16 demonstrates how this alignment is done with an event. In this example, ADC 1 initiates the trigger due to one of the the sum channels crossing the preset trigger threshold (not shown), then ADC 0 is triggered by the subsequent global trigger pulse, incurring a few hundred nanosecond delay. The rising edge of the global trigger pulse then determines a fixed point in time that is used as a reference to align the two triggers. The time synchronization offset described in Sec. 5.1.2 suggests that the cable length delays are ≈ 2 ns and thus negligible for this analysis.




Figure 5.16: Global trigger alignment of a single event in which ADC1 self-triggers, and ADC0 is triggered by the subsequent global trigger signal.

5.2.2.4 Gain calibration

Individual SiPM gains were calibrated with dedicated runs using the calibration LED integrated into each anode. To calibrate each channel, a short integration window within ~ 100 ns around the LED pulse (Q_{short}) was used to resolve the SPE and noise peaks. The calibration relating Q_{short} to 1 PE was then calculated as

$$C_{\text{short}} = \frac{1 \text{ PE}}{Q_{\text{short}}(\text{SPE}) - Q_{\text{short}}(0)} \quad (5.27)$$

where $Q_{\text{short}}(\text{SPE})$ and $Q_{\text{short}}(0)$ represent the values extracted from a multi-Gaussian fit to the Q_{short} distribution. Ultimately, a calibration for the full waveform integral was desired, so a secondary factor was extracted

$$f_{\text{long}} = \frac{\langle Q_{\text{short}} \rangle - Q_{\text{short}}(0)}{\langle Q_{\text{long}} \rangle - Q_{\text{long}}(0)} \quad (5.28)$$

where $\langle \rangle$ represents the mean and Q_{long} is the waveform integral from a fixed time-offset (listed in Table 5.2) to +500 ns. $Q_{\text{long}}(0)$ was determined by using the noise filter described in Sec. 5.2.2 to extrapolate the noise and pedestal in the pre-trigger window across the waveform. The calibration was then applied to each channels' waveforms according to

$$S_{\text{PE}} = f_{\text{long}} C_{\text{short}} S_{\text{ADC}}, \quad (5.29)$$

where S_{PE} is the response in PE per time tick and S_{ADC} is the ADC value. The LEDs were driven with an 800-Hz pulse and tuned to produce a clear SPE peak for the channels being calibrated, as shown in Fig. 5.15. A run of length ≈ 10 s was sufficient to produce $O(10^4)$ triggers, giving an expected statistical uncertainty of $\approx 2\%$. This uncertainty was confirmed using the bootstrap technique [220], yielding $<3\%$ on all channels.




Figure 5.17: Low-amplitude LED signal distribution with short (left) and long (right) integration time used to extract the individual SiPM gains.

Table 5.2: Gain calibration windows used for Module 0 LRS SPE calibration.

ADC	Channels	Offset [ns]	Q_{short} window [ns]	Q_{long} window [ns]
0	2-7	1210	+160	+500
0	9-24	1200	+100	+500
0	18-23	1210	+160	+500
0	25-30	1200	+100	+500
0	41-46	1200	+100	+500
0	50-55	1210	+160	+500
0	57-62	1200	+100	+500
1	2-7	1210	+90	+500
1	9-24	1210	+90	+500
1	18-23	1220	+160	+500
1	25-30	1210	+90	+500
1	34-39	1210	+160	+500
1	41-46	1210	+90	+500
1	50-55	1210	+160	+500
1	57-62	1200	+100	+500

This calibration was applied to tune the light simulation, discussed in Sec. 4.4. The simulated gain values were defined relative to the sample rate $\Delta t_{\text{sample}} = 10$ ns as

$$G_{\text{LRS}} = \frac{\Delta t_{\text{sample}}}{f_{\text{long}} C_{\text{short}}} \quad (5.30)$$

giving a gain value with units of ADC / (PE/ μ s). The distribution of the measured gain values is shown in Fig. 4.11.

5.3 Muon reconstruction

As discussed in Ch. 2, the primary particle used in this analysis are naturally-occurring cosmic rays. While the majority of these particles are muons, most are energetic enough to pass through the detector without stopping. Thus it is important that these through-going muons are identified and removed from the sample. In addition, cosmic rays also contain a hadronic component, consisting of primarily protons and secondary neutrons due to their longer lifetime, but also some mesons. These form a background for the analysis of stopping muons, but can be identified by their energy loss as well as their trajectory and event topology.

This section describes the reconstruction algorithm used to identify the stopping muons, along with reconstructing their vertex and trajectory. This involves a tracking algorithm

to estimate the particle trajectory, under a stopping-muon hypothesis. From this, a dQ/dx profile of the event was constructed and was used to reject backgrounds, discussed later in Ch. 6. And finally, information from the dQ/dx profile and the trajectory was used to reconstruct the stopping point within the detector.

5.3.1 Stopping-muon trajectory reconstruction

To reconstruct the trajectory and use the stopping particle's Bragg peak [221] to identify the vertex, a greedy tracking algorithm was used. First, seed points are identified by selecting track segments that originate from the exterior of the detector and cross the fiducial boundary. Only downward-going tracklets were used to avoid backgrounds from neutron events that produce a back-scattered secondary particle. Each seed point was generated by taking the charge-weighted average of hits outside of the fiducial boundary for the given tracklet. Once a list of seed points was generated, the PCA-calculated tracklet direction was used for the initial estimate of the particle trajectory. Trajectories were built step-by-step by collecting hits in the forward region relative to the current estimated position and direction. Then the current position \vec{p}_j and unit direction \vec{d}_j were then updated and the algorithm was iterated.

Figure 5.18 illustrates this procedure. Hits were collected on each step in a forward




Figure 5.18: 2D diagram illustrating the greedy tracking algorithm used to reconstruct muon decays. From the current trajectory step j , a cylindrical region (yellow) defined by the two reconstruction parameters Δ and α is projected forward along the current trajectory direction \vec{d}_j . Hits that fall within this region (red) are used to calculate the next trajectory step.

cylinder based on their position \vec{r}_i , where

$$\delta\vec{r}_i = \vec{r}_i - \vec{p}_{j-1} \quad (5.31)$$

$$d_{L,i} = \delta\vec{r}_i \cdot \vec{d}_{j-1} \quad (5.32)$$

$$d_{T,i} = |\delta\vec{r}_i - d_{L,i}\vec{d}_{j-1}| \quad (5.33)$$

and a hit i was added to the trajectory step j if $0 < d_{L,i} < \Delta$ and $d_{T,i} < \alpha\Delta$. The selection limits Δ and α were manually adjusted to balance the tracking sensitivity and jitter. For the 4.4-mm pitch and thresholds used in Module 0 data collection, $\Delta = 22$ mm (approx. 5 pixels) and $\alpha = 1/2$ provided reasonable results. The new trajectory position was calculated using the charge-weighted mean position of the associated hits, and the new trajectory direction was calculated using the difference $\vec{p}_j - \vec{p}_{j-1}$. This process was repeated until no hits could be found in the next iteration.

To avoid breaking particle trajectories that cross the disabled regions of the detector, the termination condition of the algorithm was extended if the final step fell into a disabled region of the detector. In this case, a larger search window was used $\Delta \rightarrow \beta\Delta$, $\alpha \rightarrow \gamma\alpha$ (values of $\beta = 2$ and $\gamma = 3/4$ were used for this analysis). The current position was increased by steps of $\Delta \times \vec{d}_j$ until the current position fell onto an active region of the detector or exited the active volume. If hits were found at any point of this process, the algorithm resumed the normal step-wise process.

To search for a track potentially produced by muon decay (caused by the presence of a decay electron), an iterative spherical search was used to initiate the track trajectory, starting from the last trajectory position. Hits were added to the next step if $|\delta\vec{r}_i| < k\Delta$, where k was incremented up to a maximum value. Based on the radiation length in argon, a maximum radius of 110 mm was used. If hits are found in this search, the tracking algorithm resumes.

Figures 5.19 and 5.20, show an example candidate stopping-muon trajectory reconstructed using this approach.




Figure 5.19: Stopping-muon candidate, reconstructed using the tracking algorithm described in the text. The lower images show a zoomed perspective of the same event near the stopping point. A significant gap in the data (≈ 5 cm) is successfully traversed by the tracking algorithm. The reconstructed stopping point using the dQ/dx profile (Sec. 5.3.2) is highlighted with a star.




Figure 5.20: View of the same event as in Fig. 5.19, zoomed into the region near the stopping point.

5.3.2 Stopping-muon profiled dQ/dx reconstruction

As a muon comes to rest, ionization increases, resulting in a rise in energy deposition per unit length (dE/dx) as a function of the position. The dE/dx also varies with the particle mass, so the dE/dx can be used in two ways: first, to reject non-muon backgrounds; and second, to identify and reconstruct the stopping point of the muon within the detector.

Within this analysis, the related quantity, dQ/dx , is used for the identification of the muons that stop in the detector and the reconstruction of their stopping point. The particle dE/dx is related to the dQ/dx via a recombination factor \mathcal{R} and the average energy required to ionize an electron W_e , described earlier in Sec. 3.3.1. However, additional detector effects, such as trigger deadtime and induction current, can cause additional subtleties in the reconstruction of the dE/dx from the dQ/dx . As the aim was to identify the stopping point and separate events, the dQ/dx was sufficient without a detailed reconstruction of the dE/dx .

To calculate the dQ/dx , both the trajectory displacement and the charge was reconstructed for the particle trajectory described in the previous section at each of the trajectory steps. Each hit defined by the hit-merging algorithm was associated to the nearest trajectory step. The charge sum of all associated hits was taken as the dQ for the trajectory step. The




Figure 5.21: Diagram of dx used for the stopping-muon dQ/dx calculation at each trajectory point. The outer-most box (dashed grey) indicates the boundary for the trajectory point that is set by the hits associated with the step (light blue). The region highlighted in red indicates a portion of the track that crossed a disabled pixel and thus was not included in the dx . The interpolated length of the trajectory that is used for the dx is highlighted in dark blue.

step displacement dx was then defined as the overlap of a cuboid containing the associated hits with the linearly-interpolated trajectory. The cuboid volume at each step extended out to $\pm 1/2$ of the pixel pitch beyond the center of the hit pixels furthest from the trajectory point, accounting for the finite size of the pixels. The contribution of the disabled channels that overlap with the linear interpolation was subtracted from the dx estimate. Figure 5.21 illustrates this calculation with the region highlighted red crossing a disabled channel. The dQ/dx was then calculated by taking the ratio of the charge sum to the interpolated length with the disabled-channel correction. It is important to note that the segment dx in this case is not precisely fixed, but varies slightly depending on the number of hits included in each trajectory point ($\sigma_{dx} \approx 25\%$). Figure 5.22 shows the reconstructed dQ/dx as a function of the reconstructed muon range for muon-decay events. The range extends to negative numbers to accommodate reconstruction confusion between the muon and decay electron trajectories. In general, fairly good agreement is observed between the simulation and the data in the medium-threshold sample ($\approx 2\%$ overall scale, $\approx 1\%$ variation across muon range), but some threshold-associated effects are not reproduced near the stopping point ($\approx 5\%$ effect).

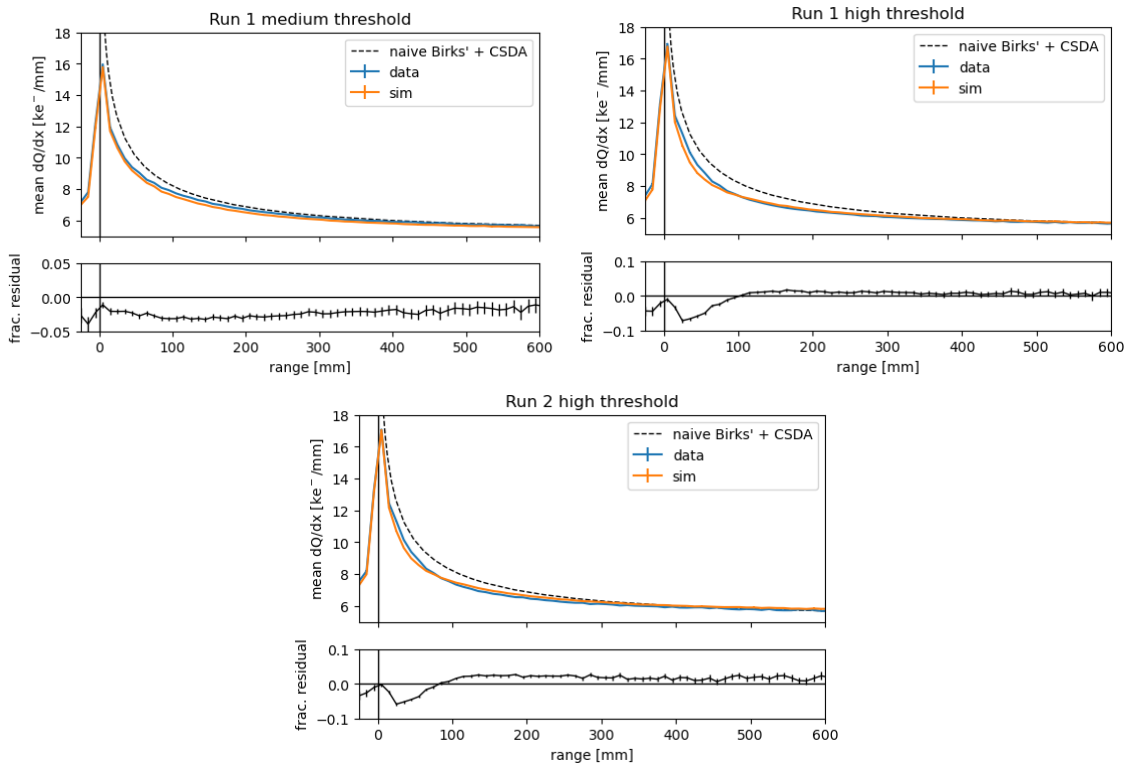


Figure 5.22: Comparison of dQ/dx as a function of the reconstructed range for muon decay events. For reference, the naive Birks' CSDA model for dQ/dx is also shown, which was used for separation of through-going and stopping muons.

After the dQ/dx was calculated, a log-likelihood score could be calculated using two proposed stopping points. The positions of the hits with the largest charge deposition associated with the trajectory points with the largest dQ/dx were selected as the proposed stopping points for each trajectory. A log-likelihood score was calculated using an L1-norm deviation of the mean dQ/dx from naive expectation using the CSDA muon range in liquid argon [212] and the Birks' recombination model. An L1-norm was used instead of a log-Landau distribution to accommodate enhanced low- dQ/dx fluctuations caused by self-trigger threshold effects and digital-analog crosstalk. The log-likelihood scores were then compared and the proposed stopping point that produced the largest log likelihood score was used as the reconstructed stopping point. Likelihood scores were also calculated under the hypothesis of a stopping proton and a MIP-like particle (i.e. no change in mean dQ/dx) and were used to remove background events from the sample. Figures 5.23 and 5.24 show the pre-cut MIP-like score and proton-like score distributions for the Module 0 data with comparison to the simulation, using the relative background contributions extracted from the fit in Ch. 6. Unfortunately, the dQ/dx only offers a modest degree of separation between through-going muons and the signal classes, due to the presence of δ -rays and dE/dx fluctuations along the through-going muon trajectory. These introduce peaks in the dQ/dx along the trajectory that can confuse the reconstruction algorithm, reducing the efficacy of this discriminator.

This reconstruction algorithm provides a reasonable stopping-point resolution with a mean absolute deviation of 8.07 ± 0.10 mm, as shown in Fig. 5.25d. The resolution is dominated by the dQ/dx fluctuations near the end of the muon track that lead to a tail, biased towards shorter residual ranges. This is illustrated by the tail shown in Fig. 5.25b, as most muons are oriented in the $-\hat{y}$ direction. A small bias is present in the drift direction (\hat{z} , Fig. 5.25c) due to the details of when the timestamp is latched relative to the self-trigger.

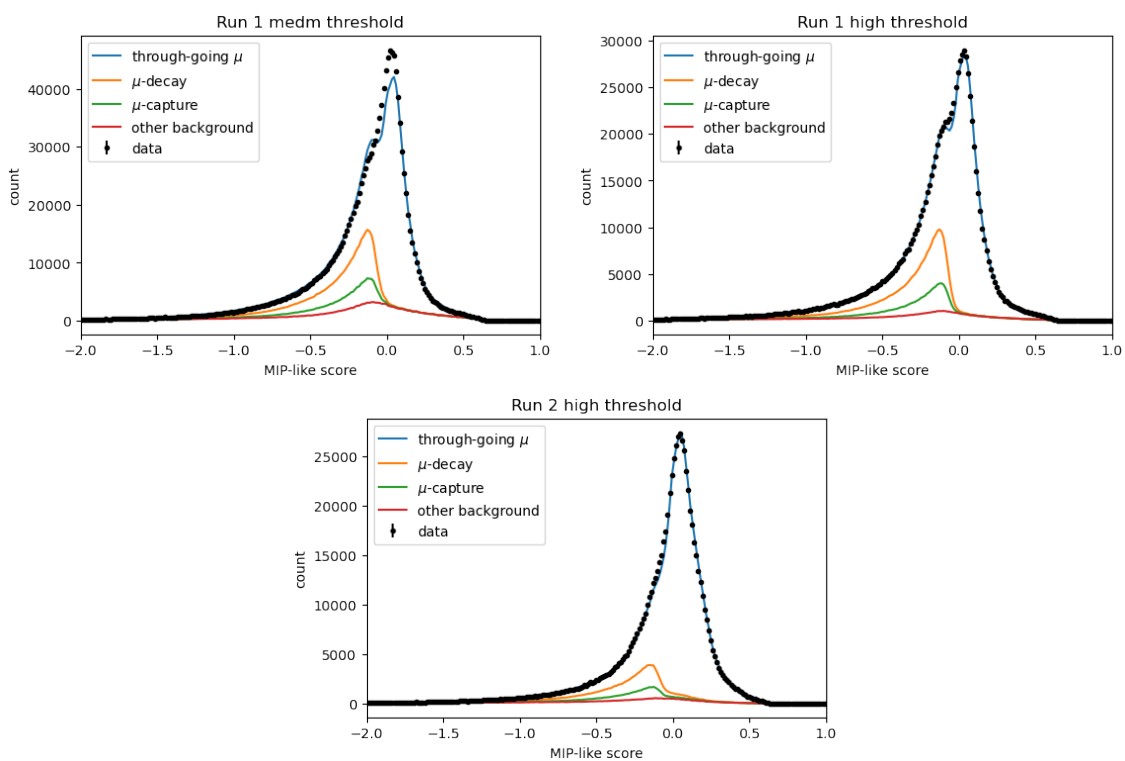


Figure 5.23: MIP-like classifier score using post-fit background contributions extrapolated from the stopping muon selection to pre-cut selection (described in Sec. 6.1.3).

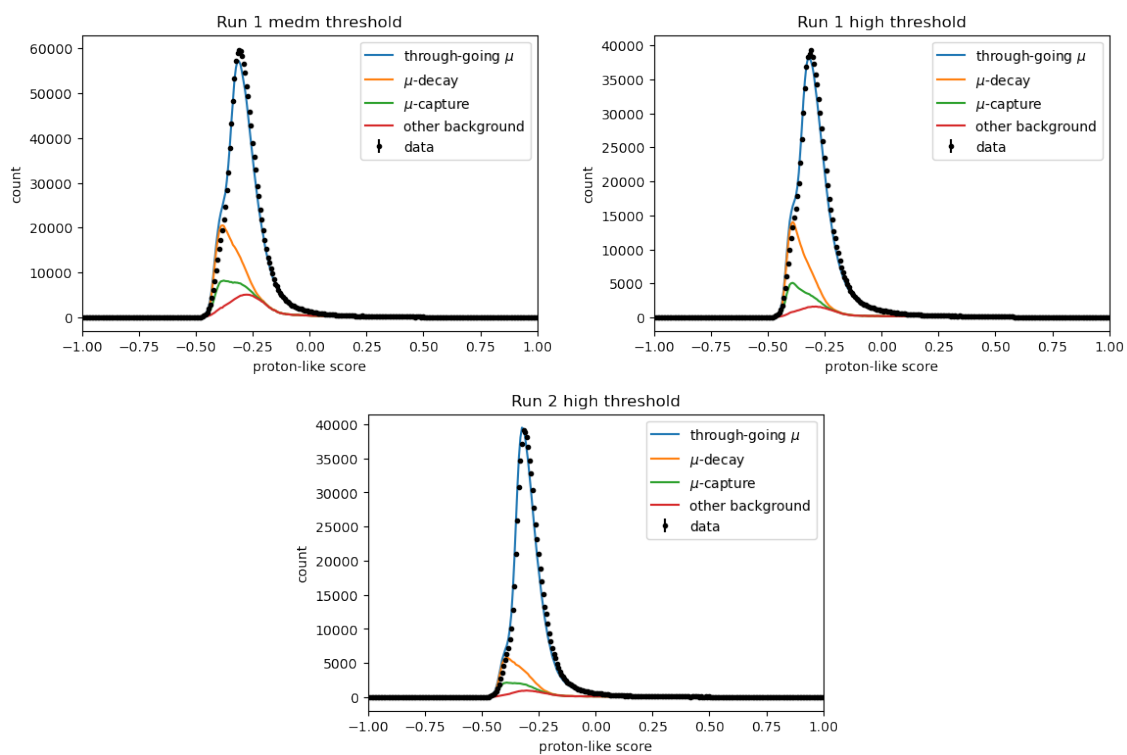


Figure 5.24: Proton-like classifier score using post-fit background contributions extrapolated from the stopping-muon selection to the pre-cut selection (described in Sec. 6.1.3).

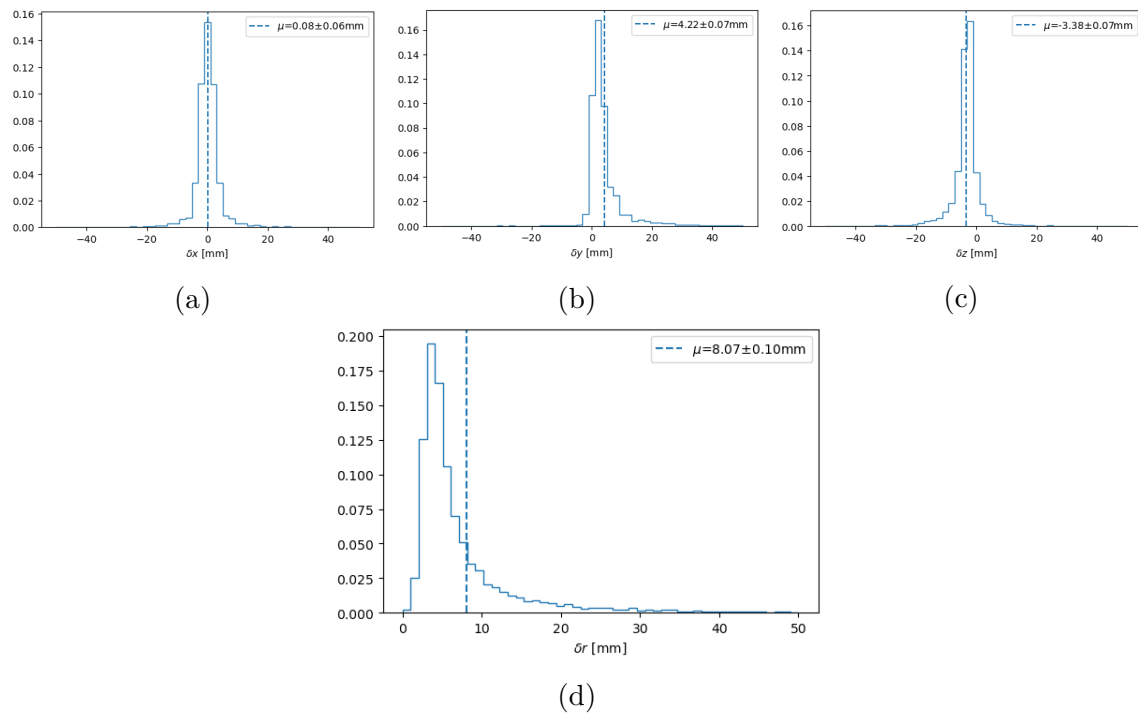


Figure 5.25: Residual vertex reconstruction error from stopping point reconstruction in each dimension (a-c) and the 3D residual (d). The residual in z is defined such that $\delta z < 0$ is biased closer to the anode and $\delta z > 0$ is biased closer to the cathode.

