

TOWARD A PRECISION ERA OF NEUTRINO OSCILLATION PHYSICS:  
LIQUID ARGON SCINTILLATION DETECTOR DEVELOPMENT FOR  
DUNE AND NEUTRINO OSCILLATION STUDIES WITH NO<sub>v</sub>A

Bruce Lynn Howard, Jr.

Submitted to the faculty of the University Graduate School  
in partial fulfillment of the requirements  
for the degree  
Doctor of Philosophy  
in the Department of Physics  
Indiana University  
November 2019

Accepted by the Graduate Faculty, Indiana University, in partial fulfillment of the requirements  
for the degree of Doctor of Philosophy.

Doctoral Committee

---

Jon Urheim, PhD

---

Stuart Mufson, PhD

---

Rex Tayloe, PhD

---

Michael Berger, PhD

October 21, 2019

Copyright ©2019

Bruce Howard

## *Acknowledgments*

First and foremost, I would like to thank my family. Mom and Dad, thanks for encouraging me to follow my desires and for providing me with help along the way. To my brother, thanks for your help as well, especially when my laptop(s) would refuse to act as desired.

To the professors I have worked with and my advisors, thank you. Thanks to Jon Urheim for being my advisor throughout my graduate school years, and for teaching me countless things, being awesome, and always having ideas. Thanks also to Stuart Mufson for working with me, helping to teach me the ropes of R&D and the passion for doing such work. To Mark Messier, though you never served in an official role as advisor or committee member, you nonetheless have also had a significant impact on my development as a scientist. Thanks for always being interested and available to chat about interesting/unexpected discoveries or NOvA more broadly. To the other members of my committee, Professors Rex Tayloe and Mike Berger, thanks for your willingness to read this thesis and to serve on my committee. Finally, thanks to my undergraduate advisor Vittorio Paolone and summer internship advisor Jyotsna Osta, who helped develop my passion for physics and guided me in becoming a neutrino physicist. To my graduate professors at IU and undergraduate professors at Pitt, thank you all for your impacts that have helped lead me to this point of writing this dissertation.

Thanks to Denver Whittington and Gavin Davies, my two main “senior” post-docs during my time at Indiana. They have both helped in so many ways, and I learned much about DUNE and NOvA, science, and life from them. I am happy to call both friends, and they are always great to chat with. In fact, this is broadly true of my fellow Indiana University neutrino family throughout my time here, and thanks to them all, especially those I have gotten to know over the years or have worked with, for example Erica Smith as another postdoc and students such as Reed Bowles, Kevin Ewart, Micah Groh, Teresa Lackey, Johnathon Jordan, Chris Macias, Maria Manrique, Ryan Murphy, Fernanda Psihas, and Justin Vasel. Furthermore, even though we did not overlap at Indiana for as long, I’m counting you in this too Michael Baird and Evan Niner. I’ve enjoyed having you as close contacts over the years as well.

Thanks also to the many other IU physics and astronomy faculty/staff I've worked with, especially on the DUNE photon detection studies: B. Adams, B. Baugh, M. Gebhard, J. Karty, M. Lang, B. Martin, J. Musser, and P. Smith. Furthermore, thanks to the front office staff of both departments for their assistance especially in matters of travel and shipping/acquisition which was incredibly helpful in being able to conduct research and present it at conferences and meetings.

Thanks to all of the various researchers, from fellow grad students to professors/scientists with whom I have had many a useful conversation related to my work. There are far too many of you to name, and yet you have been great to work with and to get to know. Much of the group of young (student and postdoc) NOvA members I have had the pleasure to work with as President of the Young NOvA organization as well. Thank you for giving me that chance to grow as a leader. To my fellow Young NOvA leadership member Diana Mendez, thank you for your consistent support of ideas and for working with me in that capacity.

Furthermore, thanks to those not mentioned above who have helped me in my research projects with DUNE and NOvA. Thank you to folks like Peter Shanahan (scientist and co-spokesperson), Alex Himmel (scientist), Matt Strait (postdoc), Karl Warburton (postdoc), Tyler Alion (graduate student), and Nitish Nayak (graduate student) for many useful/helpful conversations and/or for working together on projects, especially whilst I spent some time at Fermilab. Also thank you to the myriad people who helped make the photon detector R&D and prototype tests possible. Those not mentioned previously for other purposes in this passage but deserving acknowledgement for this are listed as follows. At Fermilab: R. Davis, C. Escobar, B. Miner, A. Pla-Dalmau, S. Pordes, B. Rebel, F. Schwartz, M. Zuckerbrot. At ANL: J. Anderson, G. Drake, A. Kreps, M. Oberling. At Colorado State: N. Buchanan, J. Jablonski, D. Warner. At Eljen Technology: C. Hurlbut.

This work was supported in part by the Trustees of Indiana University, the DOE Office of High Energy Physics through grant DE-SC0010120 to Indiana University, and grant #240296 from Brookhaven National Laboratory to Indiana University. Furthermore, Occlusion Tests discussed in the context of the NOvA oscillation analysis uses resources from the Big Red II computing cluster. This research was supported in part by Lilly Endowment, Inc., through its support for the Indiana

University Pervasive Technology Institute, and in part by the Indiana METACyt Initiative. The Indiana METACyt Initiative at IU was also supported in part by Lilly Endowment, Inc.

Bruce Lynn Howard, Jr.

TOWARD A PRECISION ERA OF NEUTRINO OSCILLATION PHYSICS: LIQUID ARGON  
SCINTILLATION DETECTOR DEVELOPMENT FOR DUNE AND NEUTRINO  
OSCILLATION STUDIES WITH NOvA

NOvA is an accelerator-based high-energy particle physics experiment that studies oscillations in a beam of mostly muon (anti-)neutrinos, looking for the disappearance of muon (anti-)neutrinos and appearance of electron (anti-)neutrinos. This search is sensitive to several key neutrino oscillation parameters. After the completion of NOvA, DUNE will continue to study neutrino oscillations with a liquid argon time-projection chamber, which provides finer granularity of tracking than NOvA. Furthermore, the beam will be more powerful, the detector fiducial volume is approximately 3 times larger than the full NOvA detector, and the detector will be situated underground. DUNE is therefore expected to usher in a precision era of long-baseline neutrino oscillation physics.

While the main signal from a liquid argon time-projection chamber is the ionization signal produced by the passage of charged particles, argon is also a copious source of scintillation. This light is useful for providing a precise event time. Scintillation light is potentially especially useful in non-beam studies in DUNE, such as atmospheric neutrinos, supernova neutrinos, and potential nucleon decay.

In this dissertation, I will discuss both projects. In addition to a study of the drift in NOvA's energy calibration, recent oscillation analysis is discussed, which finds a best fit at  $\sin^2 \theta_{23} = 0.56^{+0.04}_{-0.03}$ ,  $\Delta m_{32}^2 = 2.48^{+0.11}_{-0.06} \times 10^{-3} \text{eV}^2$ ,  $\delta_{CP}/\pi = 0.0^{+1.3}_{-0.4}$ , and prefers normal ordering at  $1.9\sigma$ . Contributions to NOvA analysis will also be highlighted, especially methods used to perform separation of neutrinos and antineutrinos in the predominantly antineutrino beam for study. This wrong-sign component is initially used to provide further understanding of the beam components. Furthermore, characterization and prototype testing of photon detectors for DUNE is discussed.

Finally, a study is performed with wrong-sign separation techniques applied to the appearance signal in both NOvA and DUNE to create two samples: one enriched in signal and the other enriched in the wrong-sign oscillation appearance. Wrong-sign in the oscillation diminishes the

sensitivity of the analysis, so this study aimed to determine if such a separation would enhance the sensitivity of the oscillation. Little benefit was found at the ultimate exposure with NOvA with the techniques currently available, but the study showed potential promise when applied to DUNE.

---

Jon Urheim, PhD

---

Stuart Mufson, PhD

---

Rex Tayloe, PhD

---

Michael Berger, PhD

## Contents

<b>1</b>	<b>The Neutrinos</b>	<b>1</b>
1.1	Neutrino Oscillations . . . . .	4
1.2	The Oscillation Landscape . . . . .	8
<b>2</b>	<b>Long-baseline Neutrino Oscillation Experiments</b>	<b>14</b>
2.1	Atmospheric neutrino studies . . . . .	15
2.1.1	Super-Kamiokande . . . . .	16
2.2	Neutrinos at the Main Injector . . . . .	20
2.2.1	Upgrades for the next generation . . . . .	23
2.3	Accelerator-driven long-baseline studies . . . . .	23
2.3.1	Main Injector Neutrino Oscillation Search (MINOS) . . . . .	24
2.3.2	Tokai 2 Kamioka (T2K) . . . . .	25
2.3.3	Hyper-Kamiokande and T2HK . . . . .	26
<b>3</b>	<b>The NuMI Off-axis <math>\nu_e</math> Appearance Experiment</b>	<b>28</b>
3.1	The Detector . . . . .	29
3.2	Energy Calibration . . . . .	31
3.3	NOvA Flux and Interactions . . . . .	36
3.4	Event classification . . . . .	41
<b>4</b>	<b><math>\nu_e</math> Appearance Oscillation Analysis with NOvA</b>	<b>44</b>
4.1	$\nu_e, \bar{\nu}_e$ appearance analysis . . . . .	45
4.2	Some contributions to the analysis . . . . .	52
<b>5</b>	<b>Studies of the Wrong-sign Component in the Antineutrino Beam</b>	<b>59</b>
5.1	Methods for characterizing WS in beam electron antineutrino sample . . . . .	60

5.2	WS component in $\bar{\nu}_e$ sample . . . . .	63
5.2.1	Preliminary results . . . . .	68
5.2.2	WS calculations using additional prong cuts . . . . .	68
5.2.3	Attempt to understand and account for bias . . . . .	73
5.3	WS in the $\bar{\nu}_\mu$ sample . . . . .	75
<b>6</b>	<b>NOvA Oscillation Results</b>	<b>79</b>
6.1	Analysis and joint fit results . . . . .	79
6.2	Kolmogorov-Smirnov style test of accumulated data . . . . .	83
6.3	Study of unexpected background . . . . .	86
<b>7</b>	<b>The Deep Underground Neutrino Experiment</b>	<b>92</b>
7.1	Detector technology . . . . .	93
7.1.1	Note on the Near Detector . . . . .	98
7.2	Physics of DUNE . . . . .	98
7.2.1	Oscillation physics . . . . .	98
7.2.2	Other physics . . . . .	100
7.3	Liquid Argon Scintillation and Detection . . . . .	103
<b>8</b>	<b>Development, Characterization, and Testing of Prototype Photon Detector System</b>	<b>109</b>
8.1	Photon detector system prototype and general component testing/characterization . . . . .	109
8.2	Laboratory-based tests of double-shift lightguide photon detection system . . . . .	119
8.2.1	Attenuation testing . . . . .	119
8.2.2	Plate testing . . . . .	122
8.2.3	Simulation for expected system performance . . . . .	126
8.3	Integrated prototype tests for research and development . . . . .	134
8.3.1	Experimental setup and operation . . . . .	134
8.3.2	Data and analysis . . . . .	137
8.3.3	Results . . . . .	139

<b>9</b>	<b>Toward a precision era in neutrino physics</b>	<b>143</b>
9.1	Oscillation analysis in NOvA and wrong-sign separation . . . . .	144
9.2	$\nu_e, \bar{\nu}_e$ appearance in DUNE . . . . .	156
9.3	Oscillation analysis in DUNE and wrong-sign separation . . . . .	159
<b>A</b>	<b>Studies on Drift in the NOvA Energy Scale Calibration</b>	<b>163</b>
A.1	The drift in calibration . . . . .	164
A.2	Supporting evidence: digging into the drift . . . . .	167
A.2.1	Digging into the drift structure . . . . .	167
A.2.2	Digging into the diblock breakdowns . . . . .	169
A.2.3	Digging into the distributions . . . . .	170
A.2.4	Digging into the drift over time . . . . .	175
A.2.5	Investigating causes . . . . .	179
A.3	Discussion . . . . .	184
<b>B</b>	<b>Studies on Timing Resolution of SiPM Readout Electronics</b>	<b>192</b>
B.1	Experimental Setup . . . . .	192
B.2	Data Analysis and Results . . . . .	196
<b>C</b>	<b>Studies on multiplexing SiPM readout</b>	<b>204</b>
<b>D</b>	<b>Further NOvA and DUNE oscillation sensitivities with wrong-sign sign separation</b>	<b>209</b>
D.1	NOvA . . . . .	209
D.2	DUNE . . . . .	217
	<b>Bibliography</b>	<b>221</b>
	<b>Curriculum Vitae</b>	

## CHAPTER 1

### The Neutrinos

The neutrino's story is interesting in its own right, and the particle has unceasingly produced important and often not entirely expected ties to fundamental physical processes and properties.

When unstable nuclei decay via the  $\beta$  process (also called the beta decay), a neutron (or proton) produces a proton and electron (or neutron and positron). Consider a nucleus originally at rest<sup>1</sup>. In this frame, the nucleus has some given energy and zero momentum. The application of conservation of energy and momentum then gives exactly one solution for the two daughter products. However, instead of experiments with beta decay yielding a single value, the energy in the beta decay turned out to be a spectrum (see for example work with Radium [1]). This led to a profound conundrum in which the principle of energy conservation was itself questioned, or at least that it might only be an approximation valid in certain regimes. In their 1956 article reporting the first direct detection of neutrino interactions [2], Reines and Cowan encapsulated the concern this view had raised by stating "One possible explanation was that the conservation laws (upon which the entire structure of modern science is built) were not valid when applied to regions of subatomic dimensions."

In 1930, Wolfgang Pauli put forth a potential second explanation in a letter to colleagues (of which a typed and translated version can be found [3]): the spectrum of energy of daughter products of beta decay could be explained if there were a third particle involved in the process. To have not been seen, it would have to be neutral. To have a beta spectrum in which the sum of the detected energy extends nearly to the parent particle/atom mass, the third particle would have to be light.

It took over 20 years from postulation to detection, but this third particle indeed exists and is the neutrino. Neutrinos abound from all manners of sources: the formation of the universe, nuclear reactions in stars, nuclear reactors, meson and muon decays, etc. However, being neutral and only weakly interacting, they have small cross-sections with material (order  $10^{-38}$  cm<sup>2</sup> per nucleon in

---

<sup>1</sup>Of course any lab frame could be boosted to the rest frame.

the few GeV region [4]) and are thus hard to detect. Frederick Reines and Clyde Cowan conducted experiments in the 1950s which finally led to a positive detection of neutrino interactions in 1956 [2]. Their experiment had sandwiched layers of water (with a Cadmium compound) and liquid scintillator that would be sensitive to the reaction  $\bar{\nu} + p \rightarrow n + e^+$  from the neutrinos originating in a nearby reactor. The positron would produce a prompt signal followed somewhat after by  $\gamma$ s from the neutron capturing on Cd, giving a clear signal on the oscilloscope with a signal rate several times higher than that of the various backgrounds [2]. Careful considerations of the results seen in this nominal condition as well as changes when varying an experimental parameter led to the conclusion that they had indeed detected neutrinos.

With this result, Pauli's "desperate" idea was found to be correct, and one of the most interesting particles was finally in view. Over the following 60+ years, the scientific community has slowly learned more and more about this particle. The experimentation of Reines and Cowan focused on neutrino interactions related to the beta decay, and their interactions produced positrons. It was initially unknown if neutrinos could couple to the muons as well.

Lederman, Schwartz, and Steinberger led an experiment at Brookhaven in 1962 showing that neutrinos could couple to muons and that these neutrinos are a distinct flavor from those coupling to electrons in the beta decay process [5], ultimately winning a Nobel Prize. They used a beam of accelerated protons and a stationary Beryllium target to create predominantly  $\pi$  mesons, which decay via  $\pi \rightarrow \mu + \nu$ . A spark chamber detector studied interactions of the resulting neutrinos, which were mostly consistent with muonic events and thus different from the electron-coupling of the inverse beta decay events in reactor experiments [5]. This provided evidence for multiple neutrino flavors. A similar process to produce neutrinos is used in modern accelerator-based experiments as well, nearly 60 years later, as described in Chapter 2.

With the observation of the tau neutrino [6], we now know there are at least 3 flavors of neutrinos, one coupled to each flavor of lepton (electron, muon, tau). Experimental evidence, especially from the width of the Z boson decay [7], suggests that there are only three [4]. Since the neutrinos interact via the weak force, reactions can proceed through the exchange of one of two mediators:

the charged  $W$  and the neutral  $Z$ . In the charged-current (CC) reaction, mediated by the  $W$ , the neutrino couples with the lepton of the given flavor. Such a reaction is the signal in the Reines and Cowan experiment noted above: the neutrinos producing positrons in the final state were electron antineutrinos ( $\bar{\nu}_e$ ). Having no charge, the  $Z$  boson mediates the neutral-current (NC) reaction, where the neutrino is present in both the initial and final states.

For neutrino interactions, many experiments over several decades have measured scattering cross-sections. These come in many varieties: CC, NC, exclusive states (specific final states, for example), inclusive (covering multiple final states, in some cases *any* final state). Cross-sections are also measured in varying media, e.g. iron, hydro-carbons, etc. These sorts of measurements are incredibly important to experiments studying other neutrino properties. Even in cases where the overall scale of the cross-section is not so important, knowledge of the products of these interactions can still be important. Energy spectrum smearing due to nuclear processes like final state interactions (FSI) in which outgoing hadrons are re-scattered within the nucleus is one example of a potentially important effect. A plot of neutrino CC interaction cross-sections in an inclusive, per-nucleon form is given in Figure 1.1 [4].

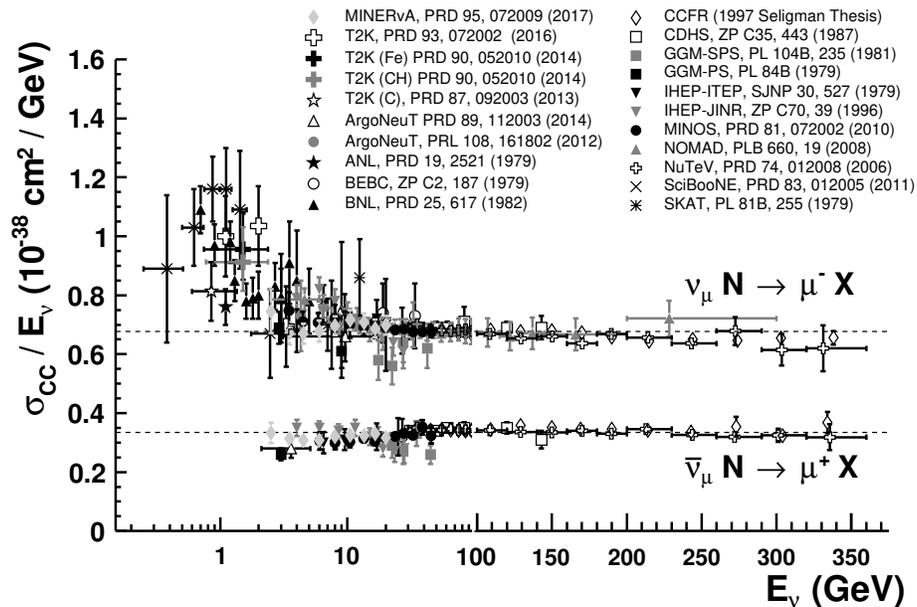


Figure 1.1: Plot of neutrino CC interaction cross-sections in an inclusive, per-nucleon form [4]. The antineutrino cross-section is lower than the neutrino cross-section.

Along with research on reactor- and accelerator-produced neutrinos, other research has focused on detecting the neutrinos from the largest local source: the sun. Solar neutrinos are produced as natural byproducts of the nuclear fusion leading to energy production. The highest flux of neutrinos comes from the main proton (hydrogen) fusion reaction  $pp \rightarrow de^+\nu$ , though these neutrinos are low in energy (a few hundred keV at most) [4]. Other nuclear processes produce neutrinos with higher energy, up to  $\mathcal{O}(10)\text{MeV}$ , for example  ${}^8\text{B} \rightarrow {}^8\text{Be}^*e^+\nu$  [4]. Raymond Davis operated the first experiment to detect neutrinos from the sun in the  ${}^8\text{B}$  process as well as another involving  ${}^7\text{Be}$ , while a number of subsequent experiments used Gallium to detect the flux of neutrinos primarily from the hydrogen fusion [8]. The Gallium experiments agreed with each other, but both these experiments and Davis’s experiment found a flux that was lower than that anticipated [8]. Though they will be discussed later in the context of atmospheric neutrino analysis, Kamiokande and Super-Kamiokande are water-based detectors that also observed a deficit of solar neutrinos [8].

It was known that the neutrinos must be light or quite possibly massless: even in Pauli’s letter this was suggested [3]. Given multiple flavors of low mass neutrinos each with the same quantum numbers, flavor state mixing is a possibility. Such mixing can occur, for example, if the flavor and mass eigenstates are distinct. Quantum mechanically, a given flavor state would be described by a linear superposition of mass states, and vice-versa. This would provide for neutrino flavor oscillation, which could explain the discrepancies noted. A neutrino created in one flavor and propagated some distance (under which the mass states are more important) may be later detected as a different flavor, probabilistically. This phenomenon is described in Section 1.1 and some original evidence for and the current state of neutrino oscillations will be discussed in Section 1.2.

## 1.1 Neutrino Oscillations

Because flavor and mass eigenstates are distinct, a flavor state is a superposition of mass states

$$\nu_{\alpha L} = \sum_{i=1}^3 U_{\alpha i} \nu_{iL}, \alpha = e, \mu, \tau \quad (1.1)$$

where  $L$  denotes the left-handed chirality of the neutrinos (antineutrinos are right-handed), and  $U$  is the so-called Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix, with elements like  $U_{e1}, U_{\mu3}$ , etc. [9]. The elements of this matrix describe the coupling of the neutrino flavors to the mass states and therefore the mixing among flavors, with elements that are conventionally parameterized as combinations of mixing angles. A common parameterization of this matrix is given as

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (1.2)$$

in Reference [9], with  $s_{ij} = \sin \theta_{ij}$  and  $c_{ij} = \cos \theta_{ij}$ .  $\delta$  is a nontrivial angle that changes sign under CP, a combination of charge conjugation (C) and spatial parity inversion (P). Specifically,  $\delta$  is the Dirac CP violating phase, the only phase if neutrinos are separate from their antiparticles. If neutrinos are their own antiparticles - if they are Majorana fermions - there are two other associated CP violating phases, neither of which are distinctly measurable in oscillation experiments.

Propagated over a distance  $L$  with energy  $E$ , the probability for detection in one flavor starting with either the same flavor (survival probability, disappearance) or another flavor (appearance probability) can be determined and characterizes the oscillation. Experiments probing oscillation via solar neutrinos are most sensitive to the “21” parameters; these are typically called the solar parameters. Reactor experiments have been sensitive especially to the parameter  $\theta_{13}$ , and atmospheric neutrino experiments provided early studies of  $\theta_{23}$ . The parameterization and style of writing the PMNS matrix in Equation 1.2 is thus especially useful in picturing the landscape. These and accelerator-driven oscillation experiments and results are the subject of Section 1.2 and Chapter 2.

In vacuum, the neutrino oscillations go as follows [9]:

$$\begin{aligned}
P(\nu_\alpha \rightarrow \nu_\alpha) &= P(\bar{\nu}_\alpha \rightarrow \bar{\nu}_\alpha) \\
&= 1 - 4|U_{\alpha 1}|^2|U_{\alpha 2}|^2 \sin^2 \Delta_{21} \\
&\quad - 4|U_{\alpha 2}|^2|U_{\alpha 3}|^2 \sin^2 \Delta_{32} \\
&\quad - 4|U_{\alpha 3}|^2|U_{\alpha 1}|^2 \sin^2 \Delta_{31}
\end{aligned} \tag{1.3}$$

for the disappearance channel, with  $\Delta_{ij} = \Delta m_{ij}^2 L/4E$ . For experiments studying  $\nu_\mu$  disappearance, the following is an approximation of the probability:

$$\begin{aligned}
P(\nu_\mu \rightarrow \nu_\mu) &\simeq 1 - 4 \cos^2 \theta_{13} \sin^2 \theta_{23} \\
&\quad [1 - \cos^2 \theta_{13} \sin^2 \theta_{23}] \sin^2 \left( \frac{1.267 \Delta m_{32}^2 L}{E_\nu} \right)
\end{aligned} \tag{1.4}$$

as provided by Reference [10]. For the appearance channel most commonly studied,

$$\begin{aligned}
P(\nu_\mu \rightarrow \nu_e) &= |2U_{\mu 3}^* U_{e 3} \sin \Delta_{31} e^{-i\Delta_{32}} + 2U_{\mu 2}^* U_{e 2} \sin \Delta_{21}|^2 \\
&\approx P_{atm} + P_{sol} + 2\sqrt{P_{atm}}\sqrt{P_{sol}} \cos(\Delta_{32} + \delta)
\end{aligned} \tag{1.5}$$

as provided by Reference [9], where

$$\begin{aligned}
\sqrt{P_{atm}} &= \sin \theta_{23} \sin 2\theta_{13} \sin \Delta_{31} \\
\sqrt{P_{sol}} &= \cos \theta_{23} \cos \theta_{13} \sin 2\theta_{12} \sin \Delta_{21} \approx \cos \theta_{23} \cos \theta_{13} \sin 2\theta_{12} \Delta_{21}
\end{aligned} \tag{1.6}$$

The CP violating phase is encoded in the interference term of the appearance probability mainly through its dependence on  $U_{e3}$ , which is why experiments looking for  $\nu_\mu, \bar{\nu}_\mu$  oscillating to  $\nu_e, \bar{\nu}_e$  are sensitive to CP violation.

In addition to the three ‘‘mixing angles’’ describing the mixing of states, encoded in the oscillation are the differences between the relevant mass states, called mass splittings. These appear in the equations for oscillation via  $\Delta m_{ij}^2$ , i.e.  $\Delta m_{21}^2$  and  $\Delta m_{31}^2$  (where  $\Delta m_{31}^2 = \Delta m_{32}^2 + \Delta m_{21}^2$ ). The

states are generally labeled conveniently such that the  $m_2$  state is more massive than the  $m_1$  state, such that  $\Delta m_{21}^2 > 0$ . It is, however, unknown if the  $m_3$  state is the heaviest or lightest, meaning the sign of  $\Delta m_{31}^2$  ( $\Delta m_{32}^2$ ) is an important unknown. This is often called the mass ordering or mass hierarchy question. As the Particle Data Group's (PDG) review on particle physics [4] points out, the hierarchy question more precisely includes the scales of the masses themselves: is  $m_3$  similar in scale to the others or is it very different? For example, the states could be such that the difference between two is much larger than the difference between one of these and the third, or it could be that  $m_1 \approx m_2 \approx m_3$ . Therefore, the sign of  $\Delta m_{31}^2$  is more accurately called the ordering rather than the hierarchy. As the current data places the masses at different scales,  $10^{-5}\text{eV}^2$  for  $\Delta m_{21}^2$  and  $10^{-3}\text{eV}^2$  for  $\Delta m_{31(32)}^2$  [4], the sign is currently the most interesting part of the question. Thus, in what follows, this will all be referred to as hierarchy. By a similar virtue to the naming of  $\Delta m_{21}^2$  the solar mass splitting,  $\Delta m_{31}^2$  is sometimes referred to as the atmospheric mass splitting  $\Delta m_{atm}^2$ .

For neutrinos traveling through a medium, for example accelerator-produced neutrinos traveling through the Earth, the neutrinos are subject to scattering off of electrons in the medium. Unlike the other species,  $\nu_e$  and  $\bar{\nu}_e$  have contributions from CC as well as NC scattering. This adds a potential to the Hamiltonian describing the oscillations, as  $V(x) = \sqrt{2}G_F N_e(x)$  with  $G_F$  being the Fermi constant and  $N_e(x)$  being the electron density [9]. The effect of matter to change the oscillation is known as the Mikheyev-Smirnov-Wolfenstein (MSW) Effect [11, 12]. A larger electron density yields a larger potential, and it was also realized that a changing density profile could lead to complicated evolution of the oscillation, e.g. an adiabatic density profile [12]. Given the right conditions, a resonant enhancement of neutrino oscillation could occur [12]: this resonant potential is  $V(n^{res}) = \cos 2\theta \frac{\Delta m^2}{2E}$  [13], where for simplicity a single mixing is considered. Such consideration is important in the case of solar neutrinos, which must propagate through the sun before traveling to Earth. For solar neutrinos above 2MeV, the resonance condition is met [13].

For accelerator-based neutrino experiments like NOvA and DUNE, where neutrinos travel under rock of the continental United States, a reasonable matter density is  $2.84\text{g/cm}^3$ . This matter potential affects neutrinos and antineutrinos in an opposite manner and depends on the hier-

arcy/ordering. The expression governing the oscillation (appearance) probability is now:

$$\begin{aligned}
P(\nu_\mu \rightarrow \nu_e) \cong & \sin^2 \theta_{23} \sin^2 2\theta_{13} \frac{\sin^2(\Delta_{31} - aL)}{\Delta_{31} - aL} \Delta_{31}^2 \\
& + \sin 2\theta_{23} \sin 2\theta_{13} \sin 2\theta_{12} \frac{\sin(\Delta_{31} - aL)}{\Delta_{31} - aL} \Delta_{31} \frac{\sin(aL)}{aL} \Delta_{21} \cos(\Delta_{31} + \delta) \\
& + \cos^2 \theta_{23} \sin^2 2\theta_{12} \frac{\sin^2(aL)}{(aL)^2} \Delta_{21}^2
\end{aligned} \tag{1.7}$$

where  $a = G_F N_e / \sqrt{2}$  and if considering antineutrinos instead,  $a \rightarrow -a$  and  $\delta \rightarrow -\delta$  [9].

The oscillation probabilities for neutrinos depend explicitly on the ratio  $L/E$ , leading to several points of maximum oscillation on an energy spectrum. Experiments that can choose the neutrino energy and/or baseline can craft parameters to be near one or more of these maximal locations. Figure 1.2 shows the probability for  $\nu_e$  appearance in a beam of  $\nu_\mu$  using Equation 1.7 at different baselines for a few current and next-generation experiments. The so-called ‘‘first oscillation maximum’’ is the one at the highest energy and is generally the most accessible with accelerator based experiments. However, note that the second oscillation maximum in DUNE occurs at an energy still approximately at 1 GeV, which can potentially be utilized.

The probability for oscillation from one flavor to another (appearance) is especially sensitive to CP violation in a way that the survival probability (say  $\nu_\mu \rightarrow \nu_\mu$ ) is not. This is particularly visible using a biprobability plot looking at the relative amounts of  $\nu_\mu \rightarrow \nu_e$  and  $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ . In practice, experimentally it is easier to make a bi-event plot that features the appearance candidates observed: the explanation is basically the same and the numbers are more meaningful in an experimental sense. An example of a biprobability plot for the NOvA experiment is given in Figure 1.3 and the description of the features is given in the caption.

## 1.2 The Oscillation Landscape

Some early experimental evidence for neutrino oscillation was noted by the Kamiokande Collaboration in 1994: not only did they find a ratio of  $\nu_\mu/\nu_e$  from atmospheric neutrinos which differed from their simulation, they found a deficit of muon-type in the upward direction with generally better data/MC agreement in the downward direction [15]. This could be interpreted as arising from  $\nu_\mu$

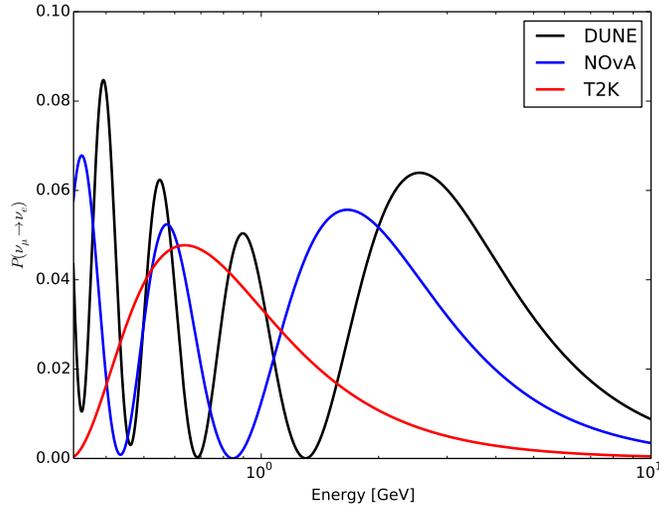


Figure 1.2: The probabilities for  $\nu_\mu$  to oscillate to  $\nu_e$  as a function of neutrino energy at different baselines. The effects of  $L/E$  can be clearly seen, as the appearance maximum occurs at a different energy at each baseline. For convenience most of the oscillation parameters are those in Table 1.2, but explicitly setting  $\delta_{CP} = 0$  and  $\sin^2 \theta_{23} = 0.5$ .

oscillating to  $\nu_e, \nu_\tau$  over the much larger distance for upward going events (they will have travelled through the Earth). Fitting this data in terms of a  $\theta$  and  $\Delta m^2$  for  $\nu_\mu$  to  $\nu_e$  and  $\nu_\mu$  to  $\nu_\tau$  oscillations gave a region of “allowed” oscillation parameters. The result seemed to point to a  $\Delta m^2$  that is approximately an order of magnitude higher than where the corresponding  $\Delta m_{32}^2$  is now understood to be [15, 16]. The IMB experiment also used an underground water detector in the 1980s and early 1990s to study potential nucleon decay and neutrinos. While the experiment provided early hints of a possible lower atmospheric  $\nu_\mu$  flavor fraction [17], some subsequent searches were more conservative [18, 19].

Strong evidence for neutrino oscillation came from Super-Kamiokande in 1998 and the Sudbury Neutrino Observatory (SNO) in the early 2000s, for which these experiments (namely leaders Takaaki Kajita and Arthur McDonald) were awarded the Nobel Prize in 2015. Super-Kamiokande, which studied atmospheric neutrinos, will be discussed in Chapter 2. SNO is a detector using heavy water with deuterium ( $^2\text{H}$ ) atoms instead of the common hydrogen. It is sensitive to the  $^8\text{B}$  solar neutrinos and measures three interaction modes

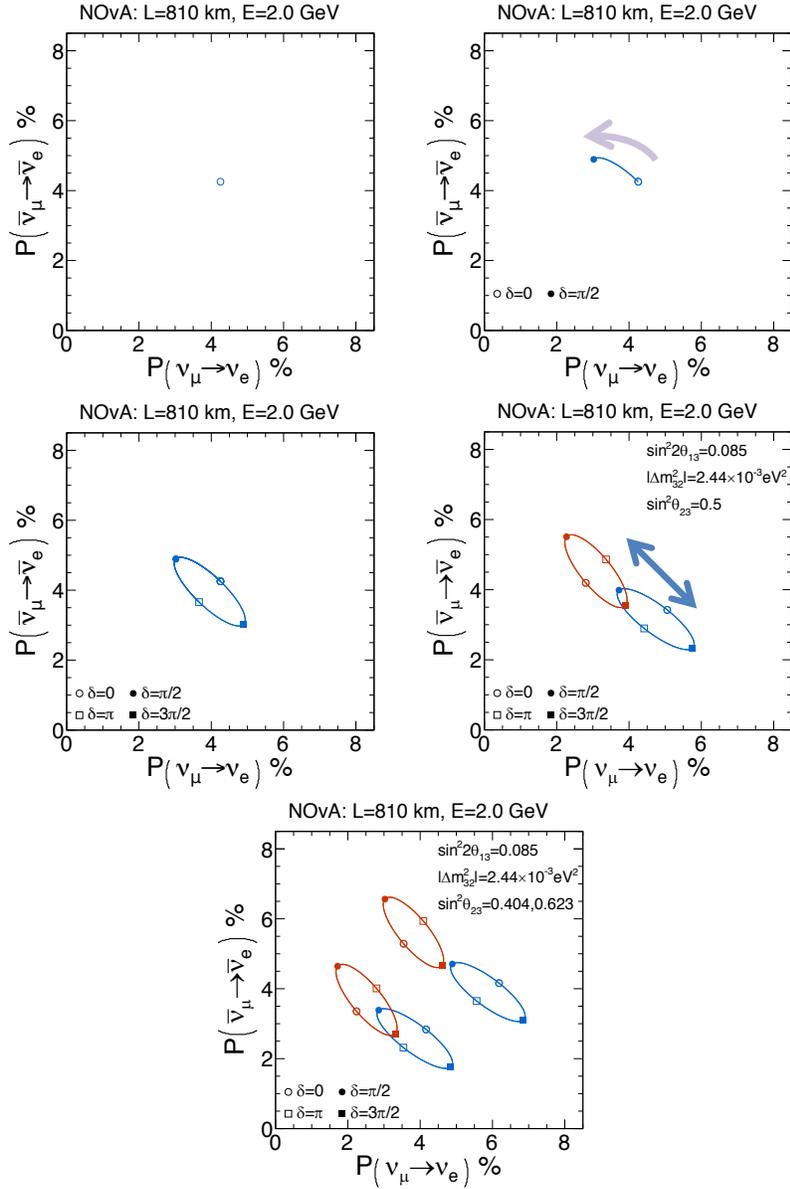


Figure 1.3: These plots show effects of parameters on the oscillation probability for muon neutrinos and antineutrinos to appear as electron neutrinos and antineutrinos, using a distance and energy relevant to NOvA to tell the story [14]. In the absence of CP violation and in vacuum, the oscillation probability is the same for both neutrinos and antineutrinos (top left). CP violation introduces an asymmetry between neutrinos and antineutrinos (top right), which over the  $2\pi$  phase space sweeps out an oval (middle left). The matter effect introduces further asymmetry between neutrinos and antineutrinos, dependent on the hierarchy (middle right). Here, blue is the normal hierarchy and red inverted. Finally, the  $\sin^2 \theta_{23}$  term causes an overall shift in probability for both neutrinos and antineutrinos, giving the overall picture in the bottom plot.

1.  $\nu_e + d \rightarrow p + p + e^-$  (CC)
2.  $\nu_x + e^- \rightarrow \nu_x + e^-$  (ES)
3.  $\nu_x + d \rightarrow p + n + \nu_x$  (NC)

where the CC process only couples to the electron neutrinos, the NC is sensitive to all three flavors the same, and elastic scattering (ES) is sensitive mostly to  $\nu_e$  but with some sensitivity to  $\nu_\mu$  and  $\nu_\tau$  also [20]. The collaboration measured the flux of neutrinos participating in these three reactions, which then allowed them to convert into a flux of flavor components

1.  $\phi_e = 1.76_{-0.005}^{+0.005}(\text{stat})_{-0.009}^{+0.009}(\text{syst})$
2.  $\phi_{\mu\tau} = 3.41_{-0.45}^{+0.45}(\text{stat})_{-0.45}^{+0.48}(\text{syst})$

where the  $\phi_{\mu\tau}$  is  $> 5\sigma$  from 0, a strong indication of flavor oscillation (recall the neutrinos are produced as  $\nu_e$ ) [20]. Furthermore, a study of the NC flux showed general agreement with the standard solar predictions [20]. Similar results to those of Super-Kamiokande were reported by other experiments, such as Soudan-2 (e.g. Reference [21]) and MACRO (e.g. Reference [22]).

The experimental landscape for neutrino oscillation and neutrino parameters (cross-section, mass, Majorana/Dirac) since the initial discoveries has blossomed. Because the major variables affecting the scale of oscillations accessible to a detector are the energy of neutrinos and the length scale, experiments using different detectors and different neutrino sources have varying sensitivity to the oscillation parameters. As an example thinking exercise, the PDG has a nice semi-quantitative example [4], as follows. The experimental location must be at least approximately an order of the oscillation length to allow the neutrinos to appreciably oscillate. The oscillation length scale associated with a given mass splitting  $\Delta m_{ij}^2$  is

$$L_{osc,ij} = 2.48m \frac{p[\text{MeV}]}{|\Delta m_{ij}^2|[\text{eV}^2]} \quad (1.8)$$

and requiring  $L_{osc,ij} \lesssim 2\pi L$  gives a general scale for the minimum  $\Delta m^2 \sim 2E_{avg}/L (> 0)$  that an experiment can probe [4]. These and the given neutrino sources are given in Table 1.1.

Source	Example Experiment	Flavor	$E_{avg}[MeV]$	L[km]	$\Delta m^2[eV^2]$
Reactor	Daya Bay	$\bar{\nu}_e$	$\sim 1$	1	$\sim 10^{-3}$
Reactor	KamLAND	$\bar{\nu}_e$	$\sim 1$	100	$\sim 10^{-5}$
Accelerator	MiniBooNE	$\nu_\mu, \bar{\nu}_\mu$	$\sim 10^3$	1	$\sim 1$
Accelerator	NOvA	$\nu_\mu, \bar{\nu}_\mu$	$\sim 10^3$	1000	$\sim 10^{-3}$
Atmospheric	Super Kamiokande	$\nu_\mu, \bar{\nu}_\mu, \nu_e, \bar{\nu}_e$	$\sim 10^3$	$10^4$	$\sim 10^{-4}$
Sun	SNO	$\nu_e$	$\sim 1$	$1.5 \times 10^8$	$\sim 10^{-11}$

Table 1.1: Ability of different sources/experiment classes to measure various  $\Delta m^2$  ranges [4], with an example experiment listed for each type.

Many of the parameters are currently fairly well measured, though very little is known about the CP violating phase, and improvement is expected in several others with more experimentation. The current state of measurements in the field is exemplified by looking at “global fits” combining the results of several experiments. Table 1.2 shows an example of such a result [4]. Other more recent fit results [23, 24] provide similar values, though with preference for  $\sin^2 \theta_{23} > 0.5$ .

Parameter	Best Fit NH	Best Fit IH	Uncertainty
$\Delta m_{21}^2 [10^{-5}eV^2]$	7.37	–	$3\sigma : 6.93 - 7.96$
$\Delta m_{31}^2 [10^{-3}eV^2]$	2.56	2.54	NH: $3\sigma : 2.45 - 2.69$ IH: $3\sigma : 2.42 - 2.66$
$\sin^2 \theta_{12}$	0.297	–	$3\sigma : 0.250 - 0.354$
$\sin^2 \theta_{23}$	0.425	0.589	NH: $3\sigma : 0.381 - 0.615$ IH: $3\sigma : 0.384 - 0.636$
$\sin^2 \theta_{13}$	0.0215	0.0216	NH: $3\sigma : 0.0190 - 0.0240$ IH: $3\sigma : 0.0190 - 0.0242$
$\delta_{CP}$	$1.38\pi$	$1.31\pi$	NH: $2\sigma : 1.0\pi - 1.9\pi$ IH: $2\sigma : 0.92\pi - 1.88\pi$

Table 1.2: Best fit neutrino oscillation parameters from several experiments combined, with ranges of uncertainty [4]. Where a value is only given in the normal hierarchy (NH), it assumed the same in both cases. In the case of the inverted hierarchy (IH), the “31” mass term is instead “23.”

Using the parameters in Table 1.2 for recent normal hierarchy best-fit parameters (but assuming  $\delta_{CP} = 0$ ) and plugging into the various trigonometric functions in the PMNS matrix as in Equation

1.2, one can get an approximation for the scale of the coefficients guiding neutrino mixing.

$$U_{PMNS} \sim \begin{pmatrix} 0.8 & 0.5 & 0.1 \\ -0.5 & 0.6 & 0.6 \\ 0.3 & -0.6 & 0.8 \end{pmatrix} \quad (1.9)$$

Interestingly, many of these parameters are sizeable and the matrix is rather non-diagonal. This is in contrast to the CKM matrix that governs quark mixing, which is close to diagonal, suggesting the physics involved may be different [25].

$$|V_{CKM}| \sim \begin{pmatrix} 1 & 0.2 & 0.004 \\ 0.2 & 1 & 0.04 \\ 0.008 & 0.04 & 1 \end{pmatrix} \quad (1.10)$$

## CHAPTER 2

### Long-baseline Neutrino Oscillation Experiments

This chapter builds on the oscillation framework presented in Chapter 1 with a specific focus on the past, present, and some of the future of oscillation studies using long baselines. Long-baseline neutrino experiments utilize a long distance between neutrino production and measurement ( $L$  in formulae such as Equation 1.7), sometimes with a specific range of neutrino energy  $E$  to match the present understanding of neutrino oscillation parameters. For example, the NuMI beamline discussed herein produces neutrinos on the order of several GeV, which for atmospheric-style oscillation (under  $\Delta m_{32}^2$ ) has maximal  $\nu_e$  appearance probability at a distance of hundreds of kilometers.

The main experiments of focus in this thesis, NOvA and DUNE, are both primarily accelerator-based long-baseline neutrino oscillation experiments. The NOvA experiment will be discussed in more detail in Chapter 3 and the upcoming DUNE project will be discussed in more detail in Chapter 7. Therefore, this chapter will describe other long-baseline experiments for context. It will also cover the accelerator and beamline used to produce neutrinos for NOvA (and MINOS) is discussed as an example of the concept of accelerator-driven (long-baseline) neutrino studies.

The long-baseline program builds upon and is facilitated by the global neutrino program. Before talking about long-baseline experiments, first consider the reactor experiment Daya Bay. Daya Bay uses reactor antineutrino inverse beta decay, similar to the Reines and Cowan experiment. However, Daya Bay uses 8 detectors at several distances from the reactors. Two reactor cores are  $\sim 365\text{m}$  from two detectors in one near site, four reactor cores are  $\sim 505\text{m}$  from two detectors in a second near site, and these six cores average  $1663\text{m}$  from four detectors at the far site [26]. With attempts to precisely constrain energy resolution, backgrounds, etc. the experiment aims to measure the spectra as a function of energy before oscillation at the near sites and after oscillation at the far sites. Using the near detectors' observations, a prediction of the far detectors' spectra is found by applying oscillation weights as a function of the mixing angle  $\theta_{13}$  and an effective mass splitting

[26]. Therefore best fitting parameters can be found. For recent Daya Bay results with over  $10^6$  neutrinos at the far site [26], the 99.7% confidence level allowed region covers  $\sim 0.075$  to  $\sim 0.096$  in  $\sin^2(2\theta_{13})$ , and the 1-dimensional  $\Delta\chi^2$  at  $3\sigma$  covers a similar but slightly tighter area than this. That  $\theta_{13}$  has been shown to be non-zero (and is as large as it is) cannot be understated for the long-baseline program. Recall Equation 1.7, in which  $\delta_{CP}$  appears in a term in the appearance probability. If  $\theta_{13} = 0$ , then this term would vanish from the appearance probability and CP violation would not be as accessible via these experiments. Furthermore, though it required large statistics and careful analysis, the precision achieved helps long-baseline appearance experiments. As will be discussed in this chapter and Chapter 4, experiments can constrain  $\theta_{13}$  in their own oscillation studies.

## 2.1 Atmospheric neutrino studies

Atmospheric neutrino studies refer to measurements performed on neutrinos produced as a result of the interactions of cosmic rays with Earth's upper atmosphere, where decays of produced hadrons (and further decays) such as

$$\begin{aligned}\pi^+ &\rightarrow \mu^+ + \nu_\mu, & \pi^- &\rightarrow \mu^- + \bar{\nu}_\mu \\ \mu^+ &\rightarrow \bar{\nu}_\mu + \nu_e + e^+, & \mu^- &\rightarrow \nu_\mu + \bar{\nu}_e + e^-\end{aligned}\tag{2.1}$$

produce neutrinos, leading to a ratio of  $\frac{\nu_\mu + \bar{\nu}_\mu}{\nu_e + \bar{\nu}_e} \approx 2$  [27]. Various factors contribute to the ratio not being exactly 2, for example rarer pion decay modes to electrons and  $\nu_e$ , the  $e + \nu_e$  channel from produced atmospheric kaons [28], or muons which do not decay in time. A strong deviation from this ratio would imply changes to the neutrino flavor states, i.e., oscillation. By utilizing a detector with pointing information, downward-going neutrinos can be distinguished from upward-going neutrinos. This is important for atmospheric neutrino studies since the neutrinos' origin points are not a known beam location but rather all over the Earth's atmosphere. In fact, downward-going events in atmospheric studies have a baseline of  $\approx 15$ -20km (they come from the sky above the detector) while upward-going events originate on essentially the other side of the planet, with a baseline of  $O(10,000)$  km [15, 27]. Therefore, by looking at events with similar energy as a function

of zenith angle, one probes oscillation as a function of baseline  $L$  in the oscillation equations. Given Table 1.1 in Chapter 1, events of order GeV at baselines of order  $10^4$  km are sensitive to mass splittings around  $10^{-4} \text{eV}^2$  or greater. The current best understanding of  $\Delta m_{31}^2$  is around  $\pm 2.5 \times 10^{-3} \text{eV}^2$  [4]. Early and much more recent Super-Kamiokande results will be discussed in this section. The earlier predecessor, Kamiokande, was mentioned in Chapter 1 for its early hints of neutrino oscillation in 1994 [15].

### 2.1.1 *Super-Kamiokande*

Super-Kamiokande is a 50 kiloton water detector (cylinder of 42m tall, 38m diameter) located in the Kamioka mine in Japan at depth 2,700 meters-water-equivalent, with a corresponding cosmic ray rate of 3 Hz [29]. The inner 22.5kt region contains the main fiducial volume, filled with ultrapure water and viewed nominally by 11,146 photomultiplier tubes [27]. The outer region of the tank is used as a veto and to alert when an event leaves the inner volume, known as a partially contained event [27]. A schematic drawing of Super-Kamiokande is found in Figure 2.1 [29]. It operates as a water Čerenkov detector, wherein photon detectors (such as photomultiplier tubes) sense particles traveling faster than the speed of light in a medium ( $v > c/n$ ). This is because these particles produce Čerenkov radiation, a cone of light that forms a ring-like image when viewed by photosensors on the cylindrical frame of the tank. The leptons produced in CC interactions are generally above Čerenkov threshold and make these Čerenkov rings; events with only one ring can then be classified as  $\mu$ -like or e-like based on the properties of the ring.

Having taken data since 1996, exposure is typically measured in kiloton-years (kton yr) equal to the volume multiplied by the years of active data. For example, the original 1998 analysis was with 33.0 kton yr of data, which is 535 days of data [27]. Super-Kamiokande has had four separate data-taking periods between 1996 and the analysis in 2018 [30]. These data-taking periods are separated where detector conditions were altered. SK-I consists of data between 1996 and 2001, after which many photomultiplier tubes were destroyed by a shock wave caused by one tube spontaneously imploding [30]. Following this event and subsequent work, SK-II took data from 2002 to 2005 with only 5,137 photomultiplier tubes: SK-III (2006-2008) and SK-IV (2008 on) have the nominal

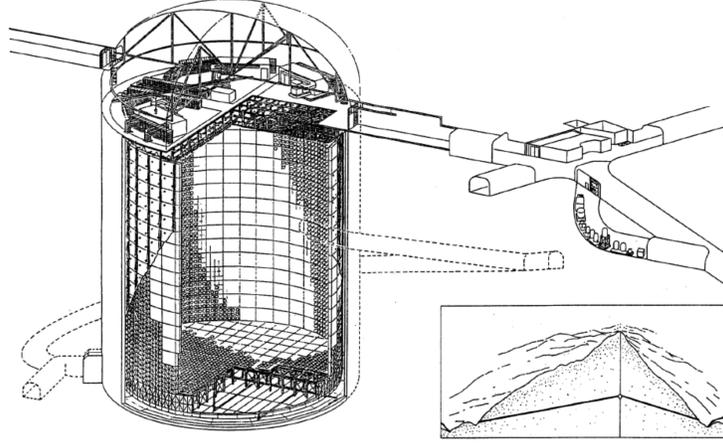


Figure 2.1: Drawing of Super-Kamiokande detector and site, beneath the mountain Ikeno-yama in Japan [29].

PMT count and are distinguished by an upgrade in the electronics [30].

### Early evidence for neutrino oscillation

With little known of neutrino oscillation and related parameters, a two neutrino approximation was adopted. Instead of the three flavors and related two mass splittings and three mixing angles, the probability for a flavor  $a$  to oscillate to  $b$  was expressed approximately as

$$P_{a \rightarrow b} = \sin^2 2\theta \sin^2\left(\frac{1.27\Delta m^2(\text{eV}^2)L(\text{km})}{E_\nu(\text{GeV})}\right) \quad (2.2)$$

with 1 mixing angle and 1 mass splitting [27]. In an era of lower precision and with less knowledge of the nature of neutrino oscillations (size of the PMNS matrix elements, for example), such an approximation held sufficiently well for the case of  $\nu_\mu$  disappearance. The samples were broken into both “sub-GeV” events and “multi-GeV” events, with the split being placed at visible energy below or above 1.330GeV. Partially contained events - that is, events that began in the inner fiducial volume but were uncontained - are almost entirely muon-like. Therefore in the multi-GeV sample, these are considered to be muon-like events. For both samples, a double ratio of muon-like to electron-like events for data compared to that from Monte Carlo, i.e.  $R = \frac{\mu/e|_{data}}{\mu/e|_{MC}}$  is found

1. sub-GeV sample:  $R = 0.63 \pm 0.03$  (stat.)  $\pm 0.05$  (syst.)
2. multi-GeV sample:  $R = 0.65 \pm 0.05$  (stat.)  $\pm 0.08$  (syst.)

where both fully-contained and partially-contained events in the multi-GeV are used [27]. In addition to this suggestive result (both are significantly low), more insight is gained by looking at the asymmetry between downward-going events (zenith angle  $0.2 < \cos \theta < 1$ ) and upward-going events (zenith angle  $-1 < \cos \theta < -0.2$ ). The asymmetry defined by  $A = (U - D)/(U + D)$  shows strong deviation in the muon-like sample from the simulation, growing as a function of energy between a few hundred MeV and several GeV [27]. Less deviation is found in the electron-like sample. The results can be explained with  $\nu_\mu \rightarrow \nu_\tau$  type oscillations with large  $\sin^2 2\theta$  ( $> 0.82$ ) and  $5 \times 10^{-4} < \Delta m^2 < 6 \times 10^{-3} \text{eV}^2$  at a 90% confidence level [27]. The contours found from this experiment provide good evidence of neutrino flavor oscillation and are shown in Figure 2.2 along with the value of  $\Delta m^2$  that is currently favored by a multitude of experiments [16].