5.4 Muon decay reconstruction

As discussed in Ch. 2, a stopping muon within Module 0 can undergo one of two processes. It can be captured on the argon nucleus, typically releasing a ~ 100 -MeV neutrino and a ~ 6 -MeV neutron, neither of which deposits an appreciable amount of energy in the detector before escaping. Or, it can decay, releasing a neutrino and an anti-neutrino along with an $O(10 \text{ MeV})$ electron. Thus the two classes of events can be separated by identifying the electron present in the muon decay. This section describes the two primary techniques used to reconstruct the muon decay via the Michel electron. First, the identification of the presence of a decay electron using the high spatial resolution of the charge information provided by the LArTPC; and second, the reconstruction of the decay time of the muon using the high temporal resolution of the scintillation light.

5.4.1 Muon decay classification

For the stopping-muon final-state classification, the aim is to produce a metric that can be used to separate capture events from decay events. For this, the analysis uses a likelihood-ratio classifier on each charge hit. Three parameters were calculated for each hit – the angle between the hit and the incoming-muon axis θ_μ , the angle between the hit and the apparent decay axis θ_d , and the distance of the hit from the stopping point d . The muon axis was simply determined by using the track segment immediately nearest to the stopping point of the muon trajectory. This accounts for multiple scattering. The decay angle was determined by searching for trajectory points that extend beyond the end of the stopping muon, i.e. having a residual range less than zero. The decay axis was determined by taking the difference between the stopping point and the first trajectory point with a residual range greater than zero, but contiguous with the muon trajectory. The classifier was "trained" on the dedicated stopping-muon simulation samples, which include only decay and capture events. A 3D histogram of θ_μ , θ_d , and d was filled for the two classification categories (muon decay or capture). The difference in the log of the normalized histograms produces a log-likelihood score for each hit, as shown in Fig. 5.26. Because normalized histograms are used, this classifier is not biased by the decay-to-capture ratio of the training sample. A classifier cut was placed at 0, where a hit is equally likely to belong to the capture or the decay samples. Hits associated with the trajectory points with a residual range greater than zero, i.e. along the muon trajectory, were excluded from both the training and the classification.

Figure 5.26 shows how the classifier boundary evolves as a function of the distance from the vertex. Within 1.3 cm of the vertex, or approximately 3 pixels, no hits are classified as decay hits. This is primarily due to the enhanced contribution of the induction signal triggering on pixels near a point with large energy deposition and due to the stopping-point resolution. Within about 4.3 cm of the vertex, an exclusion region along the muon track is introduced, particularly when the muon and the decay axis are near-parallel or near-antiparallel. Beyond this, the classifier functions effectively as a cone-like cut, with an

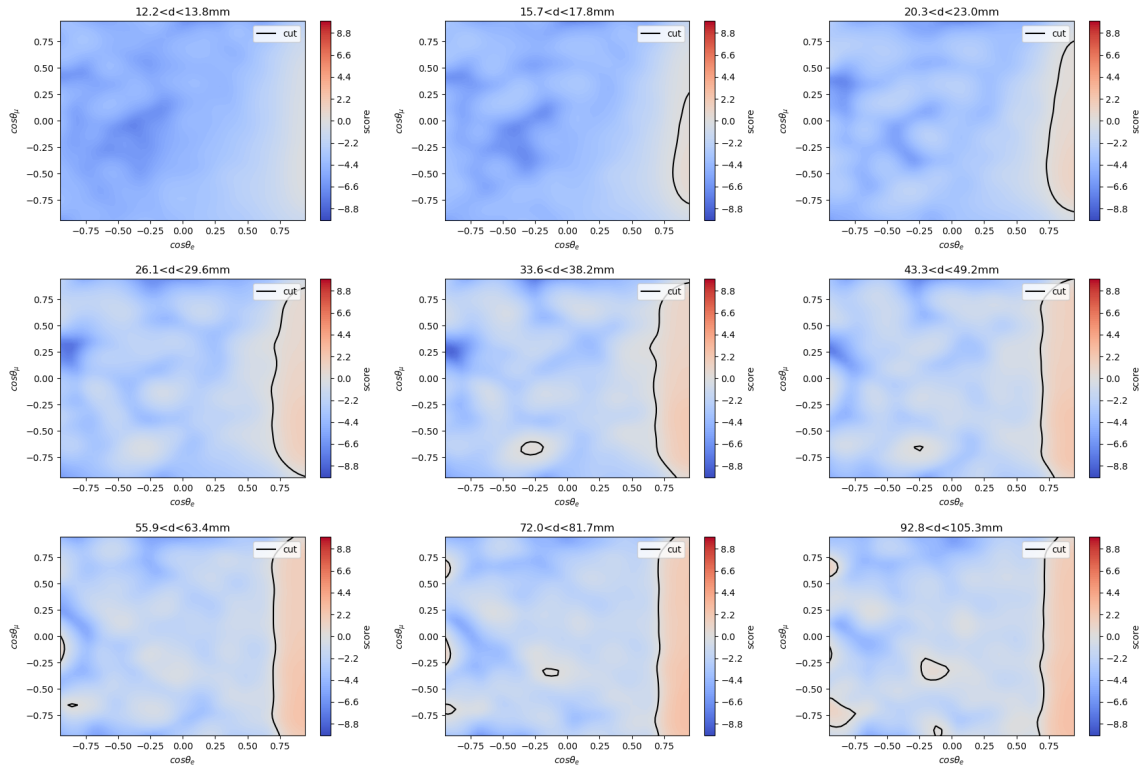


Figure 5.26: Michel likelihood score as a function of the three input parameters: x -axis is $\cos \theta$ with respect to the electron axis; y -axis is $\cos \theta$ with respect to the muon axis; and the plot grid shows the distance from the reconstructed stopping point, increasing from left-to-right, top-to-bottom.

opening angle of $\approx 45^\circ$ when the muon and decay trajectories are perpendicular. The opening angle decreases slightly when the muon and decay axes are near-parallel or near-antiparallel.

A second classifier was also constructed to discriminate decay-like topologies and background topologies. The same algorithm was used, except without the decay axis, resulting in a 2D likelihood score in the $(\cos \theta_\mu, d)$ -space, shown in Fig. 5.28. For this classifier, a training sample of CORSIKA events was used to generate the PDFs, where both capture and decay hits are used to fill the signal PDF. The behavior of the cut boundary is somewhat similar to the decay classifier – hits within ≈ 1 cm of the vertex are always classified as background hits – but it is, by definition, decay-axis agnostic. Any hits beyond ≈ 40 cm from the stopping point are classified as background hits.




Figure 5.27: Michel likelihood score as a function of the parallel and transverse directions to the muon axis. Plots from left to right show the change in the discriminator behavior as a function of the reconstructed electron direction.

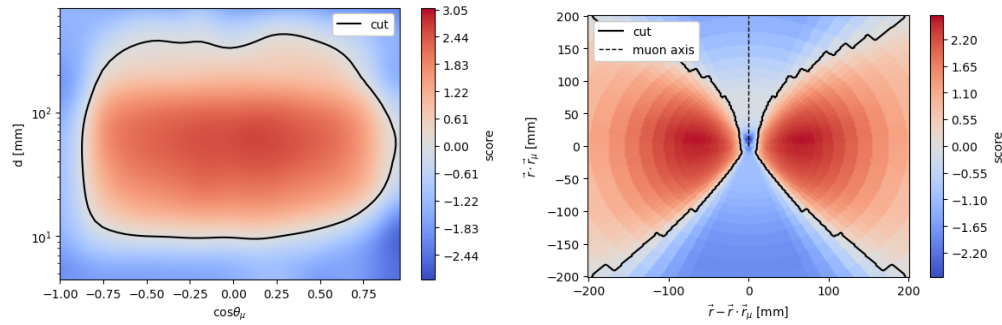


Figure 5.28: Background likelihood score as a function of the two input parameters (left) and the parallel and transverse directions to the muon axis (right). Note that the background likelihood is defined such that a decreasing score indicates an increasing likelihood of being a background-induced hit.

5.4.2 Decay time reconstruction

Due to the presence of the triplet-decay scintillation light, the light signals of the stopping muon and the subsequent decay overlap. Using coarse values of 1% for the PDE, 25% for the photo-detector coverage, 10^4 ph/MeV for the light yield, and 0.3 for the singlet light fraction, the triplet light signal of a 100-MeV muon at the timescale relevant for muon capture in argon (≈ 500 ns) is approximately 1 PE/ns. Thus the fluctuations in this triplet light within an integration period of 50 ns are of the order $\sqrt{N_{\text{PE}}} \approx 7$ PE. These triplet-light fluctuations form an irreducible background for the reconstruction of the muon decay time for events where the decay electron produces a singlet light signal of less than 7 PE. This produces an intrinsic energy threshold, where the signal-to-noise ratio (SNR) ~ 1 , of approximately 1 MeV. Similarly, the energy resolution of the LRS is dominated by the triplet light of the decay electrons of less than approximately 10 MeV (corresponding to a singlet light yield of ≈ 75 PE). Thus the challenge of the decay-time reconstruction is separating the singlet signal from the decay electron and the triplet signal of the stopping muon, which are of

similar magnitudes at low energies of the Michel electron.

To apply a decay-energy threshold that is unbiased across the decay-time spectrum, a subtraction is needed to remove the prompt-light signal that is produced by the stopping muon. A PCA technique was used to remove the known, sample-to-sample correlations from the waveform, eliminating prompt-correlated bias. In this analysis, the total light signal as a function of time was treated as a vector \vec{w}_j on each light detector j , where the vector index iterates over the sample number. This waveform vector was then projected onto the primary PCA components

$$\vec{r}_{0,j} = \vec{w}_j \quad (5.34)$$

$$\vec{r}_{i,j} = \vec{r}_{i-1,j} - (\vec{r}_{i-1,j} \cdot \vec{T}_{i-1,j}) \vec{T}_{i-1,j}, \quad (5.35)$$

where $\vec{r}_{i,j}$ is the residual vector and $\vec{T}_{i,j}$ is the unit-normalized i -th PCA component. The PCA components are derived by performing a PCA decomposition on a large number ($\sim 10^6$) of cosmic-ray waveforms. After the subtraction, the residual waveform contains only the contribution of the delayed-light signal (i.e. the decay electron) and the residual error from the muon's prompt-light signal, dominated by statistical fluctuations in the triplet light. Figure 5.29 shows a comparison of the templates for 3 of the light detectors. Since the light signal time profiles differ slightly between data and simulation, a separate set of templates was derived for the simulation samples using the same averaging routine. To account for time-walk in the trigger timing, the projection was repeated, shifting the template by up to ± 5 sample ticks. The projection that produces the smallest absolute residual $|\vec{r}_{i,j}|$ was used. Figure 5.30 shows the light signal from a muon-decay candidate before and after the prompt-light subtraction has been applied.




Figure 5.29: PCA-derived templates used to subtract the prompt-light contribution from the waveform. An oscillation is present immediately after the prompt signal in the 0th PCA component due to an imperfect impulse model used for the Wiener deconvolution. The shape of the 0th component of the PCA decomposition is very close to the mean signal shape, the shape of the 1st component is dominated by leading-edge differences between waveforms that result in a horizontal shift with respect to the trigger edge, and the 2nd component starts to account for differences in the singlet-to-triplet ratio.

To identify the presence of a delayed-light signal from the decay electron and determine the muon decay time, the prompt-light-subtracted waveforms were aligned and linearly interpolated to 10-ns sample intervals. Two significance metrics were then calculated on each waveform

$$\vec{s}_{\text{prompt}} = \sum_j \frac{w_j^2}{\sigma_j^2} \quad (5.36)$$

$$\vec{s}_{\text{delayed}} = \sum_j \frac{\bar{r}_{2,j}^2}{\sigma_j^2} \quad (5.37)$$

where j is a sum over the light detectors and σ_j is the standard deviation of the PCA residuals determined by the PCA-decomposition. A short gate from -400 ns to -200 ns relative to the global-trigger edge was used to search for the prompt-light signal and extract the prompt-light time. The prompt-light signal time was estimated using the most significant sample ($\text{argmax}(\vec{s}_{\text{prompt}})$) within this gate. The prompt-light signal amplitude was estimated by using the integral of the waveforms over the 50-ns (5 samples) interval centered around this sample and summed across all channels. After the prompt-light signal was identified, a delayed-light-signal gate from +300 ns to +20,000 ns was used to search for the delayed decay-electron signal. The delayed-light signal time was similarly estimated using the most significant sample ($\text{argmax}(\vec{s}_{\text{delayed}})$) within the delayed gate. Likewise, the delayed-light amplitude was calculated from the integral of the residual waveform within the 50-ns window around the delayed-light time. The time difference of the prompt-light time and the delayed-light time was used as the decay time. No further interpolation was deemed necessary as the 6-ns resolution provided by the time separation between the two peaks is more than sufficient to resolve the muon disappearance time of ≈ 600 ns.




Figure 5.30: Example of PCA prompt-signal subtraction on a candidate muon-decay event. The left-most frame shows the summed waveform of the SiPMs on each detector (top) and the overall signal sum (bottom). The middle frame shows the PCA fit to the prompt signal. And the right-most frame shows the residual waveform after the prompt signal has been subtracted.

For the data collected with a short-readout window, deadtime between two triggers introduces a background of muon-decay events with no light signal present and thus no reliable decay-time reconstruction. To remove these events, a light yield was estimated using the combined CRS and LRS information. As the light response of each LRS detector module depends on the solid angle of the module relative to the energy deposition, the total light yield of an event can be estimated using the coverage provided by the j -th LRS module, predicted by the CRS

$$\alpha_j = \sum_i \frac{q_i \Omega_j(\vec{x}_i)}{4\pi Q}, \quad (5.38)$$

where $\Omega_j(\vec{x})$ is the solid angle of the rectangular area of the light detector module to a point in the active volume, q_i and \vec{x}_i are the charge and reconstructed position of the CRS hits, respectively, and Q is the total charge ($\sum_i q_i$). The detector acceptance takes on a value between 0 and 0.5, corresponding to when the charge is distributed far from the light module and close to the light module, respectively. The estimated total light yield L of each event can then be calculated as

$$L = \frac{\sum_j l_j}{\sum_j \alpha_j}, \quad (5.39)$$

where l_j is the 50-ns light integral on a given detector j . Events with abnormally-low light yields are removed with high efficiency using the metric, substantially reducing the background from muon decays that occur during dead time.

Chapter 6

Capture rate measurement

As described in Chapter 2, Module 0 has sensitivity to the muon capture rate through two means: the disappearance rate of negative muons as measured by the time spectrum of muon decay, and the branching ratio of the capture process relative to the decay process. These quantities can be related to the capture rate λ_c via

$$\lambda_d = \lambda_c + R_H \lambda_0 \quad (6.1)$$

$$B = \lambda_c / \lambda_d \quad (6.2)$$

where R_H is the Huff factor described in Sec. 1.2.6, λ_0 is the muon decay rate in vacuum, and λ_d is the total disappearance rate for a stopped muon. The total disappearance rate can be determined by observing the evolution of the number of stopped muons in the detector as a function of time. The branching ratio B can be determined by assuming¹ that each negative muon has only two possible disappearance modes

$$B \equiv \frac{N_{p\mu^- \rightarrow n\nu_\mu X}}{N_{\mu^-}} \equiv 1 - \frac{N_{\mu^- \rightarrow e^- \nu_\mu \bar{\nu}_e}}{N_{\mu^-}} \quad (6.3)$$

where $N_{p\mu^- \rightarrow n\nu_\mu X}$ is the number of capture events, $N_{\mu^- \rightarrow e^- \nu_\mu \bar{\nu}_e}$ is the number of μ^- decays, and N_{μ^-} is the number of negative muons that stop in the detector.

This chapter will describe how the muon capture rate is extracted from the data through these quantities, namely by performing a maximum-likelihood fit to the distributions of the muon-decay nhit and the muon-decay time, both of which were described previously in Section 5.4. The muon-decay nhit indicates the presence (or lack thereof) of a decay electron, and so it can be used to extract the branching ratio B . Complicating the analysis, however, is that μ^+ and μ^- decays are nearly degenerate in this quantity and are challenging to separate in a non-magnetized detector. Thus an additional handle is needed to determine

¹This neglects known radiative decay modes, $\mu^- \rightarrow e^- + \nu_\mu + \bar{\nu}_e + e^+ + e^-$ and $\mu^- \rightarrow e^- + \nu_\mu + \bar{\nu}_e + \gamma$. The total branching fractions for these additional decays have been measured by other experiments [39, 40] to be less than 10^{-4} . This is far below the sensitivity of this analysis, and so the assumption of either a decay or capture final-state is reasonable.

the number μ^- in the sample. In this analysis, the complementarity of the muon-decay time distribution and the muon-decay nhit distribution is exploited to break this degeneracy. Here, a cut removing the muon-capture events from the muon-decay events is applied, and the muon-decay time distribution for this enhanced muon-decay sample is included in the maximum-likelihood fit. As a consequence, three independent parameters can be extracted from the maximum-likelihood fit: from the relative number of decays at short and long muon-decay time, the cosmic-ray muon charge ratio R ; from the slope of the muon-decay time distribution at short times, the muon disappearance rate λ_d ; and from the relative number of events at high and low muon-decay nhit, the Huff factor R_H . The muon capture rate λ_c can thus be deduced independently of R_H or λ_d .

First, Section 6.1 describes the high-efficiency, low-bias sample of muons used to populate the muon-decay nhit distribution. The high-purity sub-sample of muon decays used for the muon-decay time distribution is also described in this section. Next, Section 6.2 describes the maximum-likelihood fit and the nuisance parameters that determine the shape and amplitude of each signal and background component. A detailed description of the mechanisms that introduce backgrounds into the analysis is also included here. This section concludes with a description of how the fit errors are extracted with the bootstrap technique [220], along with a validation using a toy MC. In Section 6.3, the estimation and impact of the systematic uncertainties included in this analysis are described and summarized. Finally, the results of the fit are presented in Section 6.4, focusing on the parameters of interest: R , λ_d , R_H , and λ_c .

6.1 Event selection

In this analysis, there are two selections to remove non-muon and non-stopping muon events from the sample. The first selection focuses on providing a sample of moderate purity, but low efficiency bias between muon capture and muon decay events. This selection is used to fit the branching ratio of the decay and capture final states for negative stopping muons. The second selection is a subset of the first, targeting a sample of decay events with high purity, but low bias between positive and negative muon decays. This sample is used to determine the charge ratio of the stopping muons of the first sample, as well as, the total disappearance rate for negative muons. This section will describe the cuts used to select the events that go into these two samples. The same selection cuts are applied to the data as are the CORSIKA simulation samples.

6.1.1 Stopping muons

6.1.1.1 Implicit cuts

Prior to considering the direct event selection, there are indirect selections that are applied during the event reconstruction processing. First, the event definition described in Sec. 5.1 introduces a threshold on the activity within the detector of > 5 hits per 100 μs . This

efficiently identifies periods of time where there is activity in the detector, without incurring substantial processing overhead. However, it also rejects events with very low energies. For cosmic-ray muon events, this corresponds to an energy threshold on the entering muon of roughly > 15 MeV. Because this threshold is placed on the total activity in the detector rather than specifically targeting the muon, there is a potential to introduce a bias in the selection efficiency between captures and decays. Thus this threshold is set as low as reasonably possible without undue burden on the offline reconstruction. At the event threshold of 5 hits per 100 μs , the difference in event-building efficiency is less than 1%.

An additional implicit cut is present in the reconstruction of the stopping-muon trajectory and dQ/dx profile. In this reconstruction step, a seed point is selected for each event that is passed to the later stages of the reconstruction. Candidate seed points are selected based on the preliminary collinear-segment reconstruction. Two fiducial volumes are defined: an outer fiducial volume that encompasses the outer-most 2.2 cm of the module, and an inner fiducial volume that contains the remaining 76.6% of the detector's active volume. From the list of all collinear segments, seed points are generated using the start-point of segments originating in the outer fiducial volume with a thickness of 2.2 cm, directed downward, and that stop within the inner fiducial volume. Because the collinear segment tracking can mis-reconstruct the decay electron as part of the muon track, a fraction of decay events would be removed by requiring that the end-point of the segment ends in the inner fiducial volume. Thus, a conditional cut is placed on segments that start and end in the outer fiducial volume to improve the efficiency for these events. Specifically, track segments that start and end in the outer fiducial volume, but have a deflection of $\cos\theta < 0.75$ between any two consecutive points along the trajectory are allowed to be used for seed points. The deflection $\cos\theta$ is defined by the inner product of the trajectory direction \hat{p}_j defined in Sec. 5.3.1. This enables decay events that were reconstructed as a single trajectory to be included based on the enhanced multiple Coloumb scattering of the stopping muon or the non-negligible decay angle between the muon and the electron. A secondary veto region within 2.2 cm from the cathode rejects seed points that originate from segments that are mis-reconstructed when crossing the dead region of the cathode. Each seed point that passes this selection is reconstructed under a stopping-muon hypothesis. The seed point that produces the best dQ/dx profile score for a stopping muon is used. Overall, the selection of seed point is found to introduce an efficiency bias between simulated capture events and muon decay events of 0.8%, favoring capture events.

6.1.1.2 Quality cuts

Post-reconstruction, explicit cuts are introduced to remove backgrounds. As mentioned before in Ch. 2, only subruns with stable operation and a 500-V/cm drift field are used. An event-quality cut is placed to remove pile-up events, in which the t_0 of the stopping muon is ambiguous. This is done in a two-fold fashion. First, the starting point of the muon trajectory must be reconstructed within the nominal active volume of 310.4 mm \times 620.8 mm \times 315 mm, using the event t_0 . Additionally, if multiple light triggers are present, they are all required

to be within $15 \mu\text{s}$ of the first light trigger. This cut removes 8% of the reconstructed events, but introduces negligible bias as less than 0.1% of muon decays and captures occur outside of this region. Assuming that the pile-up events are fully uncorrelated, a flat extrapolation of the sideband into the signal region predicts that the residual background from uncorrelated events is $< 0.3\%$.

A fiducial-volume cut is placed to remove events with a reconstructed stopping point that are outside of the 2.2-cm boundary from the edges of the field cage, anode, and cathode surfaces. The cut rejects events that cannot be well-reconstructed due to their containment and have a significant background from the through-going muons.

6.1.1.3 dQ/dx cut

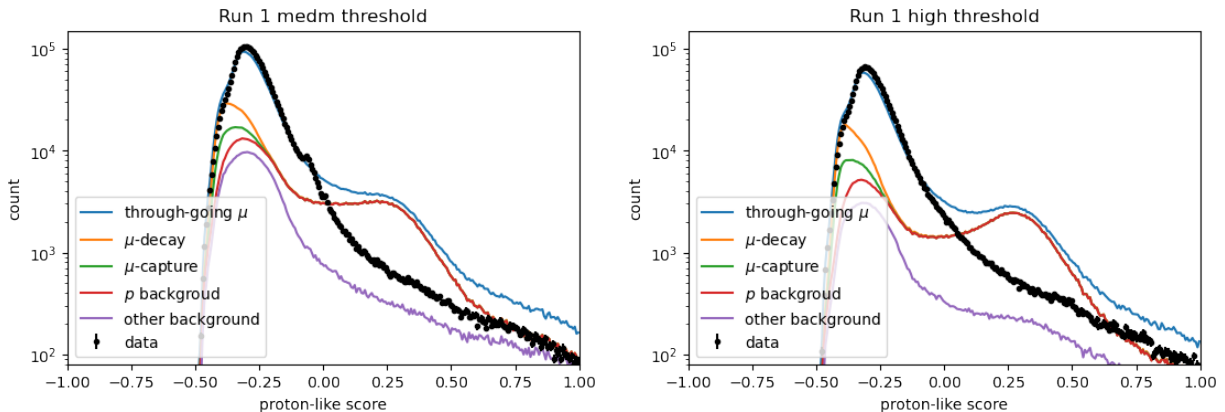


Figure 6.1: Pre-fit, pre-selection proton-likelihood score, with the simulation exhibiting a prominent peak of stopping-proton events not present in the Module 0 data.

To reject through-going muons and protons, a cut was placed on the dQ/dx profile likelihood scores described in Sec. 5.3.2. The MIP-like score was required to be less than 0.1, and the proton-like score was required to be less than -0.15 . A few discrepancies between the simulated cosmic rays and the Module 0 data are seen in the distributions of these likelihood parameters. First, the CORSIKA simulation initially over-predicted the number of proton events in the data sample to a substantial degree. Figure 6.1 shows a comparison between the observed proton likelihood score distributions and those produced by CORSIKA. The prominent proton peak is not apparent in the Module 0 data, and a conservative constraint using the number of events with a proton-likelihood-score above 0 suggests that the number of protons in the Module 0 data is $\ll 0.2$ of the CORSIKA prediction. It was determined that the most-likely cause of the CORSIKA overprediction was the use of an overly-simplistic primary flux model with a constant power-law slope. At low energies, the primary-proton spectrum is known to flatten [24], reducing the contribution of low-energy protons. For the

subsequent analysis with the CORSIKA simulation, protons were ignored as they contributed a $< 1\%$ background after reweighting by 0.2.

Second, compared to the CORSIKA simulation, the MIP and stopping-muon dQ/dx likelihood score peaks appear closer together in the data, such as is shown in Fig. 5.23. Examining the dQ/dx versus the residual range for the events in the stopping-muon sample, there is a difference in the dQ/dx width, with the simulation only reproducing 80% of the total width. This leads to a loss in the discrimination power of the dQ/dx profile likelihood score. To avoid introducing a potential bias in the decay and capture event selection efficiencies, the cuts on the dQ/dx profile likelihood scores are placed far from the signal region and preserve an efficiency of $> 99\%$.

6.1.1.4 Background veto

After applying the stopping point and dQ/dx profile likelihood score cuts, it was found that an unreasonable number of through-going muons contaminated the sample. Because such care was taken in the selection to avoid the possibility of an efficiency bias between the capture and decay events, no information about the decay electron was used. This enables a low-bias selection, but introduces a background from the through-going muons with a high dQ/dx fluctuation somewhere along their trajectory. For these events, the best-fit stopping point based on the dQ/dx profile is found at this fluctuation. If this falls within the fiducial volume, the event will likely pass the remaining selection cuts.

A naive cut on the reconstructed muon and decay-electron direction near the stopping point introduces an unreasonable systematic uncertainty in the capture efficiency. This is introduced primarily because variations in the reconstructed vertex can cause the end-point of the stopping muon to be mis-reconstructed slightly earlier along the muon trajectory. In these cases, the remainder of the muon trajectory is reconstructed as a nearly collinear decay electron, similar to the through-going muon backgrounds. Because the mis-reconstructed muon stub is preferentially reconstructed as parallel to the muon trajectory, the angular distribution of these mis-reconstructed capture events is steeply falling in the forward region with change in the efficiency of about 10% per 0.01 bin in $\cos\theta$. Thus any small mis-modeling near the end point of the muon, i.e. through pixel response functions or the recombination factor, could introduce a large effect in the selection efficiency for the muon-capture channel. Because of this, a less sensitive approach was taken.

First, events stopping in specific regions of the detector, near regions with lower sensitivity, were removed. For the through-going muon backgrounds, mis-reconstructed stopping points near these regions are particularly problematic. Because the remainder of the muon track often crosses this region, there is little information available to reject these events via the hit classifiers in Sec. 5.4. For through-going events with a mis-reconstructed stopping point far from the low-sensitivity regions, it is possible to use the remainder of the trajectory to remove these events. In particular, using the decay-like hit classifier, events with unreasonably large (> 100) decay-like hits are removed without biasing the muon capture or decay efficiencies. Additionally, the background-hit discriminator was used to create a conditional

cut on the number of background-like hits and the muon and decay-electron reconstructed axis. Events with more than 60 background-like hits or more than 15 background-like hits and a $\cos\theta < -0.992$ were also excluded. The addition of the background-like hit keeps the overall efficiency of the cut $> 99\%$ for both the capture and decay events, resulting in a capture versus decay efficiency bias of $< 1\%$.

6.1.2 Muon decay

To identify a sample of muon decays with high purity, both the CRS and LRS are used. Specific to the muon decay over capture and other backgrounds is the delayed-time structure of the entering muon and the subsequent decay with a timescale of $\sim \mu\text{s}$. First, a cut of 10 hits is placed on the number of decay-like hits, as determined by the classifier described in Sec. 5.4.1. This corresponds to a decay-electron energy of approximately > 10 MeV. Capture events and through-going muons produce a tail into the higher nhit distribution due to δ -rays originating near the end of the muon track and displaced energy arising from Compton-scattered electrons from δ -ray-induced Bremsstrahlung photons. The cut of 10 decay-like hits adequately reduces the capture-associated backgrounds by $> 99\%$ and through-going muon-associated backgrounds by $> 50\%$.

Within the LRS, mis-identified muon decays predominately arise from the intrinsic fluctuations of the LRS, particularly the noise sources that are correlated with the prompt-muon signal. Specifically, triplet light from LAr scintillation light produces a low-level photon background within the decay-time region-of-interest (ROI). This contribution is proportional to the prompt-light signal and is of the same scale as the singlet-light yield from a 10-MeV decay electron. Delayed-light emission from the wavelength shifters and the light-collection materials also contributes to the late-arriving light, albeit to a smaller degree. Measurements of late-light produced in the re-emission of TPB suggest that this late-light component has a characteristic timescale of $\tau \sim 3.5 \mu\text{s}$ but with a relative fraction of only $\approx 0.08 \pm 0.01$ [128]. Afterpulsing in the SiPMs is another contributor to the apparent light yield within the decay-time ROI. And finally, statistical fluctuations in the number of prompt-induced photoelectrons within the signal time gate can lead to an event-by-event bias in the PCA-background subtraction described in Sec. 5.4.

Due to the solid angle of the light collection modules and the continuous energy spectrum of the electron, the muon-decay signals vary in amplitude, with a portion falling below these background fluctuations. These event cannot be reconstructed reliably, and thus introduce a background in which a fluctuation in the light yield is mis-reconstructed as originating from the decay electron. Additionally, the LRS has a period of dead time just after each trigger. Decay events that occur in this dead time are always misreconstructed.

To reduce the backgrounds and mis-reconstructed events from the final sample, two cuts are placed on the light signals. First, a cut on the total amplitude of the light signal within a centered, 50-ns interval around the reconstructed decay time is placed at > 50 PE/50 ns. This aids in removing the contribution from the prompt-correlated background light. Second, a cut of > 10 PE/50 ns per decay-like hit is placed using the acceptance-corrected light

signal. This removes the mis-reconstructed muon decays that have a prominent decay-electron visible in the CRS but a small reconstructed light signal, e.g. muon decays that occur during the LRS deadtime.

6.1.3 Event selection summary

A summary of the selection efficiency and purity for each cut is provided in Tables 6.1-6.4. The total number of events and breakdown by particle species is listed in Table 6.1. Here, the number of events passing each selection cut for each data sample is provided in the column labelled 'total' along with the estimated number of background and signal contributions at each selection stage. To estimate the relative contributions, the best-fit values from the fit described in Sec. 6.2 are used. These values are only directly applicable to the fit of the post-selection sample. To extrapolate to the pre-selection particle species contributions, the efficiencies derived from simulation were used. Since there are significant uncertainties in the overall efficiency, the particle species contributions of the pre-cut selection do not necessarily add up to the total number of observed events in the pre-cut selection. So, after estimating the relative contributions for each particle species using the efficiency derived from the simulation, the total number of events are normalized for each pre-selection cut to the number of observed events of the corresponding cut

$$N_{\text{pre-cut pred., } X} = \frac{1}{\epsilon_X} N_{\text{best-fit, post-cut, } X} \quad (6.4)$$

$$N_{\text{pre-cut pred. norm., } X} = \frac{N_{\text{pre-cut pred., } X}}{\sum_X N_{\text{pre-cut pred., } X}} \quad (6.5)$$

where ϵ_X is the efficiency for a particle of species X . As a consequence, the efficiencies listed in Table 6.2 can exceed a value of 1 if the simulation under-predicts the total efficiency for any particular cut.

Table 6.1: Total number of events after each selection with the relative contributions by particle species and muon disappearance channel predicted by simulation, post-fit. Predictions are extrapolated from the background vetoed sample using the efficiency estimated with the simulation and a reweight factor determined by the fit. This extrapolation can result in an efficiency estimate > 1.00 for a particle species.

Cut	Signal (10^3)			Backgrounds (10^3)					Total	
	μ^- decay	μ capture	μ^+ decay	μ	non-fid.	μ	EM	n		π
Medium threshold, run 1 sample										
Quality	64.7	163	279	19600	284	2800	1310	30.4	24495996	
dQ/dx score	37.7	96.2	163	874	24.2	113	17.7	9.94	1336172	
Background veto	38.2	98.0	166	606	21.9	57.6	5.40	7.79	1000814	
Michel-like hits	29.7	76.3	129	101	15.3	40.0	3.33	5.33	399471	
Delayed light	24.6	0.518	108	48.6	8.26	13.1	1.13	2.67	206886	
	9.49	0.105	30.3	0.211	0.233	0.038	0.124	0.455	43026	
High threshold, run 1 sample										
Quality	29.4	73.5	126	9780	64.6	382	665	6.43	11124904	
dQ/dx score	25.8	65.8	112	653	8.04	33.0	18.9	3.18	919623	
Background veto	25.8	66.0	112	450	7.21	17.9	5.01	2.43	685665	
Michel-like hits	22.2	57.1	96.4	103	0.557	14.5	3.62	1.93	304056	
Delayed light	16.6	0.060	73.6	46.1	2.61	3.98	1.30	0.876	145081	
	6.75	0.011	21.653	0.230	0.779	0.018	0.167	0.159	29766	
High threshold, run 2 sample										
Quality	40.5	100	174	6960	81.0	396	740	8.28	8502698	
dQ/dx score	18.6	47.5	80.7	577	7.02	18.4	12.7	2.19	764621	
Background veto	17.6	44.8	76.2	353	5.58	10.2	3.57	1.58	512426	
Michel-like hits	11.3	29.1	49.0	52.3	2.83	7.38	1.84	0.983	154701	
Delayed light	8.103	0.103	35.8	23.8	1.20	1.98	0.650	0.434	72086	
	3.09	0.011	21.0	0.107	0.623	0.014	0.116	0.141	25146	

Table 6.2: Post-fit efficiency of each cut predicted with simulation.

Cut	Signal [%]		Backgrounds [%]					π	Total [%]
	μ^- decay	μ capture	μ^+ decay	μ non-fid.	μ EM	n			
Medium threshold, run 1 sample									
Quality	58.2	59.1	58.6	4.5	8.5	4.0	1.4	32.6	5.5
dQ/dx score	101.4	101.9	101.5	69.3	90.7	51.0	30.5	78.4	74.9
Background veto	77.7	77.9	77.7	16.6	69.8	69.4	61.7	68.4	39.9
Michel-like hits	82.8	0.7	83.8	48.3	53.9	32.8	34.0	50.0	51.8
Delayed light	38.5	20.3	28.0	0.4	28.3	0.3	11.0	17.1	20.8
High threshold, run 1 sample									
Quality	87.9	89.6	88.5	6.7	12.4	8.7	2.8	49.5	8.3
dQ/dx score	99.8	100.4	99.9	68.8	89.6	54.3	26.6	76.2	74.6
Background veto	86.3	86.5	86.3	22.8	77.1	80.9	72.3	79.3	44.3
Michel-like hits	74.8	0.1	76.3	44.8	46.9	27.4	35.9	45.3	47.7
Delayed light	40.6	16.7	29.4	0.5	29.8	0.5	12.8	18.0	20.5
High threshold, run 2 sample									
Quality	46.1	47.4	46.4	8.3	8.7	4.7	1.7	26.4	9.0
dQ/dx score	94.3	94.4	94.4	61.1	79.5	55.6	28.0	72.1	67.0
Background veto	64.4	64.8	64.4	14.8	50.7	72.0	51.7	62.3	30.2
Michel-like hits	71.6	0.4	73.0	45.6	42.3	26.8	35.2	44.2	46.6
Delayed light	38.1	10.7	58.8	0.4	52.1	0.7	17.8	32.5	34.9

Table 6.3: Post-fit cumulative efficiency of each cut predicted with simulation.

Cut	Signal [%]		Backgrounds [%]					π	n	Total [%]
	μ^- decay	μ capture	μ^+ decay	μ	non-fid. μ	EM				
Medium threshold, run 1 sample										
Quality	58.2	59.1	58.6	4.5	8.5	4.0	1.4	32.6	5.5	
dQ/dx score	59.1	60.2	59.5	3.1	7.7	2.1	0.4	25.6	4.1	
Background veto	45.9	46.9	46.2	0.5	5.4	1.4	0.3	17.5	1.6	
Michel-like hits	38.0	0.3	38.7	0.2	2.9	0.5	0.1	8.7	0.8	
Delayed light	14.7	0.1	10.9	< 0.1	0.8	< 0.1	< 0.1	1.5	0.2	
High threshold, run 1 sample										
Quality	87.9	89.6	88.5	6.7	12.4	8.7	2.8	49.5	8.3	
dQ/dx score	87.7	89.9	88.4	4.6	11.2	4.7	0.8	37.8	6.2	
Background veto	75.7	77.7	76.3	1.1	8.6	3.8	0.5	30.1	2.7	
Michel-like hits	56.7	0.1	58.2	0.5	4.0	1.0	0.2	13.6	1.3	
Delayed light	23.0	< 0.1	17.1	< 0.1	1.2	< 0.1	< 0.1	2.5	0.3	
High threshold, run 2 sample										
Quality	46.1	47.4	46.4	8.3	8.7	4.7	1.7	26.4	9.0	
dQ/dx score	43.4	44.7	43.8	5.1	6.9	2.6	0.5	19.0	6.0	
Background veto	28.0	29.0	28.2	0.8	3.5	1.9	0.2	11.9	1.8	
Michel-like hits	20.0	0.1	20.6	0.3	1.5	0.5	0.1	5.2	0.8	
Delayed light	7.6	< 0.1	12.1	< 0.1	0.8	< 0.1	< 0.1	1.7	0.3	

Table 6.4: Post-fit relative fraction of the sample labelled 'total' in Table 6.1 after each selection predicted with simulation

Cut	Signal [%]			Backgrounds [%]					
	μ^- decay	μ capture	μ^+ decay	μ	non-fid.	μ	EM	n	π
Medium threshold, run 1 sample									
Quality	0.3	0.7	1.1	79.9	1.2	11.4	5.4	0.1	
dQ/dx score	2.8	7.2	12.2	65.4	1.8	8.5	1.3	0.7	
Background veto	3.8	9.8	16.6	60.6	2.2	5.8	0.5	0.8	
Michel-like hits	7.4	19.1	32.3	25.2	3.8	10.0	0.8	1.3	
Delayed light	11.9	0.3	52.2	23.5	4.0	6.3	0.5	1.3	
	22.1	0.2	70.3	0.5	5.4	0.1	0.3	1.1	
High threshold, run 1 sample									
Quality	0.3	0.7	1.1	87.9	0.6	3.4	6.0	0.1	
dQ/dx score	2.8	7.2	12.2	71.0	0.9	3.6	2.0	0.3	
Background veto	3.8	9.6	16.3	65.6	1.1	2.6	0.7	0.4	
Michel-like hits	7.3	18.8	31.7	33.8	1.8	4.8	1.2	0.6	
Delayed light	11.5	< 0.1	50.7	31.7	1.8	2.7	0.9	0.6	
	22.7	< 0.1	72.7	0.8	2.6	0.1	0.6	0.5	
High threshold, run 2 sample									
Quality	0.5	1.2	2.0	81.9	1.0	4.7	8.7	0.1	
dQ/dx score	2.4	6.2	10.6	75.5	0.9	2.4	1.7	0.3	
Background veto	3.4	8.8	14.9	68.9	1.1	2.0	0.7	0.3	
Michel-like hits	7.3	18.8	31.7	33.8	1.8	4.8	1.2	0.6	
Delayed light	11.2	0.1	49.7	33.0	1.7	2.7	0.9	0.6	
	12.3	< 0.1	83.7	0.4	2.5	0.1	0.5	0.6	

6.2 Fit

To extract the Huff factor and capture rate, a binned maximum-likelihood fit is used. Data are distributed into two histograms for each data sample: a decay-nhit histogram (using the number of reconstructed decay-like hits, described in Sec. 5.4.1) and a muon decay time histogram (using the reconstructed muon decay time, described in Sec. 5.4.2). From these histograms, a likelihood function was constructed assuming that the residual of each bin is normally distributed

$$\mathcal{L}(n_i; \vec{\theta}) \equiv \log P(n_i; \vec{\theta}) = - \sum_i \frac{\left(n_i - n_{\text{pred},i}(\vec{\theta})\right)^2}{2 \left(n_i + \sigma_i^2(\vec{\theta})\right)} + \sum_j \mathcal{L}_{\text{constr.}}(\theta_j) + \text{const.} \quad (6.6)$$

where n_i is the number of events observed in bin i , $\vec{\theta}$ is the vector of parameters including both the nuisance parameters and the parameters of interest, $n_{\text{pred},i}(\vec{\theta})$ is the model prediction for the number of events in bin i for a given $\vec{\theta}$, and $\sigma_i^2(\vec{\theta})$ is the model variance of bin i (due to simulation statistics or modeling uncertainty). The additional $\mathcal{L}_{\text{constr.}}(\theta_j)$ terms are included for some of the results described in Sec. 6.4 as an external constraint on one or more parameters.

The model prediction $n_{\text{pred},i}(\vec{\theta})$ is determined by simulation, theory, or by direct measurement of a sideband sample, as listed in Table 6.5, and is described in more detail in this section. The free parameters of the fit $\vec{\theta}$ include the parameters of interest along with nuisance parameters for the relative normalization, scale, and smearing of the distributions as compared with the model prediction. A detailed accounting of these parameters is included towards the end of this section, along with a description of the toy MC used to validate the fit.

For the muon-decay nhit distribution, a bin width of 2 hits per bin was used to maintain good resolution for the transition from the capture regime ~ 0 , the decay regime ~ 50 , and the background regime ~ 100 . For the decay-time distribution, no interpolation was applied, so the reconstructed decay times are discretized with a spacing of 10 ns. Thus the bin width was required to be a multiple of 10 ns, so as not to introduce irregularities in the distributions. For the Run 1 data samples, a uniform binning of 40 ns was used. For the Run 2 data sample, the same binning as the Run 1 sample was used for decay times $< 2 \mu\text{s}$. Beyond this, the bin widths were increased logarithmically, rounding to the nearest 10 ns. This ensured that each bin has sufficient events that the residuals can be assumed to follow a normal distribution ($n_i \sim 100$).

The model prediction for the muon-decay nhit shape was determined by using the CORSIKA simulation samples described in Ch. 3. From the simulation, a template binned PDF was produced for each data sample and for each particle species and final state of interest. Nuisance parameters were incorporated to permit overall scale and smearing effects. The scale nuisance parameters were implemented using a linear interpolation of the unmodified cumulative probability distribution function (CDF) extracted from the simulation. The

Table 6.5: Summary of components that make up the model prediction in the likelihood function. The free parameters f , μ , σ , and λ represent relative normalization, scale, smearing, and exponential-decay constant, respectively. This table does not fully represent the free parameters used in the fit, which include additional constraints between the parameters. These are described in more detail at the end of Sec. 6.2.

Muon-decay nhit distribution		
Component	Source	Free parameters
μ -capture	Simulation	f, μ, σ
μ^- -decay	Simulation	f, μ, σ
μ^+ -decay	Simulation	f, μ, σ
through-going μ	Simulation	f, μ, σ
non-fid. μ decay	Simulation	f
EM-background	Simulation	f
π -background	Simulation	f
n -background	Simulation	f
Muon-decay time distribution		
μ^- -decay	Theory	f, λ
μ^+ -decay	Theory	f, λ
μ -correlated background	Sideband	f

impact on the PDF from a change in the muon-decay nhit scale was then be determined by taking the change in the CDF within a given bin

$$P_\eta(N = n|\alpha) = \Phi_\eta\left(\alpha\frac{2n+1}{2}\right) - \Phi_\eta\left(\alpha\frac{2n-1}{2}\right) \quad (6.7)$$

where Φ_η is the linear-spline CDF for particle type or final state η , a rescaling factor α , and a given muon-decay nhit bin n . This enabled an interpolation between the scale factors in a continuous fashion and without introducing additional assumptions about the underlying PDF.

Even after the event selection, backgrounds in the decay-like nhit distribution are still pronounced, with approximately 30% of the sample consisting of backgrounds. The primary contributor to the backgrounds comes from mis-reconstruction of the track endpoint in the through-going muon events. Because only approximately 3% of the muons stop in the detector, even a small number of through-going muons with mis-reconstructed end-points ($\sim 1\%$, in this case) contributes a large background to the stopping-muon sample. Improvements to the reconstruction algorithms could likely reduce this background, however care must be taken so as not to introduce a bias in the selection efficiency for muon capture and muon decay events. This limits the degree to which the decay-electron can be leveraged in the

selection. In the case of this analysis, it was preferred to use a less-performant selection with a small bias to reduce the systematic associated with the reconstruction efficiency of muon capture and decay events. The other backgrounds from π , n , electrons, and γ also contribute, but to a lesser degree. Additional information using the event topology could have been used to improve the purity of the sample from these event classes, however as they contribute less than the through-going muon background, this was not prioritized.

The number of μ^\pm decays in the detector as a function of time follows

$$\frac{dN_{\mu^\pm}}{dt} = -\lambda_{\mu^\pm} N_{\mu^\pm} \quad (6.8)$$

where $\lambda_{\mu^+} = 1/\tau_\mu$, with τ_μ being the muon lifetime, and $\lambda_{\mu^-} = \lambda_c + R_H \lambda_\mu$. Thus, an exponential distribution can be used to model the number of muons of a given type. The selection efficiency as a function of the reconstructed time was examined using simulation and was found to be consistent with a flat distribution. Thus, no relative efficiency correction was applied. The reconstructed time resolution for muon decays was estimated by simulation to be

$$\sigma = (3.79 \pm 0.05 \text{ ns}) + \left(\frac{192.0 \pm 1.9}{\mu} \cdot \text{PE} \cdot \text{ns} \right)$$

where μ is the mean delayed-light yield in units of PE. This is consistent with measurements of the LRS resolution using the Module 0 data [207]. This is less than the 10-ns binning for all signals that pass the event selection for muon decays. No dependence on the time resolution was found as a function of the decay time. Because of these factors, the migration of events between bins in the decay-time distributions is predicted to be negligibly small ($< 10^{-6}$). The reconstructed decay time does exhibit a modest (0.82 ± 0.03) ns bias in the simulation, likely due to signal deconvolution model mis-match that causes the waveform peak to be asymmetric. However this bias is constant as a function of the decay time and so does not distort the slope of the decay-time distribution. The bias also introduces a small difference in the relative efficiency of the decay-time selection window for μ^+ and μ^- decays of $< 0.1\%$. Thus the use of a simple exponential decay model was justified, without including bin-by-bin efficiency corrections or time smearing.

Backgrounds enter the selection through correlated and uncorrelated activity associated with the stopping muon. Uncorrelated backgrounds arising from multiple muon pile-up are expected to be of the $O(10^{-9})$ within the 10- μs selection window and is assumed negligible. Additionally, uncorrelated backgrounds related to low-energy radioactivity are heavily suppressed by the muon-decay nhit cut. For example, using the estimated activity of ^{39}Ar in atmospheric argon, we can anticipate approximately an intrinsic-background rate of 1 Bq/kg [222]. Thus the overall expected rate is around 600 Bq, which translates to about 0.1 decays per 200- μs drift window. The ^{39}Ar β -decay endpoint is around 600 keV, resulting in around one self-triggered channel. Thus, the muon-decay nhit cut is sufficient to remove these events, though they could cause a small bias in the muon-decay nhit cut of order +0.1 hits. For the light signal, the delayed-signal gate of 10 μs results in $< 1\%$ of the events overlapping with an ^{39}Ar decay. Considering the relatively low light yield of these

events, they will be preferentially removed by the selection cuts and thus are not expected to influence the fit. The PCB and G10 materials used to build the TPC's mechanical structure may contain a higher radioactivity (up to ~ 20 Bq/kg [223]), but only contribute at the same order of magnitude, due to their smaller contribution to the overall mass of the detector.

Activity associated with the muon-capture events was deemed to be negligible in the decay-time distributions. Neutrons produced in muon capture could introduce an intrinsic background. As the neutron propagates, low-energy n-Ar elastic scatters occur, depositing energy visible in both the CRS and LRS. However, most of this energy escapes the detector due to the small neutron cross-section. While there are significant uncertainties in the cross-section models used in the Geant4 stage of the simulation, the scatters are removed with very high efficiency by the muon-decay nhit cut. Thus, even quite large (factors of 5 or more) systematic uncertainties in the neutron cross-section would bias the fit results by less than $\sim 1\%$. Muon-capture events also produce unstable argon and chlorine isotopes, which have both fast- and slow-decay timescales. For fast-decaying intermediate states, energies do not exceed ~ 1 MeV and are thus sub-threshold [96]. For the slow-decaying states, e.g. $^{39,40}\text{Cl}$, half-lives are of the order of minutes to hours and thus have a negligible decay probability within the 10- μs LRS window or the 200- μs drift time. Overall, this leads to a negligible background ($< 0.1\%$) of the muon-capture events in the decay-time distributions.