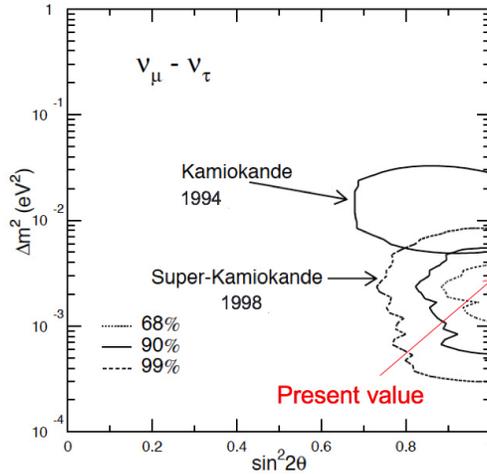


Figure 2.2: This figure shows the original results found by the Kamiokande collaboration in 1994 and the 1998 Super-Kamiokande oscillation results discussed here [16]. A more recent accepted value of the  $\Delta m^2$  value is pointed out in the figure as well. The Super-Kamiokande results match the presently accepted value which includes more experiments' measurements as well.

## Recent analysis

For more recent Super-Kamiokande analysis, as more has been learned about neutrino oscillation and matter effects and as statistical errors shrink and systematic errors become increasingly more important, Super-Kamiokande uses a full three-flavor oscillation model which includes the matter effect for neutrinos travelling various distances through the Earth. They use a simplified version of the Earth density profile with 4 layers (inner core, outer core, mantle, crust). As inferred from the lepton zenith angle for single-ring events, a neutrino will have passed through different sets of layers [30]. Due to the range of energies, path lengths, and varying density profiles, the oscillation probabilities span a wide range of possibilities. This oscillation pattern is displayed in the intricate features of a plot of zenith angle versus neutrino energy, as shown in Figure 2.3 [30]. In addition to these updates to analysis, recent analysis breaks the dataset into more categories, adding more information to the fits. For example, the effects of mass hierarchy and  $\delta_{CP}$  on the oscillation probability depend on if one is looking at neutrinos or antineutrinos (recall Equation 1.7 where the associated terms flip from positive to negative). Super-Kamiokande uses a statistical separation based on a likelihood variable to create separate  $\nu_e$  and  $\bar{\nu}_e$  enhanced samples [30].

Recent analysis in Super-Kamiokande has used 328 kton yr of data from the four running periods discussed above [30]. Since Super-Kamiokande is not sensitive to  $\Delta m_{21}^2$  and  $\theta_{12}$  in the same manner as reactor or solar neutrino experiments, they take these values as constrained by those experiments in performing their analysis. They have also performed studies in which  $\theta_{13}$  (which as was noted has been measured by reactor experiments) and T2K add additional constraints (T2K uses Super-Kamiokande as a far detector, see below). In the version of their analysis that does not use T2K or  $\theta_{13}$  as an additional constraint, Super-Kamiokande finds a slight preference for the normal hierarchy, with a best fit point in the normal hierarchy preferring  $\delta_{CP} = 4.18_{-1.66}^{+1.45}$  (on a 0 to  $2\pi$  range) and  $\Delta m_{32,31}^2 = 2.50_{-0.31}^{+0.13} \times 10^{-3} \text{eV}^2$  [30]. Adding a constraint on  $\theta_{13}$  did not change these conclusions much [30].

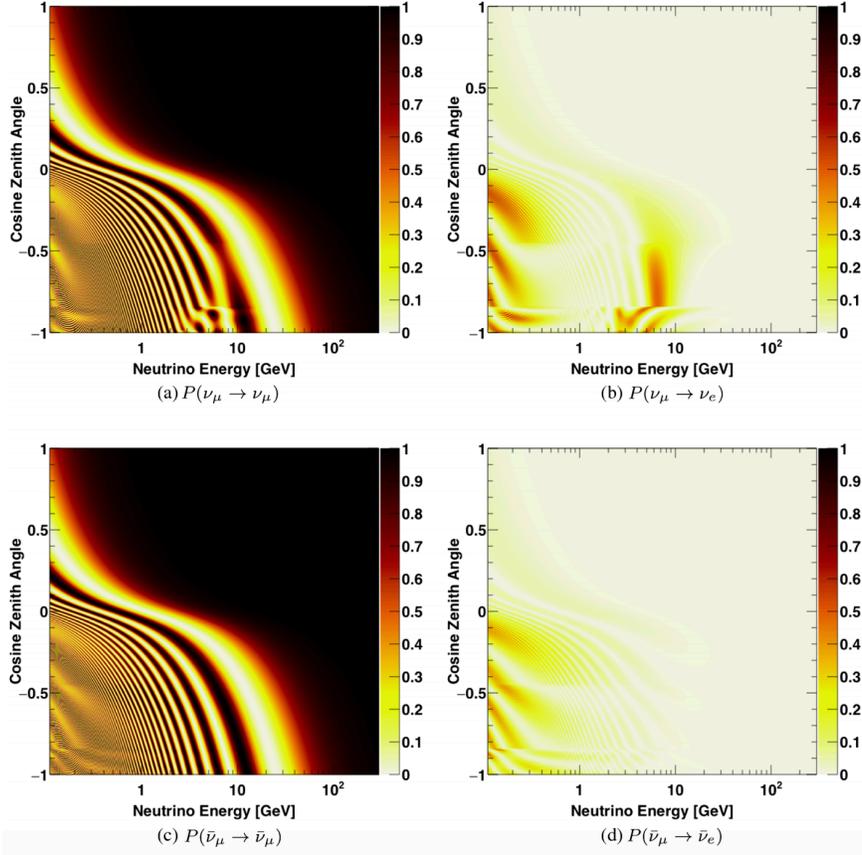


Figure 2.3: Neutrino oscillation probabilities for  $\nu_\mu$  and  $\bar{\nu}_\mu$  survival and  $\nu_e$  and  $\bar{\nu}_e$  appearance as functions of neutrino energy and zenith angle in Super-Kamiokande [30].

## 2.2 Neutrinos at the Main Injector

While studies of atmospheric neutrinos provide one path to investigating oscillations on long-baselines, neutrino production in accelerators provides an additional human-made and more customizable source. The energy of the neutrinos can be tuned based on the characteristics of the beam, and the distance to the detector(s) is also a chosen parameter<sup>1</sup>. Therefore, both L and E in the oscillation probability formula may be “chosen” to match the conditions desired.

The idea of using a beamline to produce neutrinos from meson decays was discussed in the previous chapter in relation to the discovery of  $\nu_\mu$ . In principle the concept is the same for accelerator-based neutrino experiments at present, but with many improvements. The example of the Neutrinos

<sup>1</sup>The actual freedom in this parameter may depend on borders (if the detector is to remain in one country for instance) or on the location of terrain, existing structures, etc.

at the Main Injector (NuMI) beamline, originally implemented for the MINOS oscillation experiment, will be used herein to detail accelerator-driven neutrino experimentation. NuMI is the beamline in operation for NOvA at Fermilab - a major component of this thesis. Fermilab's accelerator complex utilizes multiple stages in delivering the neutrino beam [31]:

1. H atoms are stripped of electrons, and the protons are accelerated to 400MeV in a linear accelerator (LINAC) and injected into a synchrotron known as the Booster
2. Protons in the Booster are accelerated to 8 GeV in  $1.6\mu\text{s}$  long batches and injected into another synchrotron known as the Main Injector
3. The Main Injector accumulates several (normally 6) Booster batches and accelerates the protons to 120 GeV
4. Once protons are at 120 GeV and in the right bunching, they are extracted and sent on to the NuMI target hall.

When running at 650kW there are typically  $\sim 5 \times 10^{13}$  protons per pulse<sup>2</sup>. The beamline is currently capable of running at over 700kW and exposure for analysis is often characterized by the number of protons on the target (POT). For NOvA, recent analysis has  $\sim 10^{21}$  POT in both neutrino and antineutrino mode, which is chosen by selecting the current direction in the horn (see below). The accumulated POT in both modes in NOvA is shown in Figure 2.4.

A schematic of NuMI is displayed in Figure 2.5 [31], and the components will be briefly discussed here. The NuMI target is a  $\sim 1.2\text{m}$  long series of mostly graphite target “fins” housed in a vessel with Be windows and filled with He [32, 33, 34]. 120GeV protons pass through the Be windows and strike the target, where proton interactions on the C atoms produce high multiplicities of hadrons in the final state. Most of these hadrons are pions that decay dominantly to  $\nu_\mu$  and  $\bar{\nu}_\mu$ . NuMI's target is displayed in a schematic form in Figure 2.6 [32]. Downstream of the target are two focusing horns used to correct the paths taken by mesons exiting the target. These are 3-4m

---

<sup>2</sup>These were the values obtained during a shift block I had in March 2019.

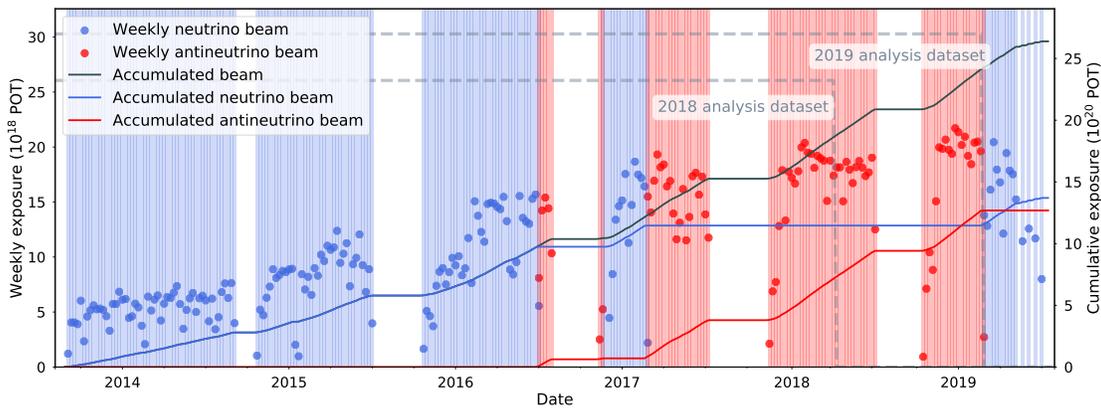


Figure 2.4: Accumulated POT in NOvA from the start of running in neutrino (blue) and antineutrino (red) mode. Each data point is for a given week: large gaps mostly correspond to the annual accelerator shutdowns.

long, curved, azimuthally-symmetric pieces of aluminum with a parabolically-changing radius as a function of position along the beam, around which large currents (up to 205kA) are pulsed [31]. This pulsed current creates a magnetic field that does not affect particles going straight along the beam direction, but will bend charged particles that left the target system at an angle. The charge sign of the pions and kaons that are focused depends on the horn current direction (generally called forward or reverse). In this way, one can select mostly positively-charged or mostly negatively-charged hadrons to make up the beam. Particles of the “wrong” sign are mostly defocused, though some will remain (especially those already traveling along the beam direction). Figure 2.7 [31] shows some of the various hadron trajectories, with hadrons being focused or defocused in the horns. In this diagram, the forward horn current (FHC) is depicted and the positively charged particles are focused while the negatively charged particles are defocused. In reverse horn current (RHC) operation, it is the opposite.

Focused particles (largely pions and kaons) are then directed toward a long decay pipe (2m diameter and 675m in length [31]) such that the particles can decay in flight and produce a boosted beam of daughter products. A guideline number is that a 10GeV pion would travel roughly 700m before decaying and would produce a 4.2GeV neutrino [31]. At the end of the decay pipe is solid material to absorb the remaining hadrons (undecayed mesons, protons which left the target without

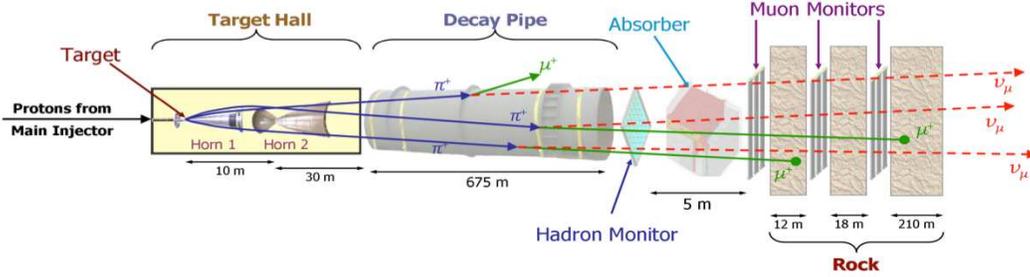


Figure 2.5: Depiction of the NuMI beamline in “forward horn current” ( $\nu_\mu$ ) operation [31]. 120 GeV protons interact with a graphite target producing hadrons (mostly  $\pi^\pm$  and  $K^\pm$ ) which decay to produce predominantly  $\nu_\mu$  and  $\bar{\nu}_\mu$ . The current running through two “horns” focuses one sign of hadron and defocuses the other, allowing for selection of predominantly  $\nu_\mu$  or  $\bar{\nu}_\mu$ . Not all hadrons with the wrong sign will be defocused, for example those going directly in line with the beam. Studying the wrong-sign component of the beam is the subject of Chapter 5.

interacting, muons, etc.), leaving just the beam of neutrinos to head toward detectors downstream. The flux at the detector depends on the location relative to the beam: NOvA is off-axis, and its flux will be discussed more in Chapter 3. The main component of the neutrino beam is  $\nu_\mu$  (FHC, focusing positively charged particles) or  $\bar{\nu}_\mu$  (RHC, focusing negatively charged particles), though there are small admixtures of  $\nu_e$  and  $\bar{\nu}_e$  from decays of  $K^\pm$ ,  $K^0$ , and muons [31].

### 2.2.1 Upgrades for the next generation

For the upcoming Deep Underground Neutrino Experiment (DUNE) project (see Chapter 7), the beam will again originate from Fermilab. However, a new beamline will be constructed to direct neutrinos to South Dakota, where the DUNE Far Detector will sit. Additionally, the Fermilab accelerator complex will be undergoing a project known as PIP-II (Proton Improvement Plan II), in which the LINAC will be replaced by a 250m long, superconducting 800MeV LINAC. The new beamline, including target, horns, and hadron absorber, will be set up such that the bulk of the neutrinos are in the desired energy range (broadly peaked between approximately 2-5GeV) and with improved power (ultimately reaching 2MW). The POT per year in DUNE is projected to be  $\sim 1.1 \times 10^{21}$  for a 120GeV proton beam.

## 2.3 Accelerator-driven long-baseline studies

The main focuses of this thesis are NOvA and DUNE, current and upcoming accelerator-driven neutrino oscillation experiments, respectively. There are a number of other accelerator-driven

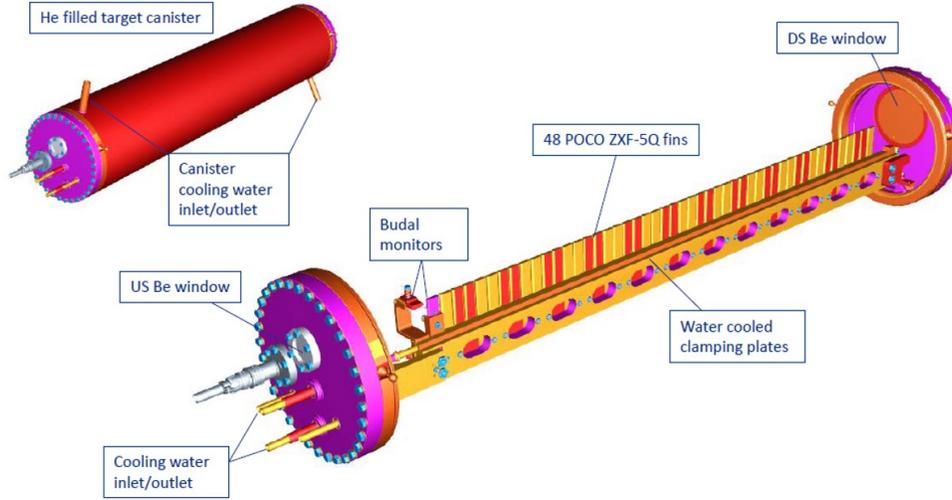


Figure 2.6: Depiction of the NuMI target system as used in the NOvA-era beam [32]. The target is housed in a He environment with Be windows to allow the initial protons to enter and protons and produced hadrons to exit. The target is  $\sim 1.2$  meters long and made of “fins” that are mostly graphite, though a few Be fins have also been tested to compare robustness [33, 34].

projects, and three major ones will be very briefly discussed here. These experiments often employ at least two detectors, with one close to the beam origin to observe the spectra before oscillations occur and one further away to measure the spectra and look for the effects of oscillations. Whereas one measures a range of baselines in one detector with atmospheric neutrino studies, multiple detectors are needed to achieve this from a single, human-made neutrino source.

### 2.3.1 Main Injector Neutrino Oscillation Search (MINOS)

MINOS is the predecessor to NOvA as a flagship neutrino oscillation experiment at Fermilab. MINOS operated two magnetized, tracking and sampling calorimeters (alternating steel and plastic scintillator) detectors located directly along the beam axis, with a Near Detector 1km from the NuMI target and a Far Detector located deep underground 735km from the beam origin, in the Soudan Underground Laboratory in northern Minnesota [35]. In its original configuration, MINOS searched for  $\nu_\mu$  and  $\bar{\nu}_\mu$  disappearance in the beam, with some additional sensitivity to  $\nu_e$  and  $\bar{\nu}_e$  appearance. Since the disappearance channel was strongest in this detector, its strongest limits were set on  $\Delta m_{32}^2$  and  $\sin^2 \theta_{23}$ . A plot of these contours including recent T2K, NOvA, and Super-Kamiokande results is given in Figure 2.8. MINOS ran for a few years after the start of NOvA

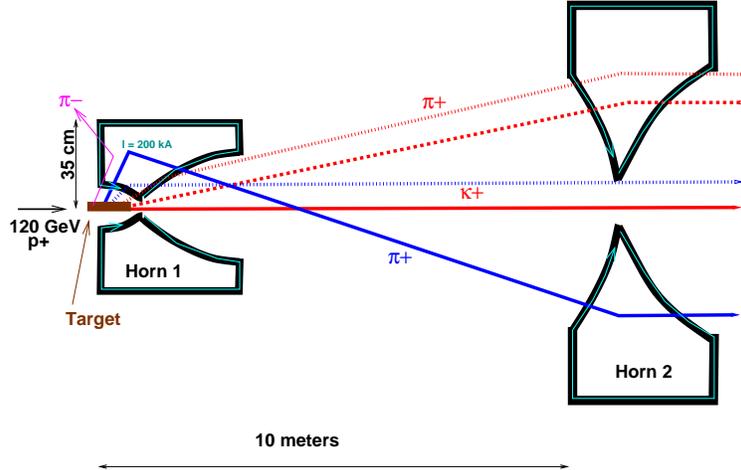


Figure 2.7: Depiction of the NuMI horn system in FHC ( $\nu_\mu$ ) operation [31], with some of the possible hadron trajectories. Note that a meson with the wrong sign ( $\pi^-$ ) is defocused, while mesons with the right sign can have their directions corrected by magnetic fields in the first and/or second horns. This focusing in turn leads to an increased neutrino flux at downstream detectors. Mesons going straight along the beam direction are unaffected by the magnetic fields.

in a higher energy configuration under the name MINOS+. Combined analysis of MINOS and MINOS+ data has set tight limits on potential oscillation to sterile flavors (separate from the 3 known active flavors) [35].

### 2.3.2 Tokai 2 Kamioka (T2K)

T2K is a currently running accelerator-driven neutrino oscillation experiment in Japan. T2K uses Super-Kamiokande as a far detector, with a beam peaked at 600MeV originating 295km away [37]. The beam is such that the Far Detector is off-axis, sculpting the neutrino energy spectrum to peak at a desired energy. This concept is discussed in Chapter 3. Two near detectors observe the neutrinos produced by the beam close to its origin. Note that both L and E are lower in T2K than in NOvA ( $\approx 2\text{GeV}$  and 810km), in such a way that both sit near the peak  $\nu_e$  appearance probability. Since the baseline is only 295km - significantly shorter than NOvA - the matter effect is diminished, as evidenced by bi-probability/bi-event plots. For a discussion of T2K and NOvA including a side-by-side comparison of bi-event plots, see Reference [16]. Situated at a shorter baseline, the Normal and Inverted Hierarchy ovals for T2K are mostly overlapping, as opposed to the more separable ovals in the NOvA version. This can be further seen in Figure 2.9, which summarizes recent results

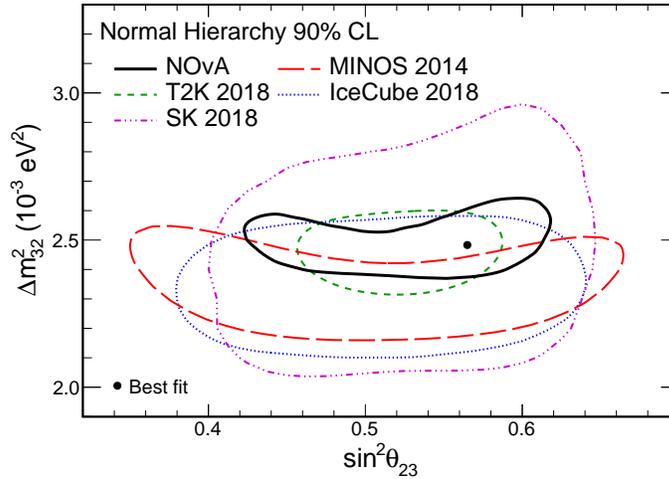


Figure 2.8: 90% confidence contours of  $\Delta m_{32}^2$  and  $\sin^2 \theta_{23}$  from several neutrino oscillation experiments [36], including MINOS, Super-Kamiokande, and T2K - all discussed in this chapter. The best fit shown is for recent NOvA results.

[38]. The observation of a larger number of neutrino ( $\nu_e$  appearance) candidate events than expected and correspondingly fewer antineutrino ( $\bar{\nu}_e$  appearance) candidates leads to a preference for nearly maximal CP violation with  $\delta_{CP}$  near  $-\pi/2$  (or  $3\pi/2$ ) in the normal hierarchy. This qualitative observation was also reported in previous analysis [37].

### 2.3.3 Hyper-Kamiokande and T2HK

Hyper-Kamiokande is the next generation upgrade to Super-Kamiokande, and T2HK will be the project using Hyper-Kamiokande as a far detector for accelerator-driven neutrino oscillation studies. Hyper-Kamiokande will employ a similar technology to Super-Kamiokande, using water Čerenkov but with improved photosensors (double the single-photon detection efficiency) and 187 kton fiducial volume (a factor of over 8 improvement) [39]. With its massive volume and a beam that will exceed 1MW, it is expected to provide precise measurements of neutrino oscillation parameters and have strong sensitivity to probe nucleon decay: the expected CP-violation sensitivity in 10 years of running is illustrated in Figure 2.10 [39]. It will run contemporaneously with DUNE, which will also provide precision measurements, as is discussed in Chapter 7.

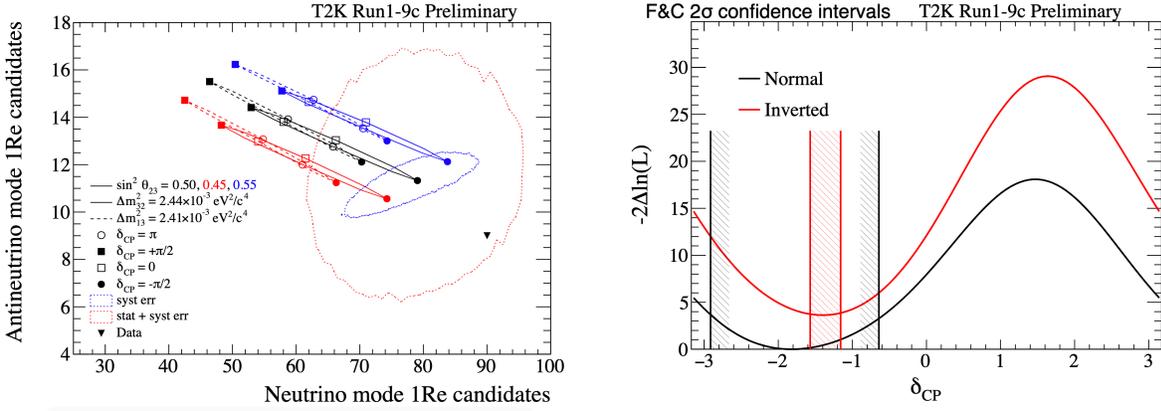


Figure 2.9: Recent analyses from T2K have a large number of neutrino candidate events and only a few antineutrino candidates [38, 37]. A bi-event plot is shown on the left [38]. Note that the Normal and Inverted Hierarchy ovals are mostly overlapping, as opposed to the more separable ovals in the NOvA version (Figure 1.3). This disparity in neutrinos and antineutrinos results in a preference for CP violation with  $\delta_{CP}$  near  $-\pi/2$  (or  $3\pi/2$ ) in the normal hierarchy [38].

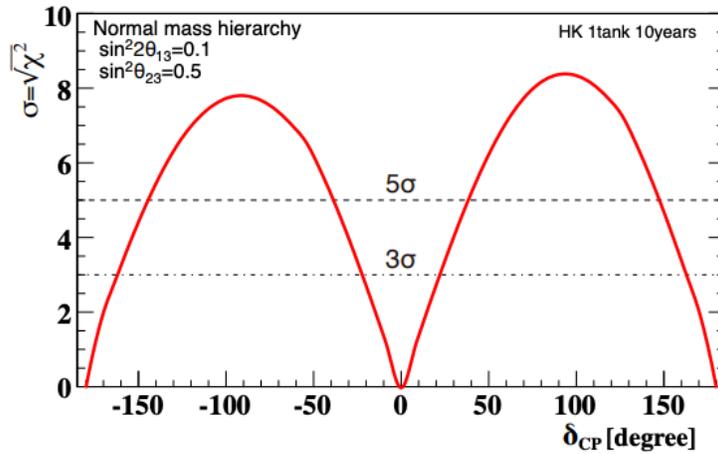


Figure 2.10: Significance to exclude CP conservation ( $\sin \delta_{CP} = 0$ ) given normal hierarchy and true  $\delta_{CP}$  in 10 years of running Hyper-Kamiokande [39]. Assuming  $\delta_{CP}$  is near  $\pm\pi/2$ , Hyper-Kamiokande is expected to surpass  $5\sigma$  sensitivity.

## CHAPTER 3

### The NuMI Off-axis $\nu_e$ Appearance Experiment

The NuMI Off-axis  $\nu_e$  Appearance (NOvA) Experiment is a current generation accelerator-driven long-baseline neutrino oscillation experiment. The experiment utilizes two functionally equivalent detectors separated by 810km to study the beam of neutrinos produced by the NuMI beamline (as discussed in Section 2.2) both before oscillations at the Near Detector (ND) on Fermilab campus, and after oscillations at the Far Detector (FD) in Ash River, Minnesota. The ND is located approximately 100m underground and its proximity to the beam origin provides a high flux for studying the spectrum produced by the beam, even with its modest size of 193t and a muon range stack at the downstream end. The FD is located on the surface and contains 14kt of material, to allow for some interactions even over 800km from the beam origin. Though there is a significant rate of cosmic activity in the detector, cuts on beam timing and other observables provide an ultimately low background, as will be discussed in chapters to come. At the baseline of 810km, the neutrino energy at which the oscillation probability  $P(\nu_\mu) \rightarrow P(\nu_\mu)$  is minimized – that is, when the  $\nu_\mu$  disappearance is maximized – is between 1-2GeV. The spectrum of neutrino energy from pion decay in the beam strongly depends on the angle at which one places a detector. This is discussed later in this chapter, and as is shown in Figure 3.6, the neutrino energy spectrum for a detector at  $\sim 14\text{mrad}$  is strongly concentrated between 1-2GeV. Therefore, the NOvA FD is placed  $14.6\text{mrad}$  from the beam axis.

The detector itself is described briefly in Section 3.1, followed by a discussion of the energy calibration in NOvA in Section 3.2. Then, the flux of neutrinos at NOvA and event classification will be discussed.

Note that NOvA is a collaborative effort, and figures with citations represent collective work by the larger collaboration. These are shown in the pursuit of a more complete narrative and context for my efforts. Figures without citations are generally produced through by me and/or through my

research efforts, though at times also in conjunction with collaborators.

### 3.1 The Detector

The NOvA detectors employ segmented liquid scintillator technology, allowing for both tracking and calorimetry. Each cell is extruded from titanium-dioxide loaded polyvinyl chloride (PVC) [40]. The cells are approximately 4cm by 6cm, with the FD cells being roughly 15.5m long and the ND cells being roughly 4m long. This gives the FD dimensions of roughly 16m x 16m transverse to the beam, and is about 60m along the beam direction (896 planes of cells). Therefore, the FD instruments 344,064 cells, each filled with the NOvA liquid scintillator blend.

The scintillator for NOvA was made in two variants. After the first two batches, the scope of the detector was downsized from 18kt to 14kt, meaning that the additional chemicals which had been purchased were now excess, so extra scintillant and wavelength shifter were added to the blend to increase light yield. The main chemicals of the scintillator are as follows [40]:

1. Mineral oil: solvent
2. 1,2,4-trimethylbenzene (Pseudocumene): primary scintillant (270-320nm), increased from 4.98% mass fraction to 5.23% in the second variant
3. 2,5-diphenyloxazole (PPO): wavelength shifter from scintillator photons to 340-380nm, increased from 0.11% mass fraction to 0.14% in second variant
4. 1,4-bis-(o-methyl-styryl)-benzene (bis-MSB): shift from PPO photons to  $\sim$ 390-440nm
5. Stadis-425: an antistatic to prevent buildup from the filling process
6. Vitamin E: antioxidant to prevent yellowing (which would reduce light yield with aging, see Section 3.2 and related appendix for more details).

In the coordinate system where  $z$  is the beam axis,  $x$  is the horizontal coordinate and  $y$  the vertical. The planes alternate between  $x$ -view (vertical cells) and  $y$ -view (horizontal cells) to allow for a full 3D reconstruction of particle tracks. In detector jargon, the planes are organized into groups of 32 called “blocks,” which naturally leads to groups of 64 being “diblocks.” In this terminology,

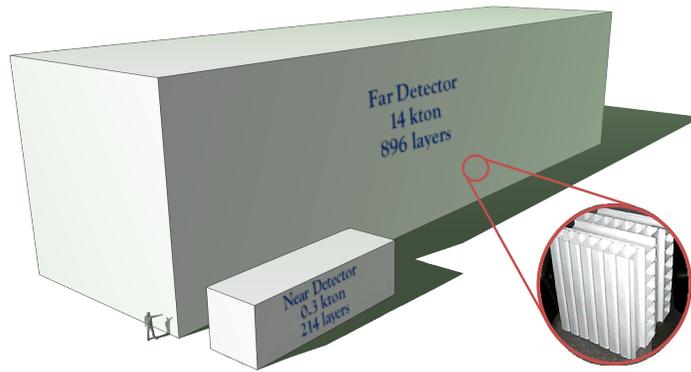


Figure 3.1: Depiction of the NOvA detectors, with a view of a few planes as well (adapted from Reference [41]).

there are 14 diblocks in the FD. The diblocks are arranged somewhat differently in the ND, due to its smaller structure and the addition of the muon catcher, leading to there being 4 diblocks. The muon catcher is a somewhat shorter and smaller region at the downstream end of the detector which interleaves planes of shorter NOvA-style cells with blocks of steel to serve as a muon range stack, significantly improving the range of muon energies the ND is sensitive to. Being so large, the FD is naturally sensitive to a wider range of muon energies and has no such catcher. A depiction of the NOvA detectors is presented in Figure 3.1.

The original scintillator blend fills the first diblock of the FD and only bits of the second; the second, somewhat brighter scintillator blend is used throughout the rest of the FD and the entire ND. Averaging out the chemicals for the modules and scintillator as used in the simulation, this “soup” has radiation length of  $\sim 35\text{cm}$ , which is about 5-6 planes in the beam direction. The stopping power is  $\sim 1.97\text{ MeV/cm}^2/\text{g}$  and  $Z/A \sim 0.55$ . This stopping power allows for the muon energy to be calculated from length and gives good particle tracking power, while the radiation length implies showers develop over a number of planes. Given that the products of  $\nu_e$  and  $\nu_\mu$  interactions are the prime interests, these characteristics are well-suited.

When charged particles are produced in interactions inside the NOvA detectors, the scintillant produces light, and the wavelength shifts lead to light that is  $\sim 390\text{-}440\text{nm}$ . Since the PVC is

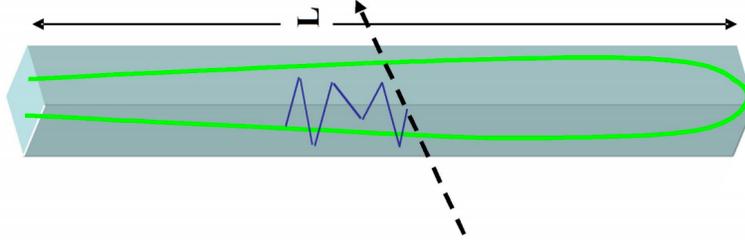


Figure 3.2: Depiction of a NOvA detector cell with the wavelength-shifting fiber loop. A charged particle traversing a cell will scintillate, producing light that is picked up by a fiber and guided to an Avalanche Photodiode (APD) at the ends (adapted from Reference [14]).

reflective, this light is allowed to reflect on cell walls. A wavelength-shifting fiber runs down the length of the cell and back, which converts the light to green and guides it out of the cell, where it is detected by an Avalanche Photodiode (APD) which produces an electronic signal with a pulse height corresponding to the amount of incident light. A NOvA cell is depicted in Figure 3.2.

Therefore, the more signal detected by the APD, the more light that was incident. In the absence of attenuation effects, the more light that was incident, the more energy that was deposited. In reality there is attenuation and other detector specific effects: a job of the calibration is to correct for these. Then, the job of calibration is determine what amount of signal corresponds to a given amount of energy deposit. This is the topic of the next section.

### 3.2 Energy Calibration

Any study of neutrino interactions or oscillations is ultimately dependent on the energy of the incident neutrino. The cross-sections are energy-dependent and oscillations are a function of  $L/E$ , with the baseline  $L=810\text{km}$  for NOvA being fixed. Therefore, a reasonably accurate accounting of the interacting neutrinos is important to properly measure oscillation parameters, cross-sections, and other quantities in the NOvA detectors. Since NOvA sees just the charged particles arising from the neutrino interactions, accurate accounting for the energies of these particles is necessary. Therefore, calibration must be performed to translate the signal read out from the cells into the correct amount of energy deposited to produce it.

The energy calibration in NOvA uses the abundance of cosmic ray muons in the detectors and

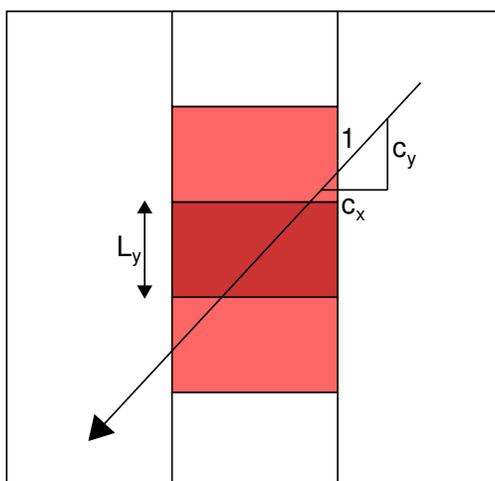


Figure 3.3: Cartoon illustrating the principle of tri-cell selection [42]: hits are required to have corresponding hits in both adjacent cells in the same plane to ensure they entered and exit through the side and did not clip a corner.

takes place in two stages. Selected hits are required to have corresponding hits on the adjacent two cells in the same plane. With this selection, called tri-cell hits, the muon is known to have entered and exited the cell through its side walls, allowing for precise calculation of the path length of the muon through the cell. Without this selection criterion, a muon could clip the corner of a cell and produce enough light to be recorded but without precise knowledge of path length. The principle of tri-cell hits are shown in Figure 3.3. The reconstructed path of the muon and the combination of hits in the three cells allows for accurate path length reconstruction and thus the  $dx$  in  $dE/dx$ .

The first stage accounts for position-dependent effects in the detector: attenuation, threshold, and shadowing. Most notably, the signal will be degraded as it travels from its origin to the readout, for example due to losses in the fiber. On the scale of the FD, with 16m-long cells and attenuation lengths of  $\sim 4\text{m}$ , such degradation can be appreciable. Further, shadowing effects are energy depositions in other parts of the detector can affect what is seen in particular cells. Finally, thresholding - the loss of low energy signals at the far end - leads to a bias towards higher energy deposits at the end of the cell. The correction factors for thresholding and shadowing are derived from the simulation, where the “true” (simulated) information is known, and is given by Equation 3.1 [42].

In this equation, the  $PE/\lambda$  term refers to the amount of light recorded from the simulated hit relative to the actual amount of simulated light (i.e. accounting for Poisson fluctuations), and the  $E_{\text{true}}/E_{\text{MIP}}$  term relates to the difference in true energy causing the recorded hit relative to the expectation for a minimum ionizing hit in the detector (roughly 1.78 MeV/cm).

$$T = \frac{PE}{\lambda} \frac{E_{\text{true}}}{E_{\text{MIP}}} \quad (3.1)$$

In fact, because the FD has been run in a lower and higher gain setting, the reduction of bias at the far end of the detector can be seen in the transition from low to high gain, as seen in Figure 3.4. These effects are accounted for on a cell-by-cell basis by fitting a functional form to the response as a function of position along the cell. After the application of the thresholding and shadowing corrections, an attenuation fit is performed. Additional effort is made to account for varying brightness of the fibers in the FD, as this was found to make some difference to the calibration and since the simulation of both detectors attempt to account for brightnesses. This accounts for degradation in signal from bulk attenuation in the middle of cells, and extra interpolations are used to handle the roll-off of signals seen at the near and far end of cells. The fit for attenuation is given in Equation 3.2 [42], where  $C$ ,  $A$ , and  $X$  are free parameters,  $L$  is the length of the cell, and  $y$  is the response. This functional form corrects the signal to an overall constant meant to represent roughly the response at the center of the cell, and this corrected unit is called PECorr (corrected PE). While this fit takes into account light from a hit going around the loop of fiber to the other end, it only uses one attenuation length. In fact, the fibers are simulated with two characteristic scales for attenuation: further thoughts about the implications of attenuation length are discussed in Appendix A.

$$y = C + A \left( \exp\left(\frac{W}{X}\right) + \exp\left(-\frac{L+W}{X}\right) \right) \quad (3.2)$$

The second energy calibration step focuses on pinning the light output of the detector to the amount of energy required to produce the signal: that is, MeV per PECorr. To accomplish this, NOvA again benefits from the abundance of cosmic rays, especially at the FD, by looking at stop-

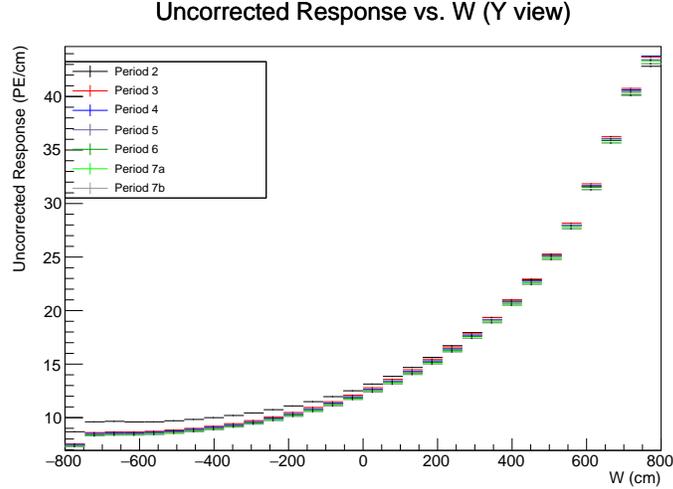


Figure 3.4: An approximate accounting of the average PE/cm recorded from hits as a function of position along cells in one view of the FD. Each color represents a different period of data-taking. The interesting point is that going from Period 2 to Period 3 a sharp decline in PE/cm per hit can be seen. This is likely due to a decrease in thresholding by increasing the gain between Period 2 and Period 3.

ping cosmic ray muons. Cells with failed attenuation fits are skipped, but beyond selection cuts, all cells in a given detector are in principle treated uniformly at this stage. That is, at the moment, each detector has one energy scale constant for a given period of data-taking, derived from a subset of the cosmic rays recorded in that period. In particular, hits in the last 1-2 meters of the muon track are used in this calibration, where the muon energy deposits are well-understood and where the energy deposits are less affected by slight inaccuracies in determining the exact stopping point of the muon (given the Bethe-Bloch curve and the non-infinitesimal size of NOvA cells). Additionally, though the intent of the attenuation calibration discussed above is to remove shape differences along the cell, this method is not perfect. There are still several regions along cells for which the recorded response in PECorr/cm is not uniform and flat. A selection of fairly flat PECorr/cm regions is taken to perform the energy scale calibration: these are  $w \in (-100, 100)$  in the ND where the readout end is near 200cm and the far end is near -200cm, and  $w \in (200, 600)$  in the FD where the readout is near 800cm and the far end is near -800cm (inexact since the cells are about 15.5m in length).

Tri-cell hits passing these selection criteria are then combined to find an average PECorr/cm and MeV/cm in the simulation, and PECorr/cm in the data. The combination of PECorr/cm and

MeV/cm (always from simulation) gives MeV/PECorr to calibrate the energy from hits in both data and simulation. Until recently, the values were fairly consistent [43]. The mean PECorr/cm is called the Muon Energy Unit (MEU), and because PECorr is set by the attenuation calibration, the MEU only varies a small amount from data period to data period: approximately 40.2-40.3 in the FD data and approximately 41.5-41.7 in the ND data. The mean energy deposition in the FD is found to be 1.786 MeV/cm in the simulation of Period 2 in the FD, 1.785 MeV/cm in the simulation of FD data periods after Period 2, and 1.794 MeV/cm in ND simulation. This gives a value of roughly  $4.3 \times 10^{-2}$  or  $4.4 \times 10^{-2}$  MeV for each PECorr detected. Efforts to check this energy scale against other tests (beam muons, protons,  $\pi^0$  mass peak reconstruction) reveal discrepancies at the order of a few percent, and a 5% systematic is taken on the energy scale to analyses.

For recent analysis, a decision was made to continue to check the energy scale calibration for any signs of change but to otherwise leave the calibration constants frozen. This decision was made due to time crunches and processing needs for the neutrino + antineutrino oscillation analysis, such as those described in Chapters 4 and 6 here. Interestingly, the calibration was found to be drifting downward over time: it turns out that this drift was masked by the fact that the attenuation calibration sets the response to a fixed scale. Since the attenuation calibration was left frozen when checks were performed on more recent data, the energy scale calibration became sensitive to these effects over time. The exact rate of the degradation in response is somewhat dependent on the diblock, especially in the FD, but in an averaged sense is  $\sim 0.78\%$  per year at the ND and  $\sim 1.36\%$  per year at the FD. Given the size of change and time of effect here and the nature of the energy calculations themselves, this effect is expected to be covered by systematics. For the second analysis of neutrino + antineutrino data (with more antineutrino data than the first), the new data were given recalibrated energy scale constants to in essence “reset” the drift. The exact nature of this drift is not wholly understood, but one could imagine scintillator degradation, fiber degradation, or electronics as potential causes. As I was involved in re-discovering this drift, a part of my research program focused on investigating the drift and its potential causes. Appendix A will go into detail on this line of study, but Figure 3.5 shows such degradation in PE/cm as a function

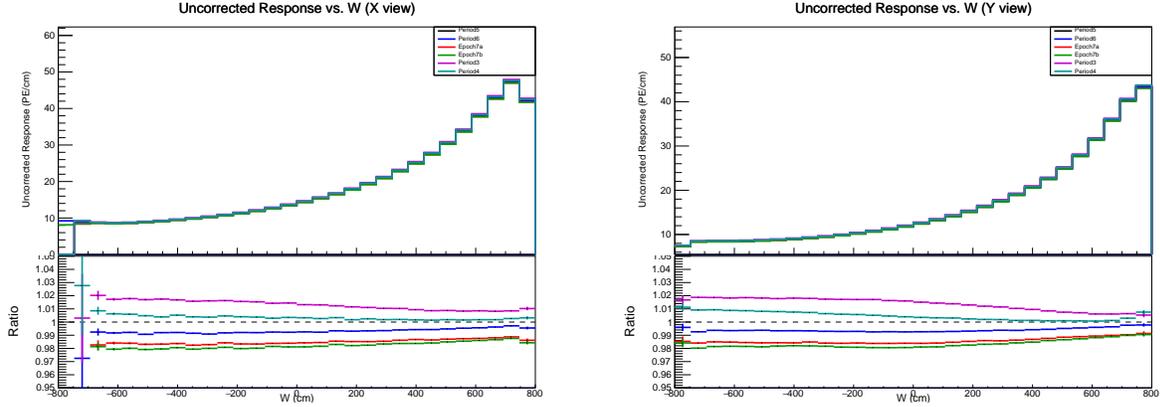


Figure 3.5: Average response in PE/cm vs position along cell in the FD. X-view is on the left and y-view on the right. Black is Period 5, and the ratios in the bottom of each plot shows each period with respect to Period 5. Blue, red, and green are more recent periods, and purple and teal are older periods. The detector shows clear signs of response degrading over time across the entire length of the cell. Perhaps the degradation is slightly less at the readout end (closer to positive end). At the readout end, the change from older periods to Period 5 and Period 5 to more recent periods appears potentially smaller than at locations further from the readout end.

of position in FD cells over a number of data-taking periods.

In addition to learning more about the drift and potential mitigation, an additional interesting lesson is that the attenuation calibration does not fully remove the spread in response among diblocks, especially at the FD. Table 3.1 shows the spread from the diblock with the lowest response to the diblock with the highest response in the FD both after the application of the attenuation calibration (left) and before (right). Clearly, the attenuation calibration does remove most of the spread, but there is still some appreciable amount of spread after the application of the attenuation calibration. This, combined with the revelation that the drift varies some from diblock to diblock (notice the spread getting larger as a function of data period), lead to the suggestion that future calibrations produce a different energy scale constant for each diblock.

### 3.3 NOvA Flux and Interactions

NOvA receives its neutrinos from the NuMI beam, as discussed in Section 2.2. With the increased (“medium”) energy configuration that the beam is now running in, the peak beam flux is at several GeV. NOvA, however, is situated 14.6 mrad off-axis, which serves to sculpt the incident neutrino flux to peak near the oscillation maximum over 810km. The predominant hadron from

	Lo MEU (Diblock)	Hi MEU (Diblock)	Diff 1-[Lo/Hi]	Lo PE/cm (Diblock)	Hi PE/cm (Diblock)	Diff 1-[Lo/Hi]
Period 2	39.91 (1)	40.26 (5)	0.0087	22.30 (3)	28.07 (9)	0.2056
Period 3	39.98 (1)	40.39 (9)	0.0101	22.14 (3)	28.21 (9)	0.2150
Period 4	39.94 (1)	40.36 (9)	0.0102	21.95 (3)	27.99 (9)	0.2158
Period 5	39.98 (1)	40.39 (9)	0.0102	21.91 (3)	27.93 (9)	0.2158
Period 6	39.73 (1)	40.18 (6)	0.0113	21.77 (3)	27.77 (9)	0.2161
Epoch 7b	39.19 (1)	39.83 (5)	0.0160	21.49 (3)	27.49 (9)	0.2184

Table 3.1: Spread in MEU [PECorr/cm] and PE/cm among FD diblocks. Period 2 is an older data-taking period made up of periods of time between 2014-2015. Each successive period is more recent, with Epoch 7b occurring from January to March 2018.

the proton-target interactions which yields neutrinos in this energy range in the NOvA detectors are pions, and the neutrinos from pion decays can be expressed by

$$E_\nu = \frac{m_\pi^2 - m_\mu^2}{2(E_\pi - p_\pi \cos \theta)} \quad (3.3)$$

where  $\theta$  is the emission angle of the neutrino relative to the pion direction [4]. Under the assumption that the pions are largely focused by the beam instrumentation,  $\theta$  then refers to the angle off-axis at which the detector is placed relative to the beam direction. Since  $p_\pi = \sqrt{E_\pi^2 - m_\pi^2}$ , the neutrino energy is dependent upon pion energy, as appears in Figure 3.6. This is a simplification for demonstration purposes, but note that the majority of neutrinos produced are within 1-2GeV for a 14mrad off-axis location, where the oscillation probability for NOvA is maximal.

To understand the flux of neutrinos in the NOvA detectors, simulations and data-driven corrections inform our knowledge. The base information for this comes from G4NuMI, which takes into account the beamline components (horn, target, etc.) as described in Section 2.2 and uses Geant4 to simulate the interactions with the FTFP\_BERT package [44], which refers to the models used to govern the interactions at various energy scales. On top of this, the fluxes are corrected using external data via a method called Package to Predict the Flux (PPFX), which is described in

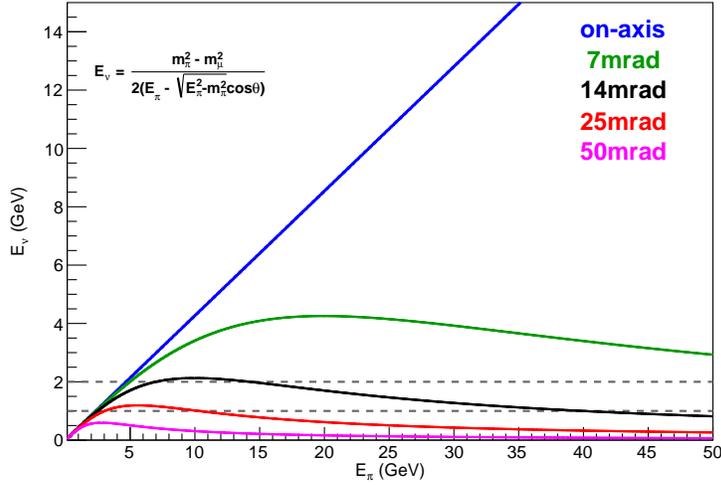


Figure 3.6: The neutrino energy produced from pion decays as a function of pion energy, at various off-axis angles. The black line represents the  $E_\nu$  spectrum at 14mrad, approximately the location of NOvA relative to the NuMI beamline. The majority of neutrinos produced are within 1-2GeV, where the oscillation probability over 810km is maximal.

[45, 44] and can be summarized as follows. Briefly, this package reweights neutrino likelihoods based on production measurements from external data, for example the  $\pi^+$  which produce neutrinos. For NOvA, this target data is used, and the weights are based on corrections for particle attenuation (whether a particle survived or interacted in traversing material) and particle interaction cross-sections. Such cross-section corrections can come directly from rescaling data available, for example from NA49, which studied hadron production at 158GeV, or from models where direct data is not available [44]. For  $\pi^+$  producing a neutrino which reaches NOvA, an example of the correction factors are presented in Figure 3.7, where the contour lines show the NOvA parent  $\pi^+$  flux in Feynman- $x$   $x_F$  (related to the longitudinal momentum and total) and the transverse momentum. The circles show the concentration of NA49 data (solid circles imply lower statistical error in NA49), and the colors indicate the correction factors. Since the reweighting is meant to correct the flux based on better data and models, the uncertainty in the flux is correspondingly reduced, leading to the ability to conduct more precise measurements. Figure 3.8 exemplifies the reduction in uncertainty (via reduced shifts) gained by utilizing the PPF weights.

In the end, this informs the flux of neutrinos which reaches the NOvA detectors, which is

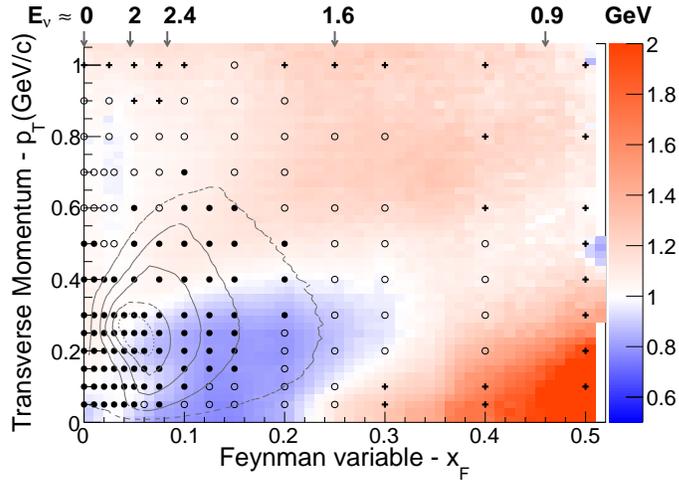


Figure 3.7: PPFX corrections scaled from NA49 data based on  $x_F$  and  $p_T$  for parent  $\pi^+$  [44]: the colors represent the correction factor and the circles represent the NA49 statistical errors. Closed circles are data points with less than 2.5% statistical error, open circles are data points with 2.5-5.0%, and crosses are data points with statistical errors greater than 5.0%. The contour lines show the concentration of  $\pi^+$  which create neutrinos in NOvA.

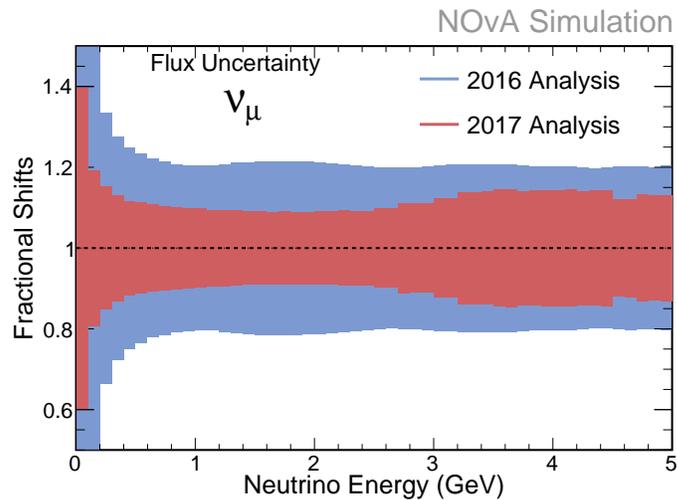


Figure 3.8: Fractional shifts from flux, hadron production, and focusing uncertainties in the  $\nu_\mu$  beam at the NOvA ND between two analyses [46]. The reduction in shifts for the more recent analysis exemplifies the utilization of PPFX.

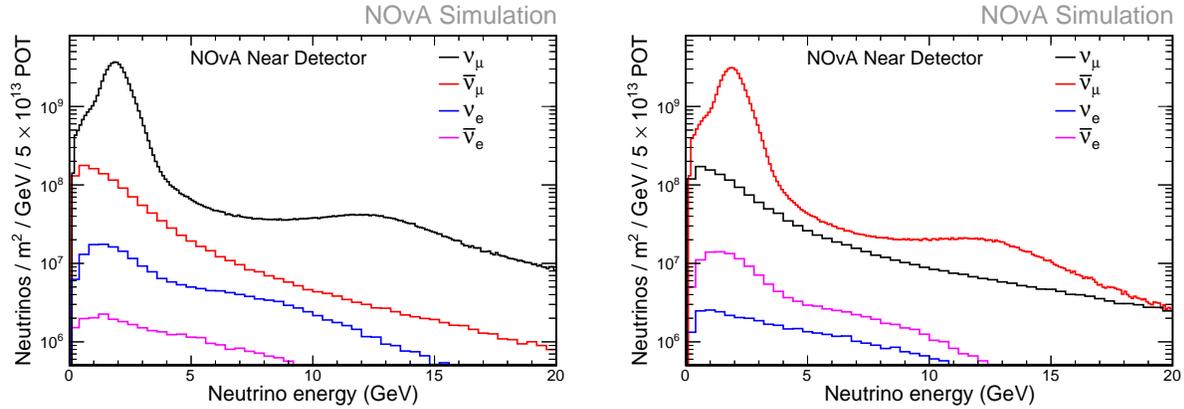


Figure 3.9: Flux at the NOvA ND after the application of PPFX corrections in both the  $\nu_\mu$  (left) and  $\bar{\nu}_\mu$  (right) beam modes [46].

shown for both the  $\nu$  (predominantly  $\nu_\mu$ ) and  $\bar{\nu}$  (predominantly  $\bar{\nu}_\mu$ ) beams in the ND after the PPFX corrections in Figure 3.9. These are the fluxes only, and do not take into account cross-sections. Note that the  $\nu_\mu$  flux in the peak in the  $\nu$  beam is similar to that of the  $\bar{\nu}_\mu$  flux in the  $\bar{\nu}$  beam. In both cases that the  $\nu_e$ ,  $\bar{\nu}_e$  fluxes are sub-dominant to both the  $\nu_\mu$  and  $\bar{\nu}_\mu$ , but appreciable enough to need to be accounted for in the oscillation analysis. Chapter 5 will deal with the wrong-sign ( $\nu$  in the  $\bar{\nu}$  beam) fraction and performing cross-checks on this important sample.

The prediction of neutrino interactions in the detectors is simulated with the event generator GENIE<sup>1</sup>. This uses cross-section knowledge to produce events of the various interaction types which are simulated in the NOvA detectors, bridging the gap between flux predictions and simulated events. NOvA uses a tuning procedure that corrects aspects of these predictions based on guidance from external constraints and ND data. A major point of emphasis in the tuning is multi-nucleon effects. In the quasi-elastic sample, nuclear charge screening is important: a random phase approximation (RPA) correction is used [47]. The same correction is also taken for resonance interactions (e.g. interactions producing the  $\Delta$  resonance), but a large systematic is taken since less is understood about this correction. Regions of phase space are corrected based on observations in NOvA ND data. Further, GENIE’s empirical meson exchange current (MEC) events (see Reference [48]) are added to the sample but in a tuned fashion to help match to our observed data. These are events

<sup>1</sup>Version 2.12.2 is the default version used for the analysis discussed here.

with two nucleons in the interactions instead of one, also known as 2p2h (two particles, two holes).

### 3.4 Event classification

Neutrinos at the FD will yield charged current interactions from unoscillated  $\nu_\mu$  and  $\bar{\nu}_\mu$ , oscillated and background  $\nu_e$  and  $\bar{\nu}_e$ , small amounts of  $\nu_\tau$  and  $\bar{\nu}_\tau$ , neutral current interactions, and backgrounds from cosmogenic events (e.g. cosmic-ray muons). To accurately study oscillations, analyses must be able to reliably classify interactions as being from one type of neutrino or another. Thus, event classification is an important topic in NOvA.

In NOvA, sets of hits with similar positions and times are grouped together into “slices” which ideally contain all the hits from an individual neutrino interaction. Various reconstruction algorithms are used to identify features of such events, some of which will be discussed later in the context of event selection. Recently, however, NOvA has developed ever-improving classification schemes by borrowing from computer vision technology and utilizing deep learning techniques. Specifically, the development of a Convolutional Neural Network called the Convolutional Visual Network (CVN) has yielded significant improvement to NOvA’s ability to classify and select electron neutrino events: the gains from the first implementation were equivalent to increasing exposure by roughly 30%. The network and implementation is described in [49, 50] so the discussion here will be at a higher level.

The network is inspired by GoogLeNet, and uses sets of convolutions in groups called inception modules which allow for the network to “learn” from extracted features with varying complexities over a number of layers. A representation of such an inception module is shown in Figure 3.10 [49]. In the implementation primarily herein, the outputs are based on final states arising from the neutrino interactions: the neutrino flavor ( $\nu_e$ ,  $\nu_\mu$ ,  $\nu_\tau$ , or NC) and the numbers of protons, charged pions, neutral pions, and neutrons in the event. There are also labels for specifying activity as cosmic. In total, 392 different labels characterize the sets of particles in the final state. The various outputs are given importance weights such that they sum to 1, so the score given to any individual output label can be thought of as the likelihood the classifier assigns to that label for the given event. These can be collapsed in various ways: Chapter 5 will use one such way for understanding

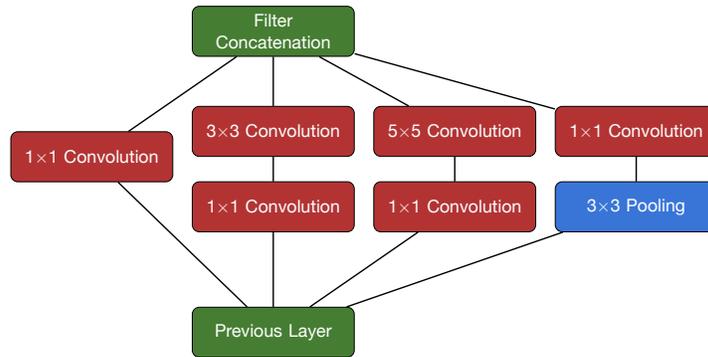


Figure 3.10: Cartoon representation of an inception module in the Convolutional Neural Network trained to classify neutrino candidates in NOvA, which performs convolutions on the input. [49]

wrong-sign events. For event classification in the selection of  $\nu_\mu$  or  $\nu_e$  events, all labels with that flavor are summed to provide the score. An example of this is shown in Figure 3.11; note that the  $\nu_e$  events rise above the background at the higher scores.

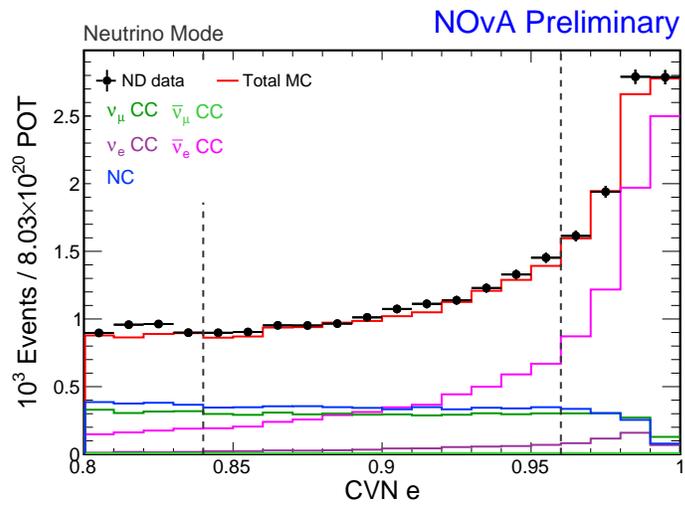


Figure 3.11: ND data and a breakdown of the simulation modes for recent analysis [51]: note that the true  $\nu_e$  increases significantly above the other neutrino interaction modes in this sample at the highest CVN electron scores. The data also increases in this region.

## CHAPTER 4

### $\nu_e$ Appearance Oscillation Analysis with NOvA

As noted in previous chapters, the disappearance analysis channel in a long-baseline experiment is particularly sensitive to the  $\Delta m_{32(31)}^2$  and  $\sin^2 \theta_{23}$  parameters. The MINOS detector was most sensitive to  $\nu_{\mu}, \bar{\nu}_{\mu}$  disappearance in the muon (anti-)neutrino beam, as was briefly discussed in Chapter 2 and displayed in Figure 2.8. The NuMI beam with an alternate configuration is also used in the NOvA experiment and again provides a  $\nu_{\mu}, \bar{\nu}_{\mu}$  beam. NOvA's result on the mass splitting is also driven by the disappearance channel, which also has strong sensitivity to  $\sin^2 \theta_{23}$ . A particular advantage to NOvA over the prior generation is the sensitivity to  $\nu_e, \bar{\nu}_e$  events, such that the appearance channel can also be probed. Recall from Equation 1.7 and the bi-probability plots that this channel has strong sensitivity to  $\delta_{CP}$ , mass hierarchy, and additional sensitivity to  $\sin^2 \theta_{23}$ .

NOvA has recently performed an analysis of three-flavor oscillations, testing  $\nu_{\mu}, \bar{\nu}_{\mu}$  disappearance and  $\nu_e, \bar{\nu}_e$  appearance at the Far Detector (FD). The first oscillation analysis of both neutrino beam and antineutrino beam data took place in 2018, but was updated to keep the same procedure and include more antineutrino data in 2019. As much of my work has been in the context of the  $\nu_e(\bar{\nu}_e)$  appearance channels, this chapter will focus on this part of the analysis. Both analyses had just under  $9 \times 10^{20}$  POT full-detector-equivalent in the neutrino beam. Full-detector equivalent POT is stated since some of the data was taken with a partial detector, especially during construction. In the antineutrino beam, the 2018 analysis exposure was nearly  $7 \times 10^{20}$  POT, and the 2019 analysis exposure was  $12.33 \times 10^{20}$ . This chapter will discuss the appearance analysis procedure and joint disappearance-appearance analysis, as well as relevant contributions of analysis work I was more involved with. The following chapter will discuss the understanding of wrong-sign ( $\nu$ ) in the antineutrino beam. Results will be discussed in Chapter 6.

Note that NOvA is a collaborative effort, and figures with citations represent collective work by the larger collaboration. These are shown in the pursuit of a more complete narrative and context

for my efforts. Figures without citations are generally produced through by me and/or through my research efforts, though at times also in conjunction with collaborators.

#### 4.1 $\nu_e, \bar{\nu}_e$ appearance analysis

Events in the NOvA detectors are formed from “slices” in time and space in a window around the beam spill time with clusters of hits in the detector. These slices have undergone reconstruction to have information such as vertices, particle clustering, and tracking. These slices will contain neutrino interactions, particles originating from cosmic activity, potential noise hits from APDs, etc. A series of selection cuts are defined for each analysis to pick the desired interaction type from these events. Since the NC interaction is insensitive to flavor, the goal is to pick CC interactions - where the signature is the lepton ( $\mu$  for  $\nu_\mu$  and  $\bar{\nu}_\mu$  and e for  $\nu_e$  and  $\bar{\nu}_e$ ).

The collaboration has developed a procedure for selecting candidate events. First, a group of operational cuts are applied to the dataset to use only events taken during “good” beam and detector conditions, for example checking the current in the focusing horns, the beam position, etc. From there, cuts are used to separate out well-reconstructed signal events from the backgrounds. A set of preselection cuts picks out events passing various quality and veto checks and further keeps events with characteristics in our region of interest, having cluster lengths and number of APD hits consistent with our signal events and energy between 1 and 4 GeV. A set of cuts on reconstructed information are useful to reject backgrounds. For one, a cut attempts to reject particles uncontained in the detector. The energy for signal events which exit cannot be accurately reconstructed, and rejecting events with particles found to be entering the detector reduces cosmic background. Another cut checks that particles near the back of the detector are inconsistent with a photon from cosmic activity entering the back of the detector. Further, since the beam enters the front of the detector while cosmic rays will be more vertical, a check on the transverse momentum fraction ( $p_t/p$ ) further reduces cosmic background. In addition, the nearest slice is examined. Finally, a particle/event identification (PID) cut is made to pick out remaining events which are consistent with  $\nu_e$  ( $\bar{\nu}_e$ ) CC signal. For this cut, the CVN electron ID (CVNe) score (as discussed in Chapter 3) retains many signal events while picking up small amounts of background. Two selection groups

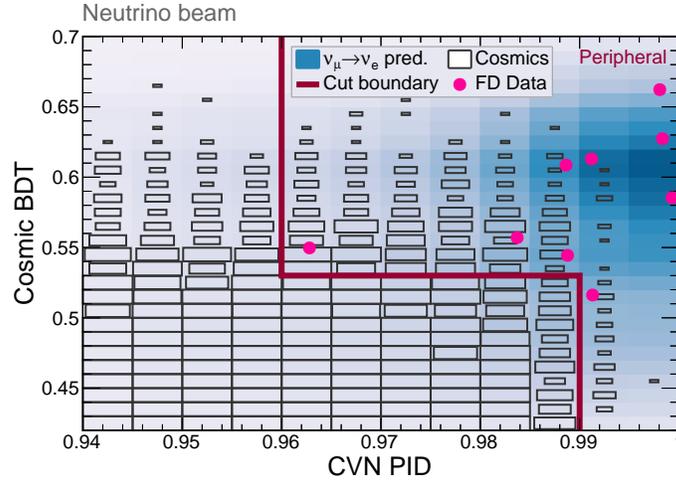


Figure 4.1: Plot of the ID scores used to select the peripheral sample for the  $\nu_e$  analysis [53]. The simulated signal prediction is given along with the cosmic background. The cut borders are given in red, and the events selected in the FD are given as well. A similar process is performed to pick up  $\bar{\nu}_e$  events in the antineutrino beam.

are made: one which requires a very high score and is pure in  $\nu_e$  ( $\bar{\nu}_e$ ) CC but not as efficient, and one for somewhat lower scores that recovers many  $\nu_e$  ( $\bar{\nu}_e$ ) events but has larger background. For this analysis, the lower PID region is defined by CVNe score between 0.84 and 0.96 for neutrino beam data (0.89 and 0.98 for antineutrino beam data), and the high PID region is made up of events scoring higher than this.

A third sample called the “peripheral” has higher selected background counts but recovers several  $\nu_e$  ( $\bar{\nu}_e$ ) signal events that would be lost to the cuts described above [52]. Events that fail the preselection but still pass the quality and veto portions and are between 0 and 4.5GeV are checked against these alternative cuts. The nearby slice is checked for cosmic rejection, and a cut is made in 2D space on CVNe vs a cosmic ID score to reject background and keep signal. This sample is condensed into one bin, rather than kept as an energy spectrum. A plot of this cut with simulated signal, cosmic background, and data events are shown for the neutrino beam in Figure 4.1.

A second CVN instance [54] is used in the energy estimate for  $\nu_e$  ( $\bar{\nu}_e$ ) events and is also used in some of the wrong-sign measurements to be presented in Chapter 5. This CVN uses additional information from the reconstructed clusters to perform particle ID on individual clusters as opposed

to event identification. Visible energy in hits from clusters identified as electromagnetic and from clusters identified as hadronic are separated and fed into a quadratic formula which reconstructs an estimate of the incident neutrino energy. This energy estimation gives resolutions of approximately 11% for neutrino beam data and approximately 9% for antineutrino beam data [55]. The reconstructed  $\nu_e$  and  $\bar{\nu}_e$  energy of candidate events then form the energy spectra used in predictions and measurements of the oscillation signal at the FD.

Assuming three flavor oscillations, the Near Detector (ND) is too close for oscillation to be visible. Events selected at the ND with the  $\nu_e$  selection cuts are therefore backgrounds, primarily the intrinsic  $\nu_e$  and  $\bar{\nu}_e$  component of the beam and misidentified  $\nu_\mu$ ,  $\bar{\nu}_\mu$ , or NC events.  $\nu_e$  and  $\bar{\nu}_e$  candidates are selected and the data and simulation compared. For data in the neutrino beam, the collaboration uses a data-driven decomposition method to probe the relative contributions of these three backgrounds. Contained and uncontained muon neutrino candidates in the ND are used to constrain the flux based on its parent ( $\pi$ ,  $K$ , etc.) [56]. Electron (anti-)neutrinos in the beam from both pion and kaon parents are important. The contained and uncontained  $\nu_\mu$  spectrum is plotted in the top of Figure 4.2, and the FD beam  $\nu_e$  based on parent particle is shown in the bottom left. In addition, differences between measured and simulated detection of Michel electrons help inform the muon (anti-)neutrino and NC backgrounds [57]. The Michel spectrum is shown in the bottom right of Figure 4.2. This decomposition is used to alter the fractions of expected backgrounds. For the ND antineutrino beam with more recent data, a simpler background decomposition is sufficient, wherein simulation-data discrepancies in the energy spectra are fixed by scaling the three components proportionally. The ND spectra agree with data after this process by definition. Figures for the neutrino beam and antineutrino beam spectra in the decomposition are given in Figure 4.3. The far to near ratio of this then informs the predicted spectrum at the FD.

While the  $\nu_e$  ( $\bar{\nu}_e$ ) selection at the ND informs the backgrounds, the signal expectation is driven by the muon (anti-)neutrino spectrum at the ND. The FD signal are the  $\nu_\mu$  and  $\bar{\nu}_\mu$  which have oscillated to  $\nu_e$  and  $\bar{\nu}_e$ . An extrapolation procedure is used to produce a prediction of expected events at the FD based on the data measurements at the ND. This procedure is graphically illustrated

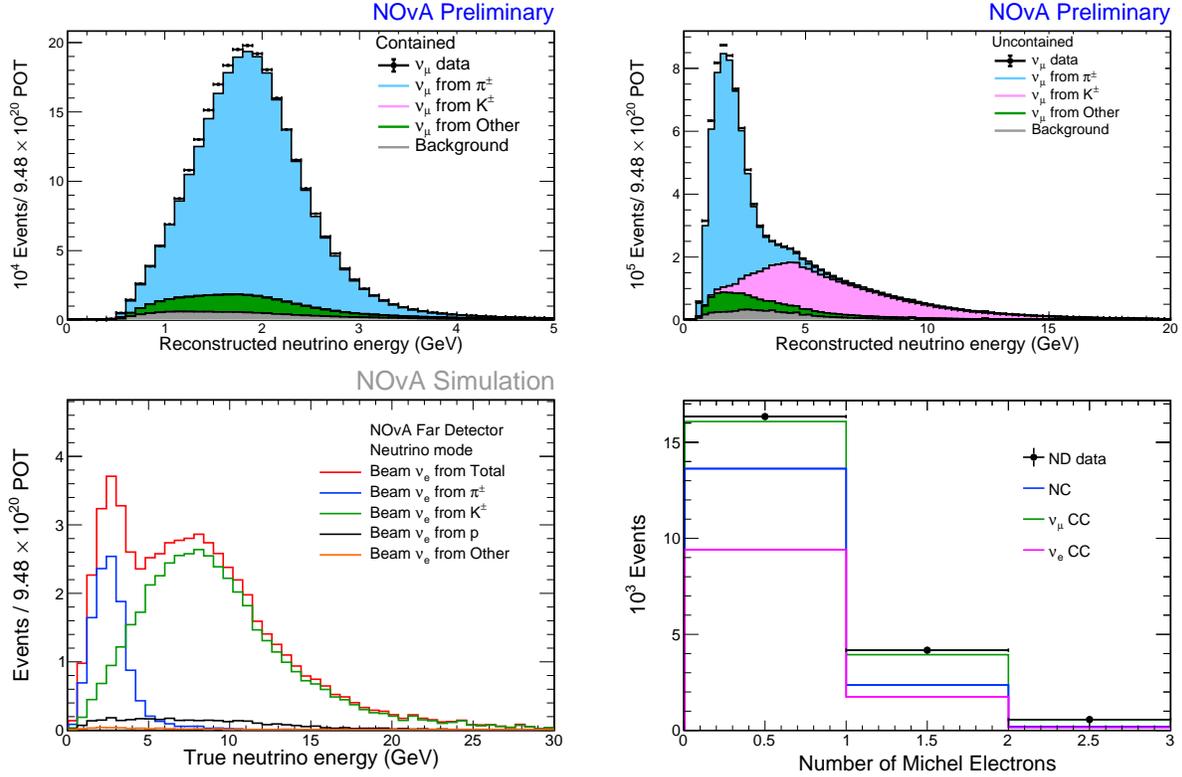


Figure 4.2: Plots describing the decomposition methods applied at the ND [58, 59]. The top two plots show the contained and uncontained  $\nu_\mu$  candidate events at the ND, with the colors representing the various parent types [58]. The contributors to the intrinsic beam  $\nu_e$  ( $\bar{\nu}_e$ ) background at the FD is shown on the bottom left, where it is noted that both pion- and kaon-produced electron (anti-)neutrinos contribute to the signal between 1 and 4 GeV [58]. This decomposition method has been discussed elsewhere, for example Reference [56]. Further, the spectrum of identified Michel electrons in data and simulation is given in the bottom right [59]. The decomposition using Michel electrons has also been discussed elsewhere, for example Reference [57].

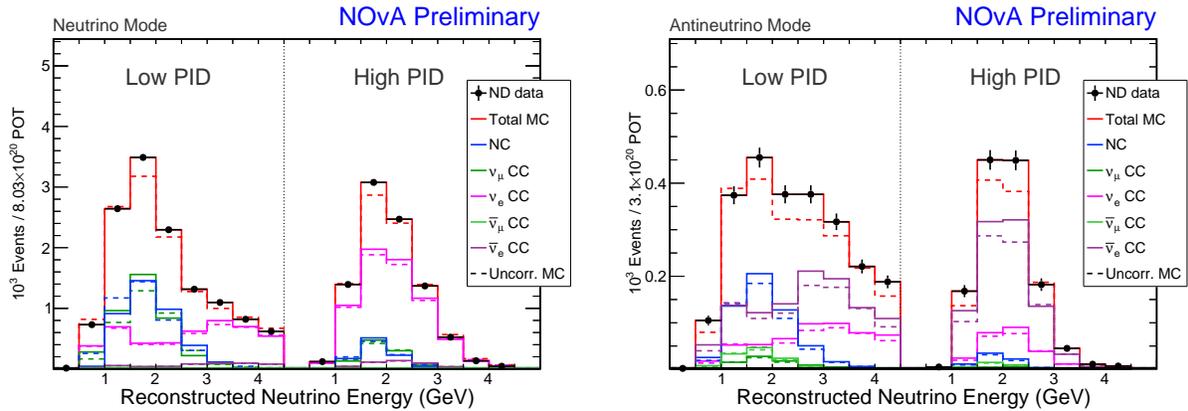


Figure 4.3: ND energy spectra for the neutrino beam (left) and antineutrino beam (right) in the decomposition [51]. Both the lower and higher PID selections are shown. The ND spectra agree with data after this process by definition.

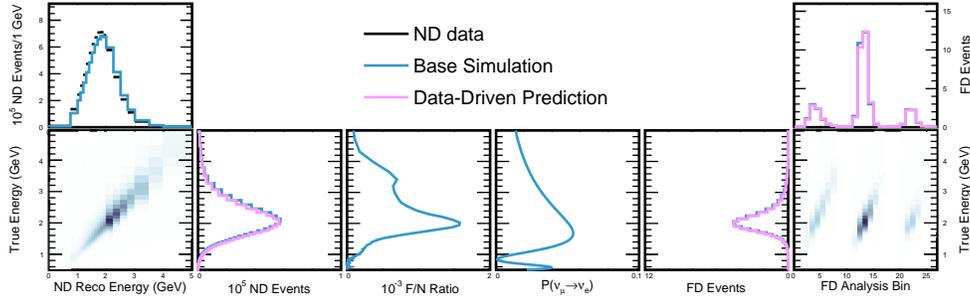


Figure 4.4: Cartoon explaining in a general sense the extrapolation procedure [60]. A small fraction of  $\nu_\mu$  and  $\bar{\nu}_\mu$  events at the ND are expected to undergo oscillation to appear as  $\nu_e$  and  $\bar{\nu}_e$  in the FD. The muon (anti-)neutrino candidate spectrum at the ND can then be modified by the ratio of acceptances between the FD and ND and a set of oscillation parameters to produce a predicted spectrum of electron (anti-)neutrinos at the FD. Varying the oscillation parameters applied varies the prediction at the FD.

in the schematic Figure 4.4. The muon (anti-)neutrino candidate spectrum at the ND can then be modified by the ratio of acceptances between the FD and ND and a set of oscillation parameters to produce a predicted spectrum of electron (anti-)neutrinos at the FD. While the energy spectra for  $\nu_\mu$  ( $\bar{\nu}_\mu$ ) at the ND and the  $\nu_e$  ( $\bar{\nu}_e$ ) at the FD are reconstructed quantities, the oscillation takes place according to true energies. The predictions save oscillatable spectra: that is spectra that store true energy information in bins as well as the reconstructed quantity. These can therefore be weighted for effects such as oscillation probabilities.

This process shows how one can form a predicted energy spectrum at the FD given data-driven techniques at the ND and simulated information. The cosmic background contribution to the prediction is determined by the rate of cosmic events selected in a timing sideband outside of the NuMI spill time during the triggers. An example prediction for the neutrino beam analysis, with data overlaid, is shown in Figure 4.5.

On top of the signal definition and background studies, an important consideration in the analysis is systematics. In the analysis framework used for NOvA (called CAFAna and discussed in a little more detail in Chapter 9), systematics can be handled in several ways. For example:

1. Systematics can be event reweighting based on event information. This can be simple such as scaling every event by 0.95, or they can be more complicated and involve truth variables

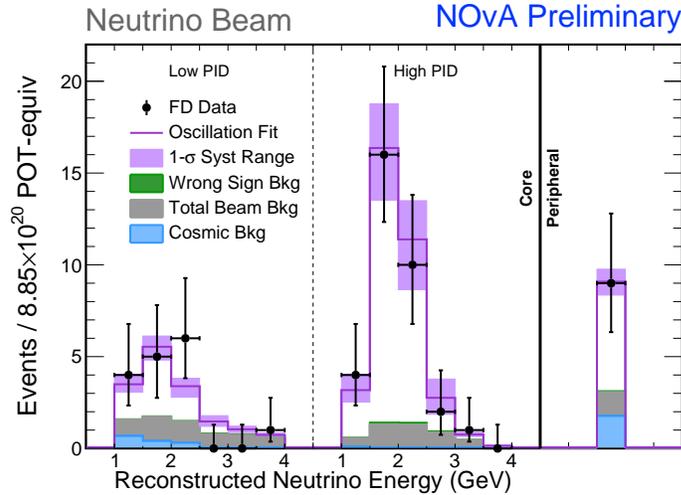


Figure 4.5: Prediction for the FD neutrino beam candidate selection [61], showing the  $\nu_e$  appearance signal, small wrong-sign  $\bar{\nu}_e$  background, cosmic background, and beam-related background - for example  $\nu_e$  and  $\bar{\nu}_e$  intrinsic to the beam. The lower PID, high PID, and peripheral sample are seen. These samples are described in the text. Overlaid is the data, which will be discussed in more detail in Chapter 6.

from simulation and/or reconstructed variable information.

2. Systematics can be alterations to the values of variables themselves. An example of this is the neutron uncertainty, which will be discussed in more detail in the text to follow.
3. Where the variable of interest is upstream of either simulation or reconstruction, systematics can also be new files produced after altering some aspect of the process. An example of this is the calibration systematic, which produces new files with updated processing after altering the energy calibration.

Extrapolated predictions are made for each considered source of systematic uncertainty in a manner as described above with variations (for example  $\pm 1\sigma, 2\sigma$  for some), such that an interpolation allows for running of systematic contributions [52]. Each considered systematic is programmed into the fit as a penalty term on the likelihood, as discussed below. The systematics considered for the analysis cover a large range of effects, including detector calibration, detector response, cross sections, neutrino flux, and differences between the ND and FD. The extrapolation procedure reduces the impact of flux and cross-section systematics (which would be  $\sim 10\text{-}20\%$  [36]), since

the flux  $\times$  cross-section is observed in a functionally similar ND before oscillation. Some specific systematics will be discussed below.

For recent NOvA results, the collaboration performs a joint analysis that combines  $\nu_e, \bar{\nu}_e$  appearance with  $\nu_\mu, \bar{\nu}_\mu$  disappearance. This is handled in the CAFAna analysis framework combining multiple spectra and constraints together in a single unit to fit together. The fit is done in log-likelihood, using the expression from [4] for binned Poisson data

$$-2 \log \mathcal{L}(\vec{\theta}) = 2 \sum_i^{N_{bins}} \left[ e_i(\vec{\theta}) - o_i + o_i \log \frac{o_i}{e_i(\vec{\theta})} \right]$$

where the set of oscillation parameters tested for significance are  $\vec{\theta}$ , the prediction at a given set of oscillation parameters in a given bin is  $e$ , and  $o$  is the observed data in a given bin (where I have adopted the notation style of Reference [62]). We call

$$\chi^2(\vec{\theta}) \equiv -2 \log \mathcal{L}(\vec{\theta})$$

Some parameters in the fit are measured elsewhere with higher precision than is capable in NOvA, so a constraint is adopted. For the joint analysis,  $\theta_{13}$  is such a constrained parameter, where reactor experiments have provided precise measurements of this. Such constrained parameter varies in the fit, but adds a penalty term to the  $\chi^2$  based on its distance from the constrained parameter as  $(x - \mu)^2/\sigma^2$ , where the constraint is  $\mu \pm \sigma$  [52]. For  $\nu_e, \bar{\nu}_e$ -only fits, a constraint is taken on  $\Delta m_{32}^2$  as well, as explored in Chapter 9. Finally, systematics penalty terms contribute as  $\delta^2/\sigma^2$ , where  $\delta$  is the fit value for the given systematic and  $\sigma$  is the systematic's error [63]. One such term exists for each systematic, and we call the set of systematic parameters  $\vec{\delta}$ . The systematic also effects

the prediction for each bin  $e_i$ . Therefore, the full fit function is

$$\chi^2(\vec{\theta}, \vec{\delta}) = 2 \sum_i^{N_{bins}} \left[ e_i(\vec{\theta}, \vec{\delta}) - o_i + o_i \log \frac{o_i}{e_i(\vec{\theta}, \vec{\delta})} \right] + \sum_j^{M_{syst}} \frac{\delta_j^2}{\sigma_j^2} + [\text{pulls from constraints}] \quad (4.1)$$

However, it is important to note that the oscillation formalism has physical boundaries (sin for instance spans 0 to 1), so the final analysis undergoes a Feldman-Cousins correction procedure [64]. NOvA runs many pseudo-experiments on a supercomputer, where a given set of parameters is tested and the significance corrected [62].

## 4.2 Some contributions to the analysis

One contribution to the analysis with added potential moving forward is studying the wrong-sign composition of the antineutrino beam:  $\nu$  in the  $\bar{\nu}$  beam. A study on the beam  $\nu_e, \bar{\nu}_e$  background is presented in Chapter 5.

Additionally, I performed POT and live-time accounting for the recent NOvA analysis. This is fairly straightforward, but will be used in Chapter 6 to perform a Kolmogorov-Smirnov (KS) type test comparing accumulated POT and selected neutrino candidates. Exposure calculated by POT and live-time do not grow at the same rate for two reasons. Beam intensities vary over time and started significantly lower than they are now. Therefore near the beginning of NOvA operations, it took longer to accumulate the same amount of POT than needed now. Furthermore, NOvA was still being constructed after the start of data taking. This means that the rate of neutrino interactions in the detector will be smaller, requiring more POT and live-time to accumulate the same number of interactions. While the POT and detector mass are responsible for the number of interacting neutrinos in the FD, the live-time and detector mass relate to the number of cosmic ray interactions appearing as backgrounds. This calculation was previously performed by colleagues and updated to include recent data. Using the efficiencies to select events with varying numbers of diblocks in the FD, a conversion from the number of diblocks (total detector mass) to an effective fiducial mass was found [65]. This scale factor is used to produce a full-detector-equivalent exposure,

which corrects the number of collected POT to an effective POT for a full detector. While this was calculated for a prior analysis, it was deemed to be sufficiently similar to reuse in the current analysis. This correction mostly applies to the neutrino beam data; for the new antineutrino data added for the recent analysis, the full detector had already been constructed. Plots of live-time and POT, both raw and equivalent, are produced to count the exposure from each data-taking period and to go into the KS-type test that will be discussed later. Table 4.1 contains the exposures in POT and live-time for the recent NOvA analysis.

Data group	$\nu, \bar{\nu}$	POT ( $\times 10^{20}$ )	Live-time (s)
Period 1	$\nu$	1.23	128.12
Period 2	$\nu$	2.64	119.18
Period 3	$\nu$	3.85	138.94
Period 4	$\bar{\nu}$	0.66	19.58
Period 5	$\nu$	1.76	51.93
Period 6	$\bar{\nu}$	3.06	76.28
Epoch 7a	$\bar{\nu}$	7.78	24.76
Epoch 7b	$\bar{\nu}$	1.86	45.75
Epoch 7c	$\bar{\nu}$	0.55	13.26
Epoch 7d	$\bar{\nu}$	2.19	53.20
Epoch 8b	$\bar{\nu}$	3.23	84.23

Table 4.1: Exposure for the recent NOvA analysis, rounded to two decimal places. The periods represent the various data-taking periods in NOvA as in Table 3.1, with each successive period being more recent.

An interesting sample to consider in the oscillation analysis is  $\nu_\tau$ . Recall from the bi-probability plot shown for NOvA in Figure 1.3 that the oscillation probabilities  $\nu_\mu \rightarrow \nu_e$  and  $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$  are around 2-7%. However, the  $\nu_\mu$  survival probability at the FD is quite small around 1.6-1.8GeV and still low at energies close to this. Disappearing  $\nu_\mu, \bar{\nu}_\mu$  that do not oscillate to  $\nu_e, \bar{\nu}_e$  instead oscillate to  $\nu_\tau, \bar{\nu}_\tau$ , assuming 3 flavor oscillations: in fact the majority of oscillations occur as  $\nu_\mu \rightarrow \nu_\tau$  and  $\bar{\nu}_\mu \rightarrow \bar{\nu}_\tau$ . However due to mass considerations, NOvA does not study  $\nu_\tau$  as an oscillation signal in the FD. In a  $\nu_\tau$  CC interaction, a  $\tau$  lepton with 1.78GeV [4] is produced, so energy and momentum conservation requires  $\sim 3.5$ GeV energy. This is well above the maximal disappearance point for  $\nu_\mu, \bar{\nu}_\mu$  and in the tail of the NOvA flux. While it does not present a usable signal, there will still be some small number of oscillated  $\nu_\tau$  and  $\bar{\nu}_\tau$  which might be selected in the  $\nu_e$  samples as background

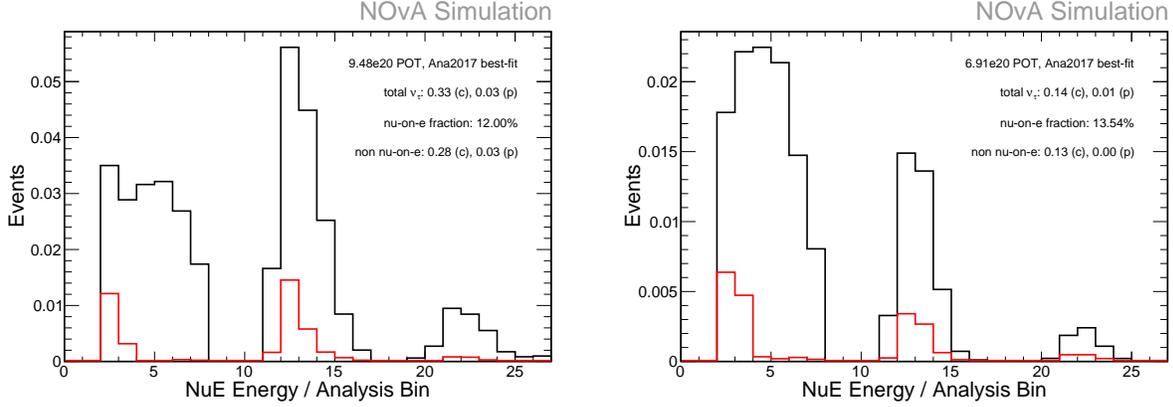


Figure 4.6: Unextrapolated predictions showing the  $\nu_\tau$  and  $\bar{\nu}_\tau$  components of the FD selection in both the neutrino beam (left) and antineutrino beam (right). The antineutrino prediction is found using the 2018 exposure. Red indicates the portion which is  $\nu$ -on- $e$  scattering. The fraction of events which is CC is used in a back-of-the-envelope calculation to set a conservative, approximate  $\nu_\tau$ ,  $\bar{\nu}_\tau$  uncertainty. The three groups in each plot represent the lower PID selection (left), high PID selection (middle), and peripheral sample (right).

and have poor energy resolution. Furthermore, that the  $\nu_\tau$  cross-section is not precisely measured provides a complicating factor. A small study was performed to look at the  $\nu_\tau$  sample in the neutrino and antineutrino beams at the FD. Unextrapolated predictions were produced using the best fit oscillation parameters of the prior NOvA analysis -  $\delta_{CP} = 1.21\pi$ ,  $\sin^2 \theta_{23} = 0.558$ ,  $\Delta m_{32}^2 = 2.44 \times 10^{-3}$ . At the 2018 analysis exposure ( $9.48 \times 10^{20}$  POT neutrino,  $6.91 \times 10^{20}$  POT antineutrino), the expected number of  $\nu_\tau$ ,  $\bar{\nu}_\tau$  are  $<1$  event:  $\sim 0.3$  events in neutrino mode and  $\sim 0.1$  events in antineutrino mode. This is approximately 0.6% of the total  $\nu_e$  candidates in the prediction and 1.3% of the  $\bar{\nu}_e$ , so we will want some sort of systematic to cover the uncertainty in the  $\nu_\tau$  cross-section, even if it is fairly rough. Figure 4.6 shows the neutrino beam and antineutrino beam selected  $\nu_\tau$ ,  $\bar{\nu}_\tau$  events that are CC or  $\nu$ -on- $e$  scattering: of the selected  $\nu_\tau$  events, roughly 86.5-88.0% are CC (not  $\nu$ -on- $e$  scattering). This information was fed back to a  $\nu_\tau$  systematic. A measurement of the  $\nu_\tau$  cross-section by OPERA [66] found a cross section of  $\langle \sigma \rangle = 1.2_{-0.5}^{+0.6} \langle \sigma \rangle_{GENIE}$ , where  $\langle \sigma \rangle_{GENIE}$  is the event generator cross-section. Taking the amount of difference with respect to GENIE and the error of the measurement in quadrature and multiplying by 0.88 for the approximate fraction of our  $\nu_\tau$  prediction that is CC, this was rounded up to 0.6 and taken as a conservative estimate for the  $\nu_\tau$  cross-section systematic [52].

Another cross-section systematic taken in the analysis is that for coherent pion production, where the neutrino interacts with a target nucleus and produces a  $\pi$  in the final state. There are CC and NC versions of this process, with interactions such as  $\nu_l + A \rightarrow l + A + \pi^+$  and  $\nu_l + A \rightarrow \nu_l + A + \pi^0$  respectively. NOvA has a set of systematics that tie to the true interaction modes, including for coherent interactions. Tests applying the coherent interaction systematics showed drastic reduction in the  $\pm 1\sigma$  cases, which is not expected. A similar quick study to the  $\nu_\tau$  study investigated the fraction of events in our predictions which are coherent. Somewhat different oscillation parameters were used in the test, but that to get a sense for the importance of this interaction mode, this is unimportant. Coherent interactions were found to be roughly 2% of the CC  $\bar{\nu}_e$  appearance signal and provided nearly 16% of the NC background to this signal. Selected NC interactions are important backgrounds, so 16% of that is at the level where we will want to have a motivated uncertainty. Rather than take the values from the coherent systematics as-is, we take a data-driven systematic based on measurements of this interaction mode, using a normalization-style systematic. A study of NC coherent pion production in NOvA with the neutrino beam found a central-value close to the GENIE prediction with errors of 16.7% [67]. One could take this result for the NC coherent uncertainty, but since an uncertainty for the antineutrino beam too is also needed, we take a more conservative 20%. For the CC coherent pion production uncertainty, a collaborator who studied CC interactions in MINERvA communicated to us that between 2-5GeV the uncertainty was calculated to be 14% for  $\nu_\mu$  and 20% for  $\bar{\nu}_\mu$  [68]. We take a conservative approach of setting the uncertainty at a flat 20%. These were implemented originally by a colleague and updated to the values here, as another example of a normalization/weight-type systematic for the analysis.

As an example of a systematic dealing with reconstructed quantities, consider neutron response in the NOvA detectors. Recall that for quasi-elastic antineutrino interactions, the hadron in the final state is a neutron. However, neutrons are neutral and therefore do not produce scintillation light in the detector. This leaves us looking at just its “daughter particles” - typically photon or proton activity. In looking at antineutrino beam data in the ND, discrepancies in data and MC were found in regions where neutrons would be expected to contribute. A colleague produced a

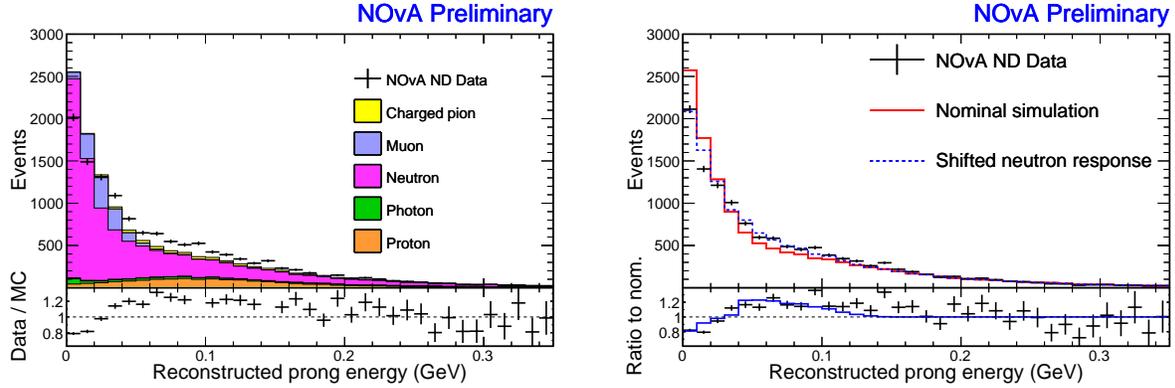


Figure 4.7: Enhanced selection of reconstructed particle prongs from neutrons in the NOvA ND [69] (see also [14]). (left) The contributions which come from neutron daughters are shown in purple, showing this is the dominant contributor to this plot. The particles making this deposition tend to be protons and photons, with most of the photons confined to low energies. (right) The application of the neutron systematic, described in the text, largely covers the discrepancy.

selection of  $\bar{\nu}_\mu$  CC candidates expected to be enhanced in neutron daughter particles that showed such discrepancy [14]. The energy of reconstructed particle prongs (reconstructed clusters) is seen in Figure 4.7. Whether this is due to production or detector modeling, a systematic to cover the incorrect energy is desirable.

When extra variables were added to the analysis files to allow for more exploration, my initial involvement in this study was to produce and look at some of the variables in a small sample of such files before a full respin of the files could be made. Ultimately, an essentially complete set of analysis files were remade with these variables in a more standard way by the Production group. A colleague found a systematic that worked in a manner deemed sufficient for this analysis, the application of which is shown in Figure 4.7. Essentially, for neutrons above a threshold, energy is probabilistically added back to the system. This is further added back directly as hadronic energy for  $\nu_\mu$  and as overall calorimetric energy for  $\nu_e$ .

This was found to shift the mean reconstructed energy for  $\nu_\mu, \bar{\nu}_\mu$  by 0.5% in the neutrino beam and 1% in the antineutrino beam. I performed a short study of the  $\bar{\nu}_e$  candidate selection with this systematic applied. Since the hadronic energy fraction is in principle a well-defined quantity which ranges from 0 to 1, the study produced spectra of “true”  $\bar{\nu}_e$  interactions<sup>1</sup> and selected  $\bar{\nu}_e$  interactions

<sup>1</sup>I use “true” interactions with scare quotes, as this sample (using the standard set of recorded information) does not

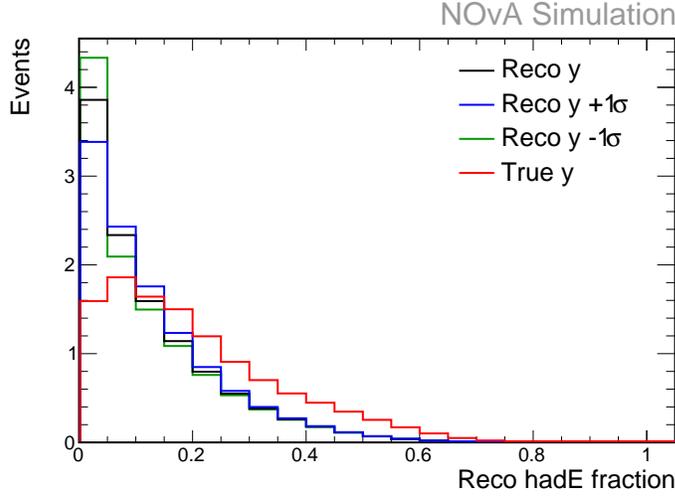


Figure 4.8: Spectra of reconstructed hadronic energy fraction for selected simulated  $\bar{\nu}_e$  CC events in the FD, with simple oscillation weights to approximate events that might appear in the oscillation signal. The effects of the neutron systematic to shift events to higher or lower reconstructed hadronic energy is visible. The true hadronic energy fraction (as opposed to reconstructed) is shown in the figure as well. While the reconstruction does not do a superb job at recreating the true fraction, the integrated event counts are what are used in this study.

as functions of hadronic energy fraction. The “true” sample is selected as slices matched to a true  $\bar{\nu}_e$  CC interaction, and the “selected” sample passes all of the previous cuts and also the FD selection cuts. Since these are at the FD and use flavor-swapped files to make the flux mostly  $\bar{\nu}_e$ ,  $\nu_e$  as opposed to  $\bar{\nu}_\mu$ ,  $\nu_\mu$ , a set of simple oscillation weights are applied to events. The selected spectra are shown in Figure 4.8, and the effects of the neutron systematic to shift events to higher or lower reconstructed hadronic energy is visible. The true hadronic energy fraction (as opposed to reconstructed) is shown in the figure as well. While the reconstruction does not do a superb job at recreating the true fraction, the integrated event counts are what are used in this study. The integral number of events in both samples and the pseudo-efficiency of selected events are shown in Table 4.2, but due to the caveat discussed in the footnote, that the number of selected events in all cases is the same is largely satisfying. The main takeaway is that the application of this systematic does not seem to drastically affect the analysis.

---

contain neutrino interactions which do not produce a slice with reconstructable information in the detector. Ideally a selection efficiency would consider any true interaction or even just any true neutrino in the flux. However, to apply the systematic which relies on manipulating reconstructed variables, I use the standard information rather than the more expansive set of true neutrino simulated information.

Sample	Selected Events	“True” Events	Pseudo-efficiency
Nominal	11.3737	17.2792	65.823%
+1 $\sigma$	11.3751	17.2791	65.8316%
-1 $\sigma$	11.371	17.2793	65.807%

Table 4.2: Event counts in the studied samples for the neutron systematic in the FD, with simple oscillation weights to approximate events that might appear in the oscillation signal. No significant problems were found for the application to the  $\bar{\nu}_e$  appearance selection.

A final study will be briefly noted in this chapter. The neutrino beam data analyzed in this analysis are the same as that in the previous analysis. However, analysis techniques have been changed, so we expect to either gain or lose some events from that which was selected previously. Therefore, an event migration study was conducted to check the core (lower PID and high PID) and peripheral samples with both the previous analysis cuts and the updated analysis cuts to understand the expectation. The study used unextrapolated predictions at the FD and used separate cosmic trigger data normalized for live-time to make a prediction about the selection of cosmic background as well. The oscillation calculation used the prior analysis results:  $\delta_{CP} = 1.21\pi$ ,  $\sin^2 \theta_{23} = 0.558$ ,  $\Delta m_{32}^2 = 2.44 \times 10^{-3}$ . Table 4.3 summarizes the results. In the core selection, a reduction of approximately 3.5 events was predicted relative to the previous results. Of these, approximately 0.9 events come from signal, with 2.6 events from background. Very small gains in signal and reductions in background were predicted for the peripheral sample.

	Both	New Only	Prev. Only	New Total	Prev. Total	Change
Core	48.93	3.45	6.98	52.38	55.91	-3.53
Appeared $\nu_e$	39.57	1.50	2.41	41.07	41.98	-0.91
$\nu$ background	8.65	1.02	3.52	9.67	12.17	-2.50
Cosmic background	0.71	0.93	1.05	1.64	1.76	-0.12
Peripheral	7.41	2.50	2.35	9.91	9.76	+0.15
Appeared $\nu_e$	5.34	1.03	0.77	6.37	6.11	+0.26
$\nu$ background	0.86	0.29	0.46	1.15	1.32	-0.17
Cosmic background	1.21	1.18	1.12	2.39	2.33	+0.06

Table 4.3: Changes in expected event counts in the FD neutrino beam samples relative to previous analysis. The Core sample contains the lower PID and high PID samples. The units for this table are “events.”

## CHAPTER 5

### Studies of the Wrong-sign Component in the Antineutrino Beam

To ultimately understand oscillation effects such as  $\delta_{CP}$ , it is necessary to understand backgrounds. The intrinsic beam-produced  $\nu_e, \bar{\nu}_e$  are one source of background, the magnitudes of which are constrained using Near Detector (ND) data. Selection efficiency could differ for  $\nu_e$  and  $\bar{\nu}_e$  however, and one can also build confidence in the simulation (Monte Carlo, MC) by checking that it has close to the right fractions of each component. Furthermore, confidence in background predictions and testing wrong-sign (WS) separation methods are useful if one wishes to apply similar methods at the Far Detector (FD) to check for enhanced sensitivity, as in Chapter 9.

Therefore, methods to identify WS events and/or produce statistically enhanced samples of WS events are of use. Of note,  $\nu$  charged-current quasi-elastic (CCQE) scattering produces visible protons at the interaction vertex with a greater fraction than  $\bar{\nu}$  scattering, as evidenced by Figure 5.1 showing major interaction types. Therefore, an identifier for protons in a given interaction can be used to create a WS enhanced sample for study. Boosted Decision Trees (BDTs) trained using kinematic information from simulated events can also be used to produce selections of WS enhanced samples. These are methods utilized in work to understand WS contamination in NOvA and are discussed in Section 5.1.