Triplet light is the most prominent cause of backgrounds within the decay-time spectrum. Typically, the light yield from the muon-decay electrons are $O(100 \text{ PE}/50 \text{ ns})$. The signal from the stopping muon is about $6\times$ larger, translating to a triplet-light signal of roughly $O(60 \text{ PE}/50 \text{ ns})$ within the decay window. Stochastic fluctuations in this light are around 7-8 PE/50 ns. Thus, the cut placed at 10 PE/50 ns is well situated to remove most of the events with poorly-reconstructed light signals. The intrinsic noise of the SiPM readout is approximately 2 PE/50 ns, and so substantially smaller than the triplet contribution. Combined, about 5% of the through-going, capture, or true muon-decay events have a mis-attributed or mis-reconstructed light signal due to these fluctuations and the LRS deadtime.

Both through-going and capture events are reasonably well-rejected using the CRS selection on the number of muon-decay-like hits, but the contribution from mis-reconstructed muon-decay events must be accounted for in the fit. To do this, a separate selection was made by inverting the MIP-likelihood score cut described in Sec. 6.1.1. This selection contains a high-purity sample of through-going muons (95%), with $\sim 1\%$ background contributions from protons, pions, and non-fiducial stopping muons, respectively, as predicted by the simulation. The light signals from these events were analyzed with the same reconstruction algorithm and selection cuts that were applied to the signal sample. The resulting empirically-derived distribution was then used as the intrinsic-background model, accounting for the effects of the triplet-light fluctuations, as well as any signal processing artifacts. The overall normalization of the background model was allowed to float during the fitting process and suggested a background contribution in line with the expectations from the simulation. This approach is non-specific to the details of the background source. Thus, it also includes any muon-associated background sources even if they were not specifically considered in this analysis.

To extract the branching fraction and charge ratio, the number of muons N_μ that stop

within the fiducial volume by charge can be partitioned in the following way

$$N_{\mu^+} = \frac{R}{1+R}N_{\mu} \quad (6.9)$$

$$N_{\mu^-} = \frac{1}{1+R}N_{\mu} \quad (6.10)$$

where $R \equiv N_{\mu^+}/N_{\mu^-}$ is the muon charge ratio. Of the negative muons, we assume they disappear via one of the two processes previously discussed, capture or DIO

$$N_{\text{capture}} = BN_{\mu^-} \quad (6.11)$$

$$N_{\text{DIO}} = (1-B)N_{\mu^-} \quad (6.12)$$

where B is the branching fraction for capture. The observable within the muon-decay nhit distribution is the relative fraction of decay or capture events, where the decay events contain both DIO from μ^- and decay-at-rest (DAR) from μ^+

$$f_{\text{decay}} \equiv 1 - f_{\text{capture}} = \frac{n_{\text{DIO}} + n_{\text{DAR}}}{n_{\text{capture}} + n_{\text{DIO}} + n_{\text{DAR}}} \quad (6.13)$$

where $n_{\text{DIO}} = \epsilon_{\text{decay}}N_{\text{DIO}}$, $n_{\text{DAR}} = \epsilon_{\text{decay}}N_{\text{DAR}}$, and $n_{\text{capture}} = \epsilon_{\text{capture}}N_{\text{capture}}$ include the effect of the event selection efficiency (ϵ_X). From simulation, $\epsilon_{\text{decay}} \equiv \epsilon_{\text{DIO}} = \epsilon_{\text{DAR}}$ holds to better than 0.1%, while $\epsilon_{\text{decay}} \approx \epsilon_{\text{capture}}$ holds at the percent level. The expressions above can be combined to express this fraction in terms of the branching ratio, the charge ratio, and the efficiency ratio as

$$f_{\text{decay}} = \left(1 + \frac{\epsilon_{\text{capture}}}{\epsilon_{\text{decay}}} \frac{B}{1-B+R}\right)^{-1} \quad (6.14)$$

by assuming that $N_{\text{DAR}} = N_{\mu^+}$.

The decay-time subsample admits a similar parameterization

$$f_{\mu^-} \equiv 1 - f_{\mu^+} = \frac{n'_{\text{DAR}}}{n'_{\text{DAR}} + n'_{\text{DIO}}} \quad (6.15)$$

$$= \frac{\epsilon_{\text{window}}(\tau_0)\epsilon_+N_{\mu^+}}{\epsilon_{\text{window}}(\tau_0)\epsilon_+N_{\mu^+} + \epsilon_{\text{window}}(\tau_d)\epsilon_-(1-B)N_{\mu^-}} \quad (6.16)$$

where $\epsilon_{\text{window}}(\tau)$ is the selection efficiency due to the time window and ϵ_{\pm} is the relative efficiency for selecting the DAR and DIO events, excluding the time window efficiency. The decay time window efficiency is readily calculated based on the exponential decay time model

$$\epsilon_{\text{window}}(\tau) = e^{-t_0/\tau} - e^{-t_1/\tau}, \quad (6.17)$$

The DIO and DAR efficiencies only differ due to the binding energy and radiative corrections to the decay electrons in low-energy (< 10 MeV) DIO events. The simulation predicts an efficiency ratio of $\epsilon_+/\epsilon_- \sim 1.03$. After some manipulation, Eq. 6.15 can be reduced to

$$f_{\mu^-} = 1 - \left(1 + \frac{\epsilon(\tau_d) \epsilon_- (1 - B)}{\epsilon(\tau_0) \epsilon_+ R} \right)^{-1}. \quad (6.18)$$

The background model was included separately within the fit. A global background fraction f_{bkg} is used as a free parameter, with the relative contributions $f_{\text{non-stop } \mu}$, $f_{\text{non-fid. } \mu}$, f_{EM} , f_{π} , and f_n such that they sum to 1. The overall contribution of each muon disappearance channel is thus $f_X(1 - f_{\text{bkg}})$, and each background is $f_{\text{bkg}}f_X/\sum_Y f_Y$.

In total, the fit parameters $\vec{\theta}$ can be grouped into two categories. The first group includes the parameters of interest and consists of the muon capture rate λ_c , the μ^+/μ^- charge ratio, and the Huff factor R_H . Also of interest for this analysis is the total disappearance rate λ_d , but this can be expressed in terms of the capture rate and the Huff factor, and so it is not a free parameter of the fit. The second group are the continuous nuisance parameters that include the backgrounds and shape modifications. Within this group, there is an overall background normalization for each of the distributions and data samples (6 parameters). The normalization of each background type was allowed to freely vary within the overall normalization (5 parameters). Because simulation determined that the through-going, non-fiducial stopping muon, and neutron backgrounds had meaningfully different normalizations in the three data samples, they were allowed to freely vary for each data sample (6 parameters). Two shape parameters were included to account for an overall scale difference in the muon-decay nhit distribution, as well as a broadening term that set the width of a Gaussian convolution kernel that was applied to the template muon-decay nhit distributions. One set of shape parameters was used for each of the muon-decay, capture, and through-going muon distributions per charge threshold setting (12 parameters). Additional systematic parameters for data loss, charge gain bias, and triplet-light decay time were included as continuous systematics (3 parameters). Other systematics were estimated independently, as described in Sec. 6.3. In total, the fit included 35 free parameters, 3 of which are of physical interest.

We minimize the negative log-likelihood to create an estimator for the model parameters

$$\hat{\vec{\theta}} \equiv \operatorname{argmin}_{\vec{\theta}}[-\mathcal{L}(n_i; \vec{\theta})]. \quad (6.19)$$

The fit errors are determined using the bootstrap approach [220]. In this, the data are used as an empirical distribution from which a random sample with the same sample size is drawn with replacement. The parameters of interest and nuisance parameters of the randomized sample are then estimated using the same maximum log-likelihood fit. A large number (> 500) of these fits are collected and used to infer the variance and correlations between the best-fit parameters.

To validate the fitting procedure, a closure test was run. The observed values n_i were replaced by the values generated from the simulation with a known value for the parameters

Table 6.6: Fit results of the toy MC fit closure test.

Parameter	True value	Asymptotic fit value	Variation in fit value	Example fit result
R	1.290	1.292 ± 0.006	± 0.04	1.35 ± 0.04
R_H	0.990	1.009 ± 0.008	± 0.05	1.03 ± 0.05
λ_c (μs^{-1})	1.273	1.257 ± 0.010	± 0.09	1.29 ± 0.10
λ_d (μs^{-1})	1.732	1.716 ± 0.012	± 0.10	1.76 ± 0.12

of interest and the nuisance parameters. The fitting procedure was applied to the test sample, and the resulting best-fit parameters compared to the true values and the distributions of the best-fit parameters were verified to produce the same variance as the bootstrap prediction.

Table 6.6 shows the results of this test. The true value for each of the parameters-of-interest was compared with the asymptotic fit value, defined as the average of 500 fits with statistical variations applied. The asymptotic fit value was found to agree to within $\approx 1\%$ of the true value of each parameter. The variation in the best-fit values was also compared to the results extracted using the bootstrap errors of a single toy MC sample. The bootstrap errors are in agreement or slightly more conservative than the true variation in the best-fit value.

6.3 Systematic effects

Table A.1 lists the systematic effects that were investigated in this measurement. For most systematics, the effect was studied by directly modifying the simulation and comparing the resulting distribution of the reconstructed decay time and the muon-decay nhit score to the simulation baseline. The change in shape for each effect X was estimated via

$$\delta_{i,X} = \frac{n_{i,\text{pred}}(\vec{\theta}_0 \pm \sigma_X) - n_{i,\text{pred}}(\vec{\theta}_0)}{n_{i,\text{pred}}(\vec{\theta}_0)} \quad (6.20)$$

where $\vec{\theta}_0$ is the nominal pre-fit values for each parameter and $\pm\sigma_X$ is one standard deviation in the effect under study. For distributions that produced a $> 0.5\%$ change in the predicted value in one of the bins, the systematic effect was either included as a nuisance parameter or studied independently. Systematics that produced $< 0.5\%$ change in the distribution shape were deemed irrelevant and not investigated further. A detailed list of the studied effects are included in Appendix A.

Where possible, the relevant systematic was included into the fit directly through a nuisance parameter. To include the systematic, a $\pm\sigma$ shift in the parameter value was simulated. A 2-parameter quadratic interpolation was then calculated for each bin using the change resulting from the $\pm\sigma$ variation. This parameterized the effect of a continuous change in the given parameter on the fit distributions, even when the underlying relationship

to the systematic was quite complicated. The systematic was then allowed to vary freely during the fit minimization.

For other systematics, a continuous variation is not relevant (e.g. using the Modified-Box or Birks' recombination model) and so, instead, the fit minimization was repeated once with the systematic effect enabled and once with the systematic effect disabled. The magnitude of the impact of the systematic was then taken to be $\vec{\sigma}_{\text{effect}} = |\vec{\theta}_{\text{best-fit, enabled}} - \vec{\theta}_{\text{best-fit, disabled}}|$. Included within this group of systematics are the cut boundaries. Each cut was independently varied by a relevant amount and then fitting procedure repeated, with the details for each cut variation listed in Appendix A.

A handful of systematics not listed in Table A.1 were investigated independently. The first potential bias arises from decay-in-flight muons. A small number of muons entering the detector decay prior to stopping. As the negative muons that decay-in-flight do not undergo muon capture, they bias the branching ratio in favor of the DIO mode, increasing R_H . Simulation suggests that only $\approx 0.7\%$ of the negative muons are impacted, resulting in a bias on R_H of $\approx 1.5\%$. However, because these events do not have a Bragg peak, they are less likely to pass our dQ/dx selection. Again using simulation, we estimate that the fraction of the negative muons in our sample that did not undergo a decay is 0.3%. Thus, decay-in-flight muons contribute a $\approx 0.6\%$ bias on the measured R_H . This effect is much smaller than other systematics and thus was neglected in the final results.




Figure 6.2: Change in best-fit value for R_H as a function of fiducial-volume cut. The overall statistical uncertainty on R_H is shown with dashed lines. The relative change in the statistical error after modifying the fiducial cut is shown with the error bars ($\sigma_{\text{rel.}}^2 = |\sigma_{\text{nominal}}^2 - \sigma_{\text{change}}^2|$) [224].

The event selection was developed specifically to avoid introducing efficiency bias between the capture and decay events. For the nominal fiducial volume, the bias was estimated using

simulation to be 2%. This bias directly impacts the measured capture rate via Eq. 6.14. Within the fit, this bias was corrected based on the value predicted by simulation. But to explore the fidelity of this correction, a change in the fiducial volume was used as a means to producing a less-biased sample. By reducing the fiducial volume to 4 cm (6 cm) from the detector boundaries, the selection bias is reduced to 1% (0.3%). This is due to the fact that the bias arises from the decay electrons that escape the active volume and are mis-reconstructed as through-going muons. A smaller fiducial volume results in a sample with fewer decay electrons that are mis-reconstructed. Thus, modifying the fiducial volume and re-fitting the data gives confidence that the efficiency bias is modelled correctly. The results for the Huff factor (which is most sensitive to the selection efficiency bias) are shown in Fig. 6.2. As can be seen, the change due to the fiducial-volume shift is much smaller than the overall statistical uncertainty and is not statistically significant. Similar results were found for the other parameters of interest. This offers some confidence that our background modelling is robust, as the reduced-fiducial-volume samples have smaller background contributions from through-going muons, and so changing the fiducial volume is sensitive to potential mis-modelling of the background.

Because the muon charge ratio is extracted primarily from the muon-decay time spectrum, any bias in the muon-decay selection efficiency due to the charge of the muon would directly bias the charge-ratio measurement and impact the measured capture rate. Based on the nominal selection, the bias was estimated by simulation to be 3% and was corrected for in the fit model. If the muon decay spectrum generated by Geant4 is unreliable, this bias could be more prominent. To investigate this, the change in the muon-decay nhit distribution for events with a long muon-decay time was compared to the nominal distribution.

A reference μ^+ sample was created by selecting only muon-decay candidates with a reconstructed decay time $> 3 \mu s$. Within this slow-decay sample, we isolate $> 99\%$ positive muons, and thus the difference in the shape of the slow-decay muons is correlated to the differences between the positive and negative muons. In particular, we expect a shift towards higher decay-nhit for the slow-decay sample due to the difference in the decay-electron energies (see for example Fig. 1.7). Overall, the simulation reproduces the same differences that are observed in the data. The Run 2 sample does not show as significant of a shift due to the 2 tiles that were disabled during this run. This causes a shift towards lower nhit that impacts positive-muon decays to a greater extent than negative-muon decays. The overall $\chi^2/\text{ndf} = 112.9 / 115$ does not suggest significant mis-modelling between the μ^+ and the μ^- decays. This analysis is somewhat limited by the number of μ^+ decays in the slow-decay sample and the relatively high threshold of the light selection cut, but it suggests no gross mis-modeling of the μ^- DIO spectrum.

Finally, in addition to neutrons, it has been observed that protons can also be emitted in muon-capture events via nuclear effects. This contributes to a charged particle tail out to 30-40 MeV for captures on Si [225], though this contributes to $< 1\%$ of the total cross-section. As a 40-MeV proton has a relatively short range (~ 3 cm), in our reconstruction, most of the track will be rejected by the muon-decay likelihood classifier. However, these events will contribute to the tail of the muon-decay nhit distribution particularly near $\sim 5 - 10$ hits.

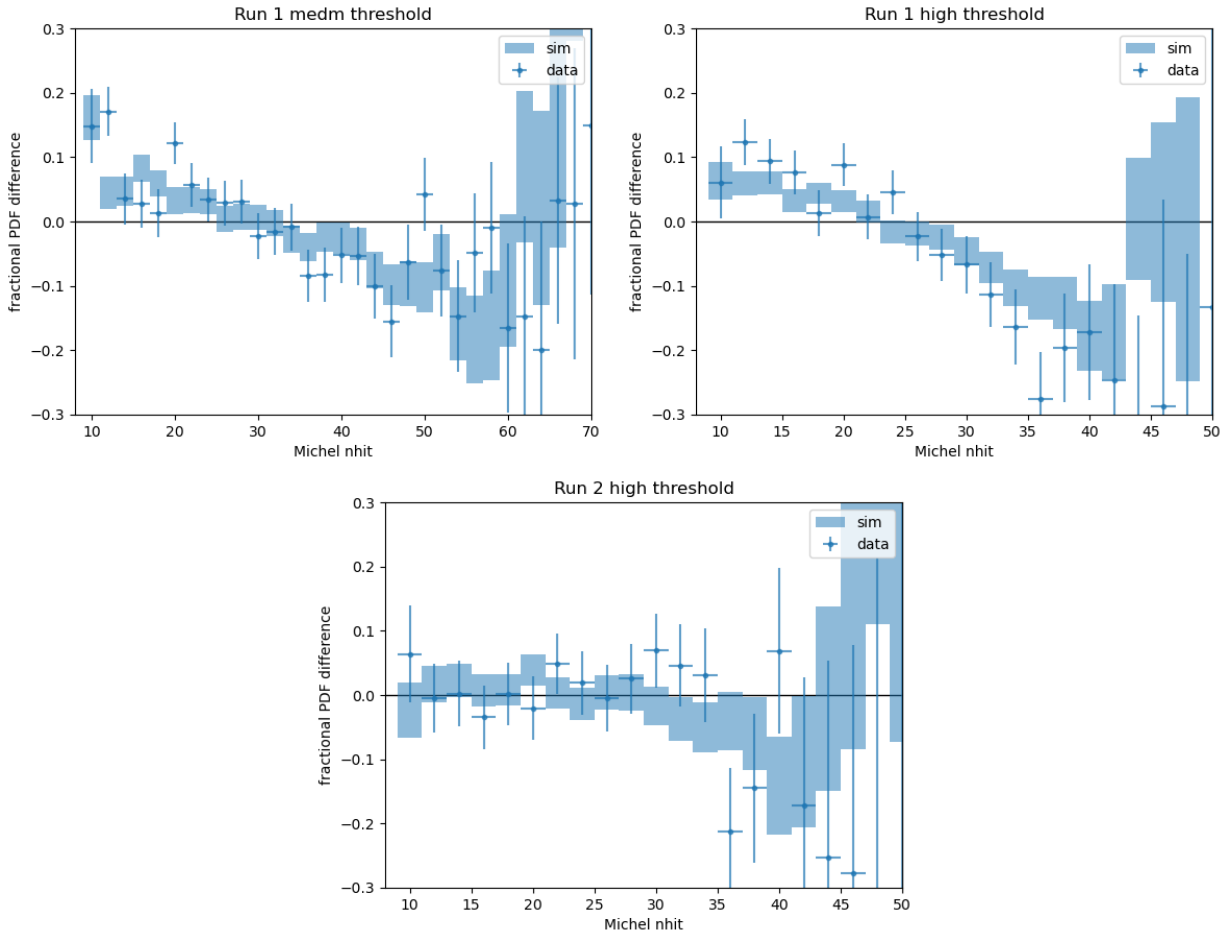


Figure 6.3: Change in the shape of the decay-like nhit distributions ($(n_{\text{nom.}} - n_{\text{slow}})/n_{\text{nom.}}$) between the nominal muon decay sample and the slow-decay samples, where $n_{\text{nom.}}$ and n_{slow} are normalized to the number of events in each sample, respectively. The corresponding χ^2/ndf are 49.76/40, 35.45/37, and 27.65/38 for the Run 1 medium threshold, Run 1 high threshold, and Run 2 high threshold samples, respectively.

As it is a small fraction of the overall cross-section, it is not expected to impact the results of this measurement by more than 1%.

Tables 6.7-6.10 summarize the relative contributions of the systematics to the final uncertainty. In general, no one systematic dominates the result fully, but depending on the parameter of interest and the applied constraint, different systematics contribute to a larger or smaller degree. For the capture rate, the uncertainty due to the change in the muon-decay nhit distribution arising from different charge threshold samples is the most important, regardless of the external constraint that is applied. The uncertainty from the light-signal threshold is of a similar magnitude, but only for the unconstrained result. The impact of

Table 6.7: Summary of measurement uncertainties for the fit with no external constraints.

	μ^+/μ^-	R_H	λ_d (μs^{-1})	λ_c (μs^{-1})
	Statistics			
	± 0.04	± 0.09	± 0.18	± 0.14
Name	Systematics			
Charge threshold	± 0.02	± 0.03	± 0.06	± 0.07
DIO/capture eff. ratio	± 0.02	± 0.03	± 0.00	± 0.01
μ^+/μ^- eff. ratio	± 0.02	± 0.03	± 0.00	± 0.01
Light threshold	± 0.02	± 0.08	± 0.11	± 0.07
Decay-time window	± 0.01	± 0.07	± 0.09	± 0.06
All others	± 0.01	± 0.03	± 0.05	± 0.04
Total	± 0.04	± 0.12	± 0.16	± 0.13
Stat. + System.	± 0.06	± 0.15	± 0.24	± 0.19

varying the decay-time reconstruction window is surprisingly pronounced and may suggest a mis-modeling of the muon-correlated background, described more in Sec. 7.2. Worth noting, is that during the systematic evaluation, the Modified-Box model produced very poor fit results ($\chi^2/\text{ndf} \sim 3$), thus our fit results strongly preferred the Birks' model. More investigation into the mechanism behind this sensitivity to the recombination model is warranted.

Table 6.8: Summary of measurement uncertainties for the fit with a λ_d constraint from Ref. [95].

	μ^+/μ^-	R_H	λ_d (μs^{-1})	λ_c (μs^{-1})
	Statistics			
	± 0.03	± 0.04	± 0.01	± 0.02
Name	Systematics			
Charge threshold	± 0.03	± 0.05	± 0.00	± 0.02
DIO/capture eff. ratio	± 0.02	± 0.02	± 0.00	± 0.01
μ^+/μ^- eff. ratio	± 0.01	± 0.02	± 0.00	± 0.01
Light threshold	± 0.02	± 0.03	± 0.00	± 0.01
Decay-time window	± 0.02	± 0.04	± 0.01	± 0.01
All others	± 0.01	± 0.02	± 0.00	± 0.01
Total	± 0.05	± 0.08	± 0.01	± 0.03
Stat. + System.	± 0.06	± 0.09	± 0.01	± 0.04

Table 6.9: Summary of measurement uncertainties for the fit with a λ_d constraint from Ref. [96].

	μ^+/μ^-	R_H	λ_d (μs^{-1})	λ_c (μs^{-1})
	Statistics			
	± 0.04	± 0.04	± 0.01	± 0.02
Name	Systematics			
Charge threshold	± 0.02	± 0.04	± 0.00	± 0.02
DIO/capture eff. ratio	± 0.02	± 0.02	± 0.00	± 0.01
μ^+/μ^- eff. ratio	± 0.02	± 0.02	± 0.00	± 0.01
Light threshold	± 0.02	± 0.03	± 0.00	± 0.01
Decay-time window	± 0.03	± 0.04	± 0.01	± 0.01
All others	± 0.01	± 0.02	± 0.00	± 0.01
Total	± 0.05	± 0.07	± 0.01	± 0.03
Stat. + System.	± 0.06	± 0.08	± 0.01	± 0.04

Table 6.10: Summary of measurement uncertainties for the fit with a $R_H = 0.99$ [63] constraint. The constraint assumes no error on R_H .

	μ^+/μ^-	R_H	λ_d (μs^{-1})	λ_c (μs^{-1})
	Statistics			
	± 0.05	± 0.00	± 0.08	± 0.08
Name	Systematics			
Charge threshold	± 0.07	± 0.00	± 0.02	± 0.07
DIO/capture eff. ratio	± 0.03	± 0.00	± 0.03	± 0.02
μ^+/μ^- eff. ratio	± 0.03	± 0.00	± 0.02	± 0.03
Light threshold	± 0.02	± 0.00	± 0.02	± 0.04
Decay-time window	± 0.01	± 0.00	± 0.01	± 0.01
All others	± 0.02	± 0.00	± 0.02	± 0.04
Total	± 0.09	± 0.00	± 0.05	± 0.10
Stat. + System.	± 0.10	± 0.00	± 0.10	± 0.13

6.4 Results

The results of the fit described in Sec. 6.2 are shown in Figures 6.5 and 6.4. The overall χ^2/ndf is calculated using the best-fit result without constraints

$$\chi^2 = \sum_i \frac{\left(n_i - n_{\text{pred},i}(\vec{\theta})\right)^2}{n_i + \sigma_{\text{pred},i}^2(\vec{\theta})} \quad (6.21)$$

$$\text{ndf} = N_{\text{bins}} - N_{\text{parameters}} \quad (6.22)$$

where N_{bins} is the number of bins across all 6 histograms (276) and $N_{\text{parameters}}$ is the number of free parameters in the fit (35). The effects of the systematics described in Sec. 6.3 are included into $\sigma_{\text{pred},i}$ using the absolute deviation of the post-fit from the best-fit value $n_{\text{pred},i}(\vec{\theta})$ for each bin i under an ensemble of fits with varied systematic parameters. The total χ^2/ndf is 229.8/237 (= 0.97), which does not suggest significant modelling issues. Examining the muon-decay time distributions, Fig. 6.4, fluctuations in the fit residuals are largely uniform. Examining the muon-decay nhit distributions, Fig. 6.5, we observe some anomalies near $\sim 5 - 10$, indicating that the fit model may be overly flexible in other bins and not fully capturing the behavior of the decay-nhit distribution in this region. Unfortunately, this area has roughly equal contributions from muon-decay, muon-capture, and through-going-muon background events, making it challenging to isolate the most-likely contributor to the modelling error.

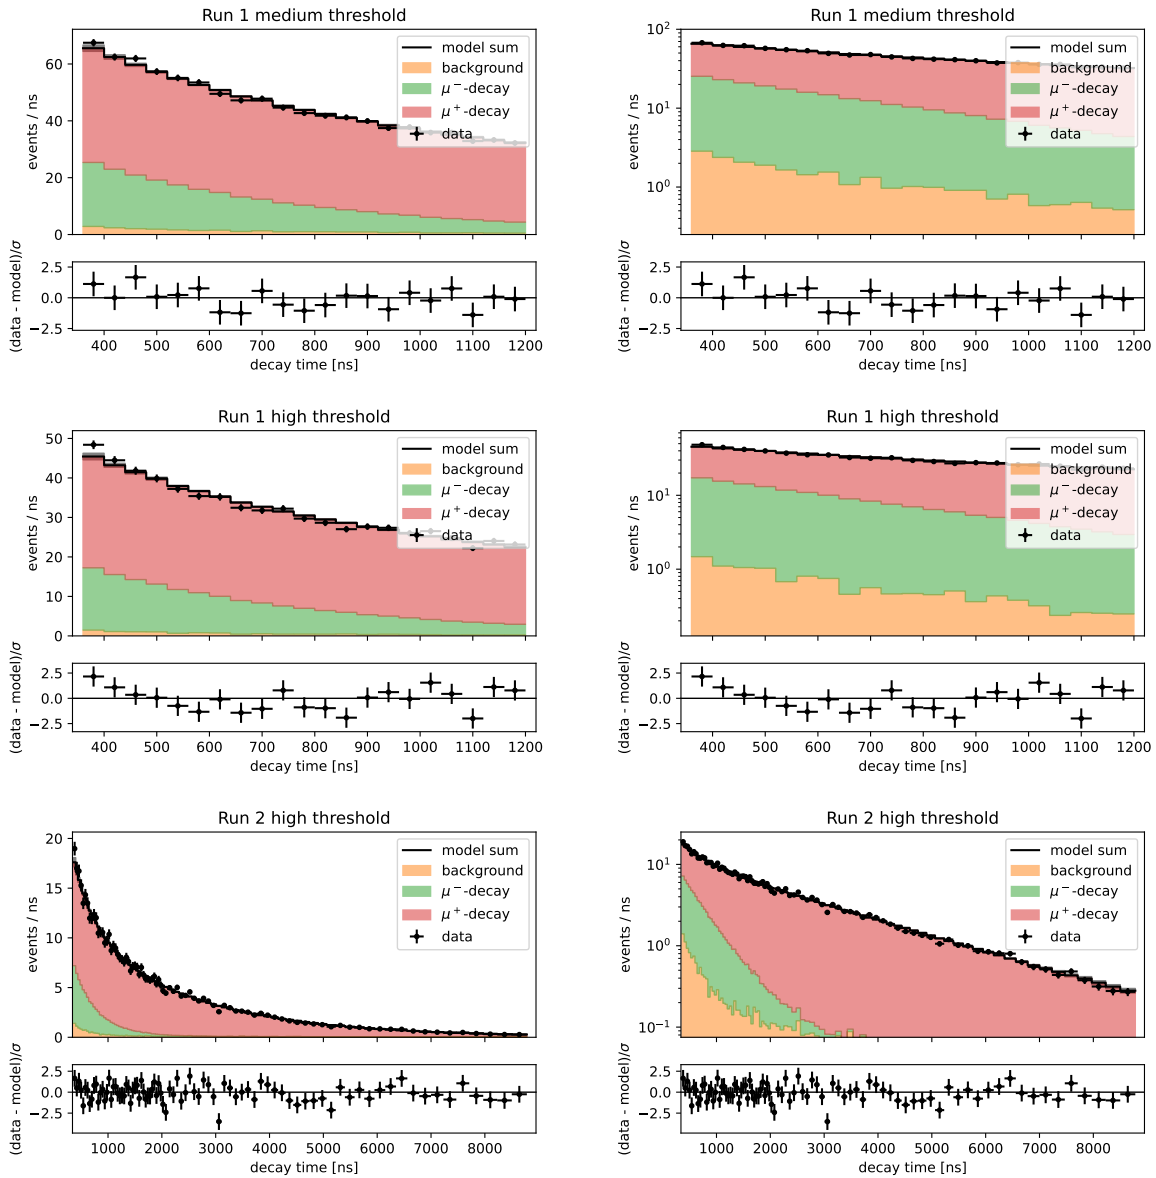


Figure 6.4: Post-fit muon-decay time distributions for each data sample. The overall χ^2/ndf for the fit, including the muon-decay n_{hit} distributions, is 229.8/237.

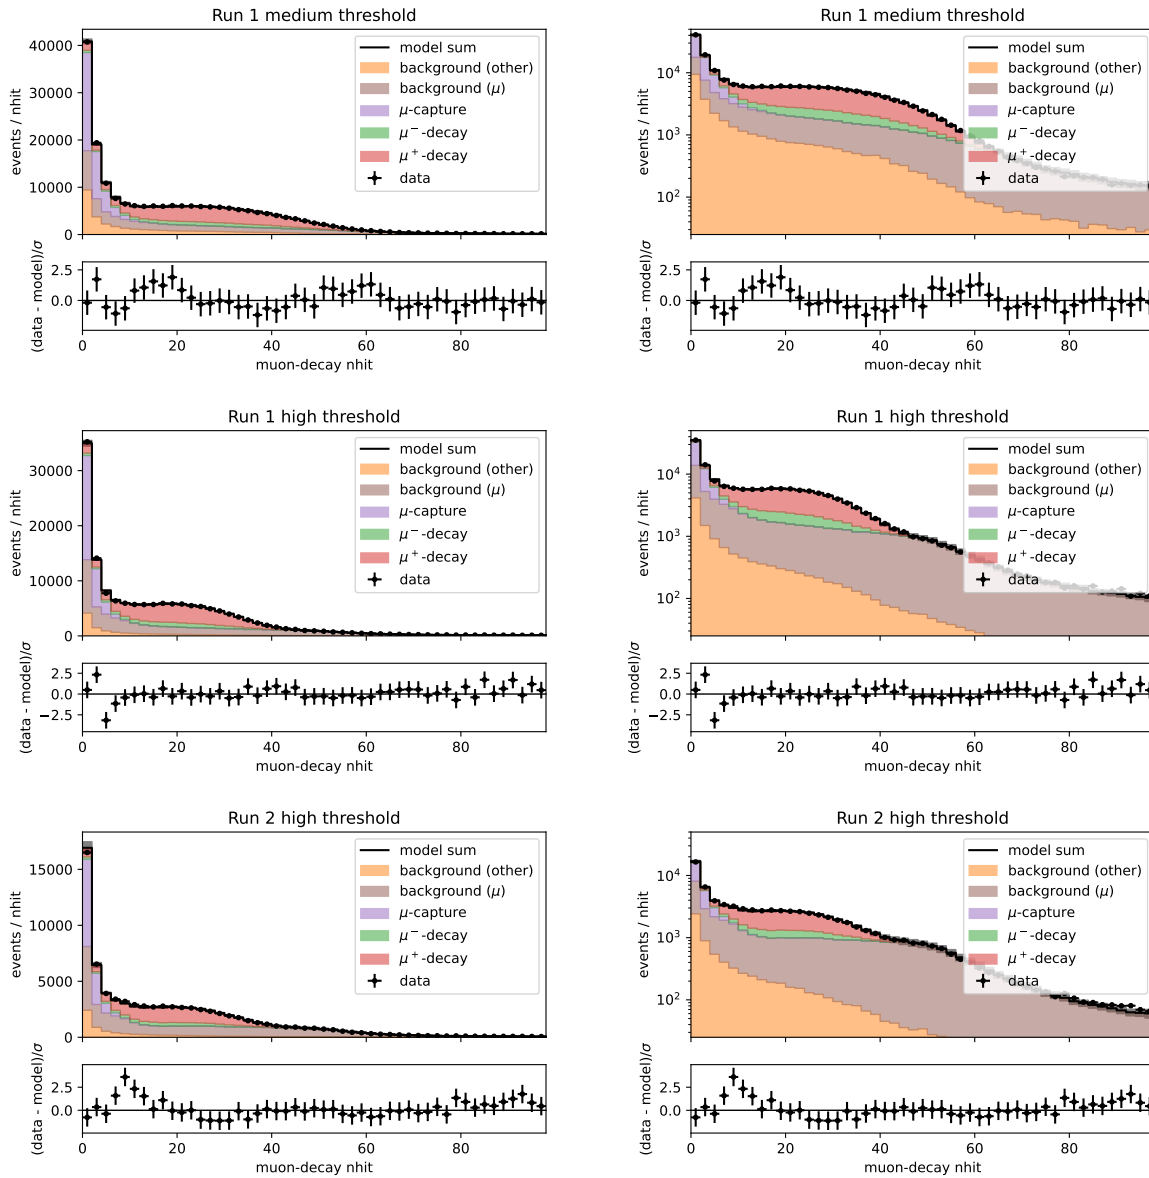


Figure 6.5: Post-fit muon-decay nhit distributions for each data sample. The overall χ^2/ndf of the fit, including the muon-decay time distributions, is 229.8/237.

Table 6.11: Fractional covariance matrix C_{rel} of parameters of interest.

	λ_d	R_H	λ_c	μ^+/μ^-
λ_d	0.0129	0.0107	0.0138	0.0014
R_H	0.0107	0.0142	0.0093	0.0022
λ_c	0.0138	0.0093	0.0156	0.0012
μ^+/μ^-	0.0014	0.0022	0.0012	0.0024

Table 6.12: Fractional covariance matrix C_{rel} of parameters of interest with a λ_d constraint from Ref. [95].

	λ_d	R_H	λ_c	μ^+/μ^-
λ_d	0.0000	-0.0001	-0.0000	0.0001
R_H	-0.0001	0.0057	-0.0001	-0.0023
λ_c	-0.0000	-0.0001	0.0009	-0.0004
μ^+/μ^-	0.0001	-0.0023	-0.0004	0.0028

Table 6.13: Fractional covariance matrix C_{rel} of parameters of interest with a λ_d constraint from Ref. [96].

	λ_d	R_H	λ_c	μ^+/μ^-
λ_d	0.0000	-0.0001	0.0000	0.0001
R_H	-0.0001	0.0059	-0.0001	-0.0026
λ_c	0.0000	-0.0001	0.0010	-0.0003
μ^+/μ^-	0.0001	-0.0026	-0.0003	0.0029

Tables 6.11-6.14 contain the fractional covariance of the four parameters of interest, where the fractional covariance is defined as

$$C_{\text{rel}} = \left(\vec{\theta}^T\right)^{-1} C \left(\vec{\theta}\right)^{-1} \quad (6.23)$$

and C is the standard covariance matrix. This form is unitless and thus has a more straightforward comparison between the elements of the matrix. Overall, the unconstrained results for the parameters of interest are highly correlated. The constrained results are also somewhat correlated but with a confounding factor removed by the constraint.

Because of the relatively small size of the Module 0 detector, only muons entering the detector with a kinetic energy of less than about 300 MeV will come to rest. The overburden and self-shielding of the liquid argon introduces approximately 350 MeV of energy loss between the nominal flux and the muons observed by the detector. Thus, the Module 0 detector is most sensitive to the low-energy distribution of the muon spectrum, between about 400 and 700 MeV. The best-fit results for the charge ratio is listed in Table 6.15. The

Table 6.14: Fractional covariance matrix C_{rel} of parameters of interest with a $R_H = 0.99$ [63] constraint.

	λ_d	R_H	λ_c	μ^+/μ^-
λ_d	0.0036	0.0000	0.0035	0.0028
R_H	0.0000	0.0000	0.0000	0.0000
λ_c	0.0035	0.0000	0.0110	-0.0006
μ^+/μ^-	0.0028	0.0000	-0.0006	0.0070

results of three of the alternative fits using constraints from theory and other experiments are also provided for comparison. Overall, the charge-ratio measurement shows relatively little sensitivity to the muon-capture rate or Huff factor constraints.

Table 6.15: Summary of the best-fit μ^+/μ^- ratio under different constraints.

Constraint	μ^+/μ^-
None	$1.22 \pm 0.04(\text{stat.}) \pm 0.04(\text{sys.})$
$\lambda_d = 1.76 \pm 0.02 \mu\text{s}^{-1}$ [95]	$1.18 \pm 0.03(\text{stat.}) \pm 0.05(\text{sys.})$
$\lambda_d = 1.62 \pm 0.02 \mu\text{s}^{-1}$ [96]	$1.16 \pm 0.04(\text{stat.}) \pm 0.05(\text{sys.})$
$R_H = 0.99$ [63]	$1.21 \pm 0.05(\text{stat.}) \pm 0.09(\text{sys.})$

Because our result is sensitive to both the muon-capture branching ratio as well as the disappearance time, it is possible to constrain the Huff factor ($\equiv \lambda_{\text{DIO}}/\lambda_{\text{DAR}}$) without a degeneracy related to the capture rate. Table 6.16 lists the value for the Huff factor in argon under different experimental constraints. As the Huff factor is a relatively small effect and is highly correlated with the disappearance rate, we have only modest sensitivity ($\approx 15\%$). Interestingly, our unconstrained fit suggests an enhancement in the muon DIO rate for argon over the theoretical calculations of 0.99 [63] and 1.1 [60]. Because the muon disappearance rate and the Huff factor are highly correlated within the fit, we gain considerable sensitivity to the Huff factor by applying a constraint on the disappearance rate from the higher-precision results of Refs. [95, 96]. We find that with this additional constraint, our results still show a slight excess, but more in line with the theoretical calculations.

Tables 6.17 and 6.18 list the best-fit of the disappearance and capture rates found in this analysis. As we are primarily sensitive to the muon disappearance rate through the lower-statistics muon decay-time spectrum, the uncertainty on the disappearance rate without a Huff factor assumption is fairly large. Including an assumption on the Huff factor using the calculations of Ref. [63], the fit settles to a value which slightly prefers the results of Ref. [96]. For the capture rate, a similar pattern follows unsurprisingly because the disappearance rate and capture rates are highly correlated.

Table 6.16: Summary of the best-fit Huff factor, R_H , under different constraints.

Constraint	R_H
None	$1.29 \pm 0.09(\text{stat.}) \pm 0.12(\text{sys.})$
$\lambda_d = 1.76 \pm 0.02 \mu\text{s}^{-1}$ [95]	$1.13 \pm 0.04(\text{stat.}) \pm 0.08(\text{sys.})$
$\lambda_d = 1.62 \pm 0.02 \mu\text{s}^{-1}$ [96]	$1.07 \pm 0.04(\text{stat.}) \pm 0.07(\text{sys.})$

Table 6.17: Summary of the best-fit disappearance rate, λ_d , under different constraints.

Constraint	$\lambda_d (\mu\text{s}^{-1})$
None	$2.11 \pm 0.18(\text{stat.}) \pm 0.16(\text{sys.})$
$\lambda_d = 1.76 \pm 0.02 \mu\text{s}^{-1}$ [95]	$1.77 \pm 0.01(\text{stat.}) \pm 0.01(\text{sys.})$
$\lambda_d = 1.62 \pm 0.02 \mu\text{s}^{-1}$ [96]	$1.64 \pm 0.01(\text{stat.}) \pm 0.01(\text{sys.})$
$R_H = 0.99$ [63]	$1.66 \pm 0.08(\text{stat.}) \pm 0.05(\text{sys.})$

Table 6.18: Summary of the best-fit capture rate, λ_c , under different constraints.

Constraint	$\lambda_c (\mu\text{s}^{-1})$
None	$1.53 \pm 0.14(\text{stat.}) \pm 0.13(\text{sys.})$
$\lambda_d = 1.76 \pm 0.02 \mu\text{s}^{-1}$ [95]	$1.25 \pm 0.02(\text{stat.}) \pm 0.03(\text{sys.})$
$\lambda_d = 1.62 \pm 0.02 \mu\text{s}^{-1}$ [96]	$1.15 \pm 0.02(\text{stat.}) \pm 0.03(\text{sys.})$
$R_H = 0.99$ [63]	$1.22 \pm 0.08(\text{stat.}) \pm 0.10(\text{sys.})$

Chapter 7

Conclusions

7.1 Summary of results

This thesis has presented an analysis of muon capture in argon using a prototype pixelated LArTPC, Module 0. This detector was operated during two run periods of 1 week each in the spring and summer of 2021, during which about 45M cosmic-ray events were collected under normal run conditions. From this dataset, a sample of approximately 625,000 stopping-muon events were selected with an estimated total efficiency of $\approx 50\%$ and purity of $\approx 60\%$. Two metrics, the *muon-decay nhit* and the *muon-decay time*, were computed for events in the sample, distinguishing muon-capture events from muon-decay events. The distribution of these metrics were then incorporated into a maximum-likelihood fit. To determine the charge-ratio of muons in the sample, a sub-sample of approximately 100,000 events excluding the muon-capture events was also defined (efficiency $\approx 15\%$, purity $\approx 95\%$) and provided to the fit. It was then possible to simultaneously estimate the cosmic-ray muon charge-ratio ($R = 1.22 \pm 0.06$), the Huff factor ($R_H = 1.29 \pm 0.15$), and the muon disappearance rate ($\lambda_d = 2.11 \pm 0.24 \mu\text{s}^{-1}$). From these, one can derive the muon capture rate ($\lambda_c = 1.53 \pm 0.19 \mu\text{s}^{-1}$). Constraints from existing measurements of the muon disappearance rate and theoretical calculations of the Huff factor were also investigated, providing more precise determinations of the parameters under different scenarios.

7.1.1 Global comparison

7.1.1.1 Cosmic-muon charge ratio

Applying cuts to the muon CSDA range corresponding to atmospheric-muon energies of 400-450 MeV, 450-550 MeV, and 550-650 MeV, the fit was repeated to determine the muon charge ratio as a function of energy, shown in Figure 7.1. Overall, our results are commensurate with the results of the BESS and CAPRICE experiments at low energies. As the latitude of Bern, Switzerland (46.9°) falls between those of Tsukuba, Japan (36.0°) and Lynn Lake, Canada (56.9°), it is not surprising that the charge ratio that we measure falls between the




Figure 7.1: Summary of published experimental measurements [165, 226, 227] of the cosmic-ray muon charge ratio plotted against muon kinetic energy and compared against this work.

two results of Ref. [165]. The high-precision charge-ratio measurement of CMS experiment saw a much higher muon energy spectrum and is not directly comparable to our result. But, it was performed at a location with similar geomagnetic characteristics and so is included as a reference for higher energies.

7.1.1.2 Huff factor

The Huff factor has not previously been measured for muon decays in argon. Comparing with previous measurements on other nuclei, Fig. 7.2, we find that our result is substantially larger. The unconstrained fit is also in tension with theoretical calculations of the Huff factor [60, 63]. However, with an additional constraint on the muon disappearance rate, the results are brought to be in-line with the theoretical prediction. This suggests that there may be additional, unaccounted-for systematic effects that are impacting our measurement of the Huff factor, which are discussed later.

7.1.1.3 Disappearance and capture rates

The disappearance rate has been measured before using beam-based measurements and, more recently, in a LArTPC, using the muon decay time spectrum. The analysis presented in this thesis uses an approach that makes use of the branching fraction of muon capture relative to muon decay. Thus, our results are complimentary to what has been performed before. Fig. 7.3 compares the results of our muon disappearance rate with those of previous experiments. Our unconstrained fit shows some tension with the high-precision results of Ref. [95, 96], though this tension is resolved using the theoretical Huff factor calculation of




Figure 7.2: Comparison of the measured Huff factor in different nuclei, Refs. [64, 71], and theoretical calculations, Refs. [60, 63]. The results of this work using different disappearance rate constraints have been offset from $Z = 18$ for clarity.




Figure 7.3: Global summary of measured μ^- disappearance rate for argon. The results from this work are compared to the experimental measurements described in Refs. [93, 94, 96, 97].

Refs. [63]. Unfortunately, our result does not have sufficient power to resolve the discrepancy between Ref. [95] and [96], though Ref. [96] is slightly preferred.

Finally, we show the muon capture rate compared with previous measurements in Fig. 7.4. Here, an assumption for the Huff factor must be made to interpret the other experimental results. A value of $R_H = 1.04 \pm 0.05$ is chosen, equally favoring the calculations of Refs. [60,

63]. The previous experiments measured λ_d only and thus

$$\lambda_c = \lambda_d - R_H \lambda_0 \quad (7.1)$$

is used to interpret their results. Overall, our results are commensurate with existing measurements. Our unconstrained result favors a slight enhancement in the muon-capture rate over the Goulard-Primakoff formula without a Huff factor constraint, but is consistent with the Goulard-Primakoff formula for $R_H = 0.99$ [63].




Figure 7.4: Global summary of the measured μ^- capture rate on argon. The results from this work are compared to the experimental measurements described in Refs. [93, 94, 96, 97] and the Goulard-Primakoff formula [56]. A Huff factor of 1.04 ± 0.05 was assumed in interpreting the previous measurements of the μ^- disappearance rate as a μ^- capture rate. The error bars were increased accordingly, as indicated.

7.2 Future work

This thesis has explored an extensive list of systematic factors that might influence the measurement of the parameters of interest. However, there is always the possibility that there is an unknown or unexplored factor that impacts the results. The tension in the Huff factor with other experimental measurements and theoretical calculations that goes away with a constraint on the muon disappearance rate could point to a missing systematic in either the decay-time spectrum or the muon-decay nhit distributions. As the muon disappearance rate is primarily constrained by the slope of the decay-time spectrum, a systematic associated with the LRS might be suspected. An enhancement in the number of events near 300 ns would cause an increase in the disappearance rate. In this vein, we notice that both of the high-threshold samples (from Run-1 and Run-2) both show a small excess in the bins near the start of the decay window. While the decay window was varied and included as one of the systematics, the correlated excess may hint at a missing background contribution associated with the change in CRS threshold. Also related to the LRS, the model used to subtract the contribution from the muon-associated backgrounds was extracted under the assumption that the sideband sample adequately reproduced the background distribution in the signal region. While the simulation suggests that the background model extracted in this fashion is reliable, there are enough differences between the light signals in the data and simulation to doubt that the 1:1 comparison is robust. The SiPMs used in Module 0 have also only been moderately calibrated and detailed measurements of cross-talk and afterpulsing have not been performed in liquid argon. The SiPM model used in the simulation for this analysis is correspondingly simplistic and may not accurately reproduce the detector response.

An alternative consideration is that apparent mis-modeling in the muon-decay nhit distribution near ~ 10 hits might be indicative of more significant issues. There are multiple ways that one could produce a model difference that is observed. Naively, the region that shows problems has roughly equal contributions from muon capture, muon decay, and backgrounds so any missing modification to the detector response to these event classes could result in a shape difference in ~ 10 hits. Alternatively, a relatively simplistic Gaussian-kernel smearing nuisance parameter was included to account for potential shape differences in the muon-decay nhit distribution. However, as the distribution is clearly asymmetric around zero, a more sophisticated model may be required. This approximation has an out-sized effect on the capture distribution, potentially leading to a bias on the Huff factor. A more problematic issue would be an unaccounted-for background that appears near 0. Due to the limited number of bins that constrain the muon-capture events, it would be relatively easy for the background to introduce a larger bias than indicated solely by the discrepancy near ~ 10 hits, as the best-fit value for the muon-capture rate would absorb these events without producing a discrepancy in the χ^2 . At this point, there is no known background that would cause this effect.