The application of these methods on the  $\bar{\nu}_e$  selection at the ND is presented in Section 5.2. While the measurement of the  $\nu_e$  fraction in the intrinsic  $\bar{\nu}_e, \nu_e$  sample provides confidence in our understanding of background and has other uses that can be explored outside the context of the current main analysis chain, the  $\bar{\nu}_\mu$  WS contamination ( $\nu_\mu$  in the  $\bar{\nu}_\mu$  beam) must be checked as well. Since the oscillation parameters depend on differences in the likelihood that a  $\nu_\mu$  oscillates to a  $\nu_e$  state as compared to the antineutrino version of this oscillation, a background from  $\nu_\mu$  oscillating to  $\nu_e$  in the sample aiming to select  $\bar{\nu}_e$  from  $\bar{\nu}_\mu$  oscillation presents an important background. Various issues made the methods described here undesirable for that sample: this and a different method

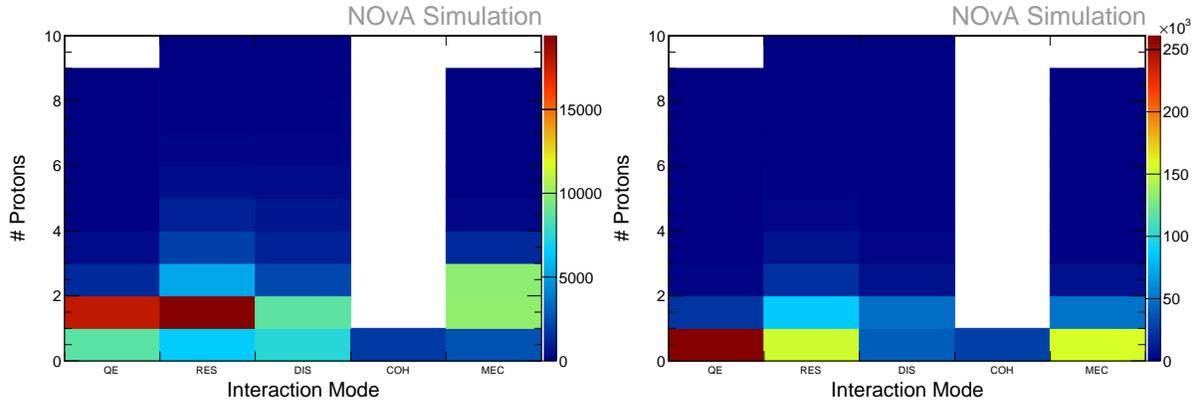


Figure 5.1: Plots of the number of primary protons with at least 40 MeV of kinetic energy in the  $\bar{\nu}$  beam simulation for various interaction modes: quasi-elastic, resonance, deep inelastic scattering, coherent, and meson exchange current. On the left is  $\nu_{\mu}$  in the  $\bar{\nu}$  beam (wrong-sign, WS) and on the right is  $\bar{\nu}_{\mu}$  in the  $\bar{\nu}$  beam (right-sign, RS).

which was applied are briefly discussed in Section 5.3.

Note that NOvA is a collaborative effort, and figures with citations represent collective work by the larger collaboration. These are shown in the pursuit of a more complete narrative and context for my efforts. Figures without citations are generally produced through by me and/or through my research efforts, though at times also in conjunction with collaborators.

### 5.1 Methods for characterizing WS in beam electron antineutrino sample

Three different methods were explored in the context of conducting WS separation in NOvA within the context of the  $\nu_e$  appearance analysis working group. One method was explored originally by a colleague, via a method using the context-enriched CVN (Prong CVN or CVN Prong) [54] discussed previously with respect to energy estimation in Chapter 4. The reconstructed clusters of hits are evaluated in a CVN network that includes information from the event and relevant to this specific cluster, and the output is then a particle identification classifier as opposed to an event classifier. For this study, the relevant CVN score is that of the cluster being identified as a proton. This identifier was found to do a reasonable job of selecting true protons, and a cut on the distance from the event vertex to the start of the proton cluster is used to reduce the impact of proton activity induced by initial neutrons away from the vertex and then detected.

A second WS separation technique using the event classification CVN is one which I spent time

developing. The event-level CVN classifier assigns scores to labels based on the particles present in the final state: in addition to the lepton information useful to the flavor identification, it has entries for the number of  $\pi^0$ ,  $\pi^\pm$ ,  $p$ ,  $n$ . To be counted toward the classification labels, a threshold of 40 MeV kinetic energy for a particle was chosen by those training CVN initially [70]. This is a small fraction of the incident neutrino energy in events analyzed for the oscillation analysis, which is typically around 2 GeV. When checking against simulated events with and without protons, I label events with true protons above 40 MeV as “having protons.” For this analysis, the CVN final state proton (FSProton) classifier score is taken as the sum of all identification labels with at least one proton. Higher scores should correspond to a higher likelihood that the event contains a proton. The labels do separate 1 proton from 2+ protons, but this method is only attempting to classify events as with or without protons, so not much effort was invested in trying to understand the distributions for 1 or 2+ protons separately. The performance of the FSProton score to separate events with and without protons in a version of  $\nu_e$ ,  $\bar{\nu}_e$  selection at the ND is shown in the top left of Figure 5.2: the event was checked for at least one primary proton with 40 MeV of kinetic energy to be considered as having a proton. For the purposes of this exploratory figure, the selection and weights applied were up-to-date at the time the plots were made, but may differ somewhat from the official selection and weights as used in the main analysis presented in Section 5.2<sup>1</sup>. Since not all  $\nu_e$  events have protons and some  $\bar{\nu}_e$  have protons, the separation of right-sign (RS) and WS is not as efficient. The FSProton score recast in terms of RS and WS is shown in the top right of Figure 5.2. The bottom plots of Figure 5.2 show efficiencies for selected events to have a FSProton score above 0.38 based on various properties in the events (number of protons and highest proton kinetic energy).

A third WS separation technique uses a BDT (Nue BDT) with multiple input variables to produce a classification score that is more  $\nu$ -like or more  $\bar{\nu}$ -like. This BDT was originally developed by a colleague and evaluates events in two categories: events with fewer than 2 reconstructed prong clusters and events with 2+ prong clusters. Both categories share 4 variables. For one, given its separation power, the FSProton score is taken as one input variable to the BDT. Second, a measure

---

<sup>1</sup>The CVN electron score for the high PID region, for example, is different from the score used for exploratory purposes here.

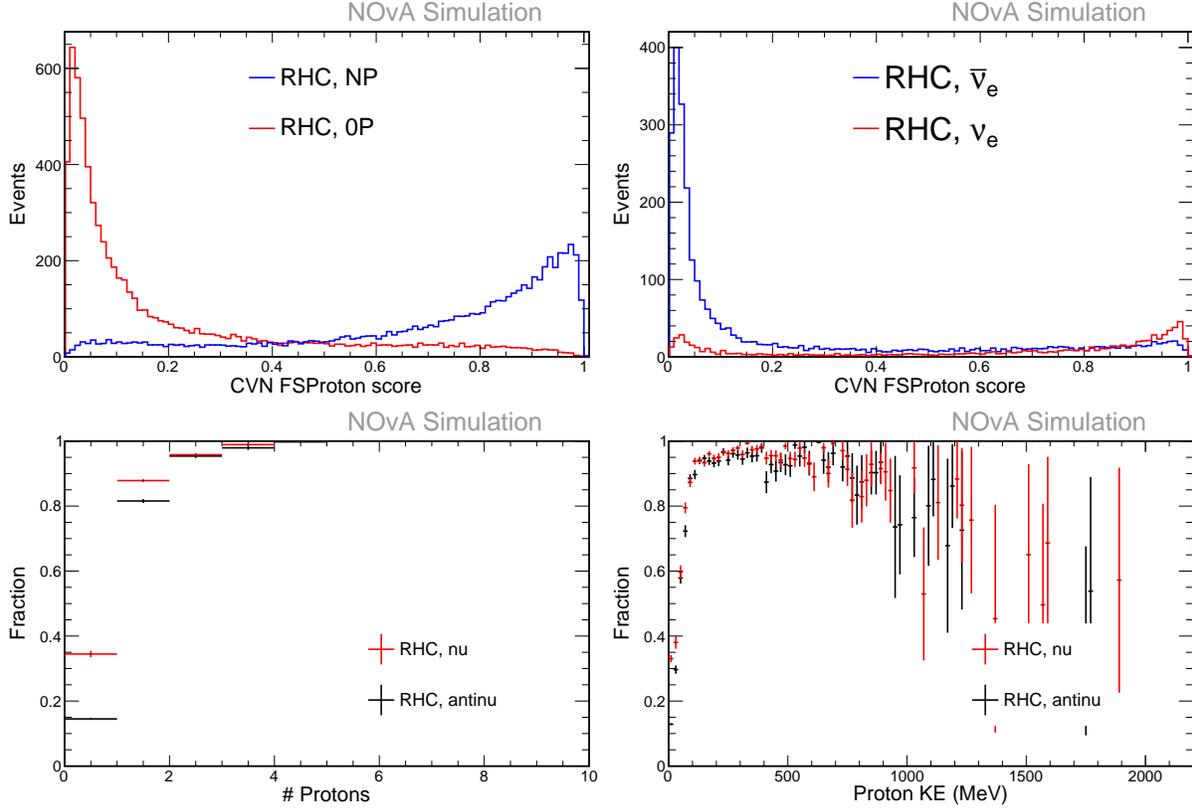


Figure 5.2: CVN FSProton score for selected  $\nu_e$ ,  $\bar{\nu}_e$  events in antineutrino beam simulation, broken into subsamples containing protons of at least 40MeV of kinetic energy (NP) and no protons (0P) above this threshold (top left) and into RS and WS with an additional requirement on CVN electron score  $> 0.97$  (top right). Fraction of selected events above some FSProton score threshold versus number of primary protons above threshold (bottom left) and versus the maximum kinetic energy primary proton (bottom right). Events which do not contain protons fill the bin containing 0MeV. Though difficult to see, the antineutrino sample in the bottom right plot is between 0.1 and 0.15 for the 0 bin. In both of the bottom plots, the  $\nu$  and  $\bar{\nu}$  samples (in the  $\bar{\nu}$  beam) differ most in the bin including events with no protons, but track each other fairly reasonably after this.

of the event’s inelasticity is used, estimating the hadronic energy fraction by calculating the fraction of the total energy not in the shower. Third, the primary (shower)  $dE/dx$  in the first plane is taken as an input. For events where the hadronic activity is not reconstructed and lost in the shower, this variable may yield sensitivity to the extra energy deposition [71]. Fourth, a measurement of the longitudinal asymmetry, or “stretch,” of an event is used as input. An example idea is that for a  $\nu_e$  CC interaction with a proton at the vertex would have the hit-weighted average position closer to the vertex, giving a potentially larger stretch value [71]. For events with at least 2 prongs, the

Prong CVN scores and distance to the prongs are taken as inputs as well. The BDTs were trained on electron (anti-)neutrinos at the FD, but ultimately applied to the ND selection for this study. An FD study is presented in Chapter 9.

## 5.2 WS component in $\bar{\nu}_e$ sample

To study the WS fractions in the  $\bar{\nu}_e$  selection in the antineutrino beam at the ND, I leveraged the above methods to create statistically enhanced WS and RS samples. Using the high PID  $\nu_e$ -appearance style selection at the ND, a score for each of the three methods for separating the events into RS-enhanced and WS-enhanced samples was found by exploring a figure-of-merit (FOM) as a function of score. The optimization for prong CVN was particularly tricky because extra cuts were included to remove events with only one prong and events with only two back-to-back prongs (more on these later). The optimization was performed with  $FOM = S/\sqrt{S+B}$ , where S is  $\nu_e$  and  $\nu_\mu$  and B is  $\bar{\nu}_e$ ,  $\bar{\nu}_\mu$ , and NC<sup>2</sup>. For the determination of WS fraction in data, define:

1. S = integrated  $\nu_e + \nu_\mu + \text{NC}(\nu)$  from the MC
2. B = integrated  $\bar{\nu}_e + \bar{\nu}_\mu + \text{NC}(\bar{\nu})$  from the MC
3. D = integrated selected data
4. RS = RS enhanced selection (events selected based on comparing given WS PID score to the score chosen to give good separation as noted above)
5. WS = WS enhanced selection (events selected based on comparing given WS PID score to the score chosen to give good separation as noted above)

such that we describe the RS- and WS-enhanced samples by

$$\begin{aligned}\alpha B_{RS} + \beta S_{RS} &= D_{RS} \\ \alpha B_{WS} + \beta S_{WS} &= D_{WS}\end{aligned}\tag{5.1}$$

---

<sup>2</sup>In determining the WS fraction the NC is treated separately depending on whether it came from a neutrino or antineutrino. This means that the FOM calculation is slightly different from the WS fraction calculation later, but the NC in the high PID selection is only about 5%.

with  $\alpha$  and  $\beta$  being RS and WS scales for the amounts in the MC, respectively. After integrating the WS- and RS-enhanced selections of data and MC, these equations have two unknowns ( $\alpha, \beta$ ) such that the scale parameters can be determined for each method. The statistical errors from the MC will have some effect, but since  $3.1 \times 10^{20}$  POT were used in ND data in the original application of this analysis, the MC had about a factor of 10 higher statistics. Therefore, the statistical errors were ignored in MC for simplicity. For this recent top-up analysis with additional data to bring the total to  $7.9 \times 10^{20}$ , this approach was not changed. Calling  $\mathcal{A}$  the matrix of MC predictions

$$\mathcal{A} = \begin{bmatrix} B_{RS} & S_{RS} \\ B_{WS} & S_{WS} \end{bmatrix} \quad (5.2)$$

then the covariance matrix is determined as

$$\mathcal{C} = \mathcal{M} \begin{bmatrix} \sigma_{data,RS}^2 & 0 \\ 0 & \sigma_{data,WS}^2 \end{bmatrix} \mathcal{M}^T \quad (5.3)$$

where  $\mathcal{M} = \mathcal{A}^{-1}$ , such that the covariance matrix is

$$\mathcal{C} = \begin{bmatrix} \mathcal{M}_{11}^2 \sigma_{data,RS}^2 + \mathcal{M}_{12}^2 \sigma_{data,WS}^2 & \mathcal{M}_{11} \mathcal{M}_{21} \sigma_{data,RS}^2 + \mathcal{M}_{12} \mathcal{M}_{22} \sigma_{data,WS}^2 \\ \mathcal{M}_{11} \mathcal{M}_{21} \sigma_{data,RS}^2 + \mathcal{M}_{12} \mathcal{M}_{22} \sigma_{data,WS}^2 & \mathcal{M}_{21}^2 \sigma_{data,RS}^2 + \mathcal{M}_{22}^2 \sigma_{data,WS}^2 \end{bmatrix} \quad (5.4)$$

using Reference [72], which borrows from the PDG (see for example [4]).

For the  $\nu_e$  selections, the lower ( $\geq 0.89, < 0.98$ ) and high ( $\geq 0.98$ ) PID selections were treated separately. One can solve for  $\alpha$  and  $\beta$  and then solve for the WS fraction as  $\beta(S_{RS} + S_{WS}) / [\beta(S_{RS} + S_{WS}) + \alpha(B_{WS} + B_{RS})]$  and the corresponding errors using the  $\mathcal{C}_{ij}$  values. However, looking at the calculation of WS fraction a little more closely, the denominator  $\beta(S_{RS} + S_{WS}) + \alpha(B_{WS} + B_{RS})$  simplifies to exactly  $D = D_{RS} + D_{WS}$ , the total number of selected data events. The error is then simplified from a more complicated formula to  $\sigma = \sqrt{\mathcal{C}_{22}^2 S^2 / D^2}$  where the error in the denominator

term is neglected, and was found to give consistently similar errors to the more complicated version.

$$f_{WS} = \frac{\beta S}{D} \pm \sqrt{C_{22}^2 \frac{S^2}{D^2}} \quad (5.5)$$

For the calculations, the central-value simulation results are used. To account for the effects of flux and cross-section systematics to alter the amount of WS, a colleague ran multiple trials with altered flux and cross-section parameters. The data-driven results are then checked against the range of possibilities from the simulation where possible.

The methods described here have not been deemed trustworthy for the  $\bar{\nu}_\mu$  sample, with one reason being data/MC distributions had shape discrepancies that were difficult to attribute to either WS or various systematics such as those relating to detector response. Before application of the methods to the  $\bar{\nu}_e$  selection, it is therefore worthwhile to consider the data/MC distributions. For the CVN FSProton and Nue BDT samples, these data/MC are shown in Figure 5.3 and show no significant issues in shape. For the CVN Prong proton score, data/MC are shown on an axis of the maximal proton score of any prong starting within 20cm of the event vertex. If no such prong exists for the event, a score of 0 is assigned. This is plotted in Figure 5.4: notice the significant data/MC discrepancy in the 0 bin.

If extra cuts are made to remove events with  $<2$  prongs or two-prong events aligned back-to-back (dot product  $\leq -0.8$ ), the remaining events in the 0 bin are found to have quite good data/MC agreement (see Figure 5.6). Such cuts could be motivated as follows. In the case of having one prong, this would be expected to be the lepton, since these are selected  $\bar{\nu}_e$  CC candidates. If only one prong is present, one therefore would expect to have no proton prongs to find. Samples of back-to-back prongs were found to contain a pathology by which the reconstruction split one particle track into two prongs [73]. Figure 5.6 shows the low CVNe selection before and after applying these prong cuts, on a linear view to more easily see the “0 bin.”

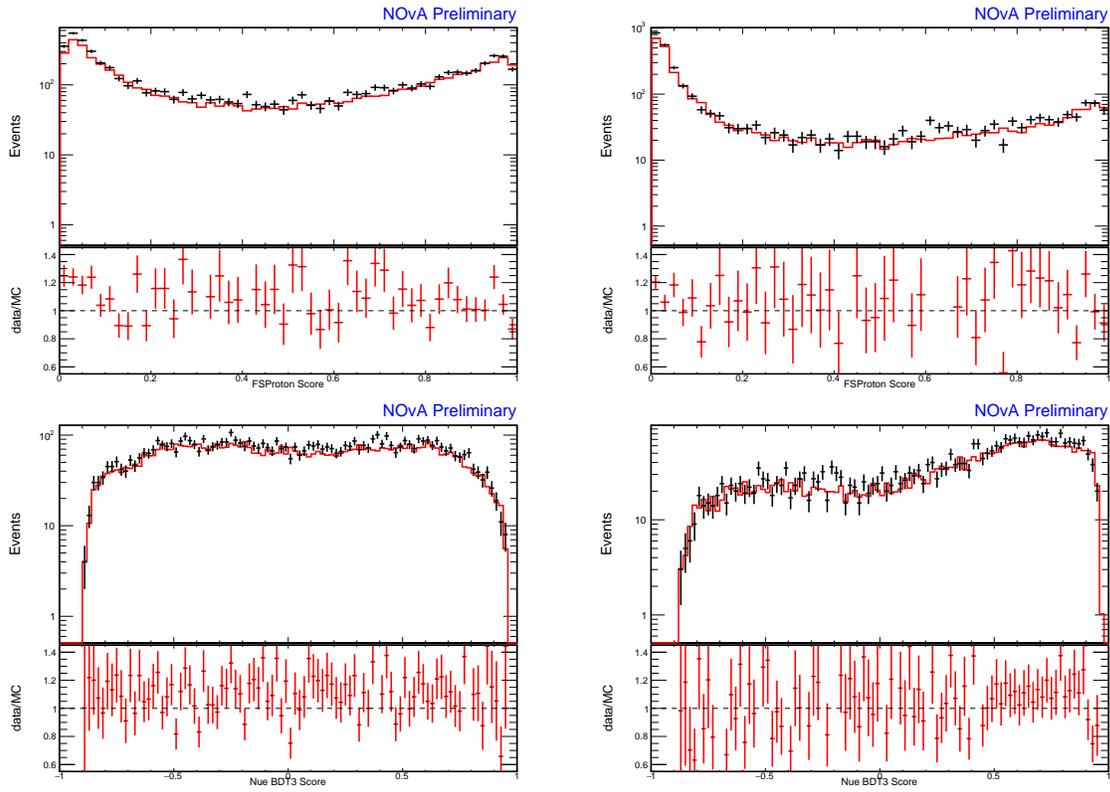


Figure 5.3: (top) Data/MC for the FSProton Score, in the low CVNe selection (left) and the high CVNe selection (right). (bottom) Data/MC for the Nue BDT3 Score, in the low CVNe selection (left) and the high CVNe selection (right). In these plots and similar plots in this chapter, the data are the black points and the simulation is presented with a line (here red) instead of points.

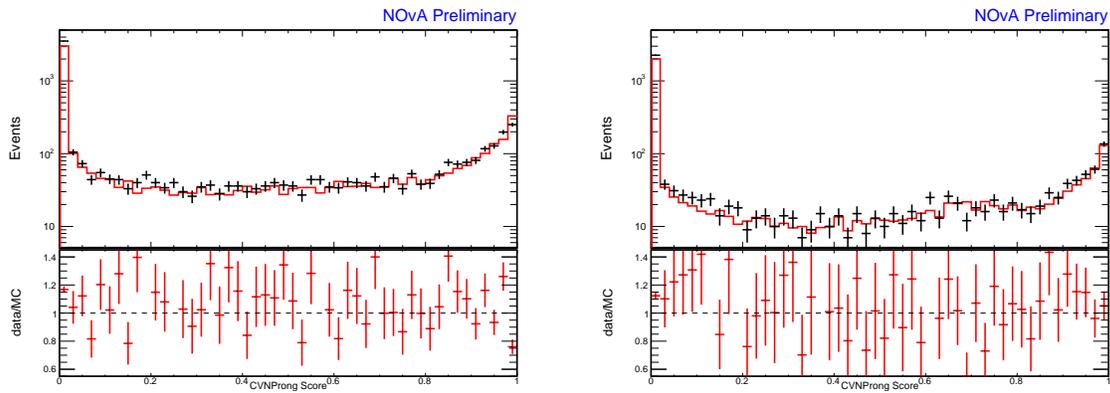


Figure 5.4: Data/MC for the CVN Prong proton Score, as described in the text, for the low CVNe selection (left) and the high CVNe selection (right). Note the especially discrepant 0 bin.

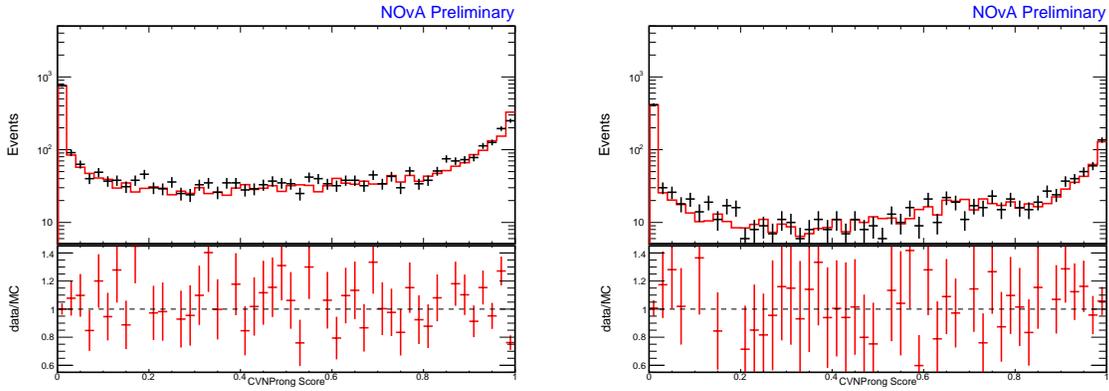


Figure 5.5: Data/MC for the CVN Prong proton Score, as described in the text with the extra added cuts, for the low CVNe selection (left) and the high CVNe selection (right). Note that the largest change is in the 0 bin and that it is now in much better agreement.

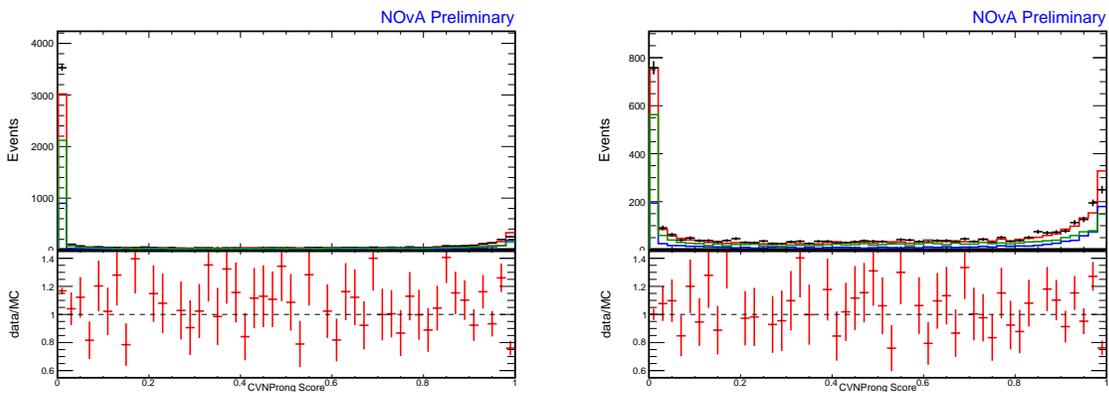


Figure 5.6: Data/MC for the CVN Prong proton Score, as described in the text, for the low CVNe selection, with no extra cuts (left) and with extra cuts on prongs (right). Note the discrepancy in the 0 bin improve with the addition of this cut.

### 5.2.1 Preliminary results

Applying the separation in the three methods gives the resulting RS and WS enhanced distributions shown in Figure 5.7. The CVN Prong selector is the least efficient at getting events into the WS enhanced selection. Note the difference in energy in the low and high CVNe selections – the high CVN selection is nicely peaked far from the energy cuts, while the low CVNe selection has a pronounced tail leading to the energy cut-off. Note that especially in the CVN Prong proton method on the low CVNe selection, the shapes of the RS and WS in the two sides of the plot look perhaps somewhat different. The data events are drawn with black crosses in the figure.

In addition to the insights gained about purity of WS in the WS enhanced region and efficiency of getting WS events into the WS enhanced region, one can note the relation of data and MC in the two regions. It is ultimately this relation which informs the calculation of  $\alpha$  and  $\beta$  (and therefore the WS fraction). In the BDT method, note that both the left (RS enhanced) and right (WS enhanced) regions have excess of data over MC, while for the CVN Prong proton score method, the left has a data excess but the right has near equal data and MC. This can be seen as a contributor to a larger swing in WS fraction<sup>3</sup>. Indeed this is what is seen in the results of calculating Equation 5.5 for these samples – see Table 5.1 and Figure 5.8.

Method	Lower PID	High PID
Event CVN Proton ID	$0.276 \pm 0.023$	$0.220 \pm 0.018$
Prong CVN Proton ID	$0.134 \pm 0.044$	$0.197 \pm 0.035$
WS BDT	$0.318 \pm 0.021$	$0.225 \pm 0.017$

Table 5.1: WS fraction in the two  $\nu_e$  selection regions, one with somewhat lower CVN electron score (lower PID) and one with high CVN electron score (high PID). The low CVNe CVN Prong proton score appears problematic, and is discussed further in the next part of the text.

### 5.2.2 WS calculations using additional prong cuts

Given justifications for additional prong cuts to filter out single-prong events and events with only two back-to-back prongs and the knowledge that the data/MC in the first bin improves drastically in the CVN Prong proton score with the addition of these cuts, the WS fraction and  $\beta$  are

<sup>3</sup>If the left proportionally controls  $\bar{\nu}_e$  more strongly, then this will drive up the  $\bar{\nu}_e$  fraction. Since data and MC are close on the right, driving up  $\bar{\nu}_e$  forces the  $\nu_e$  fraction to be driven down.

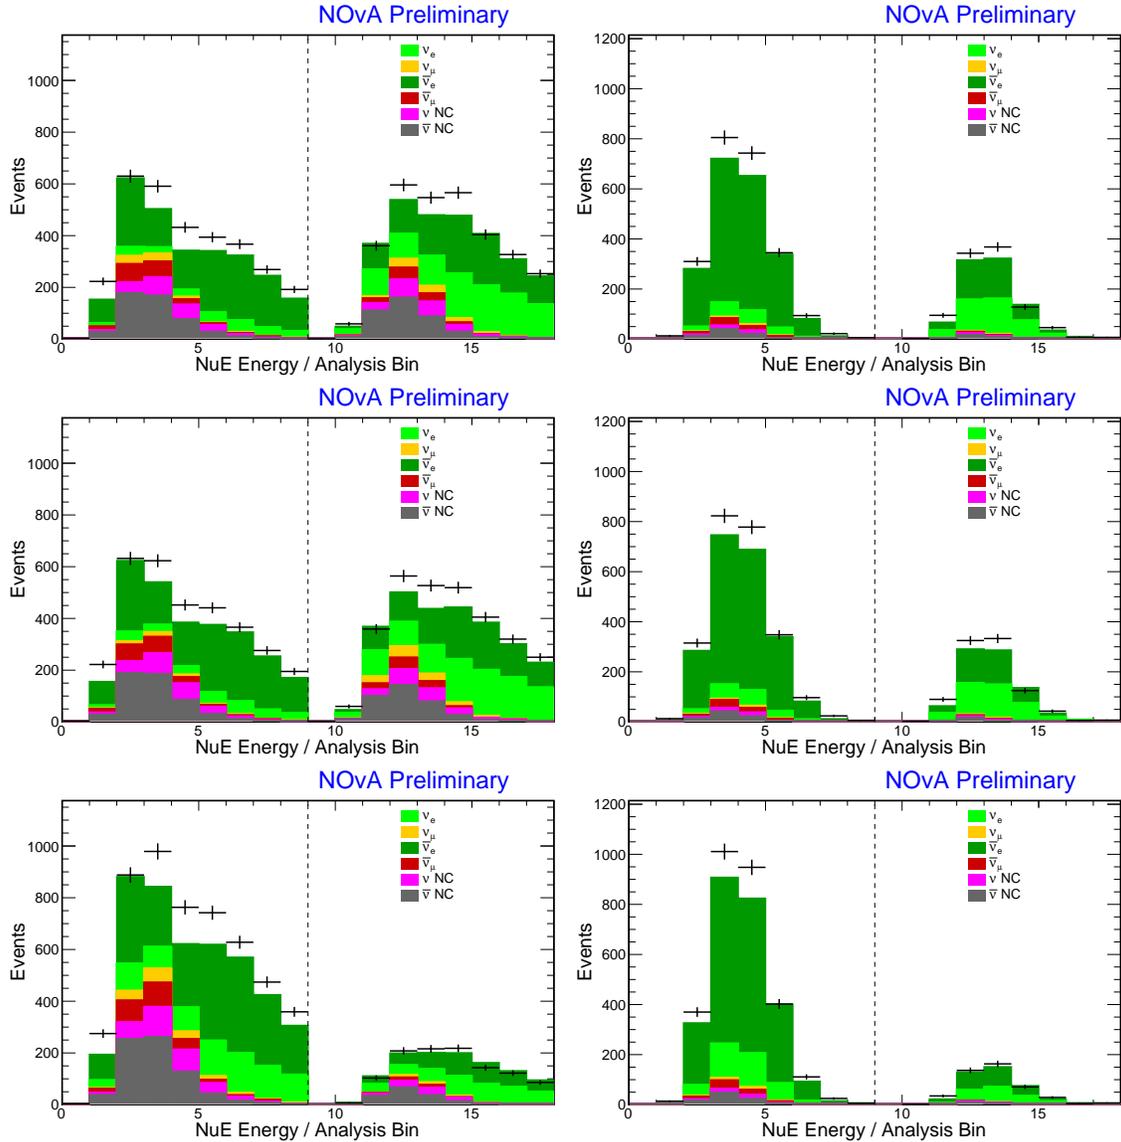


Figure 5.7: The RS enhanced (left side of plot) and WS enhanced (right side of plot) regions for the three  $\nu_e$  WS studies. In each row, the left is the lower PID selection and the right is the high PID selection. The top row is using FSProton, the middle row is using the BDT score, and the bottom row is using the CVN prong score.

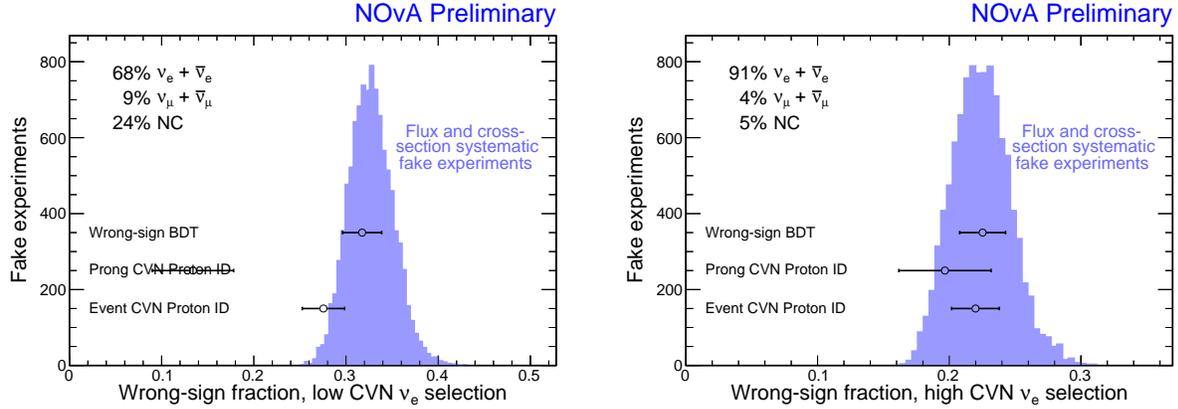


Figure 5.8: Preliminary results for  $f_{WS}$  in the lower and high PID selections in the ND. The fraction of the central-value simulation selected of each type is given in the upper left corner.

calculated in this subsample of the overall events. Since the simulation distribution was not run with additional cuts on the prongs, the central-value tuned simulation WS fractions (0.3479 in the lower PID selection and 0.2447 in the high PID selection) are used.

The RS enhanced and WS enhanced breakdowns for the three methods are presented in Figure 5.9. There are some important features to note in these plots. For one, as said before, there is a large reduction in the total number of events. However, the breakdown of these reductions is also important. Print outs of the event counts in the simulation were made, and the ratios of events before and after applying the cut in the RS enhanced and WS enhanced regions were compared. These are shown in Table 5.2. Notably, the large reduction of events in the CVNProng method comes in the RS enhanced region, while these removed events were more evenly distributed from each region of the other two methods, especially in the BDT. Furthermore, especially in the lower PID selection, the WS enhanced selection has a larger fraction of events relative to the RS enhanced selection than it nominally does. This is again especially true in the CVN Prong proton score method, where a lot of events were removed from the RS enhanced region.

The resulting WS fractions and  $\beta$  values from the calculations are given in Table 5.3, and figures to go with these results are given in Figure 5.10 comparing these results to the simulated WS fraction. In the lower PID selection, the CVN Prong proton score is again somewhat off but by significantly less than in the full analysis, off by less than  $2\sigma$ .

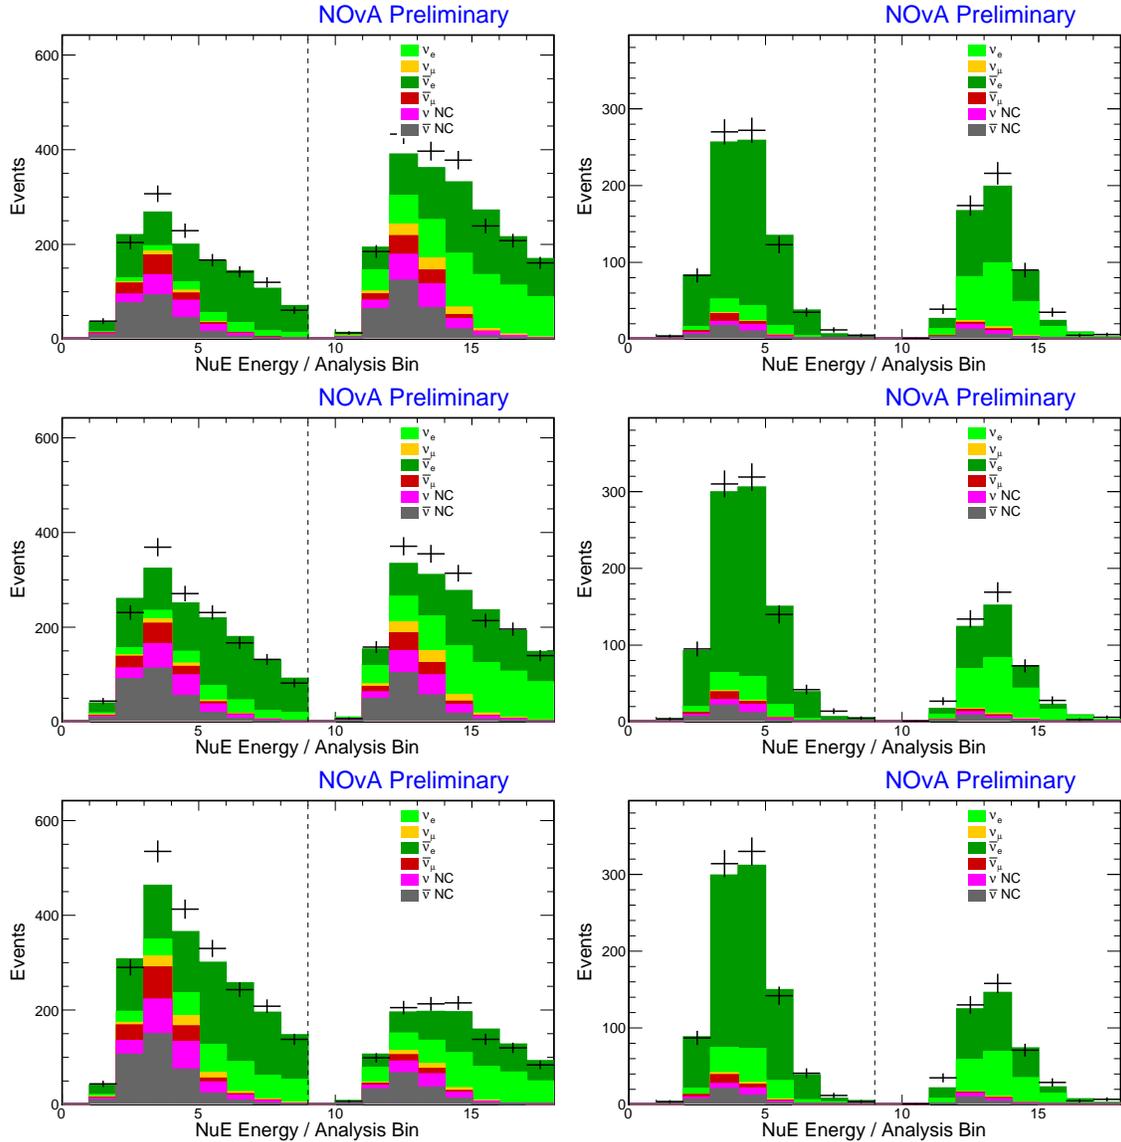


Figure 5.9: The RS enhanced (left side of plot) and WS enhanced (right side of plot) regions for the three  $\nu_e$  WS studies with the additional of cuts on prongs. In each row, the left is the low CVNe selection and the right is the high CVNe selection. The top row is using FSProton, the middle row is using the BDT score, and the bottom row is using the CVN prong score.

Simulated events after extra cuts/Nominal events		
High PID		
Method	RS enhanced	WS enhanced
Event CVN Proton ID	0.374	0.580
Prong CVN Proton ID	0.350	0.967
WS BDT	0.416	0.486
Lower PID		
Method	RS enhanced	WS enhanced
Event CVN Proton ID	0.449	0.678
Prong CVN Proton ID	0.465	0.977
WS BDT	0.524	0.612

Table 5.2: Fraction of simulated events in the given region after applying extra cuts on prongs compared to before. Note that the large reduction of events in the CVN Prong method comes in the RS enhanced region, while these removed events were more evenly distributed from each region of the other two methods, especially in the BDT.

Method	$\beta$ , Lower PID	$f$ , Lower PID	$\beta$ , High PID	$f$ , High PID
Event CVN Proton ID	1.00	$0.334 \pm 0.036$	1.20	$0.279 \pm 0.029$
Prong CVN Proton ID	0.78	$0.263 \pm 0.052$	1.21	$0.282 \pm 0.041$
WS BDT	1.13	$0.380 \pm 0.033$	1.18	$0.275 \pm 0.027$

Table 5.3: WS scales ( $\beta$ ) and resulting WS fractions ( $f$ ) in the lower PID selection and high PID selection, with the addition of extra cuts on prongs. In the lower PID selection, the CVN Prong proton score is somewhat off but by significantly less than in the full analysis.

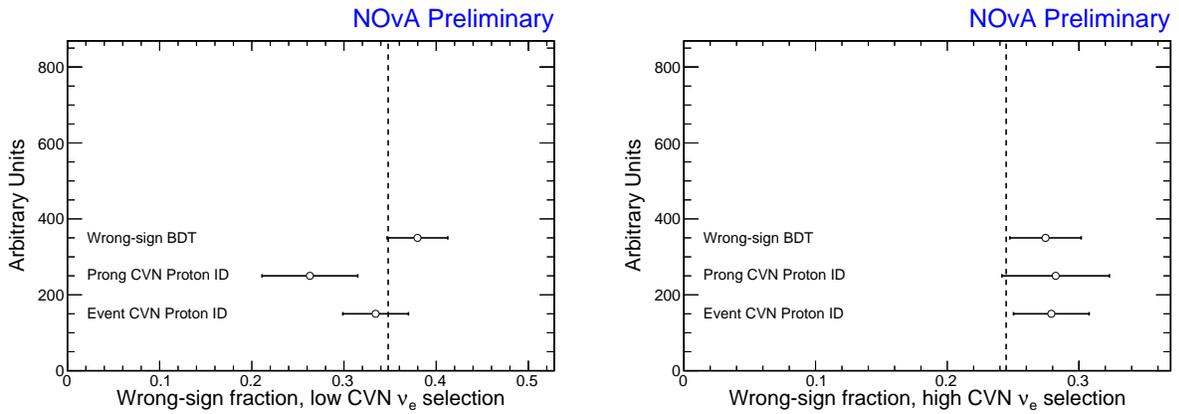


Figure 5.10: Preliminary results for the WS fraction in the lower and high PID selections in the ND with the additional cuts on prongs.

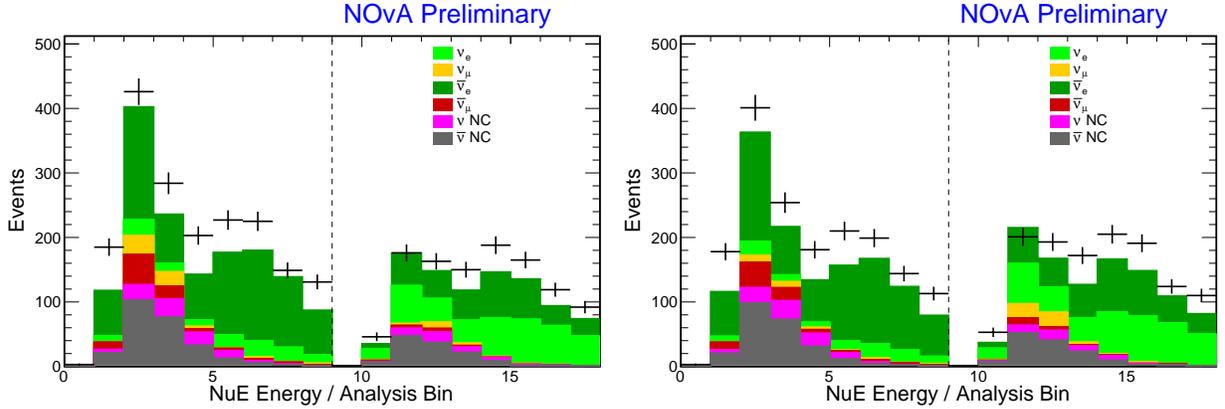


Figure 5.11: The RS enhanced (left side of plot) and WS enhanced (right side of plot) regions for events in the lower PID that fail the extra cuts on prongs: FSProton (left) and BDT (right).

### 5.2.3 Attempt to understand and account for bias

Studying just the set of events that may be more well-modeled based on the CVN prong scores may raise concerns about potentially introducing a bias to the WS calculation. In this section, such considerations are addressed. While the CVN Prong method cannot be used on the events failing the prong cuts as the RS-enhanced region is where many events are cut, the FSProton and BDT methods can be used to examine the events failing the cuts. This failing sample is then studied in the same way as the passing sample, but for just the FSProton and BDT methods. Then, an attempt is made to arrive at an alternate estimate of the overall WS fraction using the CVN Prong method, by leveraging the samples of events passing and failing the prong cuts.

First, consider the FSProton and BDT methods. The method applied is the same as was used on the full sample and on the sample of events passing the prong cuts: the only difference is that the events failing the prong cuts are used. Figure 5.11 shows the RS-enhanced and WS-enhanced breakdowns applied to this particular sample of events. Table 5.4 shows the results for the application of these methods. The WS fraction in the simulation (central-value) is 0.308, and these methods produce results that are a little low but just outside the  $1\sigma$  range. That these two methods show reasonable WS fractions is itself a result of this study.

To now make a statement about the low CVNe CVN Prong sample and potential biases, start by looking at the  $\beta/\alpha$  between the two selection samples and the various WS characterization

Method	$\beta$ , Lower PID	$f$ , Lower PID
Event CVN Proton ID	1.09	$0.277 \pm 0.029$
WS BDT	1.10	$0.278 \pm 0.026$

Table 5.4: WS scales ( $\beta$ ) and resulting WS fractions ( $f$ ) for events in the lower PID selection region that fail the extra cuts on prongs. The MC WS fraction is 0.308.

methods: these are depicted in Table 5.5. In both the FSProton and BDT methods, the ratio  $\beta/\alpha$  is smaller in the failed sample than the passing sample. While one cannot directly calculate the WS fraction or  $\beta/\alpha$  for the CVN Prong method in the failed event sample, one can be inferred from the average fractional change of  $\beta/\alpha$  from the FSProton and BDT methods. Therefore, say  $\beta/\alpha$  for the CVN Prong method in the sample of events failing the cut is  $\beta_f/\alpha_f = C\beta_p/\alpha_p = 0.559$  where  $C = 0.835$  is the average ratio of  $\beta/\alpha$  values for the other (FSProton and BDT) methods.

Method	$\beta/\alpha$ , passing cuts	$\beta/\alpha$ , failing cuts	$(\beta_f/\alpha_f)/(\beta_p/\alpha_p)$
Event CVN Proton ID	0.941	0.860	0.914
WS BDT	1.148	0.867	0.755
Prong CVN Proton ID	0.669	n/a	n/a

Table 5.5: Ratios of WS scaling to RS scaling ( $\beta/\alpha$ ) as determined for the methods in the two event samples. The subscript f is for the failing sample and the subscript p is for the passing sample. The average of the two values in the rightmost column is used to scale the CVN Prong value in the passing sample to an inferred value in the failing sample.

Taking this inferred  $\beta/\alpha$  and the counts of data events and simulated RS and WS events, one can calculate the  $\alpha$  and  $\beta$  separately. Table 5.6 shows the corresponding RS, WS, and data components for both samples. Since  $D = \alpha B + \beta S$

$$\begin{aligned}
X &= \beta_f/\alpha_f = 0.559 \\
\implies D_f &= \alpha_f B_f + \alpha_f X S_f \\
\implies \alpha_f &= \frac{D_f}{B_f + X S_f} = 1.407 \\
\implies \beta_f &= \alpha_f X = 0.787
\end{aligned} \tag{5.6}$$

Next, combining the determined values for the CVN Prong method in the sample of events

Event Type	Passing Prong Cuts	Failing Prong Cuts
WS (S)	1099.939	742.764
RS (B)	2061.81	1669.3
Total	3161.7	2412.1
Data	3282	2933

Table 5.6: The data and MC event count breakdowns for the samples of events in question.

passing prong cuts and the inferred values for the sample failing the cuts, an overall WS fraction is determined. The overall WS fraction goes as

$$f = \frac{\beta S}{\beta S + \alpha B} = \frac{\beta S}{D} \quad (5.7)$$

where the numerator is  $\beta_p S_p + \beta_f S_f = 1448.0$  and  $\alpha_p B_p + \alpha_f B_f = 4767.2$ . The sum of these two quantities agrees with the sum of the selected data events to 0.2 events, so this is considered good. Then, calculating the WS fraction as  $f = \frac{\beta_p S_p + \beta_f S_f}{D_p + D_f} = 0.233$ . Note that this is 3.0% lower than the value found using just the events passing the prong cuts: this is taken as an estimate of the bias in the calculation. To estimate an error on this WS fraction, the 5.2% error on the WS fraction in the passing sample is taken with an additional 3.0% error (100% of the difference as an estimate of bias). Taking these errors in quadrature one then arrives at  $f_{WS} = 0.233 \pm 0.060$ , which is still on the low side, but in much better agreement with the other determinations and with the simulated WS fraction, as shown in Figure 5.12.

### 5.3 WS in the $\bar{\nu}_\mu$ sample

Taking into account simply the flux, the WS fractions are similar in the two beam modes: roughly 4%  $\bar{\nu}_\mu$  in the  $\nu_\mu$  beam and roughly 6%  $\nu_\mu$  in the  $\bar{\nu}_\mu$  beam [14]. However, due mainly to cross-section differences, the WS contamination is therefore amplified in the antineutrino beam and suppressed in the neutrino beam, as depicted in Figure 5.13. The exact fraction of the sample that is WS depends on the sample, detector, and selection cuts applied. In terms of CC events in the FD, the  $\bar{\nu}_\mu$  fraction in the  $\nu$  sample is 3%, and the  $\nu_\mu$  fraction in the  $\bar{\nu}$  sample is 15% [14]. In terms of selected events at the ND and FD, the WS fractions are somewhat smaller, but still of order 10%,

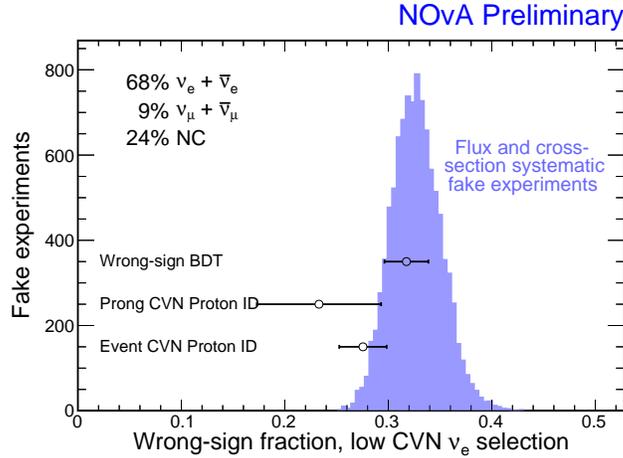


Figure 5.12: Preliminary results for the WS fraction in the lower PID selection in the ND, with the CVN Prong results replaced with the estimate determined in this section.

as shown in Figure 5.14. While the WS contamination in the  $\nu_\mu$  beam samples is negligible, the antineutrino beam WS should be considered.

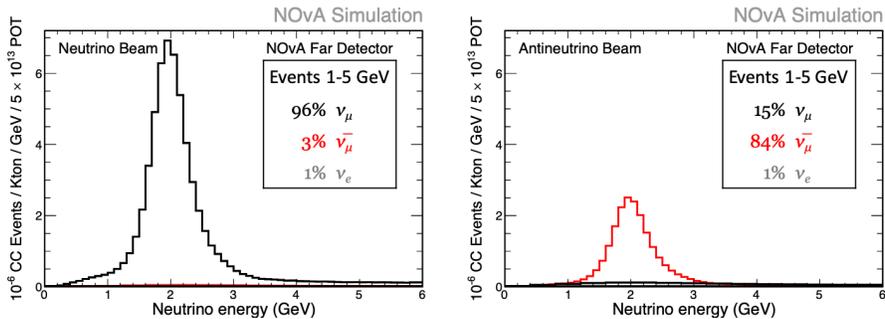


Figure 5.13: Neutrino interactions in the NOvA FD simulation [14]. Note the larger WS fraction in the antineutrino beam. This is mainly due to cross-section differences between neutrinos and antineutrinos, with some effect from the flux as well.