To further investigate the possibility of a missing systematic, additional cross-checks are likely warranted. Studies varying the cut parameters were performed, but it could be informative to investigate other sub-samples and sidebands in more detail. As an example,

the fit could be performed using a selection of only events that stop in TPC 1 (or TPC 2) or using muons with a particular orientation. This study is complicated by the fact that the muon-decay n_{hit} distribution changes as a function of the detector acceptance, as evidenced by the differences in the Run 2 sample. Thus one must also take care that these sub-samples do not contain additional biases that were not present in the original sample. Perhaps a less-fraught approach would be to examine reconstructed variable distributions of the sub-samples and sidebands using event reweighting, matched to the post-fit results.

Alternatively, since the operation of Module 0, there have been 4 additional runs using nearly identical detectors in preparation for the 2×2 program at Fermilab [228]. While the cumulative dataset from the operation of these detectors is likely only at the same order as that collected during the Module 0 runs, differences in the detector configuration may lead to improvements in the systematic uncertainties. In particular, one module is instrumented with a pixel pitch of 3.8 mm, which would modify the pixel response, reducing the importance of the far-field effects. A comparison between these datasets could provide valuable insight into the pixel response. Additionally, improvements in the LRS electronics have also reduced the pickup and noise in the readout, leading to higher-fidelity light reconstruction. It would be interesting to see how this translates into a lower threshold (and thus lower bias) in the selection of muon-decay events.

For future experiments investigating muon capture in argon, a configurable μ^\pm beam is obviously preferred over natural cosmic rays. This analysis has the undesirable property that about 75% of the included events come from backgrounds of either μ^+ or through-going muons. While the charge ratio of muon decays can be determined using the decay-time spectrum, the statistical uncertainties in this analysis are ultimately limited by this determination. Increasing the exposure is certainly a possibility, as the near-zero deadtime of Module 0 enabled the collection of $> 45\text{M}$ cosmic-ray events in just a few days. But it would be unrealistic to improve the result beyond $\sim 1\%$ with a detector of the same size. However, more important for an improved measurement will be to reduce other backgrounds, particularly those coming from through-going muons. Either a beam-based measurement, which could reduce backgrounds by isolating muons with a narrow band of energies, or a larger detector, which could reduce backgrounds by increasing the fiducial boundary, could accomplish this.

Bibliography

- [1] Charles Drummond Ellis and W. A. Wooster. “The average energy of disintegration of radium E”. In: *Proc. Roy. Soc. Lond. A* 117.776 (1927), pp. 109–123. DOI: 10.1098/rspa.1927.0168.
- [2] Charles D. Ellis and W. A. Wooster. “The continuous spectrum of β -rays”. In: *Nature* 119.2998 (1927), pp. 563–564.
- [3] Paul A. M. Dirac. “The quantum theory of the electron”. In: *Proc. Roy. Soc. Lond. A* 117 (1928), pp. 610–624. DOI: 10.1098/rspa.1928.0023.
- [4] J. Chadwick. “Possible Existence of a Neutron”. In: *Nature* 129 (1932), p. 312. DOI: 10.1038/129312a0.
- [5] E. Fermi. “Trends to a Theory of beta Radiation. (In Italian)”. In: *Nuovo Cim.* 11 (1934), pp. 1–19. DOI: 10.1007/BF02959820.
- [6] E. J. Williams and G. E. Roberts. “Evidence for Transformation of Mesotrons into Electrons”. In: *Nature* 145.3664 (1940), pp. 102–103.
- [7] Viśvapriya Mukherji. “A history of the meson theory of nuclear forces from 1935 to 1952”. In: *Archive for History of Exact Sciences* (1974), pp. 27–102.
- [8] Hideki Yukawa. “On the Interaction of Elementary Particles I”. In: *Proc. Phys. Math. Soc. Jap.* 17 (1935), pp. 48–57. DOI: 10.1143/PTPS.1.1.
- [9] S. H. Neddermeyer and C. D. Anderson. “Note on the Nature of Cosmic Ray Particles”. In: *Phys. Rev.* 51 (1937), pp. 884–886. DOI: 10.1103/PhysRev.51.884.
- [10] Paul Ehrenfest and André Fréon. “Désintégration spontanée des «mésotons», particules composant le rayonnement cosmique pénétrant”. In: *Journal de Physique et le Radium* 9.12 (1938), pp. 529–536.
- [11] Thomas H. Johnson and Martin A. Pomerantz. “The difference in the absorption of cosmic rays in air and water and the instability of the barytron”. In: *Physical Review* 55.1 (1939), p. 104.
- [12] C. G. Montgomery et al. “Slow mesons in the cosmic radiation”. In: *Physical Review* 56.7 (1939), p. 635.
- [13] S. Tomonaga and G. Araki. “Effect of the Nuclear Coulomb Field on the Capture of Slow Mesons”. In: *Physical Review* 58.1 (1940), p. 90.

- [14] M. Conversi, E. Pancini, and O. Piccioni. “On the Decay Process of Positive and Negative Mesons”. In: *Phys. Rev.* 68.9-10 (1945), p. 232. DOI: 10.1103/PhysRev.68.232.
- [15] M. Conversi, E. Pancini, and O. Piccioni. “On the Disintegration of Negative Mesons”. In: *Phys. Rev.* 71 (1947), pp. 209–210. DOI: 10.1103/PhysRev.71.209.
- [16] C. M. G. Lattes et al. “PROCESSES INVOLVING CHARGED MESONS”. In: *Nature* 159 (1947), pp. 694–697. DOI: 10.1038/159694a0.
- [17] Louis Michel. “Energy Spectrum of Secondary Electrons from $\hat{1}1/4$ -Meson Decay”. In: *Nature* 163.4155 (1949), pp. 959–960.
- [18] J. Steinberger. “On the Range of the Electrons in Meson Decay”. In: *Phys. Rev.* 74 (1948), pp. 500–501. DOI: 10.1103/PhysRev.74.500.3.
- [19] E. P. Hincks and B. Pontecorvo. “The absorption of charged particles from the 2.2-mu-sec meson decay”. In: *Phys. Rev.* 74 (1948), pp. 697–698. DOI: 10.1103/PhysRev.74.697.
- [20] E. P. Hincks and B. Pontecorvo. “Search for gamma-radiation in the 2.2-microsecond meson decay process”. In: *Phys. Rev.* 73 (1948), pp. 257–258. DOI: 10.1103/PhysRev.73.257.
- [21] T. D. Lee and Chen-Ning Yang. “Question of Parity Conservation in Weak Interactions”. In: *Phys. Rev.* 104 (1956), pp. 254–258. DOI: 10.1103/PhysRev.104.254.
- [22] C. S. Wu et al. “Experimental Test of Parity Conservation in β Decay”. In: *Phys. Rev.* 105 (1957), pp. 1413–1414. DOI: 10.1103/PhysRev.105.1413.
- [23] M. Goldhaber, L. Grodzins, and A. W. Sunyar. “Helicity of Neutrinos”. In: *Phys. Rev.* 109 (1958), pp. 1015–1017. DOI: 10.1103/PhysRev.109.1015.
- [24] R. L. Workman et al. “Review of Particle Physics”. In: *PTEP* 2022 (2022), p. 083C01. DOI: 10.1093/ptep/ptac097.
- [25] Walter F. Dudziak, Ryokichi Sagane, and James Vedder. “Positron Spectrum from the Decay of the mu Meson”. In: *Phys. Rev.* 114 (1959), pp. 336–358. DOI: 10.1103/PhysRev.114.336.
- [26] Richard J. Plano. “Momentum and asymmetry spectrum of mu-meson decay”. In: *Phys. Rev.* 119 (1960), pp. 1400–1408. DOI: 10.1103/PhysRev.119.1400.
- [27] S. E. Derenzo. “Measurement of the low-energy end of the mu-plus decay spectrum”. In: *Phys. Rev.* 181 (1969), pp. 1854–1866. DOI: 10.1103/PhysRev.181.1854.
- [28] A. Buhler et al. “A Measurement of the e^+ Polarization in Muon Decay: The e^+ Annihilation Method”. In: *Phys. Lett.* 7 (1963), pp. 368–371. DOI: 10.1016/0031-9163(63)90074-X.
- [29] F. J. Hasert et al. “Observation of Neutrino Like Interactions without Muon or Electron in the Gargamelle Neutrino Experiment”. In: *Nucl. Phys. B* 73 (1974), pp. 1–22. DOI: 10.1016/0550-3213(74)90038-8.

- [30] G. Arnison et al. “Experimental Observation of Lepton Pairs of Invariant Mass Around 95-GeV/c² at the CERN SPS Collider”. In: *Phys. Lett. B* 126 (1983), pp. 398–410. DOI: 10.1016/0370-2693(83)90188-0.
- [31] P. Bagnaia et al. “Evidence for $Z^0 \rightarrow e^+e^-$ at the CERN $\bar{p}p$ Collider”. In: *Phys. Lett. B* 129 (1983), pp. 130–140. DOI: 10.1016/0370-2693(83)90744-X.
- [32] G. Arnison et al. “Further Evidence for Charged Intermediate Vector Bosons at the SPS Collider”. In: *Phys. Lett. B* 129 (1983), pp. 273–282. DOI: 10.1016/0370-2693(83)90860-2.
- [33] M. Banner et al. “Observation of Single Isolated Electrons of High Transverse Momentum in Events with Missing Transverse Energy at the CERN anti-p p Collider”. In: *Phys. Lett. B* 122 (1983), pp. 476–485. DOI: 10.1016/0370-2693(83)91605-2.
- [34] Steven Weinberg. “A Model of Leptons”. In: *Phys. Rev. Lett.* 19 (1967), pp. 1264–1266. DOI: 10.1103/PhysRevLett.19.1264.
- [35] Serguei Chatrchyan et al. “Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC”. In: *Phys. Lett. B* 716 (2012), pp. 30–61. DOI: 10.1016/j.physletb.2012.08.021. arXiv: 1207.7235 [hep-ex].
- [36] Georges Aad et al. “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”. In: *Phys. Lett. B* 716 (2012), pp. 1–29. DOI: 10.1016/j.physletb.2012.08.020. arXiv: 1207.7214 [hep-ex].
- [37] Matthew D. Schwartz. *Quantum Field Theory and the Standard Model*. Cambridge University Press, Mar. 2014. ISBN: 978-1-107-03473-0.
- [38] Peter J. Mohr, David B. Newell, and Barry N. Taylor. “CODATA Recommended Values of the Fundamental Physical Constants: 2014”. In: *Rev. Mod. Phys.* 88.3 (2016), p. 035009. DOI: 10.1103/RevModPhys.88.035009. arXiv: 1507.07956 [physics.atom-ph].
- [39] A. M. Baldini et al. “Search for the lepton flavour violating decay $\mu^+ \rightarrow e^+\gamma$ with the full dataset of the MEG experiment”. In: *Eur. Phys. J. C* 76.8 (2016), p. 434. DOI: 10.1140/epjc/s10052-016-4271-x. arXiv: 1605.05081 [hep-ex].
- [40] Wilhelm H. Bertl et al. “Search for the Decay $\mu^+ \rightarrow e^+e^+e^-$ ”. In: *Nucl. Phys. B* 260 (1985), pp. 1–31. DOI: 10.1016/0550-3213(85)90308-6.
- [41] T. Albahri et al. “Measurement of the anomalous precession frequency of the muon in the Fermilab Muon g_2 Experiment”. In: *Phys. Rev. D* 103.7 (2021), p. 072002. DOI: 10.1103/PhysRevD.103.072002. arXiv: 2104.03247 [hep-ex].
- [42] A. Hillairet et al. “Precision muon decay measurements and improved constraints on the weak interaction”. In: *Phys. Rev. D* 85 (2012), p. 092013. DOI: 10.1103/PhysRevD.85.092013. arXiv: 1112.3606 [hep-ex].

- [43] N. Danneberg et al. “Muon decay: Measurement of the transverse polarization of the decay positrons and its implications for the Fermi coupling constant and time reversal invariance”. In: *Phys. Rev. Lett.* 94 (2005), p. 021802. DOI: 10.1103/PhysRevLett.94.021802.
- [44] T. Kinoshita and M. Nio. “Improved theory of the muonium hyperfine structure”. In: *Phys. Rev. Lett.* 72 (1994), pp. 3803–3806. DOI: 10.1103/PhysRevLett.72.3803. arXiv: hep-ph/9402260.
- [45] P. Vogel et al. “Muon Capture in Atoms, Crystals and Molecules”. In: *Nucl. Phys. A* 254 (1975), pp. 445–479. DOI: 10.1016/0375-9474(75)90230-4.
- [46] P. Vogel. “Energies of Electronic K x-Rays of Muonic Atoms”. In: *Phys. Lett. B* 58 (1975), pp. 52–54. DOI: 10.1016/0370-2693(75)90725-X.
- [47] D. F. Measday. “The nuclear physics of muon capture”. In: *Phys. Rept.* 354 (2001), pp. 243–409. DOI: 10.1016/S0370-1573(01)00012-6.
- [48] R. P. Feynman and Murray Gell-Mann. “Theory of Fermi interaction”. In: *Phys. Rev.* 109 (1958). Ed. by L. M. Brown, pp. 193–198. DOI: 10.1103/PhysRev.109.193.
- [49] Denys H. Wilkinson. “Limits to second-class nucleonic and mesonic currents”. In: *Eur. Phys. J. A* 7 (2000), pp. 307–315. DOI: 10.1007/s100500050397.
- [50] Franz Osterfeld. “Nuclear spin and isospin excitations”. In: *Rev. Mod. Phys.* 64 (1992), pp. 491–558. DOI: 10.1103/RevModPhys.64.491.
- [51] Andrzej Czarnecki, William J. Marciano, and Alberto Sirlin. “Neutron Lifetime and Axial Coupling Connection”. In: *Phys. Rev. Lett.* 120.20 (2018), p. 202002. DOI: 10.1103/PhysRevLett.120.202002. arXiv: 1802.01804 [hep-ph].
- [52] T. Cai et al. “Measurement of the axial vector form factor from antineutrino–proton scattering”. In: *Nature* 614.7946 (2023), pp. 48–53. DOI: 10.1038/s41586-022-05478-3.
- [53] V. A. Andreev et al. “Measurement of Muon Capture on the Proton to 1% Precision and Determination of the Pseudoscalar Coupling g_P ”. In: *Phys. Rev. Lett.* 110.1 (2013), p. 012504. DOI: 10.1103/PhysRevLett.110.012504. arXiv: 1210.6545 [nucl-ex].
- [54] H. Primakoff. “Theory of muon capture”. In: *Rev. Mod. Phys.* 31 (1959), pp. 802–822. DOI: 10.1103/RevModPhys.31.802.
- [55] B. Goulard and H. Primakoff. “Nuclear muon-capture sum rules and mean nuclear excitation energies”. In: *Phys. Rev. C* 10 (1974), pp. 2034–2044. DOI: 10.1103/PhysRevC.10.2034.
- [56] T. Suzuki, David F. Measday, and J. P. Roalsvig. “Total Nuclear Capture Rates for Negative Muons”. In: *Phys. Rev. C* 35 (1987), p. 2212. DOI: 10.1103/PhysRevC.35.2212.

- [57] A. Lovato, N. Rocco, and R. Schiavilla. “Muon capture in nuclei: An ab initio approach based on Green’s function Monte Carlo methods”. In: *Phys. Rev. C* 100.3 (2019), p. 035502. DOI: 10.1103/PhysRevC.100.035502. arXiv: 1903.08078 [nucl-th].
- [58] Nikolaj Thomas Zinner, Karlheinz Langanke, and Petr Vogel. “Muon capture on nuclei: Random phase approximation evaluation versus data for $6 \leq Z \leq 94$ nuclei”. In: *Phys. Rev. C* 74 (2006), p. 024326. DOI: 10.1103/PhysRevC.74.024326. arXiv: nucl-th/0606002.
- [59] T. Marketin et al. “Relativistic QRPA calculation of muon capture rates”. In: *Phys. Rev. C* 79 (2009), p. 054323. DOI: 10.1103/PhysRevC.79.054323. arXiv: 0812.1947 [nucl-th].
- [60] L. Chatterjee, T. Roy, and S. Bhattacharya. “MU+ DECAY FROM MUONIUM”. In: *Acta Phys. Polon. B* 11 (1980), pp. 635–640.
- [61] H. C. Von Baeyer and D. Leiter. “NUCLEAR RECOIL CORRECTIONS TO THE LIFETIME OF BOUND MUONS”. In: *Phys. Rev. A* 19 (1979), pp. 1371–1374. DOI: 10.1103/PhysRevA.19.1371.
- [62] Victor Gilinsky. “The decay of bound muons”. PhD thesis. Caltech, 1961.
- [63] Robert W. Huff. “Decay rate of bound muons”. In: *Annals of Physics* 16.2 (1961), pp. 288–317.
- [64] Drasko D. Yovanovitch. “Decay rates of bound negative muons”. In: *Physical Review* 117.6 (1960), p. 1580.
- [65] H. Uberall. “Decay of mu- Mesons Bound in the K Shell of Light Nuclei”. In: *Phys. Rev.* 119 (1960), pp. 365–376. DOI: 10.1103/PhysRev.119.365.
- [66] K. W. Ford and J. G. Wills. “Muonic atoms and the radial shape of the nuclear charge distribution”. In: *Phys. Rev.* 185 (1969), pp. 1429–1438. DOI: 10.1103/PhysRev.185.1429.
- [67] C. E. Porter and H. Primakofe. “The Effect of Bohr Orbit Binding on Negative mu-Meson beta-Decay”. In: *Phys. Rev.* 83 (1951), pp. 849–850. DOI: 10.1103/PhysRev.83.849.
- [68] L. M. Lederman and M. Weinrich. “Lifetime of negative muons in various materials”. In: (1956). DOI: 10.5170/CERN-1956-026.427. URL: <https://cds.cern.ch/record/1241713>.
- [69] W. A. Barrett, F. E. Holmstrom, and J. W. Keuffel. “Capture and Decay of μ - Mesons in Fe”. In: *Physical Review* 113.2 (1959), p. 661.
- [70] G. Culligan et al. “The decay rate of negative muons bound in K-orbits around iron nuclei”. In: *1st Aix en Provence International conference on Elementary Particles*. Vol. Vol.1. 1962, pp. 143–146.

- [71] IM Blair et al. “THE EFFECT OF ATOMIC BINDING ON THE DECAY RATE OF NEGATIVE MUONS”. In: *Proc. Phys. Soc.(London)*. Vol. 80. Univ. of Liverpool. 1962.
- [72] Andrzej Czarnecki et al. “Michel decay spectrum for a muon bound to a nucleus”. In: *Phys. Rev. D* 90.9 (2014), p. 093002. DOI: 10.1103/PhysRevD.90.093002. arXiv: 1406.3575 [hep-ph].
- [73] Wilhelm H. Bertl et al. “A Search for muon to electron conversion in muonic gold”. In: *Eur. Phys. J. C* 47 (2006), pp. 337–346. DOI: 10.1140/epjc/s2006-02582-x.
- [74] W. Honecker et al. “Improved limit on the branching ratio of $\mu \rightarrow e$ conversion on lead”. In: *Phys. Rev. Lett.* 76 (1996), pp. 200–203. DOI: 10.1103/PhysRevLett.76.200.
- [75] C. Dohmen et al. “Test of lepton flavor conservation in $\mu \rightarrow e$ conversion on titanium”. In: *Phys. Lett. B* 317 (1993), pp. 631–636. DOI: 10.1016/0370-2693(93)91383-X.
- [76] A. Badertscher et al. “A Search for Muon - Electron and Muon - Positron Conversion in Sulfur”. In: *Nucl. Phys. A* 377 (1982), pp. 406–440. DOI: 10.1016/0375-9474(82)90049-5.
- [77] Y. Fukuda et al. “Evidence for oscillation of atmospheric neutrinos”. In: *Phys. Rev. Lett.* 81 (1998), pp. 1562–1567. DOI: 10.1103/PhysRevLett.81.1562. arXiv: hep-ex/9807003.
- [78] Q. R. Ahmad et al. “Direct evidence for neutrino flavor transformation from neutral current interactions in the Sudbury Neutrino Observatory”. In: *Phys. Rev. Lett.* 89 (2002), p. 011301. DOI: 10.1103/PhysRevLett.89.011301. arXiv: nucl-ex/0204008.
- [79] K. Eguchi et al. “First results from KamLAND: Evidence for reactor anti-neutrino disappearance”. In: *Phys. Rev. Lett.* 90 (2003), p. 021802. DOI: 10.1103/PhysRevLett.90.021802. arXiv: hep-ex/0212021.
- [80] A. Aguilar et al. “Evidence for neutrino oscillations from the observation of $\bar{\nu}_e$ appearance in a $\bar{\nu}_\mu$ beam”. In: *Phys. Rev. D* 64 (2001), p. 112007. DOI: 10.1103/PhysRevD.64.112007. arXiv: hep-ex/0104049.
- [81] B. Pontecorvo. “Inverse beta processes and nonconservation of lepton charge”. In: *Zh. Eksp. Teor. Fiz.* 34 (1957), p. 247.
- [82] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. “Remarks on the unified model of elementary particles”. In: *Prog. Theor. Phys.* 28 (1962), pp. 870–880. DOI: 10.1143/PTP.28.870.
- [83] Ivan Esteban et al. “The fate of hints: updated global analysis of three-flavor neutrino oscillations”. In: *JHEP* 09 (2020), p. 178. DOI: 10.1007/JHEP09(2020)178. arXiv: 2007.14792 [hep-ph].

- [84] L. Wolfenstein. “Neutrino Oscillations in Matter”. In: *Phys. Rev. D* 17 (1978), pp. 2369–2374. DOI: 10.1103/PhysRevD.17.2369.
- [85] S. P. Mikheyev and A. Yu. Smirnov. “Resonance Amplification of Oscillations in Matter and Spectroscopy of Solar Neutrinos”. In: *Sov. J. Nucl. Phys.* 42 (1985), pp. 913–917.
- [86] B. Abi et al. “Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume II: DUNE Physics”. In: (Feb. 2020). arXiv: 2002.03005 [hep-ex].
- [87] V. Hewes et al. “Deep Underground Neutrino Experiment (DUNE) Near Detector Conceptual Design Report”. In: *Instruments* 5.4 (2021), p. 31. DOI: 10.3390/instruments5040031. arXiv: 2103.13910 [physics.ins-det].
- [88] J. A. Formaggio and G. P. Zeller. “From eV to EeV: Neutrino Cross Sections Across Energy Scales”. In: *Rev. Mod. Phys.* 84 (2012), pp. 1307–1341. DOI: 10.1103/RevModPhys.84.1307. arXiv: 1305.7513 [hep-ex].
- [89] P. A. Rodrigues et al. “Identification of nuclear effects in neutrino-carbon interactions at low three-momentum transfer”. In: *Phys. Rev. Lett.* 116 (2016). [Addendum: *Phys.Rev.Lett.* 121, 209902 (2018)], p. 071802. DOI: 10.1103/PhysRevLett.116.071802. arXiv: 1511.05944 [hep-ex].
- [90] K. Abe et al. “Measurement of the ν charged-current quasielastic cross section on carbon with the ND280 detector at T2K”. In: *Phys. Rev. D* 92.11 (2015), p. 112003. DOI: 10.1103/PhysRevD.92.112003. arXiv: 1411.6264 [hep-ex].
- [91] P. Abratenko et al. “First Measurement of Differential Charged Current Quasielastic-like ν_μ -Argon Scattering Cross Sections with the MicroBooNE Detector”. In: *Phys. Rev. Lett.* 125.20 (2020), p. 201803. DOI: 10.1103/PhysRevLett.125.201803. arXiv: 2006.00108 [hep-ex].
- [92] J. Nieves, Jose Enrique Amaro, and M. Valverde. “Inclusive quasi-elastic neutrino reactions”. In: *Phys. Rev. C* 70 (2004). [Erratum: *Phys.Rev.C* 72, 019902 (2005)], p. 055503. DOI: 10.1103/PhysRevC.70.055503. arXiv: nucl-th/0408005.
- [93] A. Bertin et al. “Muon capture in gaseous deuterium”. In: *Phys. Rev. D* 8 (1973), pp. 3774–3793. DOI: 10.1103/PhysRevD.8.3774.
- [94] G. Carboni et al. “Muonic X-rays and Muon Capture in Low Pressure Argon”. In: *Phys. Lett. B* 96 (1980), pp. 206–208. DOI: 10.1016/0370-2693(80)90245-2.
- [95] T. N. Mamedov et al. “Total rates of nuclear capture of negative muons in the isotopes ^{132}Xe and ^{40}Ar ”. In: *Journal of Experimental and Theoretical Physics Letters* 69 (1999), pp. 192–195.
- [96] A. V. Klinskikh et al. “Muon capture in Ar. The muon lifetime and yields of Cl isotopes”. In: *Bulletin of the Russian Academy of Sciences: Physics* 72 (2008), pp. 735–736.

- [97] W. Foreman et al. “Calorimetry for low-energy electrons using charge and light in liquid argon”. In: *Phys. Rev. D* 101.1 (2020), p. 012010. DOI: 10.1103/PhysRevD.101.012010. arXiv: 1909.07920 [physics.ins-det].
- [98] C. Rubbia et al. “Underground operation of the ICARUS T600 LAr-TPC: first results”. In: *JINST* 6 (2011), P07011. DOI: 10.1088/1748-0221/6/07/P07011. arXiv: 1106.0975 [hep-ex].
- [99] David R. Nygren. “Proposal to investigate the feasibility of a novel concept in particle detection”. In: (1974).
- [100] Georges Charpak et al. “The Use of Multiwire Proportional Counters to Select and Localize Charged Particles”. In: *Nucl. Instrum. Meth.* 62 (1968), pp. 262–268. DOI: 10.1016/0029-554X(68)90371-6.
- [101] Alfred O. Nier. “A redetermination of the relative abundances of the isotopes of carbon, nitrogen, oxygen, argon, and potassium”. In: *Physical Review* 77.6 (1950), p. 789.
- [102] C. Tegeler, Roland Span, and Wolfgang Wagner. “A new equation of state for argon covering the fluid region for temperatures from the melting line to 700 K at pressures up to 1000 MPa”. In: *Journal of Physical and Chemical Reference Data* 28.3 (1999), pp. 779–850.
- [103] Ralph L. Amey and Robert H. Cole. “Dielectric constants of liquefied noble gases and methane”. In: *The Journal of Chemical Physics* 40.1 (1964), pp. 146–148.
- [104] A. Blatter et al. “Experimental study of electric breakdowns in liquid argon at centimeter scale”. In: *JINST* 9 (2014), P04006. DOI: 10.1088/1748-0221/9/04/P04006. arXiv: 1401.6693 [physics.ins-det].
- [105] Eido Shibamura et al. “Drift velocities of electrons, saturation characteristics of ionization and W-values for conversion electrons in liquid argon, liquid argon-gas mixtures and liquid xenon”. In: *Nucl. Instrum. Meth.* 131 (1975), pp. 249–258. DOI: 10.1016/0029-554X(75)90327-4.
- [106] M. Miyajima et al. “Average energy expended per ion pair in liquid argon”. In: *Phys. Rev. A* 9 (1974), pp. 1438–1443. DOI: 10.1103/PhysRevA.9.1438.
- [107] T. Doke et al. “Estimation of Fano factors in liquid argon, krypton, xenon and xenon-doped liquid argon”. In: *Nucl. Instrum. Meth.* 134 (1976), pp. 353–357. DOI: 10.1016/0029-554X(76)90292-5.
- [108] Yichen Li et al. “Measurement of Longitudinal Electron Diffusion in Liquid Argon”. In: *Nucl. Instrum. Meth. A* 816 (2016), pp. 160–170. DOI: 10.1016/j.nima.2016.01.094. arXiv: 1508.07059 [physics.ins-det].
- [109] P. Abratenko et al. “Measurement of the longitudinal diffusion of ionization electrons in the MicroBooNE detector”. In: *JINST* 16.09 (2021), P09025. DOI: 10.1088/1748-0221/16/09/P09025. arXiv: 2104.06551 [physics.ins-det].

- [110] Eido Shibamura et al. “Ratio of diffusion coefficient to mobility for electrons in liquid argon”. In: *Phys. Rev. A* 20.6 (1979), p. 2547. DOI: 10.1103/PhysRevA.20.2547.
- [111] Bob L. Henson. “Mobility of positive ions in liquefied argon and nitrogen”. In: *Physical Review* 135.4A (1964), A1002.
- [112] T. Heindl et al. “The scintillation of liquid argon”. In: *EPL* 91.6 (2010), p. 62002. DOI: 10.1209/0295-5075/91/62002. arXiv: 1511.07718 [physics.ins-det].
- [113] M. Babicz et al. “A measurement of the group velocity of scintillation light in liquid argon”. In: *JINST* 15.09 (2020), P09009. DOI: 10.1088/1748-0221/15/09/P09009. arXiv: 2002.09346 [physics.ins-det].
- [114] C. Rubbia. “The Liquid Argon Time Projection Chamber: A New Concept for Neutrino Detectors”. In: (May 1977).
- [115] A. Hitachi, T. Doke, and A. Mozumder. “Luminescence quenching in liquid argon under charged-particle impact: Relative scintillation yield at different linear energy transfers”. In: *Phys. Rev. B* 46.18 (1992), p. 11463. DOI: 10.1103/PhysRevB.46.11463.
- [116] Akira Hitachi. “Exciton kinetics in condensed rare gases”. In: *The Journal of chemical physics* 80.2 (1984), pp. 745–748.
- [117] Akira Hitachi et al. “Effect of ionization density on the time dependence of luminescence from liquid argon and xenon”. In: *Phys. Rev. B* 27 (1983), pp. 5279–5285. DOI: 10.1103/PhysRevB.27.5279.
- [118] M. Hofmann et al. “Ion-beam excitation of liquid argon”. In: *Eur. Phys. J. C* 73.10 (2013), p. 2618. DOI: 10.1140/epjc/s10052-013-2618-0. arXiv: 1511.07721 [physics.ins-det].
- [119] R. Acciarri et al. “Effects of Nitrogen contamination in liquid Argon”. In: *JINST* 5 (2010), P06003. DOI: 10.1088/1748-0221/5/06/P06003. arXiv: 0804.1217 [nucl-ex].
- [120] R. Acciarri et al. “Oxygen contamination in liquid Argon: Combined effects on ionization electron charge and scintillation light”. In: *JINST* 5 (2010), P05003. DOI: 10.1088/1748-0221/5/05/P05003. arXiv: 0804.1222 [nucl-ex].
- [121] A. Neumeier et al. “Attenuation of vacuum ultraviolet light in liquid argon”. In: *Eur. Phys. J. C* 72 (2012), p. 2190. DOI: 10.1140/epjc/s10052-012-2190-z. arXiv: 1511.07724 [physics.ins-det].
- [122] M. Babicz et al. “Experimental study of the propagation of scintillation light in Liquid Argon”. In: *Nucl. Instrum. Meth. A* 936 (2019). Ed. by Giovanni Batignani et al., pp. 178–179. DOI: 10.1016/j.nima.2018.10.082.
- [123] Adolf Fick. “Ueber diffusion”. In: *Annalen der Physik* 170.1 (1855), pp. 59–86.

- [124] R. Acciarri et al. “A Study of Electron Recombination Using Highly Ionizing Particles in the ArgoNeuT Liquid Argon TPC”. In: *JINST* 8 (2013), P08005. DOI: 10.1088/1748-0221/8/08/P08005. arXiv: 1306.1712 [physics.ins-det].
- [125] R. Acciarri et al. “First Observation of Low Energy Electron Neutrinos in a Liquid Argon Time Projection Chamber”. In: *Phys. Rev. D* 95.7 (2017), p. 072005. DOI: 10.1103/PhysRevD.95.072005. arXiv: 1610.04102 [hep-ex].
- [126] W. H. Lippincott et al. “Scintillation time dependence and pulse shape discrimination in liquid argon”. In: *Phys. Rev. C* 78 (2008). [Erratum: *Phys.Rev.C* 81, 039901 (2010)], p. 035801. DOI: 10.1103/PhysRevC.81.039901. arXiv: 0801.1531 [nucl-ex].
- [127] R. Brunetti et al. “Warp liquid argon detector for dark matter survey”. In: *New Astron. Rev.* 49 (2005). Ed. by D. B. Cline, pp. 265–269. DOI: 10.1016/j.newar.2005.01.017. arXiv: astro-ph/0405342.
- [128] E. Segreto. “Evidence of delayed light emission of TetraPhenyl Butadiene excited by liquid Argon scintillation light”. In: *Phys. Rev. C* 91.3 (2015), p. 035503. DOI: 10.1103/PhysRevC.91.035503. arXiv: 1411.4524 [physics.ins-det].
- [129] Niccolò Gallice. “Xenon doping of liquid argon in ProtoDUNE single phase”. In: *JINST* 17.01 (2022), p. C01034. DOI: 10.1088/1748-0221/17/01/C01034. arXiv: 2111.00347 [physics.ins-det].
- [130] P. Abratenko et al. “Measurement of space charge effects in the MicroBooNE LArTPC using cosmic muons”. In: *JINST* 15.12 (2020), P12037. DOI: 10.1088/1748-0221/15/12/P12037. arXiv: 2008.09765 [physics.ins-det].
- [131] B. Abi et al. “First results on ProtoDUNE-SP liquid argon time projection chamber performance from a beam test at the CERN Neutrino Platform”. In: *JINST* 15.12 (2020), P12004. DOI: 10.1088/1748-0221/15/12/P12004. arXiv: 2007.06722 [physics.ins-det].
- [132] A. Ereditato et al. “Design and operation of ARGONTUBE: a 5 m long drift liquid argon TPC”. In: *JINST* 8 (2013), P07002. DOI: 10.1088/1748-0221/8/07/P07002. arXiv: 1304.6961 [physics.ins-det].
- [133] H. V. Souza. “Light detection with power and signal transmission over fiber”. In: *JINST* 18.02 (2023), p. C02029. DOI: 10.1088/1748-0221/18/02/C02029. arXiv: 2211.09534 [physics.ins-det].
- [134] Saba Parsa et al. “SoLAr: Solar Neutrinos in Liquid Argon”. In: *arXiv preprint arXiv:2203.07501* (2022).
- [135] Roberto Acciarri et al. “The Liquid Argon In A Testbeam (LArIAT) Experiment”. In: *JINST* 15.04 (2020), P04026. DOI: 10.1088/1748-0221/15/04/P04026. arXiv: 1911.10379 [physics.ins-det].

- [136] Elena Gramellini. “Measurement of the Negative Pion and Positive Kaon Total Hadronic Cross Sections on Argon at the LArIAT Experiment”. PhD thesis. Yale U., 2018. DOI: 10.2172/1489387.
- [137] R. Acciarri et al. “First measurement of the cross section for ν_μ and $\bar{\nu}_\mu$ induced single charged pion production on argon using ArgoNeuT”. In: *Phys. Rev. D* 98.5 (2018), p. 052002. DOI: 10.1103/PhysRevD.98.052002. arXiv: 1804.10294 [hep-ex].
- [138] C. E. Taylor et al. “The Mini-CAPTAIN liquid argon time projection chamber”. In: *Nucl. Instrum. Meth. A* 1001 (2021), p. 165131. DOI: 10.1016/j.nima.2021.165131. arXiv: 2008.11422 [physics.ins-det].
- [139] B. Bhandari et al. “First Measurement of the Total Neutron Cross Section on Argon Between 100 and 800 MeV”. In: *Phys. Rev. Lett.* 123.4 (2019), p. 042502. DOI: 10.1103/PhysRevLett.123.042502. arXiv: 1903.05276 [hep-ex].
- [140] R. Acciarri et al. “The WARp experiment”. In: *J. Phys. Conf. Ser.* 308 (2011). Ed. by Atsuto Suzuki and Koichiro Nishikawa, p. 012005. DOI: 10.1088/1742-6596/308/1/012005.
- [141] P. Agnes et al. “First Results from the DarkSide-50 Dark Matter Experiment at Laboratori Nazionali del Gran Sasso”. In: *Phys. Lett. B* 743 (2015), pp. 456–466. DOI: 10.1016/j.physletb.2015.03.012. arXiv: 1410.0653 [astro-ph.CO].
- [142] P. Abratenko et al. “Wire-cell 3D pattern recognition techniques for neutrino event reconstruction in large LArTPCs: algorithm description and quantitative evaluation with MicroBooNE simulation”. In: *JINST* 17.01 (2022), P01037. DOI: 10.1088/1748-0221/17/01/P01037. arXiv: 2110.13961 [physics.ins-det].
- [143] R. Acciarri et al. “Noise Characterization and Filtering in the MicroBooNE Liquid Argon TPC”. In: *JINST* 12.08 (2017), P08003. DOI: 10.1088/1748-0221/12/08/P08003. arXiv: 1705.07341 [physics.ins-det].
- [144] Carl Grace and Dario Gnani. *LArPix-v2 Datasheet*. Rev. 1.3. LBNL IC Group. Apr. 2021.
- [145] Jonathan Asaadi et al. “First Demonstration of a Pixelated Charge Readout for Single-Phase Liquid Argon Time Projection Chambers”. In: *Instruments* 4.1 (2020), p. 9. DOI: 10.3390/instruments4010009. arXiv: 1801.08884 [physics.ins-det].
- [146] D. A. Dwyer et al. “LArPix: Demonstration of low-power 3D pixelated charge readout for liquid argon time projection chambers”. In: *JINST* 13.10 (2018), P10007. DOI: 10.1088/1748-0221/13/10/P10007. arXiv: 1808.02969 [physics.ins-det].
- [147] Carl Grace, Daniel Dwyer, and Peter Madigan. *Hydra Network Automatic Configuration (Hydra Autoconfig) v1*. Tech. rep. Lawrence Berkeley National Lab.(LBNL), Berkeley, CA (United States), 2020.
- [148] ZeroMQ. *ZeroMQ: An open-source universal messaging library*. Tech. rep. 2020. URL: <https://zeromq.org>.

- [149] Roman Matthias Berner. “ArgonCube—A Novel Concept for Liquid Argon Time Projection Chambers”. PhD thesis. Universität Bern.
- [150] A. A. Machado and E. Segreto. “ARAPUCA a new device for liquid argon scintillation light detection”. In: *JINST* 11.02 (2016), p. C02004. DOI: 10.1088/1748-0221/11/02/C02004.
- [151] J. M. Corning et al. “Temperature-dependent fluorescence emission spectra of acrylic (PMMA) and tetraphenyl butadiene (TPB) excited with UV light”. In: *JINST* 15.03 (2020), p. C03046. DOI: 10.1088/1748-0221/15/03/C03046. arXiv: 1912.02073 [physics.ins-det].
- [152] 3M. *3MTM DICHROICTM Glass Finishes DF-PA Technical Data Sheet*. 2020.
- [153] ELJEN Technology. *Wavelength Shifting Plastics EJ-280, EJ-282, EJ-284, EJ-286*. URL: <https://eljentechnology.com/products/wavelength-shifting-plastics/ej-280-ej-282-ej-284-ej-286>.
- [154] M. Auger et al. “ArCLight—A Compact Dielectric Large-Area Photon Detector”. In: *Instruments* 2.1 (2018), p. 3. DOI: 10.3390/instruments2010003. arXiv: 1711.11409 [physics.ins-det].
- [155] N. Anfimov et al. “Development of the Light Collection Module for the Liquid Argon Time Projection Chamber (LArTPC)”. In: *JINST* 15.07 (2020), p. C07022. DOI: 10.1088/1748-0221/15/07/C07022.
- [156] F. Drielsma et al. “ArgonCube TPC High Voltage System”. 2020. URL: https://indico.fnal.gov/event/45210/contributions/195427/attachments/133913/165378/20200903_argoncube_fc.pdf.
- [157] Roman Berner et al. “First Operation of a Resistive Shell Liquid Argon Time Projection Chamber: A New Approach to Electric-Field Shaping”. In: *Instruments* 3.2 (2019), p. 28. DOI: 10.3390/instruments3020028. arXiv: 1903.11858 [physics.ins-det].
- [158] D. Heck et al. “CORSIKA: A Monte Carlo code to simulate extensive air showers”. In: (Feb. 1998).
- [159] S. Agostinelli et al. “GEANT4—a simulation toolkit”. In: *Nucl. Instrum. Meth. A* 506 (2003), pp. 250–303. DOI: 10.1016/S0168-9002(03)01368-8.
- [160] C. McGrew et al. *edep-sim: An Energy Deposition Simulation*. URL: <https://github.com/ClarkMcGrew/edep-sim/>.
- [161] M. J. Christ et al. “Cosmic-ray proton and helium spectra: Results from the JACEE Experiment”. In: *Astrophys. J.* 502 (1998), pp. 278–283. DOI: 10.1086/305882.
- [162] O. Adriani et al. “PAMELA Measurements of Cosmic-ray Proton and Helium Spectra”. In: *Science* 332 (2011), pp. 69–72. DOI: 10.1126/science.1199172. arXiv: 1103.4055 [astro-ph.HE].

- [163] Y. S. Yoon et al. “Cosmic-Ray Proton and Helium Spectra from the First CREAM Flight”. In: *Astrophys. J.* 728 (2011), p. 122. DOI: 10.1088/0004-637X/728/2/122. arXiv: 1102.2575 [astro-ph.HE].
- [164] T. Sanuki et al. “Precise measurement of cosmic ray proton and helium spectra with the BESS spectrometer”. In: *Astrophys. J.* 545 (2000), p. 1135. DOI: 10.1086/317873. arXiv: astro-ph/0002481.
- [165] Sadakazu Haino et al. “Measurements of primary and atmospheric cosmic-ray spectra with the BESS-TeV spectrometer”. In: *Phys. Lett. B* 594 (2004), pp. 35–46. DOI: 10.1016/j.physletb.2004.05.019. arXiv: astro-ph/0403704.
- [166] Thomas Hebbeker and Charles Timmermans. “A Compilation of high-energy atmospheric muon data at sea level”. In: *Astropart. Phys.* 18 (2002), pp. 107–127. DOI: 10.1016/S0927-6505(01)00180-3. arXiv: hep-ph/0102042.
- [167] M. P. De Pascale et al. “Absolute spectrum and charge ratio of cosmic ray muons in the energy region from 0.2 GeV to 100 GeV at 600 m above sea level”. In: *J. Geophys. Res.* 98.A3 (1993), pp. 3501–3507. DOI: 10.1029/92JA02672.
- [168] Suryanarayan Mondal. “Multiplicity Of Muon In 2m × 2m Detector And Charge Ratio of Cosmic Muon At Madurai”. PhD thesis. TIFR, Bhabha Atomic Res. Ctr., 2021.
- [169] B. Mitrica et al. “Investigation of cosmic ray muons with the WILLI detector compared with the predictions of theoretical models and with semi-analytical formulae”. In: *Nucl. Phys. B Proc. Suppl.* 196 (2009). Ed. by Jean-Noël Capdevielle, Ralph Engel, and Bryan Pattison, pp. 462–465. DOI: 10.1016/j.nuclphysbps.2009.10.007.
- [170] T. Djemil, R. Attallah, and J. N. Capdevielle. “Simulation of the atmospheric muon flux with CORSIKA”. In: *Int. J. Mod. Phys. A* 20 (2005). Ed. by O. Adriani et al., pp. 6950–6952. DOI: 10.1142/S0217751X05030569.
- [171] Jurgen Wentz et al. “Simulation of atmospheric muon and neutrino fluxes with CORSIKA”. In: *Phys. Rev. D* 67 (2003), p. 073020. DOI: 10.1103/PhysRevD.67.073020. arXiv: hep-ph/0301199.
- [172] Geant Collaboration et al. “Physics Reference Manual”. In: *Version: geant4 (10.6v1)* (2018).
- [173] Yung-Su Tsai. “Pair Production and Bremsstrahlung of Charged Leptons”. In: *Rev. Mod. Phys.* 46 (1974). [Erratum: *Rev. Mod. Phys.* 49, 421–423 (1977)], p. 815. DOI: 10.1103/RevModPhys.46.815.
- [174] Frank Biggs and Ruth Lighthill. *Analytical approximations for x-ray cross sections III*. Tech. rep. Sandia National Lab.(SNL-NM), Albuquerque, NM (United States), 1988.
- [175] Mihai Gavrila. “Relativistic k-shell photoeffect”. In: *Physical Review* 113.2 (1959), p. 514.

- [176] John Howard Hubbell, Heinz Albert Gimm, and I. Øverbø. “Pair, triplet, and total atomic cross sections (and mass attenuation coefficients) for 1 MeV-100 GeV photons in elements $Z= 1$ to 100”. In: *Journal of physical and chemical reference data* 9.4 (1980), pp. 1023–1148.
- [177] Oskar Klein and Yoshio Nishina. “Über die Streuung von Strahlung durch freie Elektronen nach der neuen relativistischen Quantendynamik von Dirac”. In: *Zeitschrift für Physik* 52.11-12 (1929), pp. 853–868.
- [178] W. Heitler. *The quantum theory of radiation*. Vol. 5. International Series of Monographs on Physics. Oxford: Oxford University Press, 1936.
- [179] Nimai C. Mukhopadhyay. “Nuclear Muon Capture”. In: *Phys. Rept.* 30 (1977), pp. 1–144. DOI: 10.1016/0370-1573(77)90073-4.
- [180] C. S. Wu and L. Wilets. “Muonic atoms and nuclear structure”. In: *Ann. Rev. Nucl. Part. Sci.* 19 (1969), pp. 527–606. DOI: 10.1146/annurev.ns.19.120169.002523.
- [181] E. L. Snider and G. Petrillo. “LArSoft: Toolkit for Simulation, Reconstruction and Analysis of Liquid Argon TPC Neutrino Detectors”. In: *J. Phys. Conf. Ser.* 898.4 (2017). Ed. by Richard Mount and Craig Tull, p. 042057. DOI: 10.1088/1742-6596/898/4/042057.
- [182] Adam Abed Abud et al. “Highly-parallelized simulation of a pixelated LArTPC on a GPU”. In: *JINST* 18.04 (2023), P04034. DOI: 10.1088/1748-0221/18/04/P04034. arXiv: 2212.09807 [physics.comp-ph].
- [183] Adam Abed Abud et al. “Highly-parallelized simulation of a pixelated LArTPC on a GPU”. In: *JINST* 18.04 (2023), P04034. DOI: 10.1088/1748-0221/18/04/P04034. arXiv: 2212.09807 [physics.comp-ph].
- [184] S. Huang, S. Xiao, and W. Feng. “On the energy efficiency of graphics processing units for scientific computing”. In: *2009 IEEE International Symposium on Parallel & Distributed Processing*. 2009, pp. 1–8. DOI: 10.1109/IPDPS.2009.5160980.
- [185] Guido Van Rossum, Fred L. Drake, et al. *Python reference manual*. Vol. 111. Centrum voor Wiskunde en Informatica Amsterdam, 1995.
- [186] Siu Kwan Lam, Antoine Pitrou, and Stanley Seibert. “Numba: a LLVM-based Python JIT compiler”. In: 2015. DOI: 10.1145/2833157.2833162.
- [187] Ryosuke Okuta et al. “CuPy: A NumPy-Compatible Library for NVIDIA GPU Calculations”. In: *Proceedings of Workshop on Machine Learning Systems (LearningSys) in The Thirty-first Annual Conference on Neural Information Processing Systems (NIPS)*. 2017. URL: http://learningsys.org/nips17/assets/papers/paper_16.pdf.
- [188] John Nickolls et al. “Scalable parallel programming with cuda: Is cuda the parallel programming model that application developers have been waiting for?”. In: *Queue* 6.2 (2008), pp. 40–53.

- [189] Charles R. Harris et al. “Array programming with NumPy”. In: *Nature* 585.7825 (2020), pp. 357–362. DOI: 10.1038/s41586-020-2649-2. arXiv: 2006.10256 [cs.MS].
- [190] J. B. Birks. “Scintillations from Organic Crystals: Specific Fluorescence and Relative Response to Different Radiations”. In: *Proc. Phys. Soc. A* 64 (1951), pp. 874–877. DOI: 10.1088/0370-1298/64/10/303.
- [191] S. Amoruso et al. “Study of electron recombination in liquid argon with the ICARUS TPC”. In: *Nucl. Instrum. Meth. A* 523 (2004), pp. 275–286. DOI: 10.1016/j.nima.2003.11.423.
- [192] R. T. Scalettar et al. “IONIZATION RECOMBINATION IN LIQUID ARGON AT LOW DRIFT FIELDS”. In: (July 1981).
- [193] P. Agnes et al. “Simulation of argon response and light detection in the DarkSide-50 dual phase TPC”. In: *JINST* 12.10 (2017), P10015. DOI: 10.1088/1748-0221/12/10/P10015. arXiv: 1707.05630 [physics.ins-det].
- [194] E. Aprile et al. “Signal Yields of keV Electronic Recoils and Their Discrimination from Nuclear Recoils in Liquid Xenon”. In: *Phys. Rev. D* 97.9 (2018), p. 092007. DOI: 10.1103/PhysRevD.97.092007. arXiv: 1709.10149 [astro-ph.IM].
- [195] E. Segreto. “Evidence of delayed light emission of TetraPhenyl Butadiene excited by liquid Argon scintillation light”. In: *Phys. Rev. C* 91.3 (2015), p. 035503. DOI: 10.1103/PhysRevC.91.035503. arXiv: 1411.4524 [physics.ins-det].
- [196] Patrick Koller. “A Method to Improve the Neutrino Energy Reconstruction in LArTPCs”. PhD thesis. Universität Bern, 2021.
- [197] S. Vinogradov. “Analytical models of probability distribution and excess noise factor of Solid State Photomultiplier signals with crosstalk”. In: *Nucl. Instrum. Meth. A* 695 (2012). Ed. by Philippe Bourgeois et al., pp. 247–251. DOI: 10.1016/j.nima.2011.11.086. arXiv: 1109.2014 [physics.ins-det].
- [198] Hamamatsu. *Hamamatsu MPPC®(multi-pixel photon counter) S13360 series*. 2022.
- [199] W. Shockley. “Currents to conductors induced by a moving point charge”. In: *J. Appl. Phys.* 9.10 (1938), pp. 635–636. DOI: 10.1063/1.1710367.
- [200] Daniel Douglas. “Design and Calibration of the DUNE Liquid Argon Near Detector”. PhD thesis. Michigan State U., 2022. DOI: 10.25335/5nhh-qb38.
- [201] Madigan, P. and others. “Physics performance of a PCB-based pixelated LArTPC anode (LArPix)”. In: *APS April Meeting Abstracts*. Vol. 2021. 2021, H20-002.
- [202] DuPont. *DuPont™ Kapton®XP*. 2020.
- [203] DuPont. *DuPont™ Kapton®polyimide film DR-8*. 2020.
- [204] INTERRA GLOBAL. *mSORB®4A Molecular Sieve, Type A*. 2020.
- [205] BASF. *BASF Cu-0226 S*. 2018.