An original goal was to apply the WS selection techniques discussed previously to the  $\bar{\nu}_\mu$  sample in the ND as well. However, two issues were ultimately found that led to a lack of trust in this approach. For one, as was noted in Section 5.2, data/MC distributions had interesting shape discrepancies that were difficult to attribute to either WS or various systematics such as those relating to detector response. Furthermore, the WS fraction in the  $\bar{\nu}_\mu$  sample is significantly lower than that in the  $\bar{\nu}_e$  sample, and the reduction is even more drastic if just looking at the lower hadronic energy fraction region. This has the downside of allowing small changes in the RS sample to have large

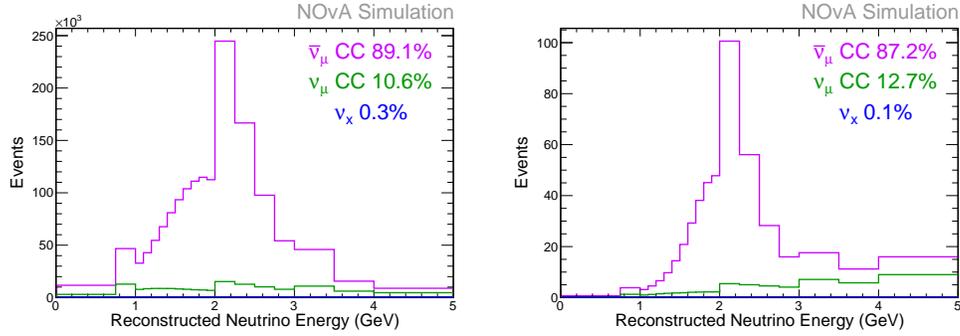


Figure 5.14: Events with the full analysis weights passing selection cuts in the  $\bar{\nu}_\mu$  candidate samples at the ND (top) and FD (bottom). The fractions of true  $\nu_\mu$  CC,  $\bar{\nu}_\mu$  CC, and other components are given in the top right of each figure.

impact on the WS samples. Therefore, the results of application of these techniques could not be fully trusted. In hopes of finding a better version of a similar method, I trained gradient-boosted BDTs on each hadronic energy fraction quantile used in the analysis in the ND. However, these too had large data/MC discrepancies, which further investigation showed was present in the input variables themselves as well. This was not understood in enough detail to decide to proceed. It is hoped that additional study outside the scope of this work can lead to the advancement in understanding necessary to proceed with this measurement.

A colleague devised a different method [74] which examined neutron captures in NOvA and used fits of this to scale the fraction of  $\nu_\mu$  in a  $\bar{\nu}_\mu$  beam selection. His work finds the WS scale to be  $1.05 \pm 0.12$  [74]. I took this number and the selected data and simulated events broken into relevant categories and formed an estimate of the WS fraction from this to be plotted against the distribution of the simulated WS fraction. This is given in Figure 5.15.

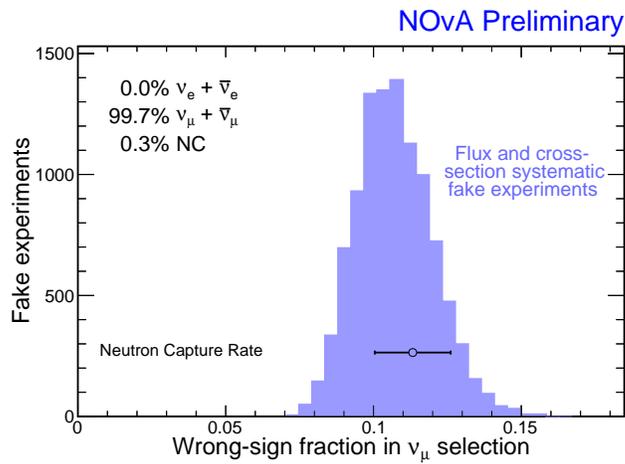


Figure 5.15: Results of a data-driven WS fraction study in the  $\bar{\nu}_\mu$  ND selection, using the WS scale derived in Reference [74] and overlaying the simulated distribution of WS fraction in the selection, allowing for systematic variation of flux and cross-section parameters. The fraction of the central-value simulation selected of each type is given in the upper left corner.

## CHAPTER 6

### NOvA Oscillation Results

Chapters 4-5 discussed elements of analysis relevant to studying 3-flavor neutrino oscillations with NOvA. The recent NOvA analysis with both neutrino and antineutrino data is also detailed in Reference [36].<sup>1</sup> This chapter will report these results with  $8.85 \times 10^{20}$  (full-detector-equivalent) POT in the neutrino beam and  $12.33 \times 10^{20}$  POT in the antineutrino beam. Following this, two aspects of the analysis results in which I had a significant role are discussed.

Note that NOvA is a collaborative effort, and figures with citations represent collective work by the larger collaboration. These are shown in the pursuit of a more complete narrative and context for my efforts. Figures without citations are generally produced through by me and/or through my research efforts, though at times also in conjunction with collaborators.

#### 6.1 Analysis and joint fit results

When the collaboration's analysis techniques were performed as discussed in Chapter 4 on the Far Detector (FD) data, 58  $\nu_e$  appearance candidates were selected from the neutrino beam, and 27  $\bar{\nu}_e$  appearance candidates were selected from the antineutrino beam. Figure 6.1 shows one such reconstructed slice from both beam samples. These events have identifiable showers, and the neutrino beam candidate has a secondary track that could be a proton. The lack of a clear secondary track in the antineutrino beam candidate could be indicative of a neutron in the final state. The  $\nu_\mu$  and  $\bar{\nu}_\mu$  disappearance studies found 113 and 102 candidates, respectively.

At the best fit predictions, the expected background in the antineutrino beam  $\bar{\nu}_e$  appearance study is  $10.3_{-0.5}^{+0.6}$ , representing a  $4.4\sigma$  measurement of  $\bar{\nu}_e$  appearance, the first such significant measurement of this at long-baseline in a  $\bar{\nu}_\mu$  beam. The expected background for appearance in the neutrino beam is  $15.0_{-0.9}^{+0.8}$  events. With respect to the previous analysis, 2 new events were gained in the core sample and 10 events were lost. This is a total change of -8 events, which is in

---

<sup>1</sup>While citations for figures in this chapters will usually point to internal documentation, some are similar to those appearing in the cited article.

the right direction of the changes expected in the migration study in Chapter 4, though larger in amount. No change was observed in the peripheral sample, completely consistent with the results of the migration study. One new event that appears to be background by eye was discovered in the  $\nu_e$  selection, and Section 6.3 discusses this.

The top two plots in Figure 6.2 show the  $\nu_\mu$  (left) and  $\bar{\nu}_\mu$  (right) disappearance spectra, with the red indicating the expectation with no oscillation. The analysis uses four quantiles of hadronic energy fraction to split the events into four kinematic regions which are then extrapolated to the FD separately: the four squares in each plot represent these four hadronic energy quantiles. The bottom two plots of Figure 6.2 show the  $\nu_e$  (left) and  $\bar{\nu}_e$  (right) appearance spectra, with the best-fit in purple and the backgrounds noted. A clear excess above background is seen. There are a projected 0.6 wrong-sign appearance events in the neutrino beam ( $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ ) and 2.2 wrong-sign appearance events in the antineutrino beam ( $\nu_\mu \rightarrow \nu_e$ ). The expected cosmic background is 3.3 events for the  $\nu_e$  candidate sample and 1.1 for the  $\bar{\nu}_e$  candidate sample.

The selected  $\nu_e$  candidates from the neutrino beam and  $\bar{\nu}_e$  candidates from the antineutrino beam are overlaid onto a bi-event plot in Figure 6.3, with two sets of oscillation parameters drawn in for both normal and inverted hierarchy. The data are more consistent with the oval for the normal hierarchy and upper octant ( $\theta_{23} > \pi/4$ ), though all four combinations of normal/inverted hierarchy and lower/upper octant ( $\{\text{NH, IH}\} \times \{\text{LO, UO}\}$ ) are considered. The solar parameters were held fixed,  $\Delta m_{21}^2 = (7.53 \pm 0.18) \times 10^{-5} \text{eV}^2$  and  $\sin^2 \theta_{12} = 0.307_{-0.012}^{+0.013}$ , and as noted in Chapter 4  $\sin^2 \theta_{13}$  was a constrained parameter. The results of the joint oscillation fit (see Equation 4.1) are shown in Figure 6.4 and written in Table 6.1 as presented in [36]. These results prefer the normal hierarchy at  $1.9\sigma$ . Furthermore, the results slightly prefer non-maximal mixing in the upper octant, but are only at the level of  $1.2\sigma$  discrepant from maximal mixing. As the  $\sin^2 \theta_{23}$  vs.  $\delta_{CP}$  contours in Figure 6.4 show, the points at  $\delta_{CP} = \pi/2$  in the inverted hierarchy are disfavored at over  $3\sigma$ .

NOvA's full projected POT exposure is  $36 \times 10^{20}$  in both beam modes, and it is expected that NOvA will run until the mid 2020s. By this time, NOvA's ultimate sensitivity will depend on the oscillation parameters. The reach in mass hierarchy is given in the left of Figure 6.5 for a range

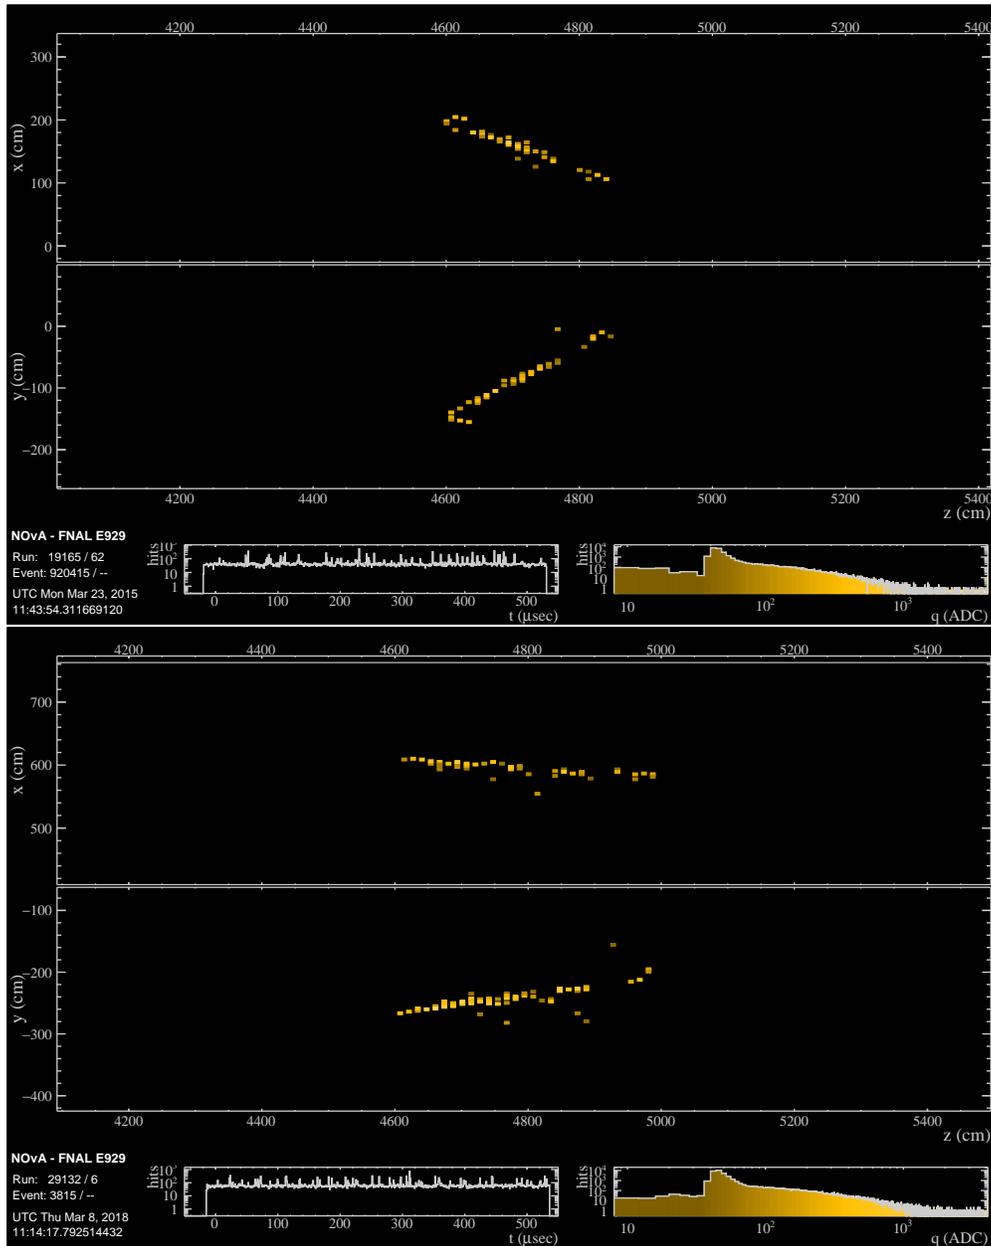


Figure 6.1: Two selected slices from candidate events from the NOvA FD used for the oscillation studies discussed in this chapter. The top event is from the neutrino beam [75], and the bottom event is from the antineutrino beam [76]. Both events have a showering primary track. The neutrino beam candidate has a secondary track that could be indicative of a proton, and the antineutrino beam candidate has no such secondary, which could be indicative of a neutron.

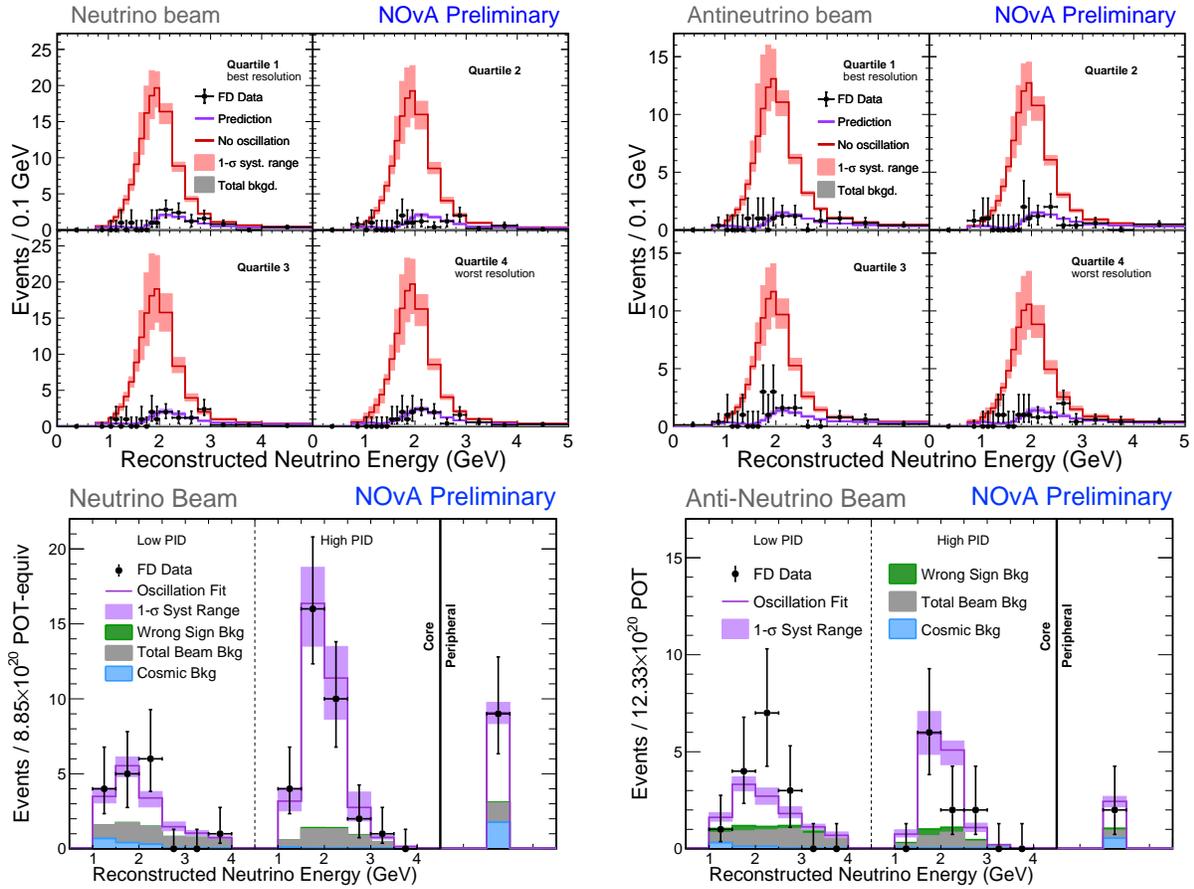


Figure 6.2: Spectra for data samples at the FD:  $\nu_\mu$  disappearance [77] (top left),  $\bar{\nu}_\mu$  disappearance [77] (top right),  $\nu_e$  appearance [61] (bottom left), and  $\bar{\nu}_e$  appearance [61] (bottom right).

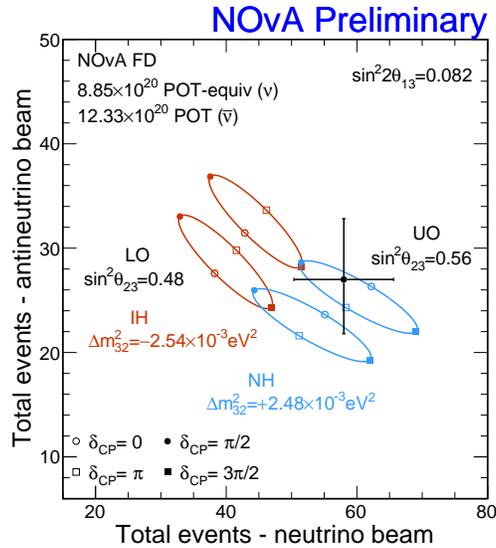


Figure 6.3: Bi-event plot showing predictions for total  $\nu_e$ ,  $\bar{\nu}_e$  candidates at various oscillation parameters given the exposure in the analysis, with data overlaid [78].

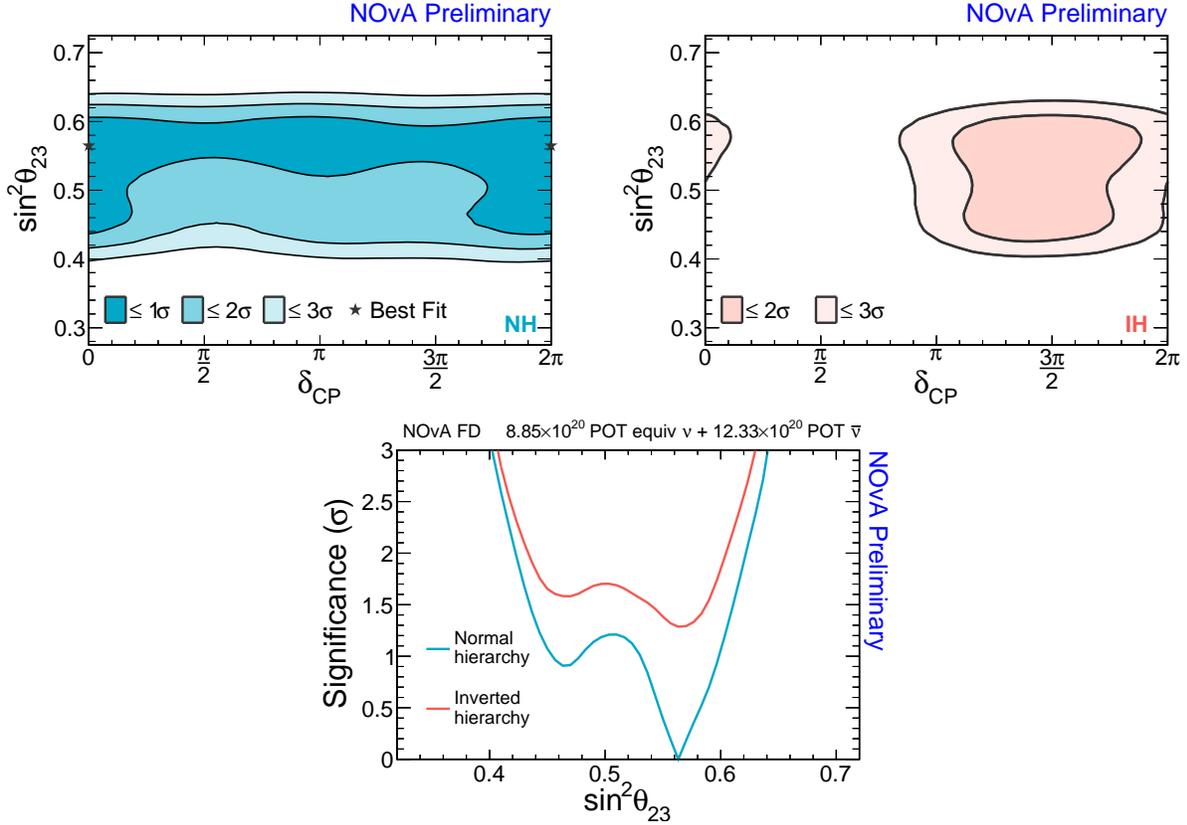


Figure 6.4: Results of the recent joint  $\nu_\mu, \bar{\nu}_\mu$  disappearance and  $\nu_e, \bar{\nu}_e$  appearance analysis in NOvA [79]. As the  $\sin^2 \theta_{23}$  vs.  $\delta_{CP}$  contours show, the current results do not constrain points in the normal hierarchy much (top left), but the points at  $\delta_{CP} = \pi/2$  in the inverted hierarchy are disfavored at over  $3\sigma$  (top right). The upper octant for  $\theta_{23}$  is slightly preferred (bottom).

of parameters, and the sensitivity to reject CP conservation for the specific case of  $\theta_{23} = \pi/4$  (maximal mixing) is given in the right of the figure. Depending on the true oscillation parameters, NOvA could reach  $2\sigma$  sensitivity to the mass hierarchy in the near future, and potentially even have  $2\sigma$  sensitivity to CP violation by the end of running. Similar discussion will continue in Chapter 9.

## 6.2 Kolmogorov-Smirnov style test of accumulated data

In Chapter 4, the accumulated exposure used for the analysis was discussed. The output of these studies are plots with POT and live-time in short bins of time over the course of NOvA data-taking. The sums appear in the exposure table there, and here the focus is on accumulated POT over time. The full-detector-equivalent versions of POT are used, as the accumulation of neutrino interactions in the FD is expected to follow this trend. For the neutrino beam, this corrects for the construction

	NH, UO	NH, LO	IH, UO	IH, LO
$\Delta m_{21}^2 / (10^{-3} \text{eV}^2)$	$+2.48^{+0.11}_{-0.06}$	+2.47	-2.54	-2.53
$\sin^2 \theta_{23}$	$0.56^{+0.04}_{-0.03}$	0.48	0.56	0.47
$\delta_{CP} / \pi$	$0.0^{+1.3}_{-0.4}$	1.9	1.5	1.4
Disfavored at	n/a	$1.6\sigma$	$1.8\sigma$	$2.0\sigma$

Table 6.1: Summary of oscillation parameters found with fits in each hierarchy (normal and inverted), octant (upper and lower) combination as in [36]. The best overall fit was found in the normal hierarchy upper octant, and the level at which each other fit was disfavored is given.

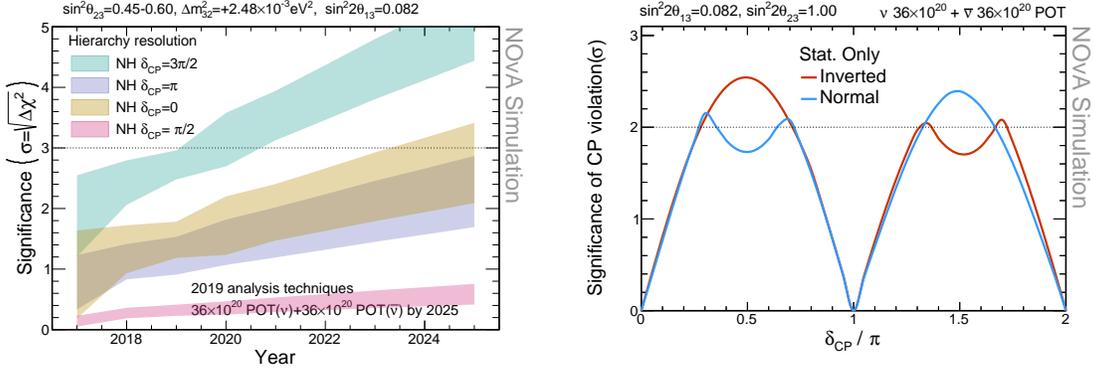


Figure 6.5: Plots of future sensitivity of NOvA analysis. (left) Reach in mass hierarchy as a function of time for many parameters [80]. (right) Sensitivity to reject CP conservation as a function of  $\delta_{CP}$  at  $36 \times 10^{20}$  POT in both beam modes [81].

period, while little change should be seen in the antineutrino beam running.

Since the neutrino accumulation rate should correspond to the “instantaneous” exposure of POT, plots of POT recorded in small time bins with respect to the whole time range represents a probability distribution function (PDF). An accumulated exposure plot is produced by replacing each bin’s contents with the total sum of POT in the given mode up to that time, a cumulative distribution function (CDF). When the analysis is performed, we save the event times at which each of the candidates were observed. These form an observed CDF.

The goodness of match between the CDF for each analysis sample and the corresponding (neutrino beam or antineutrino beam) exposure CDF is checked in a Kolmogorov-Smirnov style test. Normalizing both so that they reach 1 at the end of the analysis region, the maximum separation between a given analysis sample CDF and the exposure CDF characterizes this goodness. This separation is tested by conducting pseudo-experiments wherein the same number of events are chosen

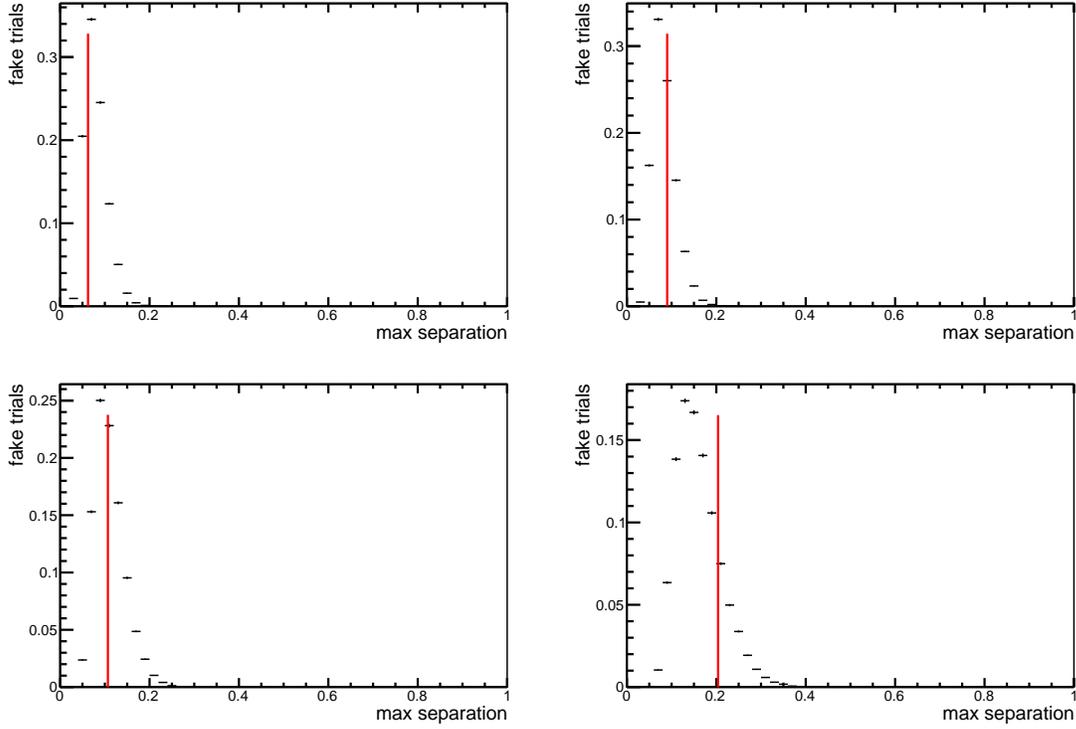


Figure 6.6: Integral-normalized distribution of the maximum separation between the accumulated POT and 10,000 pseudo-experiments with the events chosen randomly according to the exposure of the beam over time. The red lines are at the experimentally observed separation with the accumulated events over time in the  $\nu_\mu$  candidate sample (top left),  $\bar{\nu}_\mu$  candidate sample (top right),  $\nu_e$  candidate sample (bottom left),  $\bar{\nu}_e$  candidate sample (bottom right).

randomly from the exposure PDF, which we would expect our observation to follow. This randomized trial is then made into a CDF as before. Over many trials, I test how often the trial maximum separation is at least as large as our experimentally observed separation. One could plot an integral-normalized version of the maximum separation, where the probability of interest is the sum under the distribution of trials with maximum separation larger than the observation. This is shown in Figure 6.6. The fraction of performed trials with separation at least as large as the observation is equivalent to this, and is taken as the KS-style probability value. Figure 6.7 shows the accumulated exposures and selected  $\nu_e$ ,  $\bar{\nu}_e$ ,  $\nu_\mu$ , and  $\bar{\nu}_\mu$  events, along with the goodness score. The samples were found to be in reasonable enough agreement in both cases, such that the rate of accumulation of neutrino/antineutrino events is consistent with beam exposure.

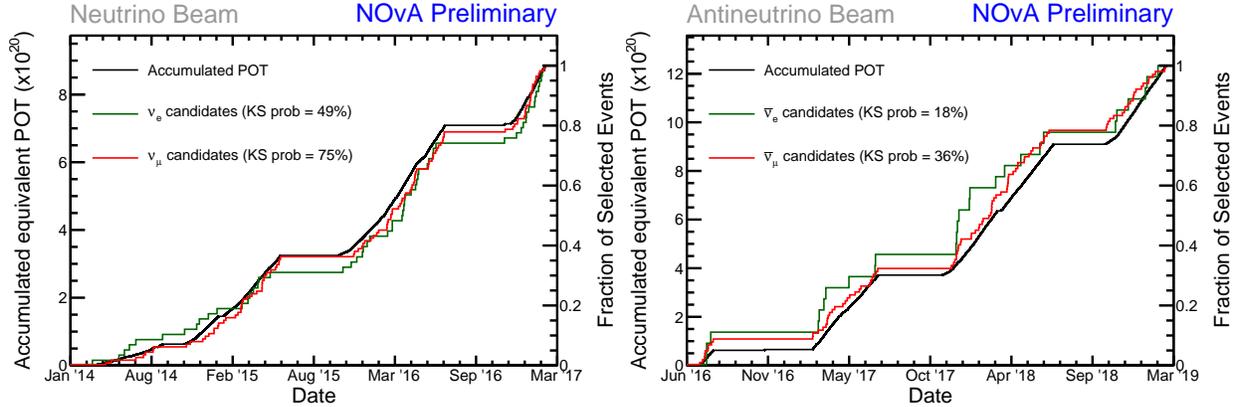


Figure 6.7: Accumulated exposure in the neutrino beam sample (left) and antineutrino beam sample (right) in terms of full-detector-equivalent POT and the  $\nu_\mu$ ,  $\nu_e$  and  $\bar{\nu}_\mu$ ,  $\bar{\nu}_e$  accumulated events (normalized to 1 at the end of the analysis time region). Probabilities as described in the text are given as a goodness of match for each sample.

### 6.3 Study of unexpected background

When the box was re-opened on the neutrino beam sample, one of the selected events seemed to be not a signal interaction. This event has no clear shower activity and does not look, by eye, to be consistent with an electron neutrino. Many hits are in a line in the x-z view with a kink near one end, and there are only a few hits, again with a kink, in the y-z view. An event display of this event is given in Figure 6.8. Interestingly, looking at event variables only raised more curiosity. This event was assigned a fairly high CVN electron score of 0.896, though it scored only 0.005 in the previous version of CVN, a large swing in scores from one version to the next. The timing of the event was somewhat after the detector had reached fully 14kt, but may have been during a period of instability. A study was undertaken to look in a little more depth at the background (cosmic) rejection portions of the selection and at the CVN scores themselves.

First, it should be noted that the transverse momentum fraction,  $p_t/p$ , for the event is 0.99. In Chapter 4, transverse momentum fraction was given as a cut variable, but 0.99 is near the theoretical maximum. The reason this event was not cut is that the check on  $p_t/p$  is performed as a function of position in the detector. For events near the top of the detector, the  $p_t/p$  cut is significantly more restrictive. For events in the centermost of the detector vertically, where these hits occur, any  $p_t/p$  is allowed. Making this cut more restrictive in maximum  $p_t/p$  allowed would reject this event.

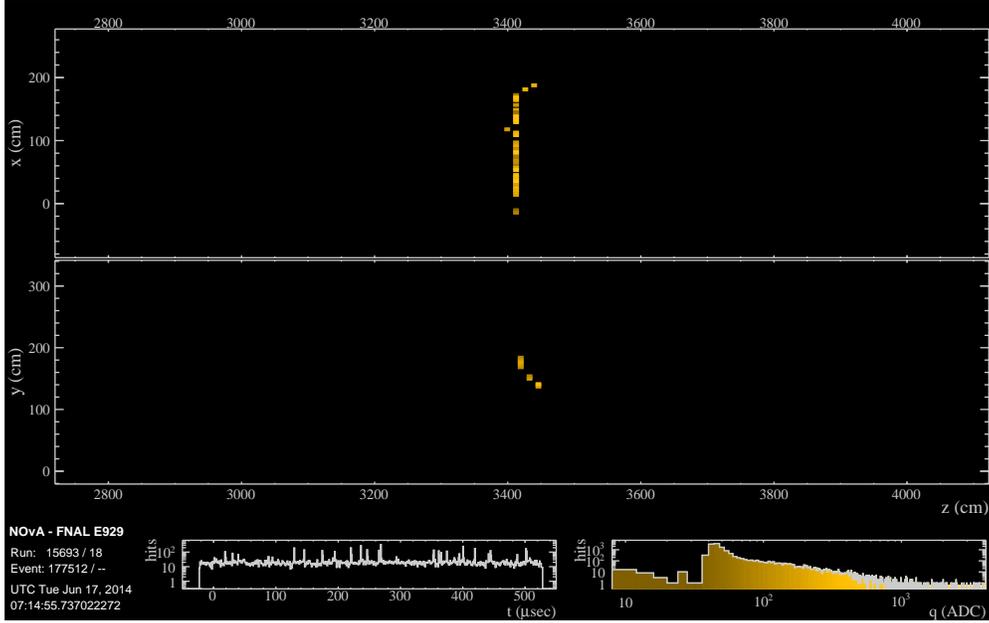


Figure 6.8: An event selected as a candidate  $\nu_e$  event from the neutrino beam sample [75]. This event has no clear shower activity and does not look, by eye, to be consistent with an electron neutrino. Many hits are in a line in the x-z view with a kink near one end, and there are only a few events, again with a kink, in the y-z view. This event launched a small study looking into selection cuts and event classification.

However, since this is an event in the blinded data sample, we wish to use another sample that we can study in more detail. The NOvA cosmic trigger at the FD is a 10Hz trigger separate from the beam trigger stream and is therefore used to study event characteristics for backgrounds which would be selected by analysis cuts. Over 120,000 seconds of live-time were studied to help understand this sample. Recall from Chapter 4 that the analysis had roughly 483.2s of live-time in the neutrino beam sample. It is the case that there is a portion of events selected with the current CVN that were not selected with the previous CVN due to changes in CVN score: in fact, a small pileup of events with nearly 0 CVN score in the previous instance of CVN shows up on the left of Figure 6.9. Secondly, these lowest scoring events (previous CVN score  $< 0.075$ ) were checked for some of the reconstruction variables. This may give some insight as to whether there is some particular population of events that were selected. Indeed, the right plot in Figure 6.9 shows that the lowest scoring events in the previous CVN are strongly concentrated near  $p_t/p$  close to 1.

Of the set of cosmic trigger events, 65 were found to be selected with current CVN score cuts

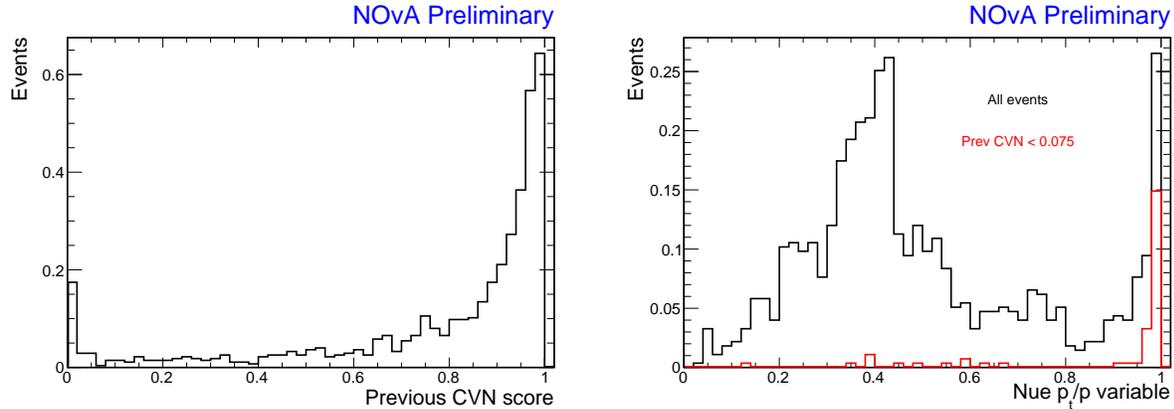


Figure 6.9: (left) A plot of the previous CVN scores for events selected with the scores and cuts in the analysis, using the cosmic trigger files. Note the pileup of events at the lowest previous CVN scores. (right) Comparing events with scores of  $<0.075$  in the previous CVN to the full set of selected events shows a striking difference in the  $p_t/p$  for these events relative to the rest.

but had scores in the previous CVN of  $< 0.075$ . In looking at the event displays and tabulated values, a number of reconstruction pathologies were noticed. For example, in some events the  $p_t/p$  was high but hits near the top of the detector were missed in the reconstruction forming prong clusters and not counted against the selection cut. Furthermore, in some of the selected cosmics the reconstruction whiffed in such a way as to make the  $p_t/p$  lower than it would seem to be from just looking at the distribution of hits. Finally, it is worth noting that the  $p_t/p$  cut only checks an event's proximity to the top of the detector. Applying similar cuts to the  $p_t/p$  cut but based on proximity to the lateral edges of the detector (x-view) would cut 51 of these 65 studied events.

While this does point to potential lessons learned, it does not say what is happening with the CVN values. There were some general differences in training noted from discussions with experts. For one, the previous CVN was trained using cosmics from the older, lower gain data, while the newer CVN used more recent, higher gain cosmic data. The event above is from the low gain period, but not all such events are low gain. Furthermore, an expert noted that the injection of cosmics to the training occurred near the end in the previous CVN, but it occurred throughout the training process in the newer CVN. Finally, preselection cuts were applied to the newer training: anything selected passes these, but maybe this makes it less likely for rare event patterns to be seen in training. It is possible that one or more of these causes this effect: hopefully such unwanted

backgrounds will be caught and/or removed in the future.

However, it is worth seeing what else we might learn. To study these events, a colleague ran a set of performance tests with CVN on approximately 50 of the events, called occlusion tests. In the occlusion tests, pixel maps are made from the event displays and calibrated energy depositions, as they would be to run CVN and compute scores. However, instead of computing the score normally, a square region of 5x5 pixels is blocked, or occluded, before the score is computed. The occluded region is moved around such that each point on the pixel maps is tested for occlusion. If an occluded feature causes large score changes, this region is likely important to the network in assigning scores. Once the colleague made the maps, many were matched with the corresponding event displays, and we could then look at these for features to assess what CVN may have found important. What we found was not one particular pathology but rather several. For one, in some events the pixel maps did not quite contain all of the hits. In particularly pathological events topologically, one could imagine that missing hits might contribute to obtaining an incorrect label.

Additionally, we found a number of recurring patterns causing large score changes in the occlusion tests. In some cases we saw a large change from the network occluding some set of hits in the x-z view. However, in other cases, the network output would change a lot by excluding hits near the end of a track. In yet other cases, the network would respond heavily to occluding high energy hits. In some cases, multiple of these conditions were met. Sometimes, the score would remain fairly stable no matter what was occluded.

Though the set of studied events consisted of a large variety of topologies and results, there were some events reminiscent of the one selected in the analysis. One such example event from this sample is shown in Figure 6.10. This event is also notable as it has hits at the side edge of the detector (-x edge). While the cuts select this background event, it would be removed by adding the x-based version of the  $p_t/p$  cuts noted above: alternatively, one could explore other cuts as well. This event has a CVN score of 0.89 in the recent iteration, but a score of  $2.6 \times 10^{-5}$  in the previous CVN. Pixel maps for this event are given in the top of Figure 6.11, with the occlusion tests results for the CVNe score in each view on the bottom. The y-z view is not particularly sensitive, but

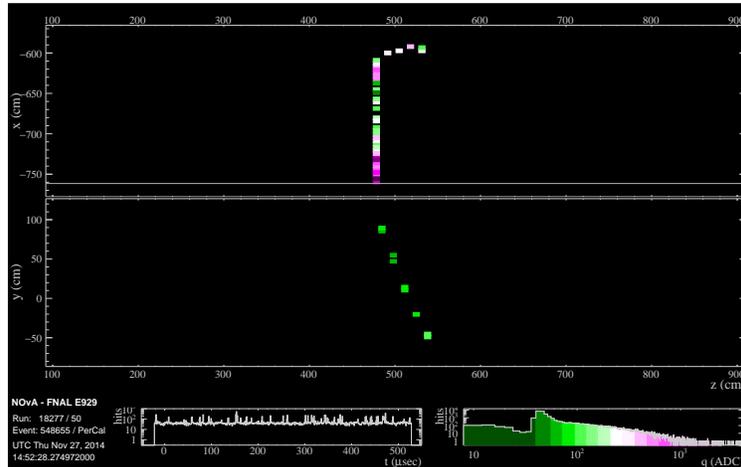


Figure 6.10: Event display for an example event picked from the cosmic sample with low score in the previous CVN but a high enough score in the recent CVN to be selected. Note the kink near one end of the track in the x-z view.

occluding hits in the x-z view near the kink causes very large ( $\sim 0.6$ ) changes in the CVNe score. The cosmic version of the score (not shown) reveals that much of what is lost in the  $\nu_e$  score is gained in the cosmic score. Potential alternative/additional cuts and other studies or cross-checks to run should be noted for future analysis, and improving/optimizing the use of and developing further understanding of CVN within the context of analysis groups remains ongoing work.

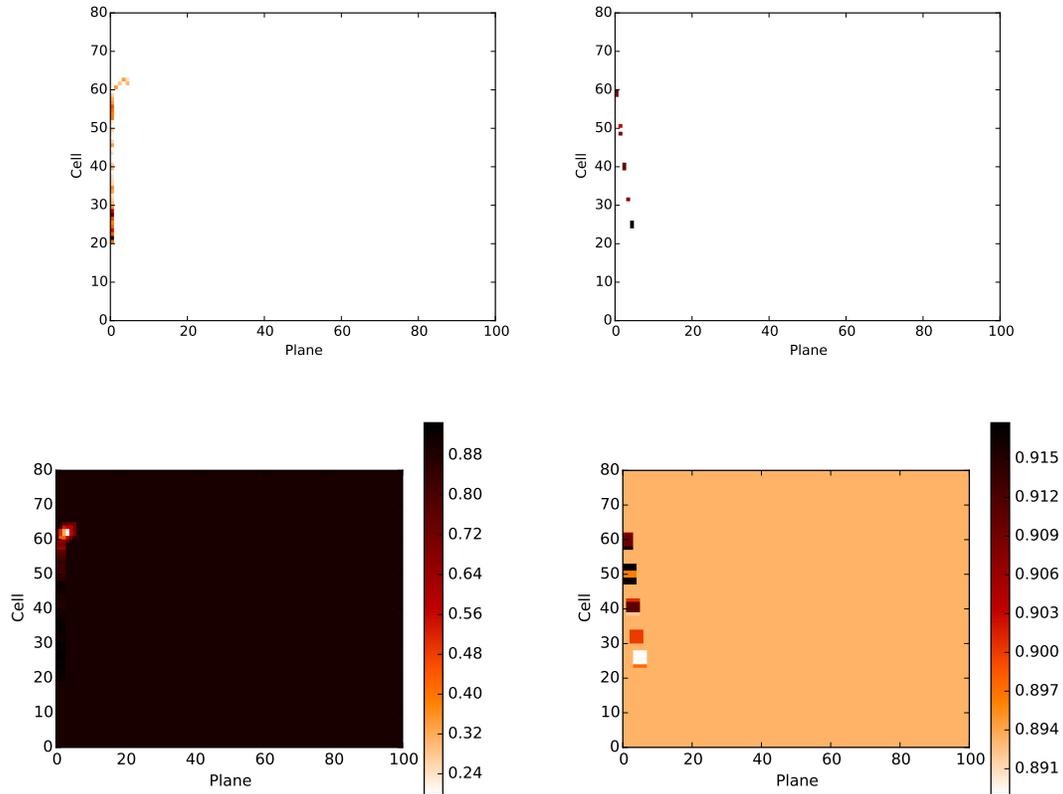


Figure 6.11: Pixel maps for the x-z view (top left) and y-z view (top right) for the event shown in Figure 6.10. The results to the network output from occluding sets of 5x5 pixels are shown for the x-z view (bottom left) and y-z view (bottom right). Occluding hits in the x-z view near the kink in the track causes very large ( $\sim 0.6$ ) changes in the current CVN electron score.

## CHAPTER 7

### The Deep Underground Neutrino Experiment

The Deep Underground Neutrino Experiment (DUNE) is an upcoming long-baseline neutrino oscillation experiment, which will use a new and improved accelerator-driven neutrino beam from Fermilab. The experiment will place a Near Detector (ND) at Fermilab to measure the spectrum, provide precision measurements of neutrino cross-sections, and perform other physics studies. The Far Detector (FD) will be located nearly a mile underground (4850 feet) in Lead, South Dakota, 1300km from Fermilab. Its depth corresponds to 4300 meters water equivalent (mwe), which means it has even more shielding from the Earth than Super-Kamiokande (about 2700mwe) [82]. Such depth of the detector will limit cosmic backgrounds, which is helpful in two ways. For one, it greatly diminishes the probability that a cosmic ray overlaps a neutrino event in the FD or that a cosmic ray is selected as a potential neutrino candidate. Such backgrounds are concerns with surface detectors like NOvA that require extra treatment and consideration as a background. Further, as will become clear in the discussion of the detector technology in Section 7.1, the DUNE detectors operate by detecting ionization signals under an electric field. The ions remaining from the ionization can cause adverse effects across the detector such as perturbing the electric field; this phenomenon is known as *space charge*. By removing much of a potentially large background such as cosmic rays, this will hopefully help mitigate space charge effects.

When completed, DUNE will measure 40kt in fiducial mass – a factor of almost three improvement over the entire mass of NOvA. DUNE will be comprised of four individual  $\sim 10$ kt fiducial mass liquid argon (LAr) time-projection chambers (TPCs). This technology will be discussed in Section 7.1, along with the two main designs that are expected to make up the four modules. Such a large, sensitive detector will be able to answer open questions in neutrino oscillation physics (with strong discovery potential for CP violation), will make precision measurements of open neutrino oscillation parameters, and opens the floor to a suite of complementary and different measurements,

as is discussed in Section 7.2. The beam servicing DUNE will be discussed in this section as well.

## 7.1 Detector technology

The four FD modules will use LArTPC technology, which is growing in popularity in neutrino physics and significantly push forward the capabilities of detector technology. Being argon-40, it provides many more nucleons as targets for neutrino interactions than carbon-12 or other materials which make up the bulk of other neutrino detectors. Further, Ar is a highly sensitive medium: traversing charged particles ionize Ar, and the medium also scintillates. Applying an electric field pulls the freed electrons away from the positively charged argon ion. In a LArTPC, this field then drifts the electrons to wire planes to be read out and provide measurements of particle tracking and energy reconstruction. The scintillation light helps to constrain the position of the particle in the direction of the drift and provides a timestamp for non-beam events. The electrons drift at a speed determined by the electric field; for DUNE the drift occurs over a few ms. However, the photons will travel at the speed of light in LAr and can therefore act as a timestamp. Since the entire region contained in the electric field is essentially active and the readout planes are typically fine-grained, one is able to obtain highly sensitive measurements of both particle positions and energy.

There are two styles of LArTPC in the discussion for DUNE: each has a description in this section. The single-phase module is a more traditional style detector where a cathode and anode plane are both submerged in LAr, with the field oriented horizontally. The electrons drift to the anode plane, where sets of wires read out induced and collected signals. The second style, the dual-phase module, is a newer iteration of the LArTPC concept, in which the cathode is at the bottom of the detector and the anode at the top. While less proven, this idea has enticing up-sides, such as a larger fraction of active area and possibly higher signal/noise ratios. Cartoon-like examples of these designs are shown in Figure 7.1 [83, 84]. The dual-phase module, on the right, shows an illustration of the detection technique as well.

The single-phase module for DUNE is shown in the representation in Figure 7.2 and to achieve the 10kt fiducial volume for one module, 17.5kt of liquid argon is required. The module is well described in [83]. The full module size is 12m high and 58m long. However, each module is itself

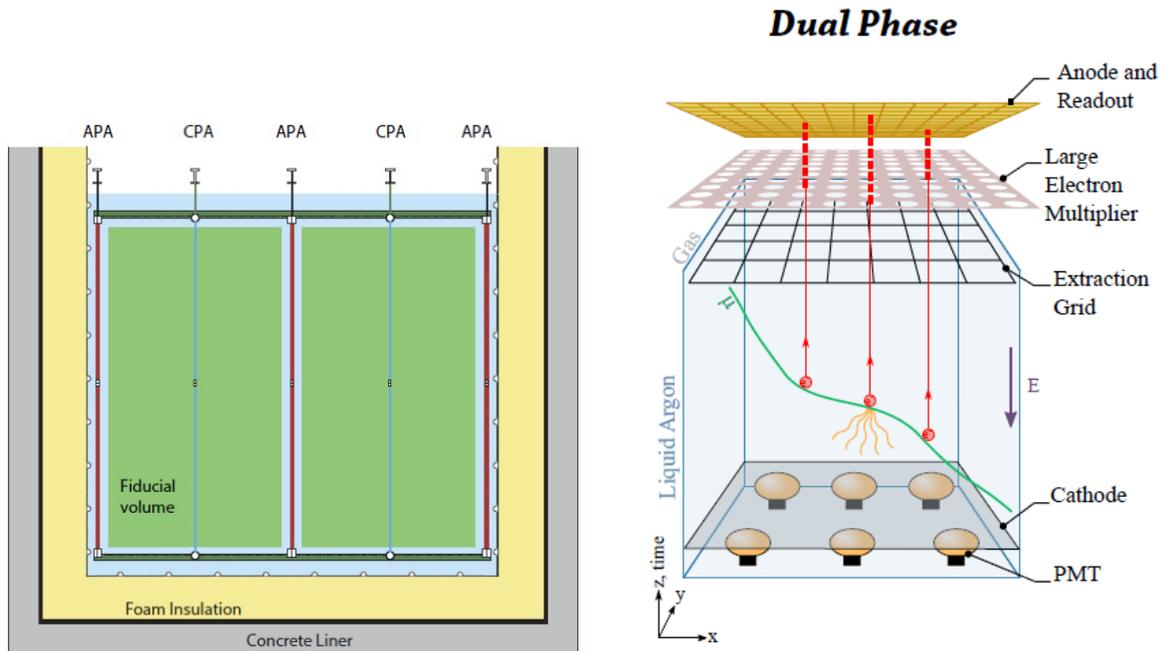


Figure 7.1: Cartoons showing depictions of a LArTPC modules (single-phase, left [83] and dual-phase, right [84]). The dual-phase cartoon further shows the concept of a traversing charged particle, with ionization and scintillation signal.

made up of several TPC volumes (2 high, 4 across, and 25 long). Since each TPC requires an anode and a cathode, clever placing of these (e.g. anode-cathode-anode) allows one to have TPCs back to back to reduce the dead space and additional material which must be added to the LAr volume. The anode for each TPC volume is roughly 2m wide and 6m tall, and the cathodes are spaced 3.53m away, such that the maximum drift distance for ionization is 3.53m. The nominal electric field will be 500V/cm, and in this field, electrons drift at approximately 1.5m/ms [83]. This means that the electron lifetime in LAr (which depends on the purity of the LAr) should be several ms to not lose signal from the far end.

The readout wire planes for the single-phase DUNE FD are depicted in Figure 7.3, with the whole thing rotated 90 degrees. Missing from the figure is a mesh of wire (in a crosshatch pattern) which serves as a grounding plane. A vertical plane of wires provides shielding such that only electrons drifting past the readout induce signals. Two planes of wires angled at  $\pm 35.7$  degrees are biased such that ionization electrons drift past these wires but induce signals to be read out. Finally,

they reach a vertical collection plane where the electrons are actually collected and read out. To achieve this effect, carefully selected bias voltages are chosen to shape the field lines to guide past the induction planes and end on the collection plane wires [83]

1. Grounding mesh: 0V
2. Grid plane: -665V
3. Induction U: -370V
4. Induction V: 0V
5. Collection Plane: 820V

Combined, these three readout planes provide the location of the ionization and particle tracking. The precision to measure the third dimension (distance of the particle from the wire planes in the drift direction) is enhanced by photon detectors helping to provide a timestamp  $t_0$  for the interaction. Since the amount of ionization signal correlates to the energy, energy can be measured precisely as well.

As noted above, the dual-phase reads the ionization electrons by extracting them vertically into a gas phase, and the photon detectors for this LArTPC design are placed at the bottom below the cathode. This means that even more of the LAr volume is in fact fiducial volume: for a 17.5kt module as in the single-phase design, approximately 12.1kt is fiducial in this design [84] as compared to the 10kt in a single-phase module. The dual-phase LArTPC design for DUNE would have 12m maximum drift height and is 60m long. Note that this maximum drift distance is 2-3 times the drift distance of a single-phase module, which places increased demands on the purity of the LAr to provide adequate time for the electrons to drift. As in the single-phase module, the nominal drift field will be 500V/cm in the LAr volume. However, in the dual-phase module, electrons must be extracted from the liquid into the gas phase. At this boundary, a 2kV/cm (four times the strength) electric field is applied [84].

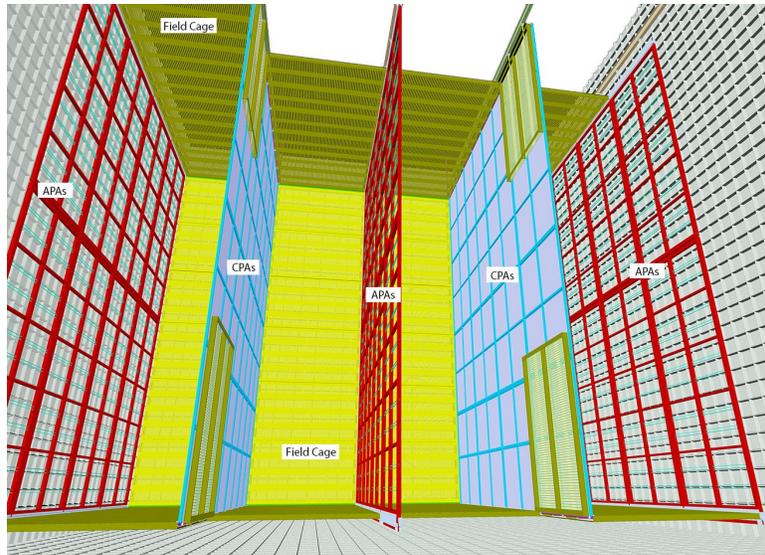


Figure 7.2: Cartoon showing a depiction of a single-phase LArTPC module [83]. The electric field goes from the cathode plane assemblies (CPAs) to the anode plane assemblies (APAs). Planes of wires at the APAs detect the ionization signal, and photon detectors in the APAs, as represented by the somewhat perceivable blueish rectangles inside the APAs, detect the scintillation signal. A single cathode plane services two sets of TPCs.

The extraction field leads electrons to components in the gas phase known as Large Electron Multipliers (LEMs). The LEMs are PCB (Printed Circuit Boards) that carry high voltages to cause electron avalanches when ionization electrons pass, creating an analog gain to the signals from the TPC. Each LEM is  $50 \times 50 \text{ cm}^2$ , with 86% active area covered with  $180 \text{ holes/cm}^2$ , each  $500 \mu\text{m}$  in diameter [84]. The field in these holes will reach up to  $35 \text{ kV/cm}$  for the  $1 \text{ mm}$  thickness, giving a signal amplification gain factor of over 20 [84]. Though the amplification process will create UV photons, the hole structure of the LEM quenches these, allowing for pure gaseous Ar to be used instead of adding a dopant [84] (maintaining LAr purity is of great concern). Signals passing the LEM are then fed to a board with gold-plated copper strips in 2D arrays, which provide the tracking and energy measurements in the dual-phase design (as opposed to wires in the single-phase design). By providing amplification in the gas signal after requiring the ionization electrons to travel some distance, this method is expected to have the ability to enhance the signal/background ratios over the levels the single-phase is capable of achieving.

The protoDUNE test experiments at CERN have been and will be testing these designs and

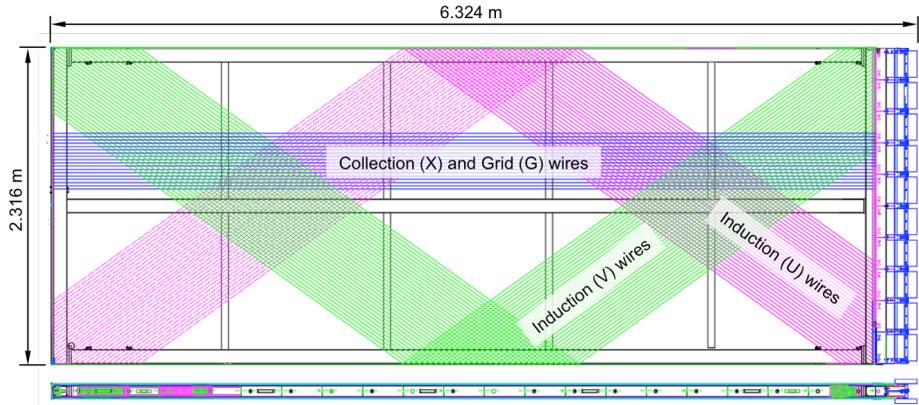


Figure 7.3: Cartoon showing a depiction of an anode plane with wire planes wrapped around it [83]. Note that in this illustration the anode is rotated 90 degrees (it is 6.324m tall and 2.316m wide). The angled U and V planes record induction signals as the electrons drift past, and the collection planes running vertically directly collect the electrons. This gives 3 planes of charge readout. Though not pictured, the photon detectors would run up and down in this figure (they are horizontal when the frame is oriented as it is in the detector.)

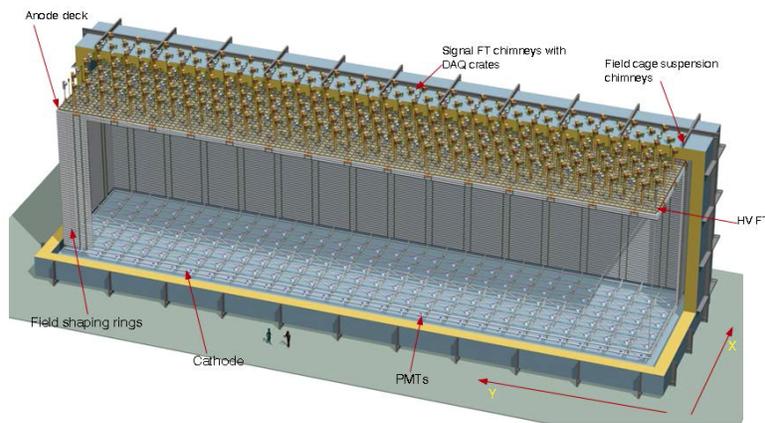


Figure 7.4: Cartoon showing a depiction of a potential dual-phase LArTPC module for DUNE [84]. The electric field goes from the cathode plane near the bottom to the anode plane in a gaseous Argon phase at the top of the detector. Photomultiplier tubes placed at the bottom of the detector measure the scintillation light.

studying the full chain of data-taking, processing, and analysis that will be required of the final detector. While these are massive undertakings - the single-phase protoDUNE is currently the largest LArTPC - they are only prototypes at the few percent level scale. Assuming both detector types test well in the protoDUNE program, it is expected that the four modules of the DUNE FD will contain at least one module of both designs.

### *7.1.1 Note on the Near Detector*

At the time of writing, the final design for the DUNE ND has not yet been determined. Philosophically, it will act much as the NOvA ND does, using the high rate of neutrino interactions close to the beam's origin to study the beam composition before oscillation and conduct physics studies related to neutrino interactions, potential sterile neutrinos, etc.

Being able to study selected spectra of neutrino interactions in a LAr medium is what will provide cancellation of errors via Far/Near ratios, as in NOvA's extrapolation process. Therefore, it is likely that a component of the DUNE Near ND will be a LArTPC. However, to properly deal with the pileup which will occur at DUNE's beam intensity, a pixelated readout for the ND is being investigated to replace the more traditional wire planes of a single-phase LArTPC.

Along with the LArTPC ND, extra calorimeter and potential gaseous argon TPC elements are being explored for the benefits they may add, including potential magnetization to sign select. Such a ND would greatly improve DUNE's understanding of the wrong-sign component over the sorts of methods explored above for NOvA (see Chapter 5). Additionally, a potential use of wrong-sign selection in the FD is explored in Chapter 9. To use this method, an understanding of the wrong-sign selection in the ND beam samples to predict signal and beam  $\nu_e$  backgrounds at the FD is expected to be beneficial.

## **7.2 Physics of DUNE**

### *7.2.1 Oscillation physics*

The primary DUNE oscillation measurements will come from a beam of neutrinos originating at Fermilab, with the upgraded beam from the PIP-II project (see Section 2.2.1). The baseline for this beam to Lead, South Dakota is 1300km, which is more than 1.5 times the baseline servicing

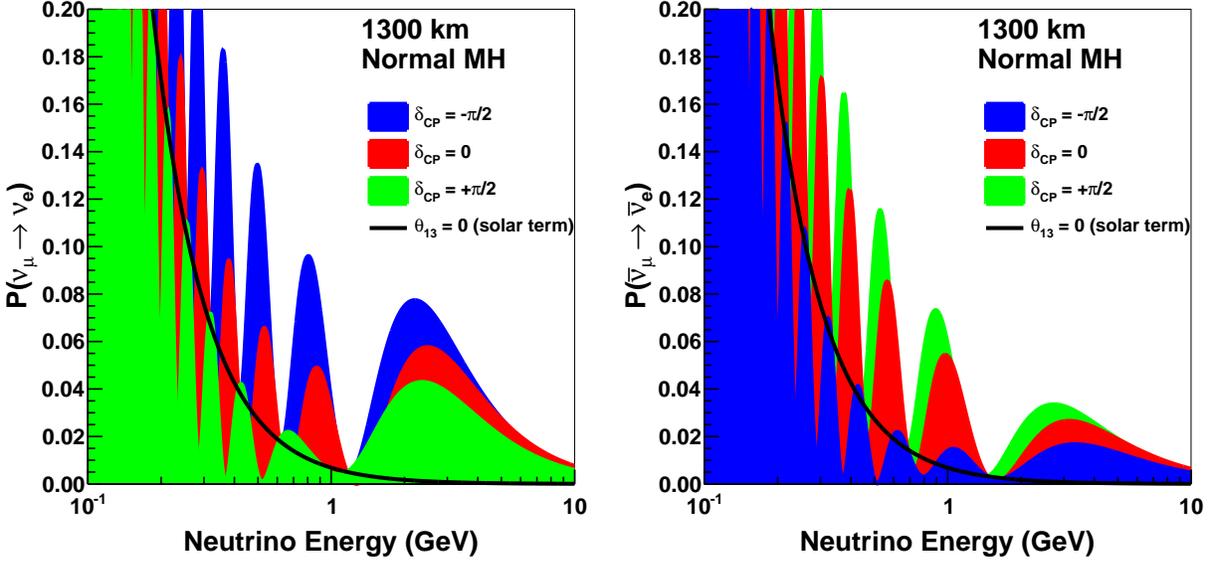


Figure 7.5: Probability for  $\nu_\mu, \bar{\nu}_\mu$  neutrinos to oscillate to  $\nu_e$  (left),  $\bar{\nu}_e$  (right) as a function of neutrino energy at the DUNE baseline (1300km), under varying assumptions for neutrino parameters [25].

NOvA. Correspondingly, the first oscillation maximum is also higher than that in NOvA. Scenarios for the oscillation maxima for  $P(\nu_\mu) \rightarrow P(\nu_e)$  and  $P(\bar{\nu}_\mu) \rightarrow P(\bar{\nu}_e)$  are shown in Figure 7.5 [25]. As DUNE is an on-axis experiment, it will have a broad-band beam (recall Figure 3.6), and the new beam from Fermilab will provide strong neutrino fluxes at the relevant energies. Figure 7.6 shows the flux intended for the beam [85].

With this new beam giving increased flux, the longer baseline providing increased disambiguation between mass hierarchy and CP effects, the highly sensitive medium and detector technology discussed above, and advances to data analysis techniques, DUNE will provide more conclusive evidence and precise measurements of neutrino oscillation parameters than is possible with the current generation. Additionally, rather than solely provide high sensitivity to one or two parameters, DUNE will enable precision to multiple parameters (e.g.  $\theta_{23}, \theta_{13}, \delta_{CP}$ ) at the same time. This benefits the idea of helping to create a more global picture of neutrino oscillations. For the most favorable oscillation parameters, DUNE will reach  $5\sigma$  sensitivity to the mass hierarchy in about a year of running, and even for the least favorable parameters, DUNE should reach this sensitivity in well under 10 years [86]. Furthermore, DUNE has a very high discovery potential for CP violation

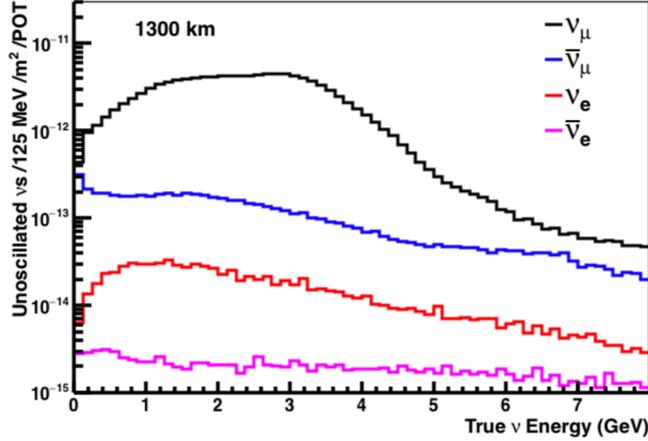


Figure 7.6: Proposed flux spectrum for the beam which will service DUNE [85]. Note the strong overlap of the neutrino flux with the oscillation probability maximum between approximately 2-3 GeV in Figure 7.5.

in the neutrino sector: it will reach  $5\sigma$  sensitivity to CP violation for 50% of possible values in about 10 years [86]. These are summarized in Figure 7.7. Finally, the sensitivity to  $\theta_{13}$  will not be especially useful outside informing the global picture at first. However, the measurement of  $\theta_{13}$  has strong potential to match the precision currently held in reactor neutrino experiments [85], which will make it an important measurement. This is depicted in Figure 7.8.

### 7.2.2 Other physics

In addition to its usefulness in studying neutrinos in the beam for oscillation properties, the large and sensitive medium of the DUNE FD will provide a test bed for non-beam physics studies as well. Three major components to this program will be supernova neutrino detection, proton decay searches, and atmospheric neutrinos.

While the largest signal from a supernova in a water Čerenkov detector comes from inverse beta decay (that is  $\bar{\nu}_e + p \rightarrow n + e^+$ ), for LAr detectors, the  $\nu_e$  CC spectrum is prominent, making the measurements that can be done with an LArTPC complementary to those that can be done with current and next generation water Čerenkov detectors (see the cross-sections for LAr in Figure 7.9) [25]. Neutrinos from supernovae are predominantly 5-50 MeV, and the DUNE FD should expect some 3000 events from a supernova at 10 kpc [86]. From these events, one can learn both about the nature of neutrinos and of the astrophysical processes going on during the supernova. The latter is

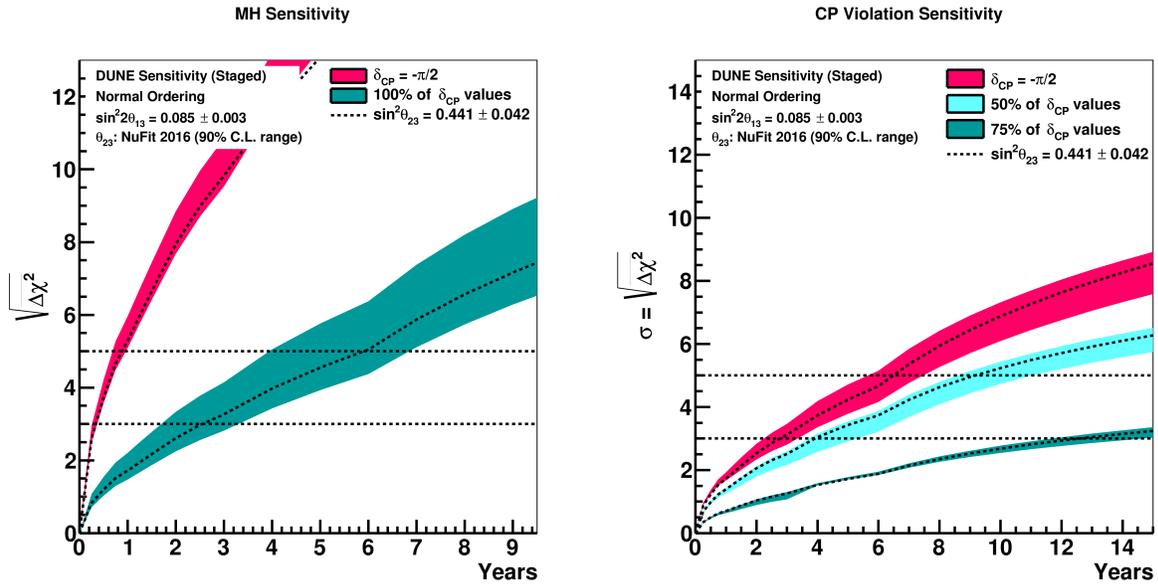


Figure 7.7: Sensitivity to mass hierarchy (left) and CP violation (right) in DUNE [86]. Note that the most favorable parameters quickly reach strong significances, and major portions of available parameters are covered on the timescale of about a decade (for mass hierarchy the entire space, for CP violation about 50%).

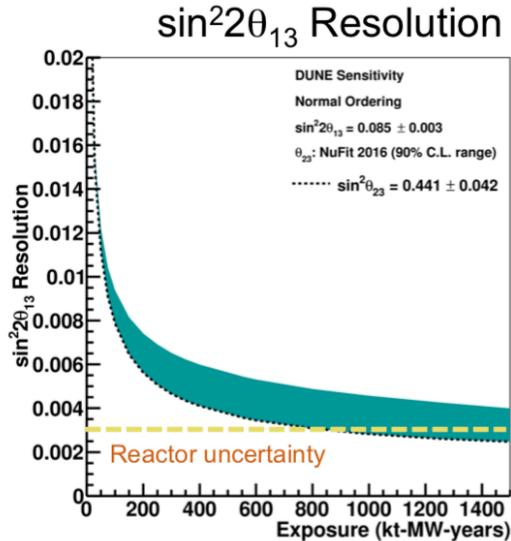


Figure 7.8: Precision of DUNE's ability to measure the  $\theta_{13}$  oscillation parameter [85]. Note that the measurement can reach the current level of precision in reactor neutrino experiments, which will be an important achievement of its own for the accelerator-based neutrino program.

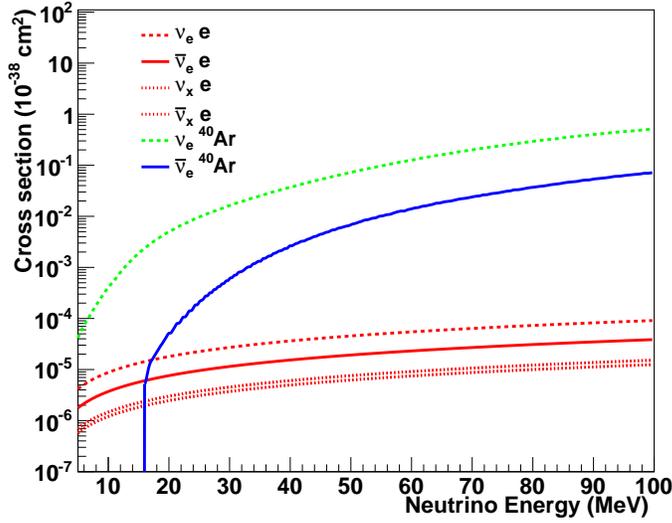


Figure 7.9: Cross-sections in LAr [25]: the CC interaction cross-section (green dashed line) in DUNE is significant for supernova  $\nu_e$  between 5-50MeV.

especially important, since for example one may find properties of shockwave effects, turbulence effects, and/or properties related to neutrino-neutrino interactions in the observable spectra; additionally, the formation of a black hole from the supernova core would be signaled by an abrupt end to events in the time-series of events [86].

A second important non-beam physics study that will be enabled by the DUNE FD is a search for nucleon decay. While water Čerenkov detectors set the current limits in many modes and are quite good at pionic modes of potential proton decay, LArTPCs are expected to provide complementary searches. LArTPCs such as DUNE are expected to be especially competitive in kaonic decay modes. The charged kaon decay mode with the highest branching ratio is  $K^+ \rightarrow \mu^+ + \nu_\mu$  ( $\sim 64\%$  of decays) [4]. Therefore this signature of a proton decay is fairly iconic in a detector with tracking capability such as DUNE: a track identifiable as a kaon (either from  $dE/dx$  information or a neural network such as a convolutional neural network), followed by a positively charged muon track with a decay electron at the end of its track. The second most common K decay mode is  $K^+ \rightarrow \pi^+ + \pi^0$  ( $\sim 21\%$ ) [4]: the  $\pi^+$  will produce a track while the  $\pi^0$  will produce showers from two photons. This is a bit of a glossy picture in that there are backgrounds that need to be well understood from, say,

atmospheric neutrinos. Additional complications arise from final state interactions in the initial kaon production which may cause it to be hard to identify. However, it is hoped that DUNE will provide competitive sensitivity to proton decay: this is an active area of research.

Finally, a third important component to the non-beam FD analysis are atmospheric neutrino studies. These are expected to be viable due to the low background environment and large detector volume, allowing for the interactions of neutrinos coming from atmospheric sources to be used to probe neutrino oscillation at long-baselines as well (akin to the measurements performed by Super-Kamiokande as discussed in Chapter 2). In this sense, the measurements should be seen as complementary to the beam-based neutrino oscillation measurements in that it creates a higher statistics sample for oscillation measurements than from the beam alone, with the same detector systematics but different spectra and acceptances.

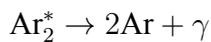
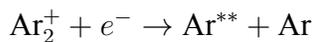
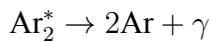
One thread common to all non-beam physics studies is that there is no inherent event time expectation. For neutrinos coming from the beam, the accelerator trigger and the distance from NuMI to Lead, South Dakota provides a fairly reasonable expectation for the time at which beam neutrinos will cross the DUNE detectors. However, no such expectation exists for these non-beam studies. Therefore, effective triggering and event location in the detectors are of critical importance. This is one important use for the scintillation light, which quickly propagates through the LAr. Therefore, detecting scintillation photons from events occurring from non-beam physics events in the LAr is useful to constrain the position of particles in the detector, and perhaps can serve as a trigger on potentially interesting events. The mechanism for scintillation of LAr and basic detection strategy and complicating factors are presented in the next section. Development and study of photon detection, with DUNE in mind, was a major aspect of my research program and is described in Chapter 8 and related appendices.

### **7.3 Liquid Argon Scintillation and Detection**

When charged particles pass through the LAr medium in the TPC, it both ionizes the LAr and induces scintillation. While in DUNE the primary signal is the ionization, the copious scintillation provides important information about the position in the drift direction and can potentially serve as

a trigger or provide extra resolution on observables.

The propagation of charged particles in the LAr medium induces excited molecular dimers of Ar atoms called Rydberg states, which in this case is a bond state containing a  $\text{Ar}_2^+$  with a shared electron [87]. A description of Rydberg states for  $\text{Xe}_2$  is given in Reference [88]. The excited dimer state then decays to the ground state (single Ar atoms), giving off 128nm light in the transition: this is the LAr scintillation. This excited state can be created directly by propagation of the charged particle (method 1 below) or be created in the recombination of ions and electrons after the ionization process has taken place (method 2 below). Adapted from the description provided in literature [89], the two mechanisms are described via



where the “heat” in method 2 is to indicate that that particular step is non-radiative and does not result in a photon.

The scintillation light released from LAr is often stated as 128nm, though values between 127nm and 129nm are sometimes given, with widths of  $\sim 8\text{-}10\text{nm}$  [89, 90]. The excited dimers occupy two states, the singlet and triplet, both of which decay to produce  $\sim 128\text{nm}$  light, but which decay on different timescales [91]. Most experiments agree that the singlet lifetime is short (something like 5-10ns) and that the triplet is much longer lived, of order  $1.5\mu\text{s}$ . Our group has measured the triplet state from scintillation induced by cosmic-ray muons to be  $1.52\mu\text{s}$  [92]. The literature covers a wider range of values for this state, from  $1.21\mu\text{s}$  to  $1.66\mu\text{s}$  (for a non-exhaustive list, see for example: [90, 93, 94, 95, 96]). Perhaps some of this is related to contaminants present in the

LAr, which will be discussed shortly and has effects on the scintillation timescales. However, it is not known that this is the cause for the variation, and the types of particles inducing the scintillation also varied from experiment to experiment.

Different ionizing particles have been found to populate these states in different amounts. This can be seen in the “fast/slow” ratios as presented in Reference [96], where the fast is mostly singlet and the triplet is expected to have the largest contribution to the slow component. In this work, gammas have a fast/slow ratio of 0.30, while alphas and neutrons have ratios of 2.6 and 3.5, respectively. Some experiments instead report the fraction of prompt, or early, light. For example, an ICARUS-like test stand reported values giving an indirect  $f_{\text{prompt}}$  measurement of  $0.24 \pm 0.08$  for cosmic-ray muons [97]. Work by the Indiana University group on measuring the timescales of scintillation using cosmic-ray muons called  $f_{\text{prompt}}$  the fraction of light collected in 133ns relative a  $10\mu\text{s}$  readout window, and found a consistent  $0.32 \pm 0.03$  [92].

The scintillation yield in LAr is quite copious. The exact yield depends on density of energy transferred to the medium (characterized as linear energy transfer, LET, essentially  $\Delta E/\Delta x$ ) [89, 98]. For particles with LET near the minimum ionizing  $dE/dx$ , the yield is  $\sim 40,000$  photons/MeV [89, 99]. The maximum yield is expected to be closer to 51,000 photons/MeV [98]. The works of Doke and colleagues to put together much information related to scintillation yield [89] gives potential explanations for why such variations in yield may occur. In the lower energy transfer region, the idea is that electrons can escape the influence of the ion and therefore not recombine till much later times, and thus the scintillation is lost from the event itself. For the higher energy density deposits, the idea is that the excited states quench themselves in a process such as  $(\text{Ar}^* + \text{Ar}^* \rightarrow \text{Ar} + \text{Ar}^+ + e^-)$  where the electrons carry away the excess energy and no scintillation is produced. A yield of 40,000 photons/MeV produces a large amount of photons for modest energy deposits.

Since one of the main mechanisms leading to scintillation light is recombination of ion states with electrons, the electric field drifting the ionization signals to the readout necessarily reduces the scintillation signal. It pulls the electrons away from the ions, so the amount of reduction in the scintillation yield is directly related to the strength of the applied field. Articles such as [89, 98]

discuss the relation of charge yield and scintillation yields, and this is summarized in plots of these versus electric field pointing out where DUNE sits at 500V/cm [100]: a reduction to  $\sim 60\%$  yield is expected. This reduces the expected light yield from  $\sim 40,000$  photons/MeV to closer to 24,000 photons/MeV, but this is still an appreciable amount of photons for moderate signals.

This scintillation yield is further affected by contaminants in the LAr volume, via quenching of scintillation light and alteration of the lifetimes. Common such contaminants affecting scintillation include nitrogen, oxygen, and water. The effects of oxygen and nitrogen on the scintillation have been studied in LAr. For concentrations less than 100ppb (atomic) of oxygen the effects of quenching are fairly negligible; for concentrations of 1ppm, there is  $\sim 38\%$  loss [94]. Nitrogen is more forgiving: concentrations up to a couple hundred ppb (atomic) of nitrogen will have little impact. The surviving fraction is still  $\sim 80\%$  at 1ppm. While the effects of water contamination have been less studied in LAr, more recent work on gaseous argon suggests that this is a more sensitive contaminant [101]. In fact, when attempting to account for the absorption lengths of contaminants in a prototype test, we found that water contamination contributes the most to the absorption length of the three, even at the lowest concentrations (see Chapter 8). However, at the level of a few ppb, it is not expected to significantly quench the scintillation yield [101].

For DUNE, filtering the LAr to reduce contaminants will take care of water and oxygen, removing these below the level of concern for photon detection by necessity (the TPC requires more stringent control of these contaminants). The filtering will not remove nitrogen, so the LAr purchased should ideally have low enough nitrogen to not produce significant effects on the scintillation.

Photon detection in LAr, and the design and testing of prototypes, is an interesting field of study due to several challenges:

1. 128nm scintillation is in the Vacuum UV (VUV) and is therefore hard to detect, especially with conventional photon detection devices. Typically, wavelength shifters are needed to couple practical photon detectors with such a system.
2. The system cannot interfere with the operations of the TPC and should avoid taking up fiducial volume since the LArTPC is the primary detection element.

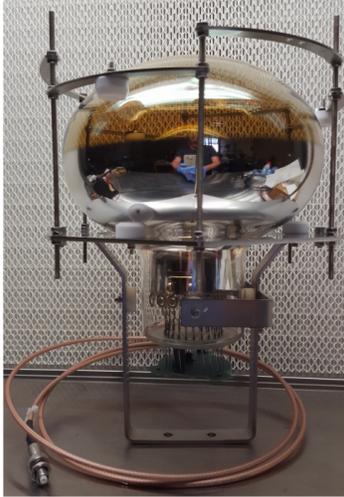


Figure 7.10: Example of a prototype PMT that would be coated with a wavelength shifter like tetraphenyl-butadiene (TPB) and placed in a dual-phase module to detect scintillation photons [84].

3. To fully take advantage of the system and achieve the proper resolution in the drift direction, timing of the system needs to be  $\leq 1\mu\text{s}$ .

For the dual-phase design, which is not broken into many back-to-back TPC volumes as in the single-phase, the path forward is much clearer. It will use wavelength-shifter coated photomultiplier tubes (PMTs) placed below the cathode plane at the bottom of the cryostat, in the LAr. Such a system is depicted in Figure 7.10 and is a more traditional photon detection system for a LArTPC.

Since the TPC volumes are placed back to back in the single-phase design and there is a field cage surrounding the TPC volumes to guide the electric field, it is not feasible to put PMTs on the top, bottom, or sides of the cryostat to detect light. To mitigate this situation and not take up active volume, it is desirable to place the photon detectors *inside* the APA frames, between the wire planes for each TPC volume. That is, a photon detector can be located between the collection wires of back-to-back TPC volumes and be sensitive to both volumes. The natural solution here is a system which uses wavelength shifters and a long, thin collection and transport piece to carry photons to a photosensor. Two such designs are the so-called lightguide bar designs, as pictured on the top row of Figure 7.11. The blue colored cartoon is a bar that is dipped (“dipped bar”) in a wavelength shifter like tetraphenyl-butadiene (TPB) to convert the scintillation light from VUV to

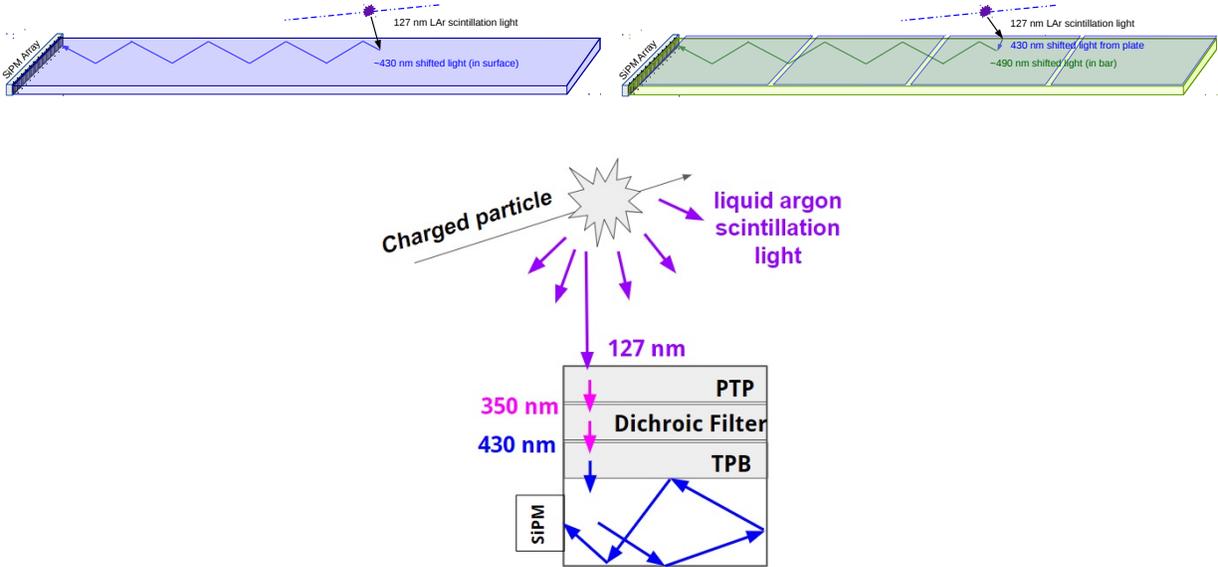


Figure 7.11: Examples of prototype photon detection systems for a single-phase LArTPC module for DUNE [83]. These work on the principle of converting the signals from VUV to higher wavelengths and then trapping or guiding converted photons to a silicon photomultiplier (SiPM) readout. One such model is a bar dipped in a wavelength shifter (top left), one requires two shifts but decouples the conversion and transportation (top right), and one uses dichroic filters to trap converted light in a box-like structure with SiPMs (bottom).

near-visible ( $\sim 430\text{nm}$ ). The light then travels via total internal reflection in the bar to an array of silicon photomultipliers (SiPM) at the end. A second variant of this type of design (“double-shift”) uses TPB-coated plates in front of and behind a bar produced commercially by Eljen Technology, doped with a second wavelength shifter as depicted in the top right of Figure 7.11. Light produced on the plates which make it to the bar can then be shifted again to  $\sim 490\text{nm}$  and then also be trapped by total internal reflection as in the dipped bar. While this requires two shifts, it separates the conversion from the transport, and allows one to optimize the initial conversion without affecting characteristics of the bar. A third design is the ARAPUCA, a “light-trap” which uses dichroic filters to trap wavelength-shifted light to be detected by SiPMs.

The development and testing of a photon detection system is discussed in the next chapter and related appendices, both from the general standpoint of individual components and from the standpoint of integrated prototype testing. The group at Indiana University pushed the development and testing of the double-shift style photon detection system, as will be discussed.

## CHAPTER 8

### **Development, Characterization, and Testing of Prototype Photon Detector System**

This chapter builds on the discussion of scintillation photon detection in DUNE that began in the last section of Chapter 7. Specifically, this chapter will focus on the development, testing, and characterization of the double-shift style photon detector (PD) system for the single-phase DUNE Far Detector. However, the prototype systems share many commonalities, and where such broader impact is appropriate, it will be noted here. Many of the results of this chapter have been published in Reference [102], and several figures are those used in said reference.

The system itself is described in more and finer detail than previously in Section 8.1. In order to ensure that the system is characterized and well understood, an extensive research and development program exploring the system and several prototype designs was undertaken. Projects like ProtoDUNE and future tests will study full-size prototype devices. The work described herein is somewhat different. For one, much work was done on studying individual components in laboratory testing. Various instruments and test stands are appropriate to study and test such components in the laboratory. Some relevant work to the system description and characterization will be presented in Section 8.1. Development and characterization effort more useful for the double-shift prototype PD system will be discussed in Section 8.2. Furthermore, larger facilities such as the LAr testing facilities at Fermilab provide controlled environments for slightly larger tests of integrated prototypes. A study conducted in the Blanche dewar facility at Fermilab with a 30 inch double-shift prototype device will be discussed in Section 8.3.

#### **8.1 Photon detector system prototype and general component testing/characterization**

A depiction of the double-shift style of technology is presented in Figure 8.1. The version with 4 plates is akin to that tested in the Fermilab facilities, as discussed in Section 8.3. For clarity, plates are only shown on one side of the light guide, but in reality plates line both sides. This makes the photon detector sensitive to scintillation from the TPC volume on either side of the anode in the

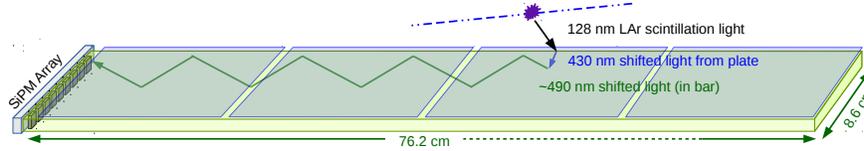


Figure 8.1: Cartoon of the double-shift style lightguide photon detection system [102]. Incident VUV photons from the LAr scintillation are converted to nearly visible wavelengths by a wavelength-shifter on acrylic plates placed in front of and behind a lightguide bar. Some of the converted photons are then incident on the lightguide bar, where a second wavelength shift and total internal reflection lead photons to an array of SiPMs at the end to be detected.

DUNE TPC layout. A full-size device would be approximately 2.1m long and have 6 such plates, each plate a bit longer than those in the prototype test.

The plates are spray coated with a solution containing tetraphenyl-butadiene (TPB) in a fume hood with a High Volume Low Pressure (HVLP) sprayer. The solution contains 5g of scintillation grade ( $\geq 99\%$ ) TPB for every 1000g of dichloromethane (DCM), which dissolves the TPB. Such a spray device is a common industrial tool. The plates for prototypes to date have been sprayed by a lab technician, though if this photon detection system were to go forward, this is a step in the procedure which could be industrialized with expected enhancement in yield and uniformity. More information on the production process is given elsewhere [102].

Incident VUV photons from the LAr scintillation are then absorbed and converted by the TPB to nearly visible wavelengths ( $\sim 430\text{nm}$ ). Light which is emitted by the TPB molecules through the plate can then be incident on the lightguide bar. Since there is a gap between the plates and the lightguide bar, there will be LAr filling the space between. This means there is the possibility for total internal reflection of the  $\sim 430\text{nm}$  light at the back surface of the plates – this will be accounted for in simulations to be described shortly.

An emission spectrum for converted light from coated plates was taken by Anna Pla-Dalmau at Fermilab, using a Hitachi F-4500 fluorescence spectrophotometer. This device records the output spectrum from the plate at a 90 degree angle from the incident direction, using 200nm light to activate the TPB: this wavelength has been shown to produce an equivalent output spectrum to incident VUV photons to a very high degree [103]. The results for three plate samples were combined to

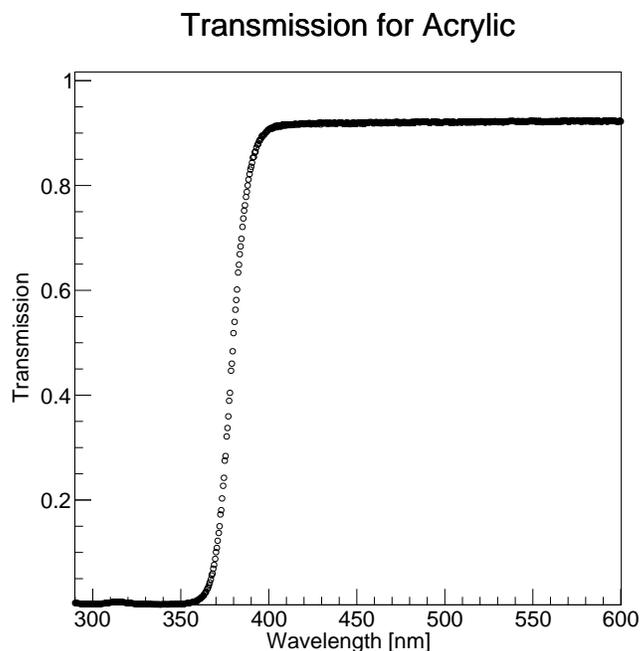


Figure 8.2: Transmission measurements for uncoated 1/16-inch acrylic plates, which are then coated with TPB to be used as the initial wavelength-shifting component of the double-shift light-guide photon detection system. For the peak region of TPB emission around 430nm, the transmission is above 90%.

account for any plate-to-plate variations.

The acrylic used for the plates was purchased from McMaster-Carr and is  $\frac{1}{16}$ -inch thick. Measurements of transmission for converted light through the acrylic were made by giving uncoated acrylic samples to John Karty in IU Chemistry, who uses a Cary 50 UV-VIS spectrophotometer to study the transmission as a function of wavelength. Figure 8.2 shows the transmission percentage of the acrylic in the main region of interest. Interestingly there is a bump in the acrylic transmission around 250nm (outside the range of the figure), but this is significantly above the LAr scintillation peak and below the emission and detection wavelengths. The top plot in Figure 8.4 shows the product of the TPB emission and acrylic transmission in red. For the dipped bar, which uses TPB as the only wavelength-shifter, the emission should be approximately 430nm as well.

The lightguide bar used in the double-shift system is polystyrene, with EJ-280 wavelength-shifter dopant converting to  $\sim 490$ nm, from Eljen. Full-length prototypes are approximately 2.1m long, 8.6cm wide, and 0.6cm thick. The prototype tested as depicted in Figure 8.1 was the same

width and thickness, but 76.2cm in length. Since these are commercially made and minimally altered in putting together the PD system, they are expected to have sufficiently long attenuation lengths ( $>2\text{m}$ ). Attenuation length was probed as part of the testing and system characterization (see Section 8.2.1). Eljen has provided absorption coefficients and emission spectra as functions of wavelength, and these are plotted in black in the top and in red on the bottom of Figure 8.4, respectively. The output of the TPB plates is well-matched to the absorption of the EJ-280.

Photons reaching the end of the bar are then incident on an array of silicon photo-multipliers (SiPMs), which are solid state photon detection devices. A good basic description of SiPMs can be found [104] from SensL, a SiPM manufacturer, and is described here. When photons are incident on a silicon chip, an electron-hole pair is formed. Applying a sufficient voltage across the chip with a depletion region from a junction formed by p and n type semiconductor materials causes the drifting e/hole pair to accelerate enough to produce an avalanche of electrons in the device. The lowest sufficient avalanche voltage is called the breakdown voltage, and the avalanche produces a measurable current from the device. A quenching resistor is added in series with the chip to limit the current, such that during avalanche, the voltage seen by the chip drops below breakdown and the avalanche stops. Once the avalanche (current) stops, the device then charges back up to the chosen voltage (at or above breakdown), waiting for another photon. Such a single chip then only serves as an on-off type detector – one or many incident photons will produce the same result. However, putting many such devices together on a single unit, one adds the ability to count photons as well: this is precisely an SiPM. Each such device on the SiPM unit is called a microcell. If 5 photons are detected by the device, the output is expected to be five times that of a single microcell. Cross-talk and afterpulsing are mechanisms by which one microcell may either trap some left-overs from the avalanche and fire again in quick succession or may leak some of the products of the avalanche to another microcell and cause it to fire. In this manner, one photon may produce a signal that corresponds to two or more microcells. This process leads to a characteristic quickly rising and

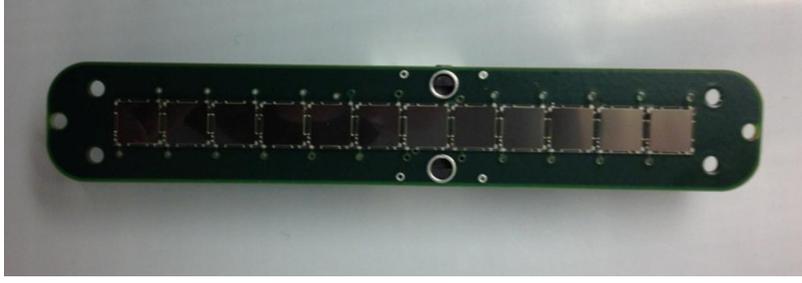


Figure 8.3: A set of twelve 6mm x 6mm SiPMs on an electronics board, equivalent to the SiPM arrays used in protoDUNE placed at the end of lightguide bars [105]. The boards and SiPMs used in prototypes described in this text are similar in idea, though different and containing varying numbers of SiPMs.

slowly falling signal. The recovery tail is that of an RC circuit with the time constant [104]

$$\tau_{RC} = C_d(R_q + R_s N) \quad (8.1)$$

with  $C_d$  as the capacitance of the microcell,  $R_q$  the resistance of the microcell's quenching resistor,  $N$  the number of microcells, and  $R_s$  any resistance in series.

For the testing reported here, both SensL B-Series and C-Series SiPMs have been used: these are similar devices. All of the prototypes mentioned for the single-phase LArTPC use SiPMs to detect wavelength-shifted photons. For the prototype testing in Section 8.3, the C-series SiPMs were used: MicroFC-60035-SMT SiPMs. The device active area is 6mm x 6mm, making it well-matched to the dimensions of the light guide. These SiPMs have 18,980 microcells, each  $35\mu\text{m} \times 35\mu\text{m}$ , for a fill factor of 64%. The SiPM cross-talk, photon detection efficiency, gain, and other properties are functions of the bias voltage applied. The properties of 7 SiPMs used in the prototype testing with bias voltage of 25.5V were tested in liquid nitrogen (LN2) and are reported in Table 8.1. LN2 is desirable for this testing as it does not scintillate but has a similar temperature to LAr ( $\sim 10\text{K}$  cooler than LAr). A set of 12 SiPMs in a setup such as that used in the protoDUNE test experiment is shown in Figure 8.3 [105].

The individual photoelectron (PE) peaks from the devices are typically very distinguishable. The digitizer used is described later, but for now it is only necessary to know that it records wave-

forms and can self-trigger when SiPMs fire. To separate out triggers on background fluctuations, events determined to have  $<0.5\text{PE}$  are typically discarded. The SiPM dark noise is defined by the rate of events having over  $0.5\text{PE}$  in amplitude. Similarly, since this is dark noise in LN2, any event should have  $1\text{PE}$ . Events with more signal indicate cross-talk. The cross-talk probability is simply the ratio of events with over  $1.5\text{PE}$  to overall events with over  $0.5\text{PE}$ . Labelling all delayed cross-talk or afterpulsing as simply “afterpulsing,” one can explore this in dark noise by looking for a second pulse delayed from the first but fitting within the readout window. Anything after that would simply be seen as an extra noise pulse. From this study, afterpulsing was estimated to be  $\lesssim 1\%$ , and is therefore considered negligible. Furthermore, the gain is determined by looking at the integrated charge collected from the mean  $1\text{PE}$  signal and dividing by the charge of the electron ( $\sim 1.6 \times 10^{-7} \text{pC}$ ). Such charge measurement is described below by equation 8.2. The breakdown voltages of the SiPMs were determined by finding the gain at a series of bias voltages between  $24.5\text{V}$  and  $29.5\text{V}$  and fitting a straight line: the breakdown voltage is where the line crosses  $\text{gain}=0$ . The results of such tests show that the nominal bias voltage of  $25.5\text{V}$  is about  $5\text{V}$  above breakdown. The photon detection efficiency (PDE) of the C-series SiPM at this overvoltage is provided by SensL and plotted in black in the bottom of Figure 8.4. Note that the PDE is better matched to the output of TPB directly but is still reasonably well-matched to EJ-280 emission.

SiPM	Noise [Hz]	Cross-Talk Probability	Gain ( $\times 10^6$ )	Break. Volt. [V]
0	14	0.20	4.6	20.5
1	11	0.19	4.5	20.6
2	13	0.19	4.6	20.5
3	13	0.20	4.6	20.5
4	14	0.19	4.6	20.4
5	11	0.19	4.5	20.6
6	14	0.19	4.5	20.6
Mean	13	0.19	4.6	20.5

Table 8.1: Properties of SensL C-series SiPMs, determined from dark noise in LN2 at  $V = 25.5\text{V}$ .

There are some important considerations with respect to SiPMs in the single-phase photon detectors. While the versions of SensL SiPMs used in the course of testing were generally fairly

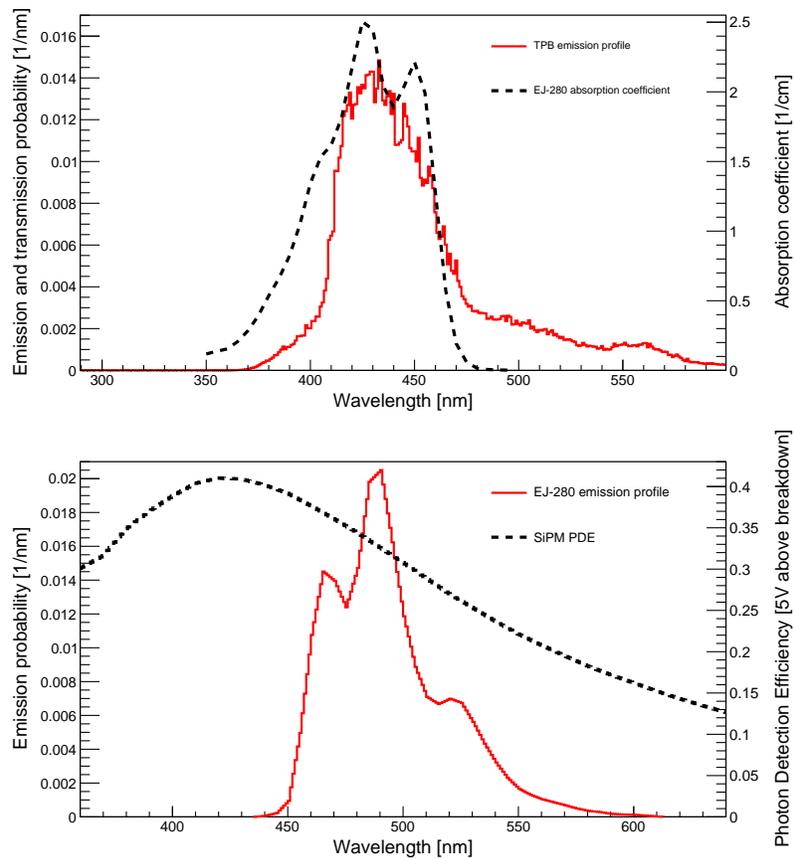


Figure 8.4: Wavelength-shifting and detection properties of the double-shift lightguide bar style photon detector. The top plot shows the transmission, in red, of converted light through the plates (the product of plate emission and transmission) as described in the text and the absorption coefficient for light in the bar in black. The bottom plot shows the emission from the wavelength-shifter in the lightguide bar in red and the photon detection efficiency of the SiPMs used in black. In general, there is good overlap between these various steps of emission, absorption, and detection.

robust, a change in specifications of the devices resulted in newly purchased versions of these devices cracking at intolerable rates in LN2 (for more on SiPM comparisons, see [83]). However, SiPMs are produced by several companies, and a few (Hamamatsu and Fondazione Bruno Kessler) have signaled an interest in working with DUNE to produce viable SiPMs [83]. In the ideal scenario, a device would have minimal cross-talk and after-pulsing, and perhaps in the case of the double-shift style PD, be a better match with higher PDE at the EJ-280 output.

If a photon is detected by the SiPM - that is, it results in an avalanche and output stream of electrons - the signal is sent over a cable, across a feed-through at the warm/cold interface, and to a processor outside the detector. A system used to date is the SiPM Signal Processor (SSP), which was developed at Argonne National Lab. The device has 12 channels and is capable of providing a bias voltage and 14-bit readout over each channel. In the device, the conversion factor is  $6.60 \times 10^{-3} \text{mV/ADC}$  [92]. Furthermore, since the device's impedance is  $100\Omega$ , the summed ADC of an SiPM pulse can be converted to charge (as used for gain) via the following conversion [92]

$$c [\text{pC}] = \frac{s}{10^{-12}} \times 6.60 \times 10^{-6} [\text{V/ADC}] \times \frac{1}{100} [1/\Omega] \times 6.67 \times 10^{-9} [\text{s}] \quad (8.2)$$

where  $s$  is the summed ADC in the pulse.

The digitizing sample rate is 150MHz, corresponding to 6.67ns between samples, though the device is capable of running at different clock speeds, such as twice the NOvA timing system rate, 125MHz [106]. A readout window in the nominal setting is capable of lasting approximately  $13.6\mu\text{s}$ . In our prototype tests such as those to be discussed in Section 8.3, we typically read out  $13\mu\text{s}$ , with  $2\mu\text{s}$  of baseline before the trigger and  $11\mu\text{s}$  after. This is sufficient to record almost all of the singlet and triplet emission, recalling that the lifetimes are a few ns (singlet) and  $\sim 1.5\mu\text{s}$  (triplet). The working goal in the creation of the SSP is 30ns timing resolution [107]. A digitized waveform of two SiPM pulses with the sharp rise and recovery tail is seen in Figure 8.5.

We performed tests to characterize the timing resolution of the readout (technically the SiPMs and SSP together) by sending pulsed signals from an LED separated by  $5\mu\text{s}$  and studying the differ-

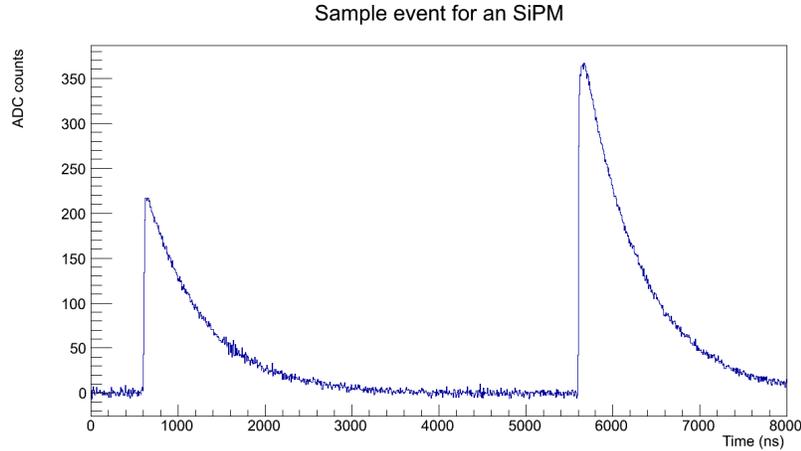


Figure 8.5: A sample waveform digitized by the Argonne SSP, showing two pulses from a single SiPM. Note the characteristic sharp rise and recovery tail.

ence in reconstructed pulse times. This study is described in detail in Appendix B and the results are summarized in Figure 8.6. Figure 8.5 shows an example event from this study. Though the timing resolution is better for stronger signals, even for modest signals the timing resolution is several ns. For scintillation signals, there will be smearing effects from effects such as the scintillation lifetime, reflections off of cryostat components, etc. However, the SSP itself more than meets its goal to help ensure the best system is in place.