- [206] P. Cennini et al. “Argon purification in the liquid phase”. In: *Nucl. Instrum. Meth. A* 333 (1993), pp. 567–570. DOI: 10.1016/0168-9002(93)91209-6.
- [207] B. Abi et al. “Performance of a modular ton-scale pixel-readout liquid argon Time Projection Chamber”. In: (2024). In preparation.
- [208] M. Adamowski et al. “The Liquid Argon Purity Demonstrator”. In: *JINST* 9 (2014), P07005. DOI: 10.1088/1748-0221/9/07/P07005. arXiv: 1403.7236 [physics.ins-det].
- [209] David Montanari et al. “Performance and Results of the LBNE 35 Ton Membrane Cryostat Prototype”. In: *Phys. Procedia* 67 (2015). Ed. by H. J. M. ter Brake, H. H. J. ten Kate, and S. Vanapalli, pp. 308–313. DOI: 10.1016/j.phpro.2015.06.092.
- [210] J. Linsley. “US Standard Atmosphere”. In: *US Govt. Printing Office, Washington* (1962).
- [211] Mengyun Guan et al. “A parametrization of the cosmic-ray muon flux at sea-level”. In: (Sept. 2015). arXiv: 1509.06176 [hep-ex].
- [212] Donald E. Groom, Nikolai V. Mokhov, and Sergei I. Striganov. “Muon stopping power and range tables 10-MeV to 100-TeV”. In: *Atom. Data Nucl. Data Tabl.* 78 (2001), pp. 183–356. DOI: 10.1006/adnd.2001.0861.
- [213] The HDF Group. *Hierarchical Data Format, version 5*. 1997. URL: <https://www.hdfgroup.org/HDF5/>.
- [214] Laura Dominé and Kazuhiro Terao. “Scalable deep convolutional neural networks for sparse, locally dense liquid argon time projection chamber data”. In: *Phys. Rev. D* 102.1 (2020), p. 012005. DOI: 10.1103/PhysRevD.102.012005. arXiv: 1903.05663 [hep-ex].
- [215] Xin Qian et al. “Three-dimensional Imaging for Large LArTPCs”. In: *JINST* 13.05 (2018), P05032. DOI: 10.1088/1748-0221/13/05/P05032. arXiv: 1803.04850 [physics.ins-det].
- [216] Martin Ester et al. “A density-based algorithm for discovering clusters in large spatial databases with noise”. In: *KDD*. Vol. 96. 34. 1996, pp. 226–231.
- [217] Martin A. Fischler and Robert C. Bolles. “Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography”. In: *Communications of the ACM* 24.6 (1981), pp. 381–395.
- [218] Stefan Gundacker and Arjan Heering. “The silicon-photomultiplier: fundamentals and applications of a modern solid-state photon detector”. In: *Phys. Med. Biol.* 65.17 (2020), 17TR01. DOI: 10.1088/1361-6560/ab7b2d.
- [219] Norbert Wiener. *Extrapolation, interpolation, and smoothing of stationary time series: with engineering applications*. Vol. 113. 21.
- [220] B. Efron. “Bootstrap Methods: Another Look at the Jackknife”. In: *Annals Statist.* 7.1 (1979), pp. 1–26. DOI: 10.1214/aos/1176344552.

- [221] William Henry Bragg and RLXXIV Kleeman. “LXXIV. On the ionization curves of radium”. In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 8.48 (1904), pp. 726–738.
- [222] P. Benetti et al. “Measurement of the specific activity of ar-39 in natural argon”. In: *Nucl. Instrum. Meth. A* 574 (2007), pp. 83–88. DOI: 10.1016/j.nima.2007.01.106. arXiv: astro-ph/0603131.
- [223] G. Heusser. “Low-radioactivity background techniques”. In: *Ann. Rev. Nucl. Part. Sci.* 45 (1995), pp. 543–590. DOI: 10.1146/annurev.ns.45.120195.002551.
- [224] Olaf Behnke et al. *Data analysis in high energy physics: a practical guide to statistical methods*. John Wiley & Sons, 2013.
- [225] Stanley E. Sobottka and Edward L. Wills. “Energy Spectrum of Charged Particles Emitted Following Muon Capture in Si28”. In: *Phys. Rev. Lett.* 20.12 (1968), p. 596. DOI: 10.1103/PhysRevLett.20.596.
- [226] J. Kremer et al. “Measurements of ground-level muons at two geomagnetic locations”. In: *Phys. Rev. Lett.* 83 (1999), pp. 4241–4244. DOI: 10.1103/PhysRevLett.83.4241.
- [227] Vardan Khachatryan et al. “Measurement of the Charge Ratio of Atmospheric Muons with the CMS Detector”. In: *Phys. Lett. B* 692 (2010), pp. 83–104. DOI: 10.1016/j.physletb.2010.07.033. arXiv: 1005.5332 [hep-ex].
- [228] Michele Weber. “Toward the FDR: Design Development and Prototyping”. In: *DUNE ND-LAr Preliminary Design Review*. 2022.

Appendix A

Systematics summary

Table A.1 summarizes all of the systematic effects investigated for this thesis, along with the determined relevance for the measurement. Effects that were deemed irrelevant produced less than a 0.5% change in the shape of the fitted distributions, determined by simulation.

Table A.1: Summary of systematics considered for this measurement.

Name	Description	Scale	Deemed relevant	Technique
Through-going μ background	Total number of through-going muon background events in muon-decay nhit distribution.	$\sim 25\%$	Yes	Free-floating nuisance parameter for each data sample.
Non-fid. μ background	Total number of stopping muons that are mis-reconstructed inside the fiducial boundary.	$\sim 4\%$	Yes	Free-floating nuisance parameter for each data sample.
EM backgrounds	Total number of electron and gamma showers that are mis-reconstructed as stopping muons.	$\sim 10\%$	Yes	Free-floating nuisance parameter.
n backgrounds	Total number of neutron-induced events that are mis-reconstructed as stopping muons.	$\sim 1\%$	Yes	Free-floating nuisance parameter.
π backgrounds	Total number of pion-induced events that are mis-reconstructed as stopping muons.	$\sim 1\%$	Yes	Free-floating nuisance parameter.

(Continued on next page.)

Table A.1: Summary of systematics considered for this measurement (continued).

Name	Description	Scale	Deemed relevant	Technique
p backgrounds	Total number of proton events that are mis-reconstructed as stopping muons.	< 1%	No	N/A
Muon-decay nhit scale	Overall scale for muon-decay nhit count.	N/A	Yes	Free-floating nuisance parameter for each data sample.
Muon-decay nhit smear	Broadening term for muon-decay nhit count.	N/A	Yes	Free-floating nuisance parameter for each data sample.
LRS LUT tune	Impact of using Sec. 4.4 LUT modification.	N/A	No	N/A
Energy deposition in LRS modules	Light produced within LRS modules is not simulated.		No	N/A
Correlated light gain bias	Overall LRS $dE/dx \rightarrow$ light yield (LY) bias.	$\pm 5\%$	No	N/A
Uncorrelated light gain bias	Bias in each light detector LRS $dE/dx \rightarrow$ LY.	$\pm 5\%$	No	N/A
Correlated charge gain bias	Overall CRS $dE/dx \rightarrow$ charge yield (QY) bias.	$\pm 5\%$	Yes	$\pm \sigma$ nuisance parameter
Uncorrelated charge gain bias	Increase channel-to-channel $dE/dx \rightarrow$ QY variation.	+5%	Yes	Re-fit with different value.
CRS self-trigger threshold	Changes in template distributions due to self-trigger threshold.	N/A	Yes	Use medium threshold templates to fit high threshold data and vice-versa.
Recombination model	Choice of Birks' recombination over Modified Box Model.	N/A	Yes	Re-fit with a different value.
CRS noise	Uncertainty in single-channel noise.	4%	No	N/A
Electron lifetime	Uncertainty in electron lifetime correction.	$\pm 500 \mu s$	No	N/A
Cathode voltage	Uncertainty in applied drift field.	0.5%	No	N/A
Singlet light yield	Uncertainty in singlet-to-triplet light yield ratio.	± 0.075	No	N/A

(Continued on next page.)

Table A.1: Summary of systematics considered for this measurement (continued).

Name	Description	Scale	Deemed relevant	Technique
Triplet time constant	Uncertainty in triplet-decay time constant.	$\pm 30\mu\text{s}$	No	N/A
CORSIKA μ -flux model	Discrepancy between CSDA muon energy between CORSIKA and the data.	$\sim 5\%$	Yes	Re-weight simulation events and re-fit.
Data loss	Global bit errors result in data loss.	$< 3\%$	Yes	Free-floating nuisance parameter.
Fid. volume	Semi-arbitrary choice in fiducial volume.	+2.2 cm, +4.4 cm	Yes	Re-fit with different value.
Decay-time gate	Semi-arbitrary choice in decay-time gate window start time.	+40 ns, +80 ns	Yes	Re-fit with different value.
LRS amplitude cut	Semi-arbitrary choice in light amplitude cut.	± 10 PE	Yes	Re-fit with different value.
LRS/CRS yield cut	Semi-arbitrary choice in light yield cut.	± 5 PE/hit	Yes	Re-fit with different value.
Muon-decay nhit cut	Semi-arbitrary choice in muon-decay nhit cut to select muon decays.	± 2	Yes	Re-fit with different value.
Charge ratio bias	Bias in decay-time distribution due to difference in low-energy e^\pm spectrum.	3%	Yes	Re-fit without correction.
Efficiency ratio bias	Selection bias between capture and decay events	2%.	Yes	Re-fit without correction.

Appendix B

Event gallery

This appendix contains a collection of Module 0 events and events generated by the simulation described in Chapter 3. Each event display shows the positions of the reconstructed hits, colored by the output of the muon-decay electron likelihood score described in Section 5.4.1. Blue indicates muon-like and red indicates electron-like. Two perspectives of the detector are shown: x, y on the left and z, y on the right, where y is vertical, z is the drift axis, and x points in the direction that yields a right-handed coordinate system. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left. Note that because the anode planes lie on top of each other with this perspective, it is possible to observe hits in disabled regions, as the disabled channels are not equivalent in each TPC. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM. Events from the Module 0 dataset are labelled by the subrun ID, event number (as generated sequentially using the algorithm described in Section 5.1), as well as the low-precision unix and high-precision GPS PPS timestamps.

B.1 Stopping muon sample

The events shown in Figures B.1-B.6 pass the event selection to be labelled as stopping muons. They have been randomly selected from the Module 0 data described in their captions.




Figure B.1: Stopping-muon candidate events from the medium-threshold Module 0 run 1 data sample (batch 1). Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM.




Figure B.2: Stopping-muon candidate events from the medium-threshold Module 0 run 1 data sample (batch 2). Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM.




Figure B.3: Stopping-muon candidate events from the high-threshold Module 0 run 1 data sample (batch 1). Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM.




Figure B.4: Stopping-muon candidate events from the high-threshold Module 0 run 1 data sample (batch 2). Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM.




Figure B.5: Stopping-muon candidate events from the high-threshold Module 0 run 2 data sample (batch 1). Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM.

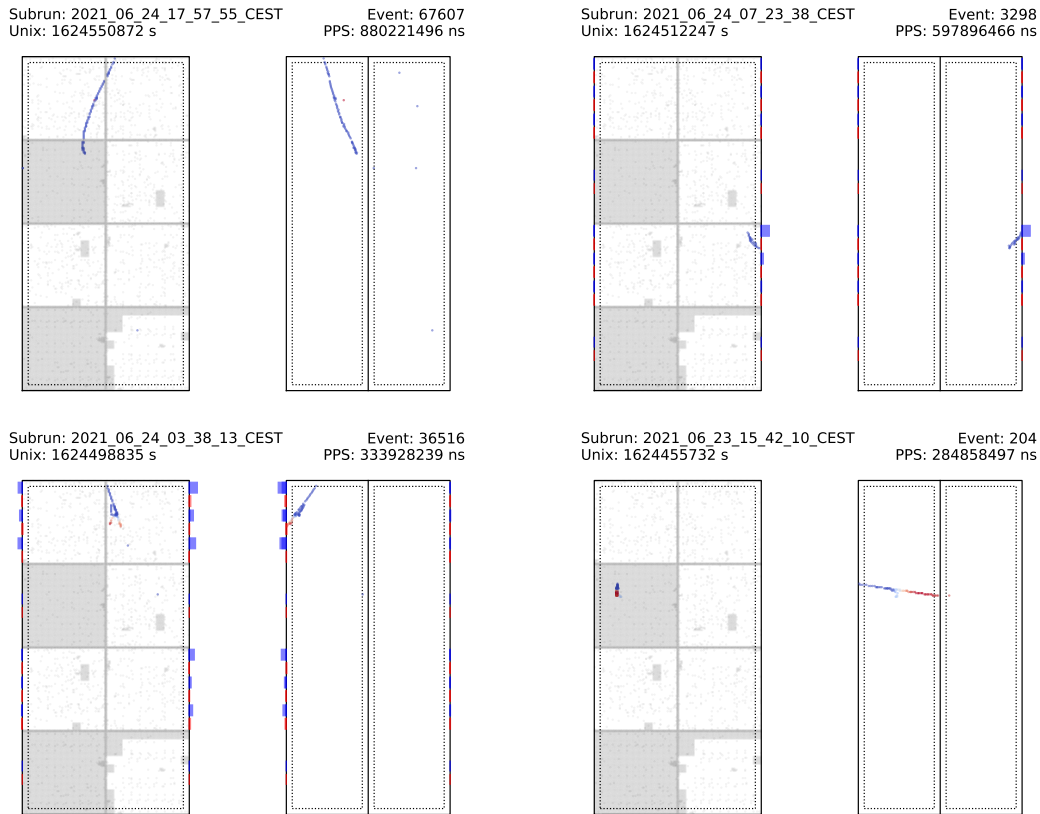


Figure B.6: Stopping-muon candidate events from the high-threshold Module 0 run 2 data sample (batch 2). Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM.

B.2 Muon-decay sample

The events shown in Figures B.7-B.9 pass the event selection to be labelled as muon-decay events. They have been randomly selected from the Module 0 data described in their captions. The light signals of the ArcLight detectors are not shown as they were not used to tag or reconstruct events in this analysis.




Figure B.7: Muon-decay candidate events from the medium-threshold Module 0 run 1 data sample. Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM.




Figure B.8: Muon-decay candidate events from the high-threshold Module 0 run 1 data sample. Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM.

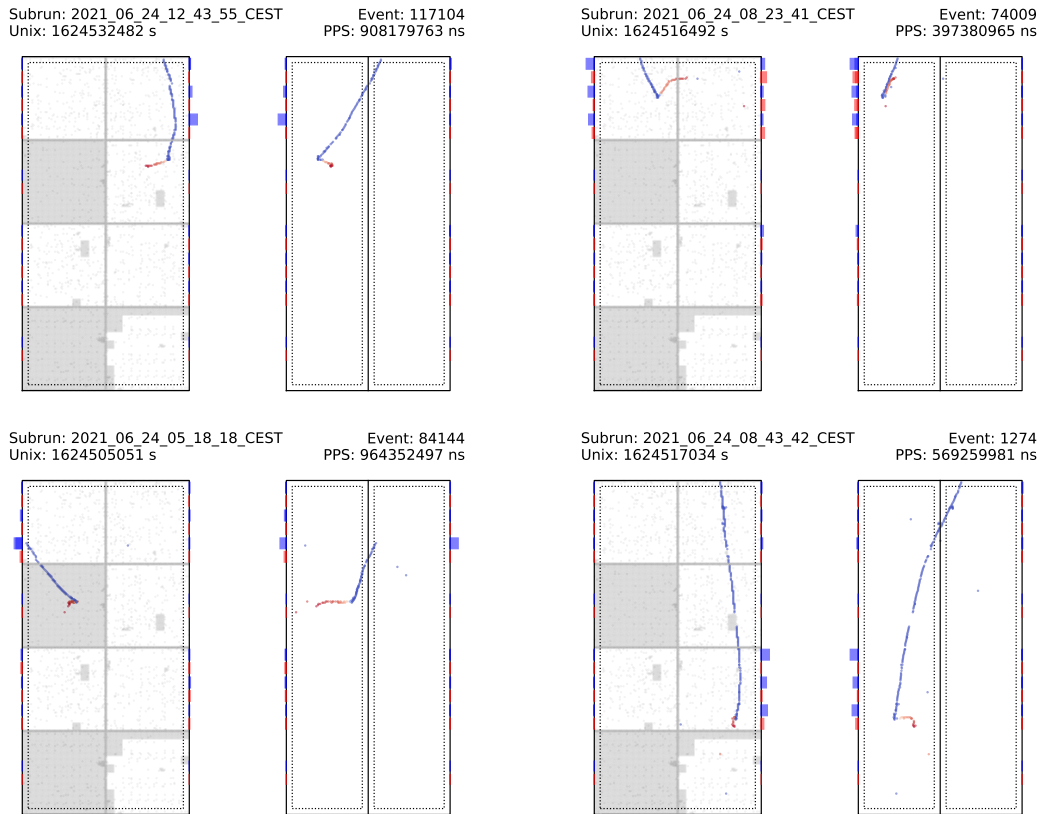


Figure B.9: Muon-decay candidate events from the high-threshold Module 0 run 2 data sample. Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM.

B.3 CORSIKA muon sample

The events shown in Figures B.10-B.15 pass the event selection to be labelled as stopping muons. They have been randomly selected from the CORSIKA-generated data described in their captions.




Figure B.10: Stopping-muon candidate events from the CORSIKA-generated medium-threshold Module 0 run 1 sample (batch 1). Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM. All four of these events are stopping μ^+ with a subsequent decay.




Figure B.11: Stopping-muon candidate events from the CORSIKA-generated medium-threshold Module 0 run 1 sample (batch 2). Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM. From left-to-right, top-to-bottom, these events are a μ -capture event, a μ^+ -decay event, a μ -capture event, and an EM-shower background event.




Figure B.12: Stopping-muon candidate events from the CORSIKA-generated high-threshold Module 0 run 1 sample (batch 1). Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM. From left-to-right, top-to-bottom, these events are a through-going μ background event, a μ^+ -decay event, a μ^+ -decay event, and a μ^- -decay event.




Figure B.13: Stopping-muon candidate events from the CORSIKA-generated high-threshold Module 0 run 1 sample (batch 2). Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM. From left-to-right, top-to-bottom, these events are a μ^+ -decay event, a μ^- -capture event, a μ^+ -decay event, and a μ^+ -decay event.




Figure B.14: Stopping-muon candidate events from the CORSIKA-generated high-threshold Module 0 run 2 sample (batch 1). Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM. From left-to-right, top-to-bottom, these events are a μ^- -decay event, μ^+ -decay event, a μ^+ -decay event, and a μ^+ -decay event.




Figure B.15: Stopping-muon candidate events from the CORSIKA-generated high-threshold Module 0 run 2 sample (batch 2). Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM. From left-to-right, top-to-bottom, these events are a μ^- -decay event, an EM-shower background event (with a pile-up μ), an EM-shower background event, and a μ^+ -decay event.

B.4 CORSIKA muon-decay sample

The events shown in Figures B.16-B.18 pass the event selection to be labelled as muon-decay events. They have been randomly selected from the CORSIKA-generated data described in their captions.

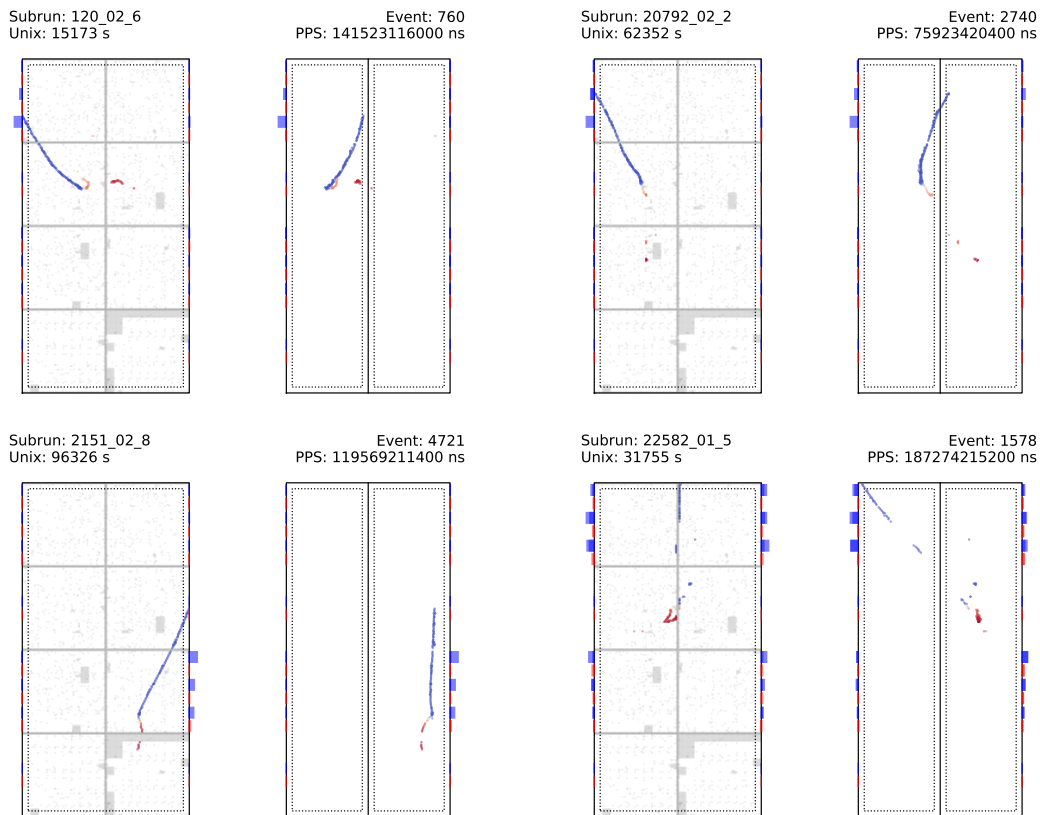


Figure B.16: Muon-decay candidate events from the CORSIKA-generated medium-threshold Module 0 run 1 sample. Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM. From left-to-right, top-to-bottom, these events are a μ^+ -decay event, a μ^+ -decay event, a μ^- -decay event, and a μ^- -decay event.




Figure B.17: Muon-decay candidate events from the CORSIKA-generated high-threshold Module 0 run 1 sample. Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM. From left-to-right, top-to-bottom, these events are a μ^- -decay event, a μ^+ -decay event, a μ^- -decay event, and a μ^+ -decay event.




Figure B.18: Muon-decay candidate events from the CORSIKA-generated high-threshold Module 0 run 2 sample. Blue points represent muon-like hits and red points indicates electron-like hits. Two perspectives of the detector are shown: x, y on the left and z, y on the right. Outlines of the detector are indicated (solid-black line) along with the fiducial volume used to select stopping muons (dotted-black line). The disabled channels on each anode plane are projected into the x, y perspective on the left and colored gray. Bar indicators from the outside of the detector show the reconstructed prompt (blue) and delayed (red) light signal for each LCM. From left-to-right, top-to-bottom, these events are a μ^+ -decay event, μ^+ -decay event, a μ^- -decay event, and a μ^+ -decay event.