Electronics channels in such a system are expensive, given the timing and ADC capabilities of the device and the digitization of waveforms. This becomes especially true when considering that a Far Detector module would contain tens of thousands of SiPMs. In most photon detection prototypes made till now, the SiPMs are only along one end of the lightguide bar. Since light will escape both ends of the lightguide bar, one way to enhance the photon detection efficiency of these systems is to employ SiPMs at both ends (double-ended readout). This would be expected to double the overall efficiency of the system, but also doubles the number of SiPMs needed.

For this reason, schemes which reduce the cost of photon detector channels are being explored. Of course, one way would be to change the system to one with lower design specifications – maybe one which only records certain information from an event rather than saving a whole waveform. Another method to reduce the channel and cable count is to multiplex several SiPMs to readout over

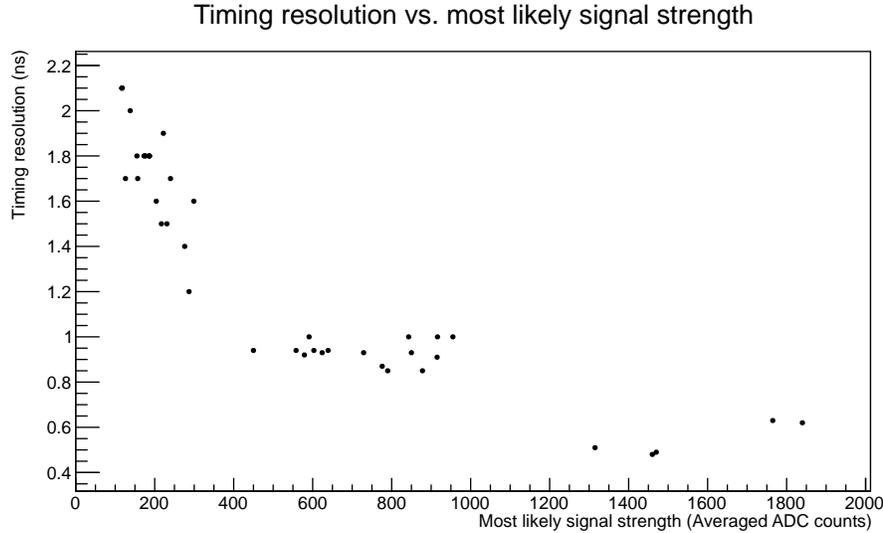


Figure 8.6: The timing resolution of signals in the PD system electronics, the SiPMs and the Argonne SiPM Signal Processor (SSP) is dependent upon the amount of photons. Each point is the resolution of a data channel, as reported in Appendix B. The timing resolution is on the order of a few ns for modest signals.

one channel, referred to as “ganging.” There are two strategies to ganging: passive and active. In a passive ganging scheme, the devices are connected in parallel. In an active ganging scheme, the devices are summed and fed through the same signal, for example via an operational amplifier on the SiPM board. Both schemes have the advantage of enabling more SiPMs to be used on any given device (for example in a double-ended readout), though each has drawbacks as well. When the SiPMs are passively ganged, the increased RC constant leads to a lower amplitude SiPM pulse with a longer recovery time. For the SensL SiPMs, this means that gangs of a few SiPMs will not seriously compromise detection, but gangs of 6 or more SiPMs on a single channel might mean low level signals are washed out. Different SiPM types may behave differently when ganged, so many options will need to be explored. An active ganging scheme may be seen as ideal, since with the right tuning of parameters one may expect to not degrade the signal. However, the addition of an op-amp to the system presents a potential source of noise, shaping effects in the waveforms, and/or heat. Some early work on understanding ganged SiPMs is presented briefly in Appendix C, though more work has been done and more will need to be done.

## 8.2 Laboratory-based tests of double-shift lightguide photon detection system

One specific prototype design for a single-phase detector, the double-shift lightguide PD system, was heavily pushed forward in research and development by the Indiana University group. Several of the tests described in the last section and related appendices highlight laboratory-based testing. Further laboratory tests provided characterization of the components specific to the double-shift PD system, namely the lightguide bar and the wavelength-shifting plates. Some of the tests of these components are described herein, as well as a C++ ray-tracing simulation which uses this information to calculate an expected performance of the system.

### 8.2.1 Attenuation testing

To accurately characterize the performance of the PD system, it is important to understand the components. The polystyrene lightguide bars with EJ-280 are purchased from Eljen Technologies, and largely left unaltered. This is because the goal of separating the initial VUV wavelength shift from the transport is to preserve the attenuation length of the system, such that a 2m-long lightguide bar will still have modest efficiency at the far end. To test this, we employ a medium-sized dewar filled with LAr, which will hold our 30-inch prototype bars.

The bar was held in a stand that employed an array of SiPMs across the bottom and allowed a holder to slide along the face of the bar. A lab technician operates the holder with an  $^{241}\text{Am}$  source that produces  $\alpha$ s that stop quickly in the LAr, producing scintillation. In front of the source, a 1-inch wide piece of a TPB-coated plate was placed. This then rides along with the  $\alpha$  source, acting as a relatively constant source of  $\sim 430\text{nm}$  photons to test the response of the bar as a function of distance from the readout end. The lab tech moves the source along the bar in 1-inch increments for 25 inches, though the exact starting point of the  $\alpha$  source is not known *a priori*. It is likely that it sits near the readout end offset by holder spacings, but since it is in a closed-off environment, it is not known for sure that it is not slightly above this. In the comparison to the simulation which comes later, the starting point was allowed to float some. The apparatus is depicted in Figure 8.7.

The collected information is processed and analyzed to get information on the attenuation length and characteristics of the bar. Four SiPMs located near the middle of the bar were used to record

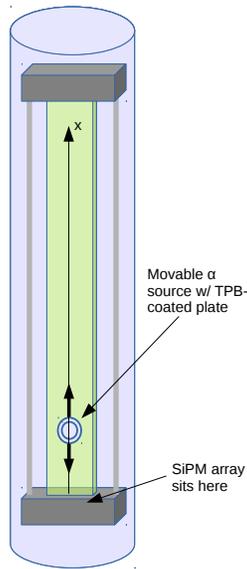


Figure 8.7: The experimental setup for the wavelength-shifting bar attenuation property scan. The IU-dewar can reasonably fit a 30-inch prototype bar as used in the testing in Section 8.3. An  $\alpha$  source with a piece from a wavelength-shifting plate is moved along the face of the bar. This tests the bar's response as a function of distance from the readout end.

the scintillation signals using an SSP. Since the SSP was run in self-triggered mode, the SiPMs will record other signals above threshold as well, for example scintillation signals from cosmics or SiPM dark noise. The summed charge in the waveform is measured from each waveform (see Equation 8.2). Approximately 148,000 waveforms from each SiPM is included in the analysis, and the results from one of the four SiPMs at two positions along the bar are shown in Figure 8.8.

There are two pieces of the distributions which are immediately visible: a tail at low integrated charge which decreases quickly above the threshold, and a Gaussian-like distribution which is separated from this tail. A fit was performed in ROOT as the sum of an exponential (for the background scintillation in the tail) and a Gaussian (for the  $\alpha$ s). The mean of the Gaussian is then taken to be the response of the bar as a function of position along the bar. The fits are superposed in the figure, and the vertical dotted lines are the means of the Gaussian distributions at the two locations. The blue dashed curve is at a higher position on the bar (farther from the readout), and therefore has lower signal due to attenuation.

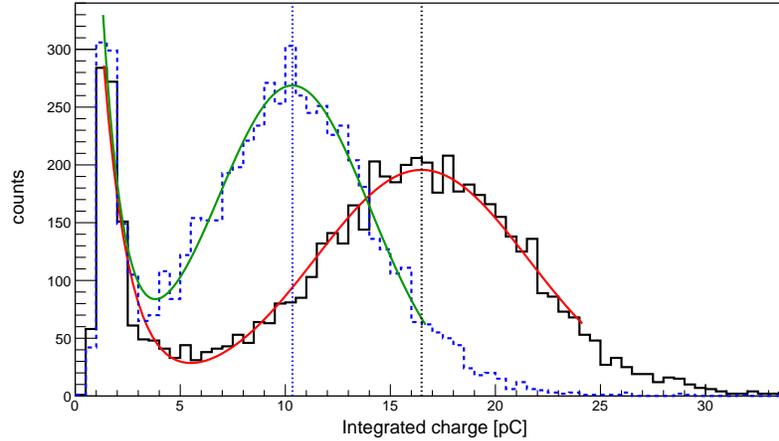


Figure 8.8: Example of the response of the bar to signals from the alpha source at two locations along the bar, in a histogram of the integrated charge collected from waveforms. At both positions a clear peak can be seen, with a tail near 0 resulting largely from residual scintillation backgrounds in the LAr. The black solid curve is at a position closer to the readout than the blue dashed curve: the attenuation resulting from the longer paths in the latter case results in fewer detected photons on average. To quantify this effect, the peak of a Gaussian fit to the signal region is used. The values are shown by the vertical dashed lines.

At each point, the mean of the Gaussian for the four SiPMs were averaged in a weighted mean to provide the signal as a function of distance from the readout end of the bar. The errors on each point were taken as the standard deviation of the four values. The data are characterized by the black crosses (the vertical lines are error bars) in Figure 8.9. At this point in the discussion the y-axis in the figure is fairly meaningless, as it requires information from Section 8.2.3. However, the data have been normalized in a way to preserve the shape and compare with simulation. What is important to note is that two characteristic fall-off scales are seen in describing the curve, both in data and in the simulation, and can be attributed in part to combinations of light directly incident on the SiPMs and light which undergoes reflections off the sides of the bar before hitting a SiPM.

Ultimately, the simulation will be used to provide a functional description of the attenuation, but the long tail is found to correspond to a  $>2\text{m}$  long attenuation length. The initial fall-off is accentuated by only producing and detecting photons in a somewhat central location on the bar. In a fuller description, the short attenuation length is a few cm.

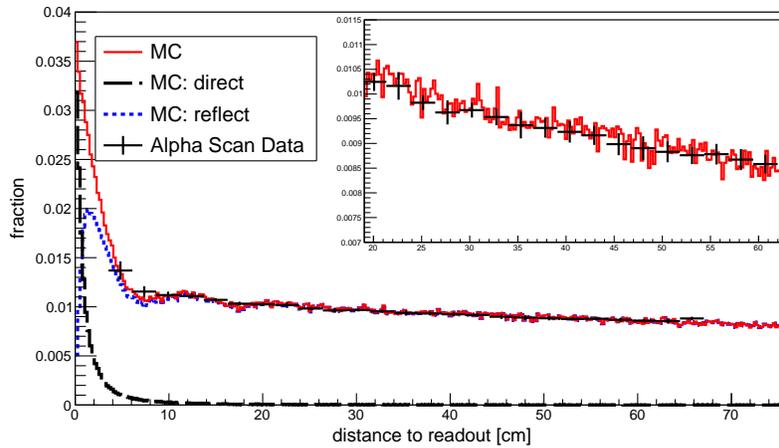


Figure 8.9: Results of the attenuation length testing from the alpha scan are shown in the crosses in the figure. The test started over an inch up the bar, but the quick drop in the first few bins can be seen, clearly followed by a much slower fall-off along the bulk of the bar. These two characteristic length scales are also seen in the simulation in red (further broken into direct and reflected components). Such simulation is discussed more specifically in Section 8.2.3. The inset shows a zoomed in region from  $\sim 20$ cm to  $\sim 60$ cm to show the agreement in this flatter region. Such a bar would have a  $\geq 2$ m bulk attenuation length.

## 8.2.2 Plate testing

A VUV monochromator provides a light source with the ability for an operator to choose a wavelength and was used to characterize the relative performance of plates and to obtain estimates of TPB-coated plate efficiencies. For our testing we used a McPherson model 234/302 VUV monochromator, which utilizes a deuterium lamp with usable output in the VUV region. A grating in the monochromator is used to select a wavelength from the spectrum and send it on to a sample chamber, where a piece of a wavelength shifting plate could be illuminated. A VUV-sensitive photodiode examines the output of the lamp incident in the sample chamber at the selected wavelength, and measurements of a plate sample taken on the same day (since the lamp brightness varies as a function of time between cleanings) with an SiPM behind the sample characterizes the plate's output in converted light. The concept of this measurement is illustrated in Figure 8.10. A Keithley picoammeter was used to measure the currents output by these photosensors, for which the errors are small compared to the typical readout level. In a purely relative comparison, the brightest plates would have the highest ratio of SiPM readout to VUV photodiode readout. The plates were

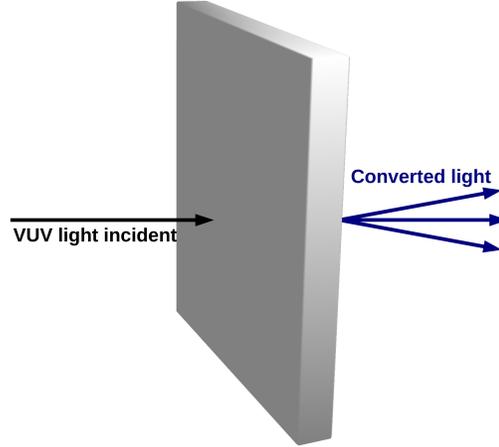


Figure 8.10: Cartoon explaining the process of measuring the efficiency of a wavelength shifting plate in a VUV monochromator. Photons from a deuterium lamp are selected at a given wavelength and sent to a sample chamber. 128nm light is sent to the sample chamber, where it is incident on a sample cut from a wavelength shifting plate. Shifted light leaving the back of the plate is then incident on a SiPM. Comparisons of the SiPM signals relative to a measurement taken of the lamp itself with a VUV-sensitive photodiode gives the efficiency of the plate.

originally characterized by testing sample pieces cut from extra plastic on the plates.

A technician in our group made many plate samples and tested them in the VUV monochromator using the prescribed techniques that we developed. We analyzed this data and selected the best of these for the Blanche experiment. Once the experiment at Blanche (see Section 8.3) was completed, more destructive tests were performed. In this process, our group ran a series of tests to understand the wavelength-shifting efficiency. Three pieces of the actual plate surfaces for the four main plates used in prototype testing at Blanche were cut out and tested in the monochromator.

The plate efficiency is given by the equation

$$\epsilon_{\text{TPB}} = \frac{1}{f_{\text{geo}}} \frac{I^{\text{SiPM}}}{I^{\text{VUV}}} \frac{\mathcal{R}^{\text{VUV}}}{\mathcal{R}^{\text{SiPM}}} \quad (8.3)$$

where  $I^{\text{SiPM}}/I^{\text{VUV}}$  is essentially the ratio described above for the relative tests. The other terms are the corrections necessary to convert this ratio into a measurement of efficiency. In order to convert the current to a number of photons, the  $\mathcal{R}$  terms give functions of the responsivity  $R$  (Amps/Watt)

and photon energies relating photons to the current. For the photodiode readout,

$$\mathcal{R}^{\text{VUV}} = R^{\text{VUV}}(128\text{nm}) \frac{hc}{\lambda(128\text{nm})} = 2.44 \times 10^{-19} [\text{As}/\gamma] \quad (8.4)$$

gives this relation, where  $R^{\text{VUV}}(128\text{nm}) = 0.157\text{A/W}$  is read from a responsivity curve provided for Opto Diode AXUV-100g devices. Conversely, measurements of the SiPM current examine the output spectrum, meaning a wide range of wavelengths are relevant. Integrating over the information for the TPB output as measured for our plates (much like in Figure 8.4) and the SiPM responsivity as provided by SensL for 5V over breakdown gives the relation

$$\mathcal{R}^{\text{SiPM}} = \int d\lambda_{\text{vis}} \mathcal{P}_{\text{TPB}}(\lambda_{\text{vis}}) R^{\text{SiPM}}(\lambda_{\text{vis}}) \frac{hc}{\lambda_{\text{vis}}} = 4.09 \times 10^{-13} [\text{As}/\gamma] \quad (8.5)$$

where the integral is between 250nm and 600nm where the relevant data overlap in range. Estimates of 5% were adopted as fairly standard uncertainties on the responsivities of the photosensors, and a quadrature sum of these (7.1%) was taken as an error on the efficiency measurements.

Since the studies aim to characterize the overall efficiency of the plate, a geometrical factor  $f_{\text{geo}}$  corrects the observation based on the acceptance of the readout system. A series of tests and experimentation was used to provide an estimate for this geometrical correction factor. Tests were performed as the above measurements but with apertures placed just in front of the VUV photodiode to reduce the amount of sensitive area, with the readout from a 4mm diameter aperture providing the measurement used. Furthermore, tests with half of an aperture blocked (vertically and horizontally separately) determined where on the photodiode the light seemed to fall. Combining these, it was determined that the light fell in the upper left quadrant. A rectangular pattern of incident light in this quadrant using the aperture tests for guidance with respect to dimensions was chosen for the estimation, and a Monte Carlo simulation assuming Lambertian emission of photons calculated the fraction of photons emitted that would fall on the SiPM. The determined  $f_{\text{geo}} = 0.31$ , and to allow for some of the estimation which went into the calculation, a relative error of 5% was taken.

As noted, 3 pieces (herein called A, B, and C) cut from the plates after use in Blanche were

measured to perform the efficiency measurements. The overall efficiency of the plate was taken to be the average of the three measurements. A relative error in the uniformity of the plate (“plate”) was estimated by the variation in the samples for a given plate: standard deviations divided by the means. Error associated with measurements using the VUV monochromator (“VUV”) were estimated by variation from multiple trials with a somewhat similar configuration to be 5.7%. This is taken to be a reasonable estimate of the repeatability uncertainty. Original measurements were taken with a NIST-calibrated VUV photodiode, but as it was found to have degraded, a new VUV photodiode was used. In measurements with the new photodiode, plate 3 was found to have results which made little sense compared to those of plates 1, 2, and 4. It is assumed that something affected plate 3 between measurements, and instead a correction factor was applied to the previous measurements of plate 3 based on other measurements with an estimated uncertainty of 5.9% (“corr”). The measurements are summarized in Table 8.2.

Plate	Sample	$(I^{\text{SiPM}}/I^{\text{VUV}}/10^5)$ (128nm)	$\epsilon_{\text{TPB}}$
1	A	3.27	0.63
	B	2.87	0.55
	C	3.39	0.65
mean $\epsilon_{\text{TPB}} = 0.61 \pm 8.7\%$ (plate)			
2	A	3.51	0.68
	B	3.08	0.59
	C	3.51	0.68
mean $\epsilon_{\text{TPB}} = 0.65 \pm 7.4\%$ (plate)			
3	A	3.28	0.63
	B	3.66	0.71
	C	3.29	0.63
mean $\epsilon_{\text{TPB}} = 0.66 \pm 6.3\%$ (plate) $\pm 5.9\%$ (corr)			
4	A	2.54	0.49
	B	2.70	0.52
	C	2.33	0.45
mean $\epsilon_{\text{TPB}} = 0.49 \pm 7.2\%$ (plate)			

Table 8.2: Results of tests of wavelength-shifting efficiency of TPB-coated plates used in the prototype experiment discussed in Section 8.3. In addition to the uncertainties in each row are 5.7% uncertainties in VUV repeatability and 7.1% response uncertainties. A 5% uncertainty is taken on the estimate of  $f_{\text{geo}}$ .

### 8.2.3 *Simulation for expected system performance*

It is desirable to have a simulation which will describe the expected performance of the system as a whole. A preliminary simulation was worked on by Denver Whittington (IU postdoctoral scientist at the time) and Johnathon (Lowery) Jordan (IU undergraduate at the time) as a ray tracing program that tracked photons created at the top of a long rectangular prism, to understand the attenuation behavior of the dip-coated lightguide bars for an alternative PD system model.

I upgraded and expanded the simulation to essentially contain the entire double-shift lightguide PD system, using laboratory tests and vendor information and yielding a performance expectation. It retains the ray-tracing nature of the original simulation, but now starts photons on the top surface of TPB-coated plates above the light guide and tracks them through various processes of emission/absorption, ultimately arriving at a prediction of performance including SiPM readout. The implementation is meant to describe the system essentially as shown in Figure 8.1. Two versions of the simulation were used, one which is set up similar to the  $\alpha$  scan and gives the correct scale to use for attenuation parameters. The second version gives an expectation of detector efficiency and is compared to the Blanche results discussed in Section 8.3.

As noted above, the simulation starts with the wavelength shifting plates. As the main point is to create a calculation of detection efficiency as a function of position along the PD system, we distribute the photons uniformly over the surface and keep track of where they started. In the full detector simulation, roughly  $1 \times 10^9$  are created all along the plate surfaces, as depicted in Figure 8.11. For the alpha scan simulation, the photons are constrained to a region nearer the center of the bar to approximate the effect moving the  $\alpha$  source up and down along the bar: for this simulation,  $5 \times 10^8$  photons are used instead, since the area to cover is smaller.

The photons are considered to be incident normally on the plate surface. This is an approximation as in reality photons would be incident at angles on both sides of normal depending on its position relative to the scintillation source. For the prototype testing with cosmic-ray muon tracks in LAr as described in Section 8.3, it was found that the major density of photons is incident essentially where the track crosses the PD system's face. Therefore, normal incidence in the ray-tracing

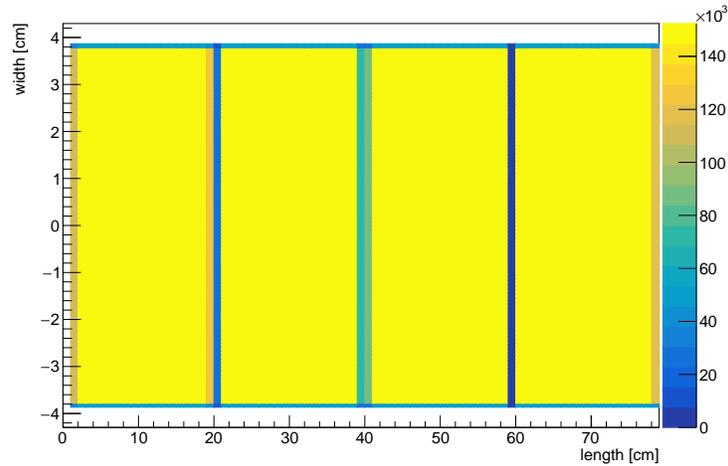


Figure 8.11: Generation point of photons on the surfaces of the wavelength-shifting plates. The x-axis is distance from the readout end, and the y-axis is the width of the plates ( $\sim 7.6$ cm). Note that the color is uniform over the bulk of plate surfaces. The gaps in regions of the x-axis between yellow regions are gaps between the four plate surfaces due to holding structures.

program is taken as an approximation. In fact, the distributions are seen to be peaked and somewhat Gaussian-like. For each photon, a wavelength is chosen randomly from the TPB emission spectrum described earlier in this chapter (see Figure 8.4), and the direction is picked randomly from a Lambertian (cosine) distribution [103]. If directed away from the light guide, that photon is considered lost. An overall scale factor equivalent to the plate efficiency will be used later to account for the fact that not every incident VUV photon produces a converted photon.

If directed toward the light guide, the simulation continues. The transmission probability for a 1/16-inch plate (described earlier) determines if the photon will reach the light guide boundary, again by drawing randomly and comparing to this percentage. For an efficiency of  $X$ , a random number between 0 and 1 that is  $\leq X$  is accepted. If it reaches the boundary, the angle is compared relative to the critical angle for total internal reflection, using a typical index of refraction for acrylic of 1.49 and 1.23 for LAr. If it is reflected, the photons are tracked via ray tracing and further checked at boundaries. In principle, most of these photons are likely lost; the ends of each plate are covered by opaque holders and so photons reaching these ends are considered lost.

For photons which do exit through the bottom of the wavelength shifting plate heading toward

the lightguide bar, their directions are recalculated using Snell's law, and intersection points with the plane of the top of the lightguide bar are computed. If that point falls on the light guide, Snell's law again is used to determine the trajectory in the polystyrene bar. The index of refraction for an Eljen polystyrene bar is taken as 1.6 (approximated from information provided by Eljen). The fraction of photons created at a position on the plates which are incident on the bar is shown in Figure 8.12. The theoretical maximum would be 50% since half the photons are expected to be sent backward away from the system. Including photons lost in transit, for example by total internal reflection in the plates, the number of photons which make it to the lightguide bar is reduced further to  $\sim 35\%$ . The gap between the plate and lightguide bar is set to an estimated measurement taken of the prototype device used in Section 8.3. For the  $\alpha$  scan simulation, this distance is set instead to an approximation of the distance from the bar to the plate sample.

If the photon is not incident on the top of the lightguide bar, it is considered lost and the simulation moves to the next photon. Using a random number drawn from an exponential given the absorption length in EJ-280 at the incident photon's wavelength (see Figure 8.4), an absorption path length for this specific photon is determined. If this length is longer than the path to the given lightguide bar boundary the photon is incident upon, it would not be absorbed in the light guide and is considered lost. If it is shorter, then the photon's capture point is where that path length is reached. Given the EJ-280 efficiency stated by Eljen of 86% ( $\pm 1.7\%$ ), a random number is used to determine if the EJ-280 will re-emit the photon.

In principle, since the SiPMs are sensitive to the photons created by the initial wavelength-shift by TPB, photons which are not absorbed in the lightguide bar but would otherwise be incident on the SiPM could still be detected. The geometry is such that most of the effect should be in the bins close to the readout end. Contributions from this process are not captured in this simulation. However, since most of the tracks pass further along the PD system in the prototype test described in Section 8.3, a relatively small amount of scintillation photons are expected to be incident nearest the readout end. Therefore, this and other simplifications/approximations affecting the first several centimeters/the first plate are not expected to matter much for the overall conclusions. Likewise,

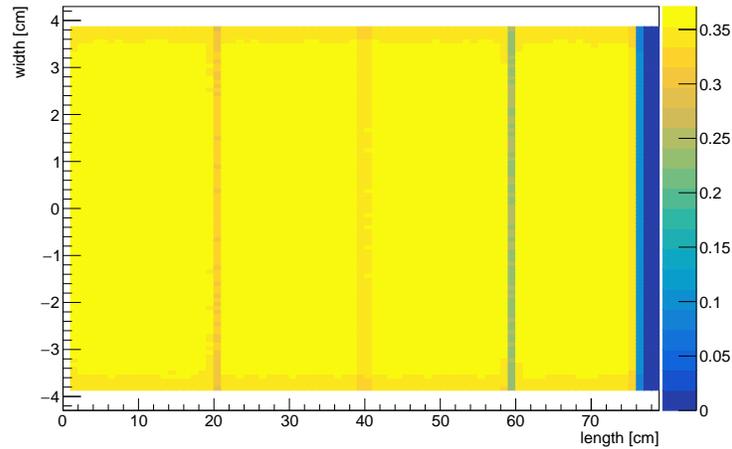


Figure 8.12: The fraction of created photons from a given point on the plate surface which make it to the light guide. The theoretical maximum would be 50% since half of the photons would be sent backward away from the system. Including photons lost in transit, for example by total internal reflection in the plates, the number of photons which make it to the lightguide bar is reduced further to  $\sim 35\%$ .

more recent work by others to produce a Geant4 model of this PD system found that our ray-tracing simulation underpredicts the efficiency relative to the Geant4 model [108, 109], but mostly contained to the first several centimeters<sup>1</sup>. It is suspected this direct component could be a contributing cause of the disagreement.

For photons that pass the re-emission criteria in the lightguide bar, a new wavelength is chosen randomly from the EJ-280 emission spectrum with direction chosen from uniform solid angle. Any time a photon is re-emitted it is required to have a wavelength equal or longer to the absorbed photon (for energy conservation). Photons are then tracked through the light guide in a ray-tracing method similar to that used in the plates and to that of the original simulation this builds from. If the photons are incident on a lightguide bar surface at or greater than the critical angle, they are reflected; otherwise, the photon is considered lost. A parameter called the “survival parameter” ( $P = 0.9988$ ) is checked at reflections to allow for the loss of photons due to some unspecified physical process

<sup>1</sup>Reference [109] is an email exchange after the author of the Geant4-based model talk received more details and the ray-tracing results. The simulation was rerun and the updated results overlaid on the ray-tracing results. The agreement was worse in the first several centimeters than in the original talk, but after this portion of the first plate the two were in a similar regime. As said in the main text, only a small fraction of light is believed to be incident this far up the detector module, so the simplified version used when producing these results is deemed sufficient.

at boundaries. In this sense, the simulation is allowed to account for bulk attenuation parameters not covered by other photon loss methods during transit down the bar. This parameter is chosen to provide good match to data from the  $\alpha$  scan as discussed shortly.

Photons are tracked until they are either lost or escape out the readout end of the lightguide bar. Since there is overlap in the emission and absorption spectra of EJ-280, it is possible for photons to be reabsorbed/re-emitted by the EJ-280. A distance is picked at random from an exponential for the path length at which this reabsorption occurs, based on the absorption coefficients discussed earlier. The same efficiency, random wavelength selection at an equal or higher wavelength, direction, etc. govern this process.

Figure 8.13 describes the major contributions of reasons that photons are lost as they traverse through the bar. “Destruction at reflection” refers to photons lost at a boundary instead of reflecting (survival probability is 0.9988) to account for attenuation. Another process that is considered for the more complete picture here is when photons are reabsorbed by the EJ-280 and then emitted away from the readout. Additionally, photons could be lost to the wavelength-shifter inefficiency. Finally, the bulk of photons appear to be lost after a first encounter with a boundary allows for reflection, but a subsequent encounter does not. Since the bar is a three-dimensional object, it is not such that if it is trapped once it will be trapped at the next boundary.

Recall that in the baseline design, a row of SiPMs sits at one end of the lightguide bar. If photons reach the readout end and escape - the angle is less than the critical angle for total internal reflection - they can be detected by the SiPMs. The wavelength corresponds to a particular photon detection efficiency in the SiPM, and this is then checked against a random number to determine if that photon is indeed detected. In the general simulation meant to describe the baseline system, 12  $6 \times 6 \text{mm}^2$  SiPMs in an optimized layout will cover most of the surface (84%). Due to this large coverage and near symmetry, the overall number of photons escaping the readout end is weighted by 84% to account for the SiPM coverage. In the simulation of the  $\alpha$  scan, since only four SiPMs were used in a more central region of the lightguide bar, approximate locations of the SiPMs were entered into the simulation. Rather than an overall weight, the position of the exiting photons were

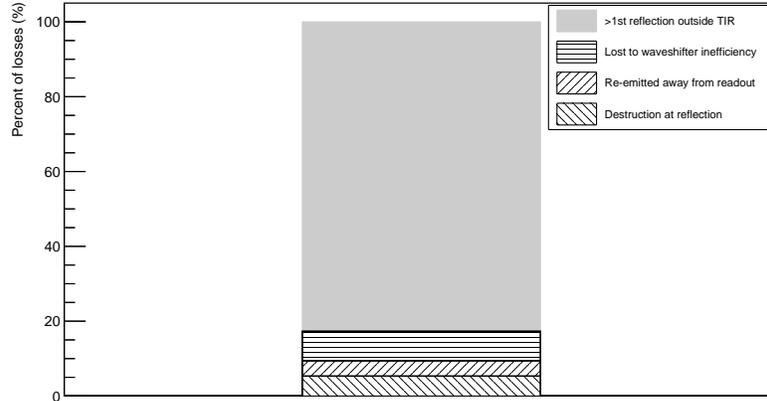


Figure 8.13: The relative contributions of major processes leading to loss of photons as they travel in the lightguide bar towards the readout. “Destruction at reflection” refers to photons lost at a boundary instead of reflecting (survival probability is 0.9988) to account for attenuation. Another process that is considered for the more complete picture here is when photons are reabsorbed by the EJ-280 and then emitted away from the readout. Additionally, photons could be lost to the wavelength-shifter inefficiency. Finally, the bulk of photons appear to be lost after a first encounter with a boundary allows for reflection, but a subsequent encounter does not. Since the bar is a three-dimensional object, photons trapped once need not be trapped at the next boundary.

checked against approximated SiPM positions to determine if the photon was incident on an SiPM.

One major goal of the simulation is to understand the attenuation effects in the lightguide bar. The first step is comparing the simulation of the  $\alpha$  scan to the data collected and picking a survival probability for photons at reflections. The transported fraction of photons in the simulation is determined by comparing the fraction of photons detected as a function of the position in the *light guide* at which the originally converted photon (from the TPB plate) was absorbed. The fall-off is essentially the attenuation. The results were compared to the  $\alpha$  scan study in shape by normalizing the plots such that the shared range had the same area. The starting point of the  $\alpha$  source with respect to the readout end in data was not precisely known, but by allowing the starting position to vary a little in a given test, the starting point of 4.83cm for the chosen survival probability gave a good fit. This starting point is a bit less than 2 inches, which is reasonable given the size of the movable source and holder. The survival probability is chosen to be  $P=0.9988$ , which gives  $\chi^2/dof = 17.8/25$  to match the data and simulation curves. A fuller description of this is given in Reference [102].

The comparison was shown earlier in discussing the  $\alpha$  scan in Figure 8.8. The bump in the simulation is possibly there in the data too. It can likely be explained by reflection modes in the total internal reflection criterion. For photons created in exactly the center of the lightguide bar and detected by a point-like readout in the center of the lightguide bar, reflected photons must be a certain distance to strike a wall and then hit SiPMs given that there is a critical angle. This process will turn on at some distance and the light received will increase. Smearing will be caused by neither the photons nor the SiPMs being constrained to exactly the center and by the ranges of different path lengths to the top and side walls.

Overall attenuation effects as light travels along the lightguide bar are characterized by a transport function. To determine the transport function to explain the full prototype, the same sort of plot as the alpha scan comparison is made but for the nominal simulation. The fraction as a function of position along the bar appears to fall off with two characteristic scales, so a double exponential fit is made with the form

$$f(x) = Ae^{-x/x_{sh}} + Be^{-x/x_{lng}} \quad (8.6)$$

for the tested bar. In Equation 8.6  $x_{sh}$  is the shorter attenuation length,  $x_{lng}$  is the longer attenuation length, and A and B are the amplitudes characterizing the relative contributions of the two exponentials. The results are given in Table 8.3, and the plot illustrating this is given in Figure 8.14. When used, A and B are normalized by dividing by the sum to provide a function which is equal to 1 at  $x=0$  (the readout end). The large  $\chi^2/dof = 1.5$  from the fit is believed to be due to the fact that only statistical errors from the simulation were used<sup>2</sup> and to the fact that a double exponential is a simplification of the fall-off. Given the data and efficiency test performed in experimentation, such a model was deemed justified. The “long” attenuation length recovered from the fit is indeed longer than 2m, which is reasonable given the designed full-scale length.

	A/(A+B)	$x_{sh}$ [cm]	B/(A+B)	$x_{lng}$ [cm]	$(\chi^2/dof)_{sim}$
transport func.	0.29	4.3	0.71	225	448/296

Table 8.3: Transport function fit parameters characterizing attenuation along the light guide.

<sup>2</sup>These are small for the number of simulated photons.

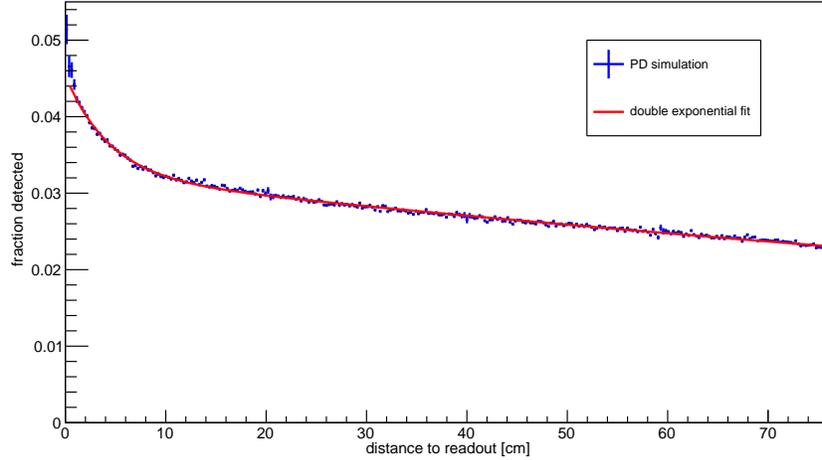


Figure 8.14: The fraction of photons detected at which the photon was initially absorbed/re-emitted in the bar as a function of distance to the readout. This is used to get the transport function, an interpretation of the effect of attenuation along the bar. For the alpha scan test described previously, the simulation results of the transport function are shown in Figure 8.9.

Finally, multiplying the fraction of photons detected by the plate efficiencies provides the computation of the efficiency of the PD prototype as a function of position for an incident photon. For the prototype PD system used in this work, four plates are lined up along a side of the lightguide bar. The plates are put between spacers that block some surface area, so these areas are masked. The plot of efficiency is shown in Figure 8.15, where the errors are taken from the determination of the plate efficiency and the statistics of the simulation, which are again small. As shown in Table 8.2, the four plates have different efficiencies, hence the discontinuity across gaps on the plot. Due to these varying plate efficiencies, using this plot alone makes it difficult to give an overall efficiency calculation. The approach taken was to average the efficiency of plates 1-3 and use these in the simulation to determine efficiency: plate 4 was dropped from the average as one that could ideally be rejected in a full-scale production. The resulting efficiency curve was then fit with a free parameter times the transport function, where the free parameter is then an estimate for the efficiency at the readout end. The resulting absolute efficiency model for the photon detection system prototype was determined to be

$$\epsilon = 0.0048 \times [0.29\exp(-x/4.3) + 0.71\exp(-x/225)] \quad (8.7)$$

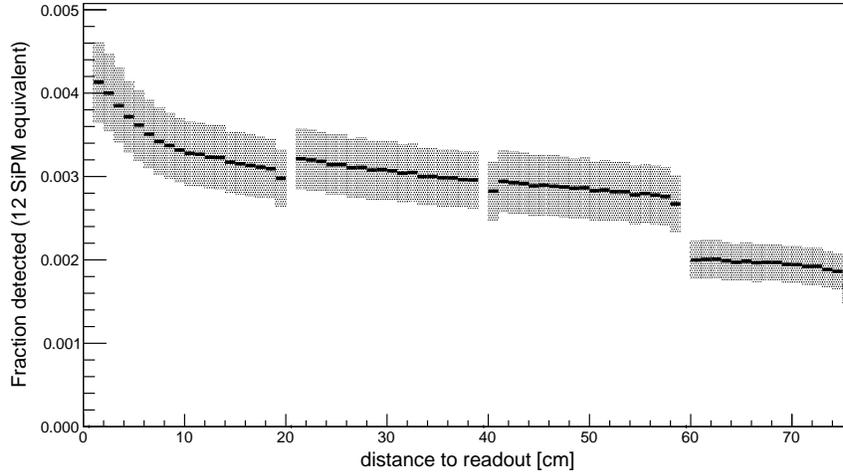


Figure 8.15: The prediction of the efficiency of the prototype double-shift lightguide bar PD system, as a function of distance to the readout end at which an initial VUV photon hits the wavelength-shifting plate. The jumps between section of the fraction of photons detected are due to the four different wavelength-shifting plates with their four individual efficiencies, as shown in Table 8.2.

Averaged over the length of a 76.2cm prototype, the efficiency is  $3.0 \times 10^{-3}$ . Instead taking an approximate full-scale of 2m, the average efficiency is estimated to be  $2.2 \times 10^{-3}$ .

### 8.3 Integrated prototype tests for research and development

As has been alluded to multiple times in the discussions above, an experimental program was conducted using the LAr test facilities at the Proton Assembly Building (PAB) at Fermilab, with the “Blanche” dewar facility. This section will detail the experiment and analysis of data taken to provide an estimate of the photon detection system efficiency. This work has been reported in Reference [102]: the description will largely mirror the description there, with gaps or abbreviated descriptions often pointing there for more detail.

#### 8.3.1 Experimental setup and operation

For the tests in Blanche, two double-shift PD prototypes (PD1 and PD2) were mounted in a frame suspended from the dewar lid, each  $\sim 76$ cm in length. Both sides of PD1 and PD2 had four TPB-coated plates placed in holders along the length of the lightguide bars, a total of 16 plates. Two hodoscope arrays borrowed from the CREST experiment (see for example Reference [110]) were placed on either side of the Blanche dewar to trigger on through-going cosmic rays. Each array is a set of barium-fluoride crystals each monitored by a PMT arranged in an  $8 \times 8$  pattern. Two

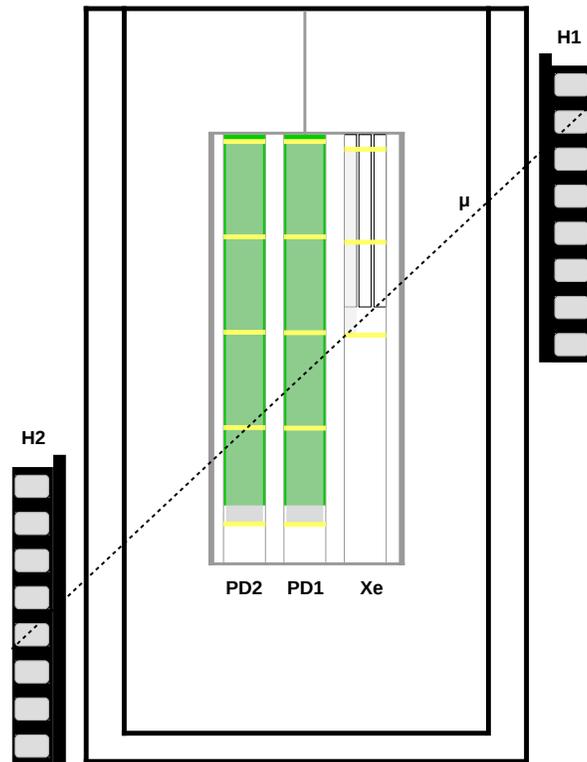


Figure 8.16: A representation of the experimental setup as used in the experiment at Blanche. PD1 and PD2 are prototypes of the double-shift style photon detector this chapter has focused on. The portion labeled “Xe” is a set of shorter dipped-bar style photon detectors for use in a Xe-doped LAr test taking place after the efficiency test detailed here. H1 and H2 are the hodoscope arrays used to trigger on through-going cosmic rays: a representative muon track hitting a detector in each array is drawn on the figure.

scintillator paddles were placed behind each array to provide an additional rejection to extraneous (non-muon) activity, specifically targeting reduced sensitivity to photons. Requiring signals from both hodoscope arrays and sets of scintillator paddles forms a four-fold coincidence trigger for through-going cosmic-ray muons. The scintillation signals were read out using an SSP for hodoscope triggers. Single track candidate events are selected from this sample by requiring only one PMT in each hodoscope array to have fired along with the corresponding scintillator paddles. The setup for the Blanche experiment is depicted in Figure 8.16.

The data collection for the prototype efficiency test in Blanche took place between September 14 and September 27, 2016. Our estimate is that the typical volume of LAr in Blanche over the

course of the experiment was 570 liters. After the setup was completed, the dewar was evacuated by a turbo pump and back-filled with Ar gas to reduce the initial levels of contamination from air and other residual gases. Filtration on the input line removes much of the remaining O<sub>2</sub> and H<sub>2</sub>O contamination from the ultra-high-purity LAr used for the fill, and contamination monitors allow further monitoring. O<sub>2</sub>, N<sub>2</sub>, and H<sub>2</sub>O monitors were connected to Blanche for various periods over the course of the experiment. Concentrations of O<sub>2</sub> contaminations were found to be ~30ppb. H<sub>2</sub>O contamination was found to be ~5ppb or less. The N<sub>2</sub> monitor provided readings of ~80-90ppb. All of these contamination levels are found to be within acceptable levels for scintillation studies, as discussed in Chapter 7.

The dewar was held at an internal pressure of 8 psig to attempt to keep outside contamination from entering the system during the experiment. A liquid nitrogen condenser helped to keep the LAr level consistent by taking gaseous Ar in the space above the liquid and condensing it. However, there was still some boil off, and occasional top-offs with fresh LAr were necessary. The level occasionally dipped below the top of the mounting frame before a top-off, but data taken at varying liquid levels suggests the data was consistent across the experimental period.

One half of the dewar was referred to as the “front,” with the other being called the “back.” By selecting specifically single track candidate triggers corresponding to the “front” hodoscope PMTs or the “back” hodoscope PMTs, tracks on each side of the photon detectors could be chosen. In principle this meant that there would be 4 data sets (front tracks for PD1 and PD2, and back tracks for PD1 and PD2). However, data from both sides of PD2 were excluded for a number of reasons. One important reason is that the described attenuation studies with an  $\alpha$  source on this bar provided short ( $\lesssim 1\text{m}$ ) attenuation lengths. Data from the back of PD1 were further excluded, as a plate fell completely out of the holding structure. This left the data taken for tracks on the front of PD1, in the best condition of the four, which form the basis of the analysis discussed here. These four plates are those in Table 8.2 with the efficiencies.

### 8.3.2 Data and analysis

Calibrating the SiPM readout is performed in a manner similar to the description of looking at PE peaks provided at various points in this document. The SSPs were run in a “free run” mode, in which signals above a given threshold are collected and analyzed with minimal cuts. Since each SiPM was read out separately for this test, a large number of free run triggers were collected from each SiPM and histogrammed for ADC counts. The separation between PE peaks then give a calibration constant to convert summed ADC (counts) to PE for each SiPM.

The scintillation detected by eight SiPMs on PD1 are analyzed for hodoscope-triggered events. For each trigger, each SiPM waveform is pedestal subtracted and summed from  $0.7\mu\text{s}$  before the trigger to  $11\mu\text{s}$  after, giving a summed ADC for the scintillation readout. The calibration is then applied, giving a summed PE detected by each SiPM for a given trigger. Adding together these 8 PE counts then gives the total number  $N_{det}$  of PE detected by PD1 for a given trigger.

Track-like trigger candidates from the front of PD1 were preselected, and two further cuts were added to reduce the effect of spurious triggers with no actual traversing particle in the dewar. First, both PD1 and PD2 were required to collect at least 10PE for a given event. The set of all preselected triggers for PD1 is given in Figure 8.17: the signal peak is well above 10PE. Secondly, events with relatively high energy deposition in the hodoscope were cut: this was made to exclude events less likely to have been triggered by minimum ionizing muons. Resulting distributions of signal for the remaining and removed samples are visible in Figure 8.17. Rejected events appear consistent with a background (a long-tailed distribution peaking at or near 0 PE). In total, 11,223 tracks passed the selection cuts.

A ray-tracing Monte Carlo simulation for scintillation from tracks in a dewar was originally developed by Denver Whittington and Johnathon (Lowery) Jordan of IU and used in prior testing with the LAr facilities. The track-trajectory is given based on the hit PMTs in the hodoscopes, and the assumption of minimum ionizing muons depositing 40,000 photons/MeV was used. This gives the number of photons for a given track trigger, and photons are thrown from this line segment and traced around the dewar. The dewar geometry is estimated to be a cylinder, which is largely true

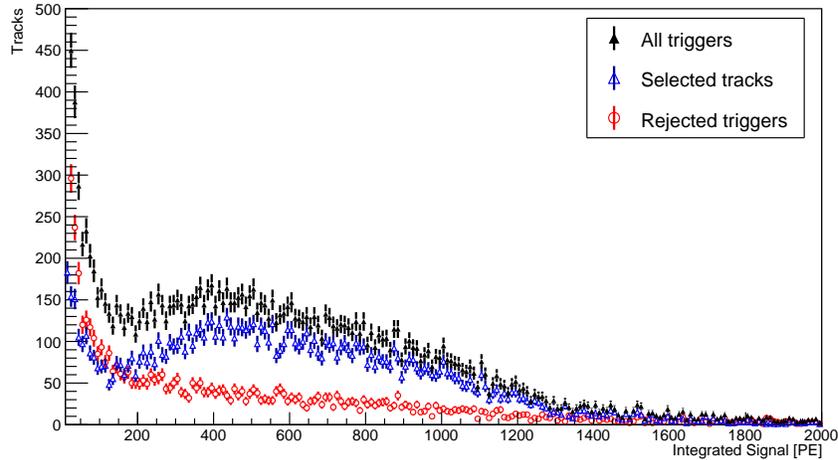


Figure 8.17: Integrated signals (in photoelectrons, PE) for all preselected track-like triggers in the Blanche experiment for PD1. The blue, open triangles show the integrated signals for triggered events passing further cuts added to reject spurious triggers. Rejected events are shown in the red, open circles.

but neglects the bowl shape at the bottom, below the frame. Scintillation photons can be lost by leaving the LAr volume or by being absorbed rather than reflected: the reflection coefficients<sup>3</sup> are taken from work reported by ICARUS [111]. The end result of the simulation is a distribution of photons falling along the PD system, reported in a histogram of photons versus position along the system.

For this experiment, the code was updated and made to include knowledge of photons hitting the front side or back side. Two major additional features were added to this simulation: Rayleigh scattering and photoabsorption. The probability for a photon to undergo Rayleigh scattering is based on the characteristic scattering length for a given medium at a given wavelength. For LAr scintillation, this characteristic length is not settled, so the simulation was run with three different lengths from two experimental groups [112, 113, 114], with the median length taken as nominal for this analysis (1.1m [112]) and others (0.66m [113] and 1.63m [114]) used to determine uncertainties. A random number drawn from an exponential with the characteristic length was chosen, giving the path length at which a given photon would scatter. The procedure handling scattering

<sup>3</sup>25% reflectance, half specular/half diffuse

was adopted from Geant4's G4OpRayleigh routine<sup>4</sup> with Geant4-specific commands converted to ROOT commands. Photoabsorption lengths were chosen in the same way as Rayleigh scattering. A characteristic absorption length for each of the 3 discussed contaminants was determined using

$$\lambda_{abs} = 1/n_c\sigma = 1/[\chi n_{LAr}\sigma] \quad (8.8)$$

with  $n_c$  being the contaminant number density,  $\chi$  being the contaminant fraction,  $n_{LAr}$  being the number density of LAr atoms, and  $\sigma$  being the absorption cross-section for 128nm photons [102]. The characteristic length at 5ppb for H<sub>2</sub>O, 11.9m, provided the shortest characteristic length and was chosen for this process, as described in Reference [102]. The simulation calculates the VUV scintillation photons falling on the plates. Secondary processes happening to photons after they fall on the plates are not considered. For example, reflected  $\sim 430$ nm photons from the plate conversion hitting exposed bar (gaps or where the plate fell off on the back) or hitting plates and undergoing re-absorption/re-emission are neglected. For each selected track, the number of expected photons at the readout ( $N_{exp}$ ) is computed by multiplying each distance bin in the simulation histogram by the transport function (Equation 8.6).

### 8.3.3 Results

While  $N_{exp}$  describes the number of photons one would expect to reach the readout from a given track, the number of detected PE described above ( $N_{det}$ ) represents the number of photons that actually were detected at the readout<sup>5</sup>. Therefore, the fraction  $N_{det}/N_{exp}$  characterizes the efficiency of the prototype at the readout end. For each selected track in the PD1 front dataset discussed in Section 8.3.2, this ratio is calculated.

A few important errors to consider with this ratio are statistics and calibration. The statistical errors are a few tenths of a percent, and the error in ADC-to-PE calibration as used in  $N_{det}$  was found to be  $\approx 3.5\%$  by averaging the standard deviations of peak separations for the SiPMs in the calibration data. To handle this uncertainty, a Monte Carlo method was adopted. The ratio was

---

<sup>4</sup>As implemented in version 4.10.03

<sup>5</sup>The number of PE detected is not exactly the number of photons detected due mostly to cross-talk, but we will apply a correction factor to the final number to account for this.

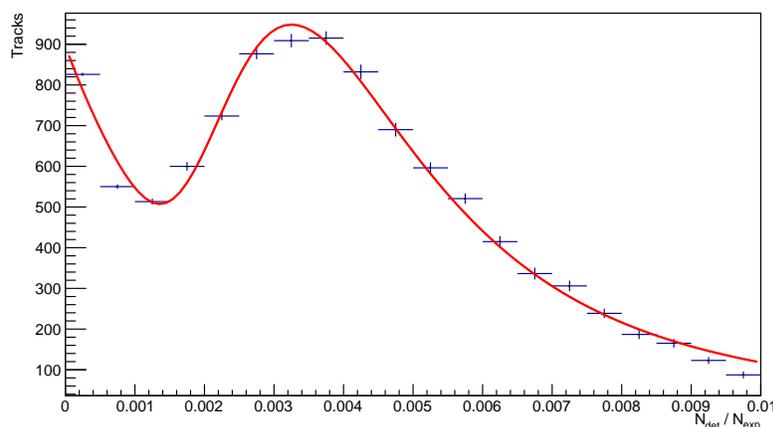


Figure 8.18: Photon detector efficiency is characterized by the ratio of number of detected PE to the expected number of photons, with corrections applied as discussed in the text. This figure shows this ratio for PD1 data and simulation, averaged over 100 trials using the Monte Carlo method as discussed in the text. The most probable value is found to be  $N_{det}/N_{exp} = 0.0035 \pm 0.3\%PE/photon$ . The y-axis is somewhat zero-suppressed and shows the main region of interest in more detail.

filled for each track into 100 separate histograms each taking a random number pulled from the normal distribution based on the quadratic sum of these errors. The 100 histograms were then averaged to produce a final  $N_{det}/N_{exp}$  curve as shown in Figure 8.18, where the errors are equal to the standard deviation of contents in a given bin across the 100 histograms. Each of the 100 histograms was fit in ROOT with a model that combines an exponential background term with a Landau function to characterize the cosmic-ray muon signal. The most probable value (MPV) for the signal is the parameter of interest, found by averaging the peaks of the Landau in the fit for the 100 trials:  $0.0035 \pm 0.3\%PE/photon$ . The average fit function is the red line in Figure 8.18.

Further major sources of uncertainty depend on the scintillation yield, Rayleigh scattering length, and the reflectivity of the steel. The uncertainty in the scintillation yield is 10% [99], and since this is a direct scale factor for  $N_{exp}$  this is taken as is. For the other two uncertainties, a similar analysis was performed as above with altered parameters but with some simplifications. This value is then compared to the nominal to provide a relative uncertainty. For the reflectivity of the steel to the scintillation photons, a conservative approach was taken: the nominal reflection probability was varied by 50% each way along with testing the effects of making all reflections specular

or diffuse. The uncertainty was estimated by considering these variations. The set of considered systematic uncertainties are summarized in Table 8.4.

Systematic	Uncertainty
Scintillation yield	$\pm 10\%$
Rayleigh scattering	$+2.3\%, -3.1\%$
Average MPV from fits	$\pm 0.3\%$
Scintillation reflection	$\pm 7\%$

Table 8.4: Summary of systematic uncertainties on the efficiency calculated in the Blanche analysis.

To convert  $N_{det}/N_{exp}$  to an efficiency for a full DUNE model (photon/photon), two corrections are made. First, the Blanche measurement was made with 8 SiPMs, while a full-scale DUNE model is supposed to use 12 SiPMs across the end of the light guide. Therefore, we apply a factor 12/8, giving  $N_{det}/N_{exp} = 0.0053\text{PE/photon}$ . Finally, a correction is made to convert the PE read out to a number of photons detected. Since the afterpulsing rate is low, we neglect afterpulsing and correct for the cross-talk. Prescriptions for understanding cross-talk adding up to 4 additional PE from a single photon [115] use the overall cross-talk probability, which we gave in Table 8.1, and the geometry of the sensor. Though it is possible to have more than 4 additional PE from cross-talk for a given detected photon, the range considered covers  $>99.7\%$  of the probability, so this taken to be adequate. A Monte Carlo calculation uses the probabilities from the prescription over many individual photon throws and arrives at 1.25-1.26PE/photon for multiple models. Therefore, we convert the ratio  $N_{det}/N_{exp}$  to photons detected per photon expected, and arrive at

$$\epsilon_{\text{readout}} = N_{det}/N_{exp} = 0.0041$$

The efficiency from the Blanche experiment is given in Figure 8.19, where the efficiency is 0.0041 at  $x=0$  and uses the transport function (Equation 8.6) to describe the fall-off, with gaps added to match the plate holders as in the simulation. Errors correspond to those given in Table 8.4.

A weighted average of the plate efficiencies in Table 8.2 was calculated based on the incident photons in the Blanche simulation for analyzed tracks. This efficiency was then taken as the plate efficiency for all plates in the photon detector simulation. The resulting expectation of efficiency

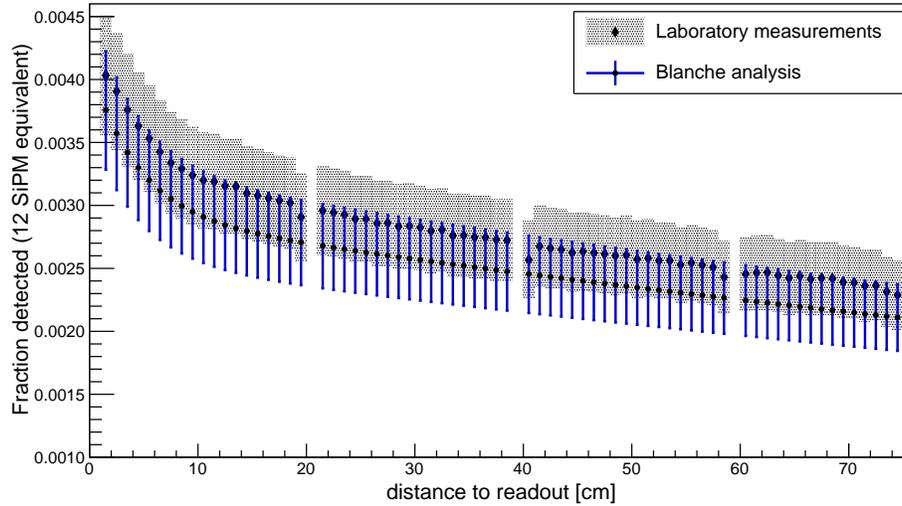


Figure 8.19: Comparison of photon detection efficiency as a function of distance from the readout end for a simulation as discussed in Section 8.2.3 and the Blanche test using PD1. The Blanche result uses the transport function from the simulation, so the shapes should be very similar, but the overall scales suggest consistency between the two determinations of efficiency. The Blanche result assumes minimum ionizing energy deposit at 40,000 photons/MeV.

from the lab measurements and simulation is given in gray in Figure 8.19. The comparison in this figure suggests that the determination of efficiency from laboratory measurements and from the prototype test in Blanche are in reasonable agreement. Taken one step further, the free parameter that would scale the Blanche result up or down is the overall scintillation yield: the agreement of these two methods suggests that the Blanche measurement is consistent with this scintillation yield.

## CHAPTER 9

### Toward a precision era in neutrino physics

Chapters 4 and 6 discussed recent analysis of neutrino and antineutrino oscillation with NOvA. This chapter will examine potential future improvements to this analysis by NOvA and ultimately in the case of DUNE, beyond the analysis techniques already in place and demonstrated. Specific focus will be given to the treatment of wrong-sign background in the context of  $\nu_e$ ,  $\bar{\nu}_e$  appearance analysis. After NOvA, DUNE will perform a similar analysis, with many improved features leading to precision measurements, as noted in Chapter 7.

While the focus of this chapter is on oscillations of beam neutrinos, it should be noted that measurements using atmospheric neutrinos can enhance the overall sensitivity of NOvA and DUNE, adding an additional sample yielding even more interactions to study with the Far Detector (FD), with different baselines and potentially different energy ranges. DUNE especially has the advantage of being larger and underground. In both cases, the flux is not filtered through current-bearing horns to choose  $\bar{\nu}$  or  $\nu$ . Here, the ability to separate into samples that are  $\nu$  and  $\bar{\nu}$  enhanced can provide additional sensitivity in the oscillation fit, especially to parameters such as the CP violating phase that underscore  $\nu$ ,  $\bar{\nu}$  differences. Super-Kamiokande's recent atmospheric neutrino analysis indeed utilizes a statistical separation technique [30]. Chapter 5 provides an investigation of wrong-sign ( $\nu$ ) in the NOvA  $\bar{\nu}$  beam and gives a number of tools used to create samples enhanced in wrong-sign and right-sign. Atmospheric neutrino oscillation analysis in DUNE or NOvA could explore the tools provided in Chapter 5 in the case of NOvA or Section 9.2 from this chapter in the case of DUNE.

The focus of this chapter is to determine whether statistical separation would be useful in the case of 3-flavor oscillation studies using the NOvA and DUNE beams. The  $\nu_\mu$  and  $\bar{\nu}_\mu$  will oscillate to produce appearing  $\nu_e$  and  $\bar{\nu}_e$  in the FD, and this is what we refer to as the  $\nu_e$ -analysis appearance signal. The overall rates of these oscillations inform  $\theta_{13}$  and  $\theta_{23}$ , and relative differences between

the  $\nu_\mu \rightarrow \nu_e$  and  $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$  processes inform mass hierarchy and CP-violation measurements. This was described earlier by using a bi-probability or bi-event plot (Figure 1.3). The wrong-sign contamination of the antineutrino beam therefore presents a background that weakens the sensitivity to measurements of parameters, especially mass hierarchy and  $\delta_{CP}$ . Figure 9.1 shows such a bi-event plot with two sets of  $\theta_{23}$  and the normal (NH) and inverted (IH) hierarchies in NOvA.

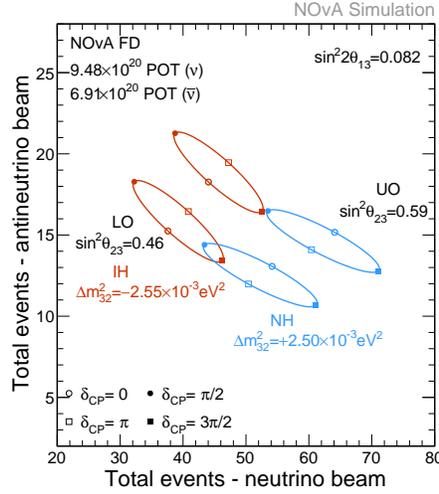


Figure 9.1: Bi-event plot for  $\nu_e$ ,  $\bar{\nu}_e$  appearance analysis in NOvA [116]. This is taken from the initial study with less antineutrino beam data, before the top-up, as the chosen  $\theta_{23}$  values are used in the story told here. Here,  $\theta_{13}$  is taken from reactor constraints. DUNE will ultimately aim to measure this parameter as well.

## 9.1 Oscillation analysis in NOvA and wrong-sign separation

First, consider the impact that the wrong  $\nu$  fraction in a  $\bar{\nu}$  beam can have. Note that that the following two sets of points have nearly identical  $\nu_e$  appearance candidates but differ primarily in the number of  $\bar{\nu}_e$  appearance candidates: ( $\sin^2 \theta_{23} = 0.59$ ,  $\delta_{CP} = \pi/2$ , NH) and ( $\sin^2 \theta_{23} = 0.46$ ,  $\delta_{CP} = 0$ , NH). Let us consider the case where the wrong-sign  $\nu_\mu$  (and therefore appearing  $\nu_e$ ) in the beam is uncertain. Nominally, the wrong-sign contamination in the neutrino beam (in this case  $\bar{\nu}_\mu$  in the  $\nu_\mu$  beam) is only a few percent, so even a large relative error in this fraction would not necessarily skew the result. However, the fraction of  $\nu_\mu$  in the  $\bar{\nu}_\mu$  beam is closer to 15%. Using selected events in the Near Detector (ND) and variations of simulated flux and cross-section parameters within estimated uncertainties, the relative uncertainties are estimated to be approximately 10%. For the

sake of argument and to illustrate the story, consider a case where a very large deviation from expectations in the wrong-sign appearance background is possible.

For simplicity, in this chapter I use unextrapolated predictions for FD neutrino energy spectra, where no correction to the FD signals is made for ND observations. The considerations of this chapter are not to provide a precise *data* analysis result for NOvA, where the ND data measurements provide important corrections before looking at the FD data. Since these are exploratory tests making generally relative comparisons of samples, the central-value tune-weighted simulation is therefore deemed sufficient. Furthermore, in DUNE no oscillation data or ND data yet exist. FD predictions for NOvA in the case where  $\sin^2 \theta_{23} = 0.59$ ,  $\delta_{CP} = \pi/2$  in the NH are shown for the neutrino and antineutrino beams in Figure 9.2. The two sets of bins reflect the two main selections, the lower PID region and the high PID region: the extra peripheral sample is neglected in this chapter. Each set of bins covers 0-4.5GeV in bins of 0.5GeV. Differences in the total counts with respect to the bi-event plot will mostly be from leaving out the peripheral bin, with some change also from not correcting for ND data. Higher exposure will lead to increased event yields. Altering the wrong-sign fraction to 50% nominal produces the  $\bar{\nu}_e$  candidate prediction in the left of Figure 9.3. Note that the oscillated  $\bar{\nu}_e$  signal remains the same, but the overall number of selected events (candidates) is lower. If instead the oscillation parameters are taken to be  $\sin^2 \theta_{23} = 0.46$ ,  $\delta_{CP} = 0$  in the NH but with 4 times the nominal amount of wrong-sign events, then the overall number of selected events is almost the same. The right plot in Figure 9.3 shows this example. Although the oscillation parameters are vastly different (including both having and not having CP violation) and have different expected  $\bar{\nu}_e$  appearance signal yields, one finds the same number of total candidates.

Since there are different amounts of right-sign and wrong-sign in each case, forming a right-sign-enhanced and wrong-sign-enhanced FD prediction similar to the ND work presented in Chapter 5 may break this degeneracy. The wrong-sign  $\nu_e$  BDT, as discussed in Chapter 5, was reused with new thresholds for FD studies picked using the FD simulation in a similar manner to that used in the aforementioned chapter. Indeed, using the wrong-sign  $\nu_e$  BDT (Figure 9.4) shows that the two cases then produce different event counts in the regions. It is important to remember that the

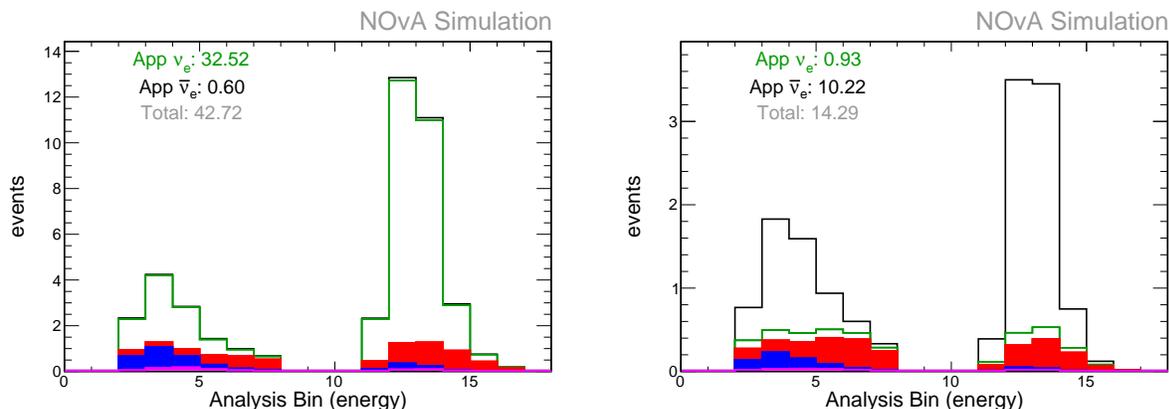


Figure 9.2: Unextrapolated FD predictions with  $\sin^2 \theta_{23} = 0.59$ ,  $\delta_{CP} = \pi/2$  in the NH for the neutrino beam (left) and antineutrino beam (right). Differences in the total counts with respect to the bi-event plot will mostly be from leaving out the peripheral bin, with some change also from not correcting for ND data. Higher exposure will lead to increased event yields.

wrong-sign is not expected to be off by such large factors, the bi-event plot is not how the analysis is conducted, and a full fit of spectra may be able to pick up on some shape difference, but this is meant to illustrate a potential impact.

Secondly, we test separation for its ability to improve the analysis. For example, creating a  $\nu$  enhanced sample leaves a  $\bar{\nu}$  enhanced sample that ideally provides a lower background measurement of the antineutrino oscillation. For NOvA, the full analysis and measurements of oscillation parameters are done by conducting fits to extract parameters, nominally using Common Analysis Format (CAF) files and CAF Analysis (CAFAna) framework.

In CAFAna, one creates “experiments” from created predictions of spectra at the FD or from parameter constraints where desired, for example if the parameter has been better constrained elsewhere than is possible with NOvA. Multiple of these are combined together to form a “multi-experiment” with which one can then perform fits and draw sensitivity slices and/or contours. For the studies performed here, only  $\nu_e$  and  $\bar{\nu}_e$  appearance is considered, ignoring the  $\nu_\mu$  and  $\bar{\nu}_\mu$  disappearance. All cases use a neutrino beam fake data sample from simulation with the lower and high PID selection. From the antineutrino beam simulation, fake data spectra are made with the nominal lower and high PID selection, but separately also with selection broken into right-sign-enhanced and wrong-sign-enhanced groupings with 4 sets of bins instead of 2. Examples of such a prediction

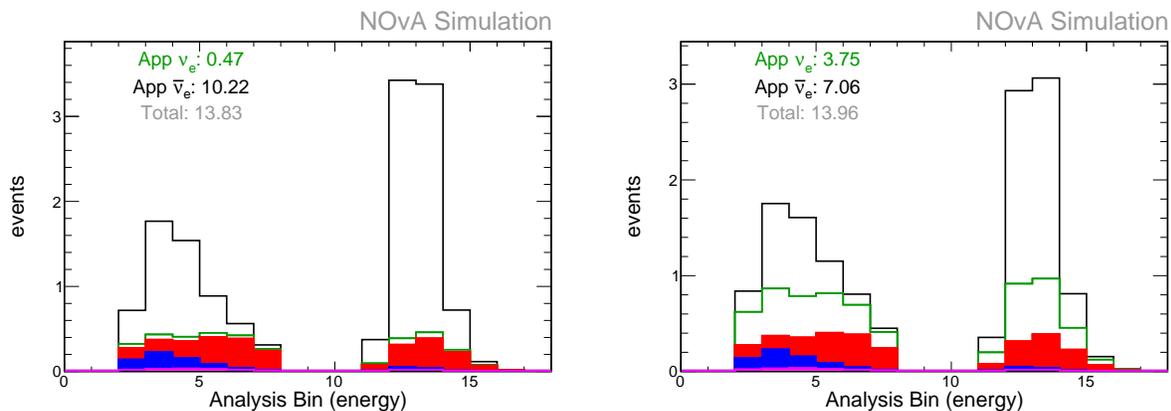


Figure 9.3: Unextrapolated FD predictions with  $\sin^2 \theta_{23} = 0.59$ ,  $\delta_{CP} = \pi/2$  in the NH for the antineutrino beam as in Figure 9.2 but with half the  $\nu_e$  appearance (left) and with  $\sin^2 \theta_{23} = 0.46$ ,  $\delta_{CP} = 0$  in the NH but with 4 times the nominal  $\nu_e$  appearance (right). Note that even though the oscillation parameters are very different - one conserves CP and one has maximal CP-violation - the total selected event counts would be the same. As discussed in the text, this is purely hypothetical, but illustrates the story that wrong-sign can make the measurement more complicated.

are given in Figure 9.4, though unlike in that case the wrong-sign weighting is unaltered. Since this only considers  $\nu_e$  and  $\bar{\nu}_e$ , which has very little sensitivity to the  $\Delta m_{32}^2$  parameter, a constraint is taken:  $(2.45 \pm 0.05) \times 10^{-3} \text{eV}^2$  (NH) and  $(2.52 \pm 0.05) \times 10^{-3} \text{eV}^2$  (IH) [117]. Additionally, as noted previously, reactor experiments have strong constraints on  $\theta_{13}$ , so  $\sin^2 2\theta_{13} = 0.82 \pm 0.004$  [117] is added for this. The likelihood function in the fit and aspects of selection and analysis are discussed in Section 4.1.

For the fake data used in the fit, I use the neutrino beam and antineutrino beam predictions (with or without wrong-sign splitting as needed) and generate a spectrum with a set of oscillation parameters. The default oscillation calculator is set to  $\theta_{23} = \pi/4$ ,  $\delta_{CP} = 0$ ,  $\Delta m_{32}^2 = 2.45 \times 10^{-3}$ . From there, the parameters can be altered as desired to test various constraints. The experiments are fit by defining a set of parameters to fit over and the seeds to test to fit over. For the fits described here, the following parameters are fit:  $\delta_{CP}$ ,  $\sin^2 2\theta_{13}$ ,  $\Delta m_{32}^2$ , and  $\sin^2 \theta_{23}$ . For one,  $\Delta m_{32}^2$  is seeded at  $-2.5 \times 10^{-3}$  for fits in the IH and  $2.5 \times 10^{-3}$  for fits in the NH. In investigating a slice on a variable, the profile is performed and then the oscillation calculator set such that  $\Delta m_{32}^2 = -\Delta m_{32}^2$ , and the profile repeated to explore the other hierarchy. Furthermore,  $\sin^2 \theta_{23}$  is seeded at 0.45, 0.5, and 0.55 to allow it to search for the best octant in the fit. In official NOvA analysis plots, it is

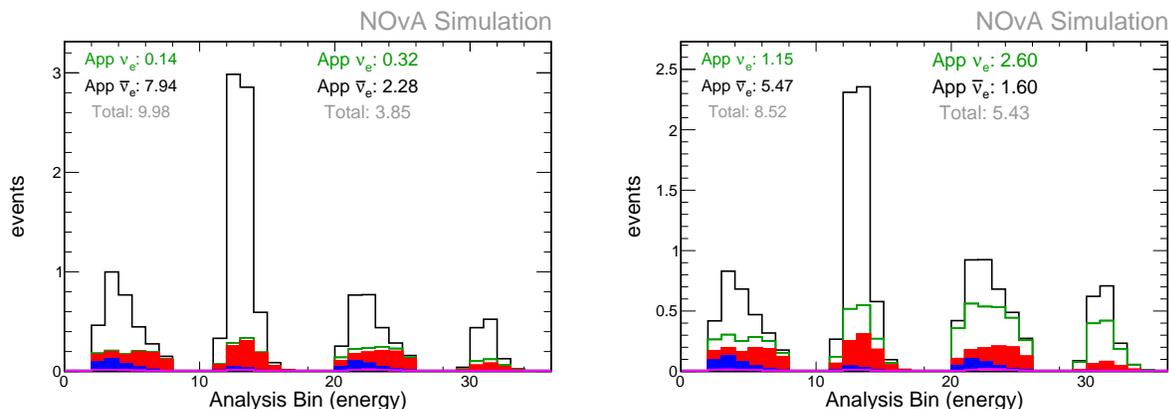


Figure 9.4: Applying the  $\nu_e$  BDT creates separate right-sign and wrong-sign enhanced regions. (left)  $\sin^2 \theta_{23} = 0.59$ ,  $\delta_{CP} = \pi/2$  in the NH for the antineutrino beam but with half the  $\nu_e$  and (right)  $\sin^2 \theta_{23} = 0.46$ ,  $\delta_{CP} = 0$  in the NH but with 4 times the nominal  $\nu_e$ . Within each plot, the leftmost two curves are the lower and high PID right-sign enhanced selection and the rightmost two bins are the lower and high PID wrong-sign enhanced selection, respectively. Note that while the total event counts are almost the same, the selected events in the wrong- and right-sign enhanced regions differ for these two cases.

not uncommon to see a slice with the possible octants overlaid, but in this section I just show each hierarchy once. Finally  $\delta_{CP}$  is seeded at  $0\pi$ ,  $0.5\pi$ ,  $1.0\pi$ , and  $1.5\pi$ . A fit is performed to obtain the minimal  $\chi^2$  and then slices or contours can be made. These are made by scanning over one variable to map out sensitivity in the case of a slice or two variables to map out likelihood surfaces in the case of contour. Other fit parameters are then profiled over, with seeds if given any. For most of this section, I show  $\delta_{CP}$  slices to detail sensitivity to the best fit in that parameter. The development of a fitting procedure and code was guided first by a colleague with who originally wanted to look at added sensitivity to the analysis from the wrong-sign methods we tested [118, 119] and then further by tutorials for the NOvA fit (for example [120, 121]) and related demonstration code and discussion with the authors.

As an example result of this fit, take the result from the recent analysis of NOvA neutrino and antineutrino beam data: the resulting  $\delta_{CP}$  slice is on the left side of Figure 9.5. The plot on the right was produced by the fitting method described here, given the true parameters of the best fit ( $\sin^2 \theta_{23} = 0.56$  and  $\delta_{CP} = 0$ ) and the data exposure in the analysis. The best fit  $\Delta m_{32}^2$  was not adopted for this plot, as  $2.48 \times 10^{-3}$  is very close to  $2.45 \times 10^{-3}$  and should have little effect on this

slice. The dashed curves denote different conditions in each plot and should not be compared. In the plot on the right, the IH sensitivity around  $\delta_{CP} = \pi/2$  is somewhat lower and the NH sensitivity appears to be slightly higher than on the left, but the overall shapes and qualitative results are very similar. One should not expect the curves to match exactly given the several differences in making the two. Among the differences, the fits I have run are not Feldman-Cousins corrected, though the official analysis results are. Furthermore, the fits incorporate statistical errors only and do not take into account systematics, nor do they use the peripheral bin. This simplified analysis reproduces the full analysis used for published NOvA results with sufficient validity to justify studies of the impact of wrong-sign/right-sign separation.

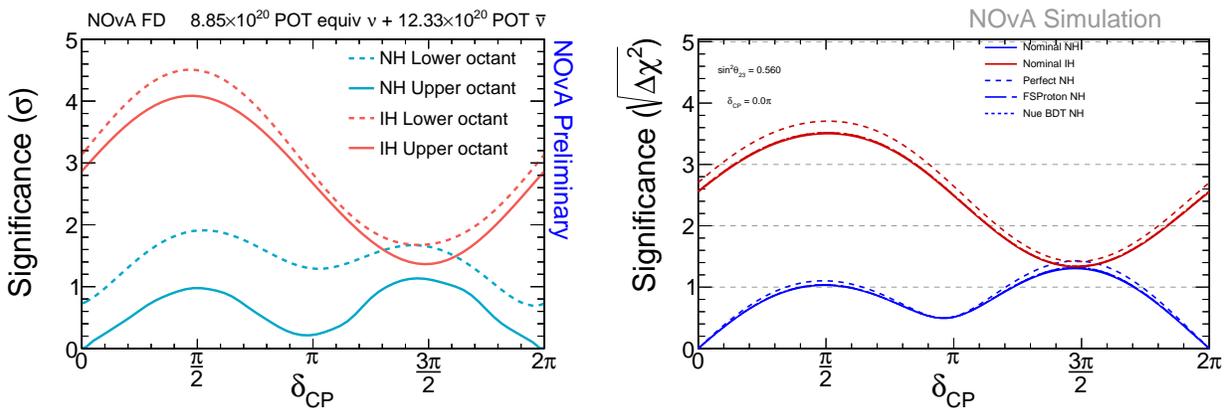


Figure 9.5: Left: the  $\delta_{CP}$  slice result from the recent NOvA analysis using neutrino beam and antineutrino beam data, including Feldman-Cousins corrections [79]. Right: a fit as described in the text with true  $\delta_{CP} = 0$  and  $\sin^2 \theta_{23} = 0.56$ . Recall that the fits produced for this chapter do not fit each octant but rather have one slice in the NH and one slice in the IH. The dashed curves denote different conditions in each plot and should not be compared. In the plot on the right, the IH sensitivity around  $\pi/2$  is somewhat lower and the NH sensitivity appears to be slightly higher than on the left, but the overall shapes and qualitative results are very similar. Differences could be due to many factors as described in the text.

A goal of this section is to then explore if the results could be improved by using wrong-sign separation techniques to produce 4 samples for the antineutrino beam appearance experiment (right-sign-enhanced lower PID, right-sign-enhanced high PID, wrong-sign-enhanced lower PID, wrong-sign-enhanced high PID). For this study, two of the tools developed in the context of Chapter 5 are used. The CVN final state proton score (FSProton) uses the output likelihood for the event

to have a proton in the final state based on CVN, the event level classifier neural network used in NOvA. Furthermore, the wrong-sign  $\nu_e$  BDT using input variables based on CVN information and shower/reconstruction information. To create right-sign-enhanced and wrong-sign-enhanced samples in the antineutrino beam selection, events were grouped based on falling above or below certain threshold values. Two wrong-sign separation methods discussed in Chapter 5 (BDT and CVN final-state proton score) were used in these studies, with thresholds picked for the FD simulation using similar methods to the one described in that chapter. From the studies of Chapter 5, the ratios of background beam  $\nu_e$  to beam  $\bar{\nu}_e$ , which form backgrounds in each of the split samples, are believed to be close to correctly modeled. This is important to build confidence in the background samples in each subsample of the split antineutrino beam data.

These separated antineutrino predictions and fake data were then also fit and compared to the sensitivity achieved in the nominal case in both hierarchies at the given oscillation parameters. These appear as dashed lines in the right of Figure 9.5, and another set of dashed lines uses perfect wrong-sign separation based on simulation truth information to produce the fake data and predictions. While systematics involved in such separation may lead to additional smearing and reduced sensitivity from what is ultimately shown in these plots, these statistics-only fits are instructive in showing whether such techniques could be beneficial. Given the exposure in hand, the current techniques do little to benefit the analysis if the true oscillation parameters are  $\delta_{CP} = 0$  and  $\sin^2 \theta_{23} = 0.56$ . However, perfect wrong-sign separation shows there is potential for advancement. The story is much the same with other values of  $\delta_{CP}$ . Figure 9.6 shows slices for true  $\delta_{CP} = 0, \pi/2, \pi, 3\pi/2$  at the recent analysis exposure for  $\sin^2 \theta_{23} = 0.56$ . Figure 9.7 shows the difference in  $\chi^2$  between the sensitivity curve using the wrong-sign BDT and the nominal, taken by squaring the  $\sqrt{\Delta\chi^2}$  values at each point and subtracting them. Though small, the gains are indeed non-zero.

We only pick one true set of oscillation parameters and exposures at a time, but the results could vary as a function of parameters and exposure. To handle this, fits and slices are made in a large loop over parameters and exposure beyond the values chosen above. The following permutations

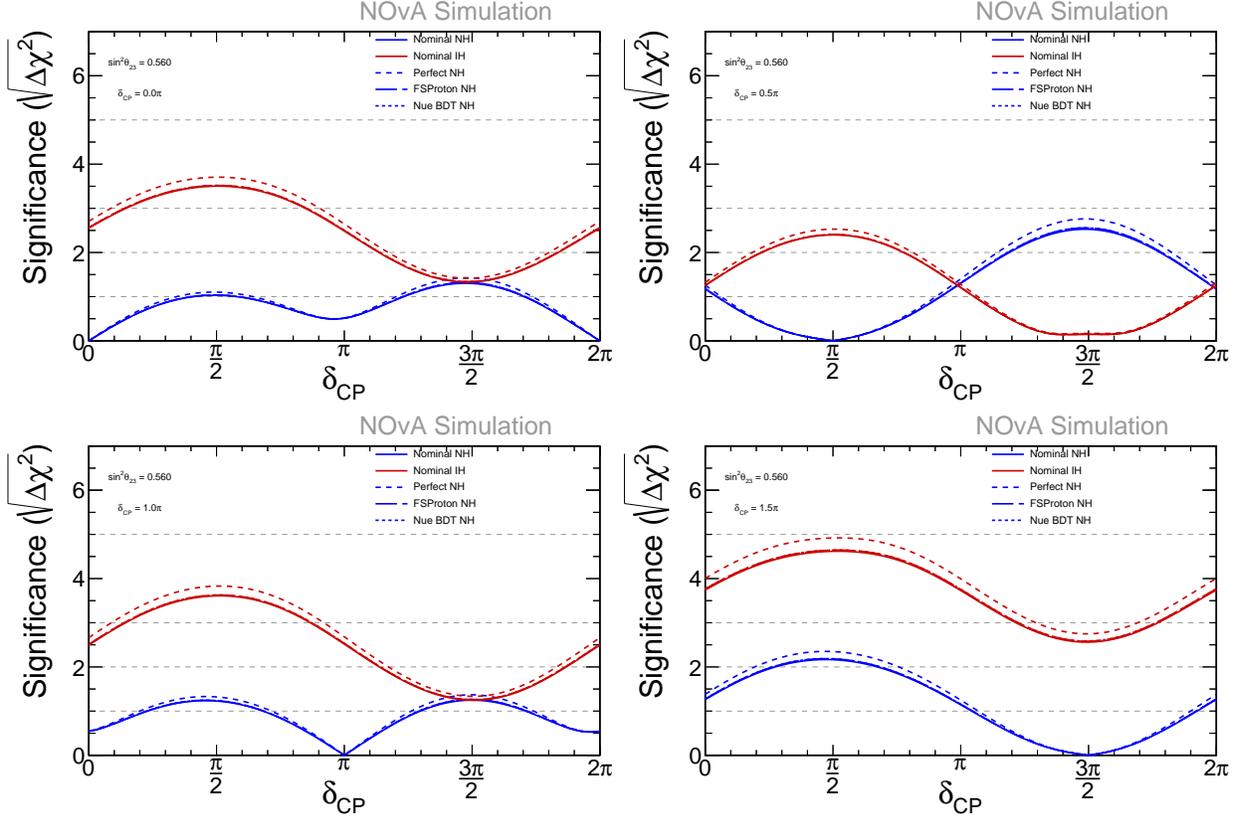


Figure 9.6: Slices for true  $\delta_{CP} = 0$  (top left),  $\pi/2$  (top right),  $\pi$  (bottom left),  $3\pi/2$  (bottom right) with  $\sin^2 \theta_{23} = 0.56$  at the recent analysis exposure for  $\sin^2 \theta_{23} = 0.56$ . At the current exposure (POT collected), the current techniques do little to benefit the analysis. However, perfect wrong-sign separation technique shows some potential for advancement.

of true parameters are considered:

1. In addition to the exposure levels of recent analysis, POT (both beams given this amount):  
 $15 \times 10^{20}, 20 \times 10^{20}, 25 \times 10^{20}, 30 \times 10^{20}, 36 \times 10^{20}$
2.  $\delta_{CP} = 0$  to  $2\pi$  in increments of  $0.05\pi$  (41 values)
3.  $\sin^2 \theta_{23} = 0.48, 0.5, 0.56$

where the recent analysis POT were  $9.48 \times 10^{20}$  (neutrino beam, not taking full-detector-equivalent) and  $12.33 \times 10^{20}$  (antineutrino beam). The values of  $\sin^2 \theta_{23}$  were chosen to be maximal mixing along with the points in the upper and lower octant where the recent analysis found a best fit. The projected total POT expected for NOvA is approximately  $36 \times 10^{20}$  in both beams by the time

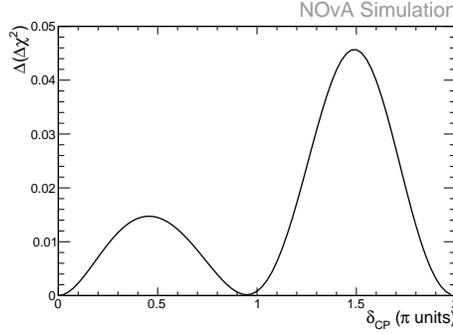


Figure 9.7: Difference in  $\chi^2$  between the sensitivity curve using the wrong-sign BDT and the nominal, taken by squaring the  $\sqrt{\Delta\chi^2}$  values at each point in the graph and subtracting them, for the recent analysis exposure with true  $\delta_{CP} = 0$  and  $\sin^2 \theta_{23} = 0.56$ .

the accelerator turns off to prepare for DUNE around the mid 2020s. This exercise is performed assuming NH as truth for convenience. Two sets of plots are useful in this procedure. The slices inform how an individual set of true parameters fares in terms of sensitivity. To collate this information and boil it into CP violation sensitivity, each slice is checked for the minimum sensitivity at  $\delta_{CP} = 0, \pi$  in each hierarchy, with the minimum of these said to be the sensitivity with which one might reject CP violation. Of course, if  $\delta_{CP} = 0, \pi$  then CP is conserved and not violated. A test was also run checking a single profile but seeding that profile in both hierarchies such that it found the minimum from both up front. Checking the plot appearing here for  $\sin^2 \theta_{23} = 0.5$  at  $36 \times 10^{20}$  POT in both beams showed that the CP conservation rejection plots find the same result.

Figure 9.8 combines the four most interesting  $\delta_{CP}$  points (conservation and maximal violation) in the upper octant ( $\sin^2 \theta_{23} = 0.56$ ) at the full projected exposure of  $36 \times 10^{20}$  POT in both beams. Appendix D shows further plots of the development of these curves with increasing POT and some plots with the same POT but differing  $\sin^2 \theta_{23}$ . One can see that the extra sensitivity gained by performing wrong-sign separation does increase with increased exposure - this is expected - but that the current forms of wrong-sign separation still do not provide much improvement<sup>1</sup>.

<sup>1</sup>One will note a few places where the minimization/profiling has inadvertently thrown discontinuous values, notably in the IH portion of a few of the plots, especially around  $\pi/2$  for the plots with true  $\delta_{CP} = 3\pi/2$ . It is important to note that this point is maximally incorrect: the neutrino vs. antineutrino expectations are on the opposite side of the biprobability plots. A look at debug level plots of the profile variables show that  $\sin^2 \theta_{23}$  has trouble profiling in this fit. While it gets close to the true value for  $\delta_{CP} = 3\pi/2$  in the NH, for the troubling point of  $\delta_{CP} = \pi/2$  in the IH,  $\sin^2 \theta_{23}$  gets pulled all the way to 1. This is the end of its physically allowed space. However, since this does not appear to be occurring at CP-conserving points for use in the sensitivity curves and since this is seemingly happening

The results of exploring these slices for the three  $\sin^2 \theta_{23}$  values at  $36 \times 10^{20}$  POT in both beams is summarized in Figure 9.9. These plots should be read as: “If the true  $\delta_{CP}$  is <x-value>, the sensitivity at the given exposure to reject CP conservation is <y-value>.” Again note that the current wrong-sign separation techniques do not provide much improvement to the analysis even at full planned exposure. However, the green dashed lines for perfect wrong-sign separation indicate there is some room for potential improvement to the ultimate sensitivity of NOvA.

The dip in the rejection sensitivity plots around  $\delta_{CP} = \pi/2$  is due to the degeneracy of CP-violation and mass hierarchy in NOvA. This is visible in the biprobability plots, where the NH and IH ovals nearly overlap at NH  $\delta_{CP} = \pi/2$  and IH  $\delta_{CP} = 3\pi/2$ . These two values are hard to distinguish from each other, and Figure 9.8 shows that even at larger exposures, one obtains a fairly low  $\Delta\chi^2$  at IH  $\delta_{CP} = 3\pi/2$  when the true parameters are NH and  $\delta_{CP} = \pi/2$ . Since the minimum sensitivity at  $\delta_{CP} = 0, \pi$  from both NH and IH are considered for a given data point and since the IH significances are a little lower, it therefore tends to pull the overall rejection sensitivity down at  $\pi/2$  in the NH. Further information is given in Appendix D.

---

in areas with already significant sensitivity, these are deemed okay for the exploratory purposes of this study.

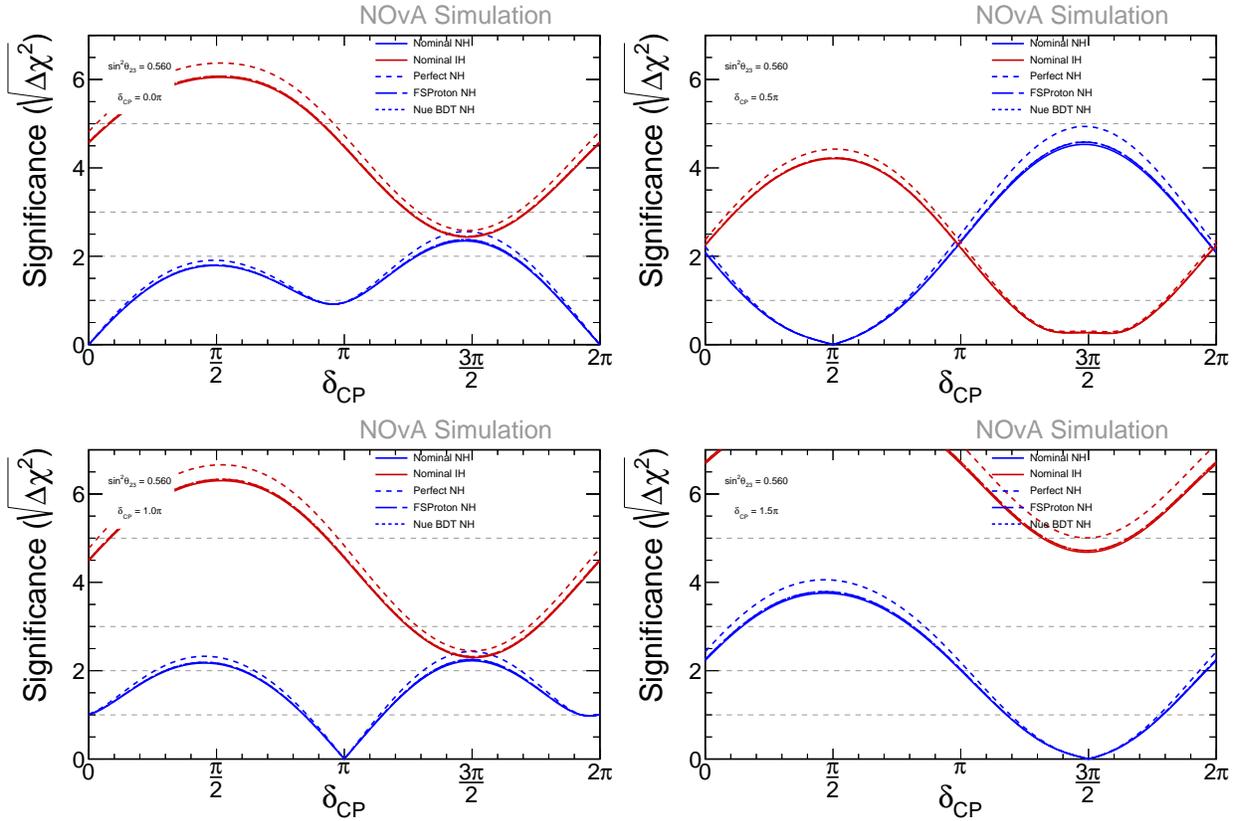


Figure 9.8: Slices for true  $\delta_{CP} = 0$  (top left),  $\pi/2$  (top right),  $\pi$  (bottom left),  $3\pi/2$  (bottom right) with  $\sin^2\theta_{23} = 0.56$  at a full  $36 \times 10^{20}$  POT in both beam modes. The current wrong-sign separation techniques still do little to benefit the analysis. However, perfect wrong-sign separation technique shows potential for advancement.

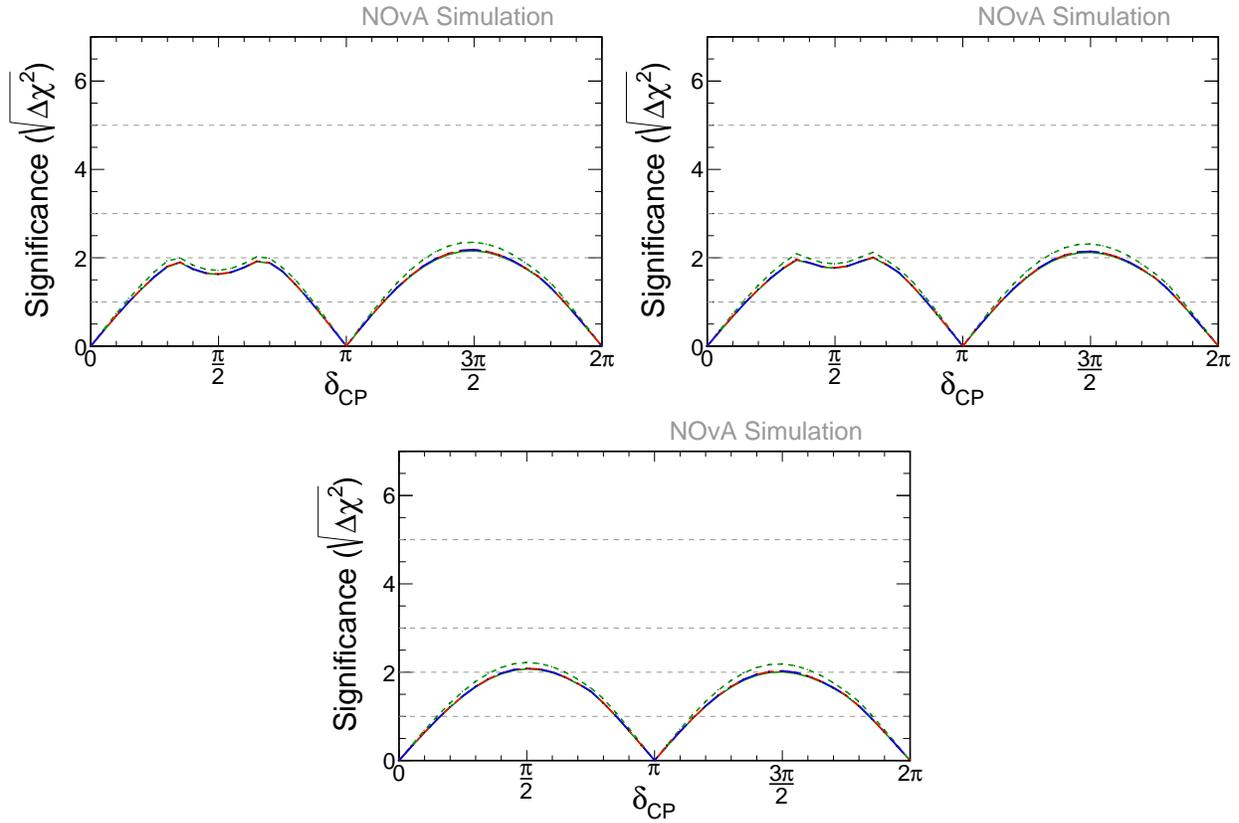


Figure 9.9: CP conservation rejection curves for the three considered values of  $\sin^2 \theta_{23}$  at  $36 \times 10^{20}$  POT in both beam modes (0.48 (top left), 0.5 (top right), 0.56 (bottom)). These plots should be read as: “If the true  $\delta_{CP}$  is  $\langle x\text{-value} \rangle$ , the sensitivity at the given exposure to reject CP conservation is  $\langle y\text{-value} \rangle$ .” Again note that the current wrong-sign separation techniques do not provide much improvement to the analysis even at full planned exposure. However, the green dashed lines for perfect wrong-sign separation indicate there is some room for potential improvement to the ultimate sensitivity of NOvA.

## 9.2 $\nu_e, \bar{\nu}_e$ appearance in DUNE

Having considered the idea of wrong-sign separation and future analysis in NOvA, now consider DUNE. The DUNE detectors were discussed in Chapter 7 and will utilize a fiducial mass of 40kt of liquid argon to detect neutrino oscillations, significantly larger than the 14kt total size of NOvA. Since the DUNE FD will be situated at 1300km, nearly 500km further than NOvA, the degeneracy between delta CP and the matter effect is broken. If one were to draw the biprobability plots for DUNE, this would have the effect of separating the ovals for the NH and IH, such that there would no longer be overlap. As noted in Chapter 7, DUNE will therefore measure the correct hierarchy with high significance. In this chapter, we therefore only show fits in DUNE for the NH, where we again assume the true hierarchy to be NH for our oscillation fake data.

Analysis methods for DUNE are still being developed and improved, but at present the analysis chain uses a similar CAFAna framework. This means that a study very similar to the one described above can be conducted with simulated neutrino interactions in DUNE as well. Since simulating a full 40kt of LAr is computationally challenging, these simulations were run with a small fiducial volume 2 TPCs wide, 2 TPCs tall, and 6 TPCs long. To properly account for this, a scale of 40/1.13 is applied to the exposure. Cuts on simulation truth information are used to check that the interaction is within the fiducial volume for oscillation studies. DUNE has a version of CVN that originally borrowed from NOvA but has since grown into its own network. The DUNE CVN has multiple outputs, three of which [122] are useful in this study:

1. The output used in the main analysis translates to relative scores for  $\nu_\mu, \nu_e, \nu_\tau, \text{NC}$
2. A binary score that returns either  $\nu$ -like or  $\bar{\nu}$ -like
3. A return that gives scores for 0 protons, 1 proton, 2 protons, or 3+ protons

where the latter two will be used as wrong-sign separation techniques to explore here<sup>2</sup>. In addition to the CVN PID cut (item 1 in the list above) and the fiducial volume cut, a cut is placed on neutrinos

---

<sup>2</sup>The exact details of the CVN are not so critical at the level of study described here. Follow up studies exploring systematics would consider model-dependence or any potential biases in the CVN methods. For now, the separation is performed based on the simulated score results, which are shown in Figures 9.11 and 9.12. That the shown variables provide reasonable separation is sufficient to be used in this study.

for reconstructed energy between 0 and 10GeV, as this is the analysis binning range chosen for oscillation studies. Events are weighted to reflect a set of cross-section weights. Predictions are made for these, as in NOvA, though the default oscillation calculator uses a slightly older but similar value for  $\theta_{13}$ . A set of nominal predictions is shown in Figure 9.10 for both the neutrino beam and antineutrino beam in DUNE. It has been normalized to events per 0.25GeV to preserve the spectrum shape though the bins are variable in size.

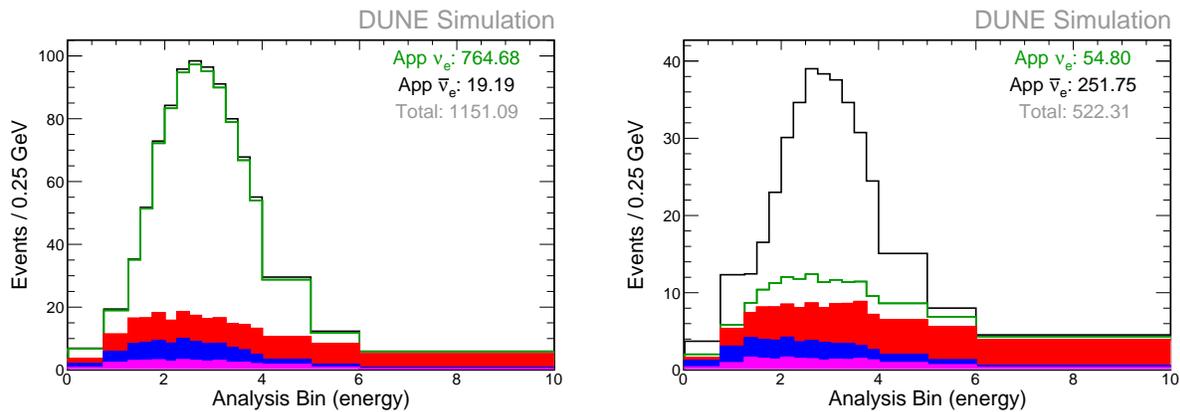


Figure 9.10: Unextrapolated FD predictions with some default oscillation parameters in the NH for the neutrino beam (left) and antineutrino beam (right) in DUNE.

Though the full answer depends on the oscillation parameters, there is enough information in this figure to see that the wrong-sign background is still negligible for the neutrino beam ( $\sim 1.7\%$  of total events) and considerably higher in the antineutrino beam ( $\sim 10.5\%$  of total events). Two preliminary wrong-sign separation techniques using the CVN scores can be explored in a manner similar to that in the last section. One is simply the binary output from the CVN, attempting to classify events as  $\nu$ -like or  $\bar{\nu}$ -like. Plots were made using the “nue-swap” FD files, wherein the FD flux was swapped to have predominantly  $\bar{\nu}_e$  instead of  $\bar{\nu}_\mu$ . A default oscillation calculator calculates the oscillation probability at the given energy and is applied as a weight: in essence this sculpts the flux to roughly approximate the oscillation signature. Figure 9.11 shows the antineutrino and neutrino selections from this sample. This score indeed does an adequate job of separating neutrinos and antineutrinos. In the fits to follow, two values of this score will be used as separators: 0.2 (referred to as AntiN2 in plots) and 0.5 (referred to as AntiN5 in plots). A second variable used to

test wrong-sign separation in this study is the sum of the scores with 1+ protons from the proton labels. This is done to roughly mimic the FSProton score used in NOvA studies and is called CVN Proton (or CVNPr) for these studies. Figure 9.12 shows that the variable is good at separating events with protons and those without and provides some separation for wrong-sign events as well.

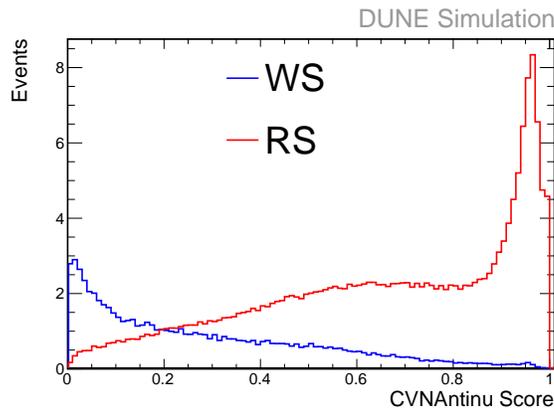


Figure 9.11: A spectrum of the CVN antineutrino score for events selected with an additional cut on true antineutrino (right-sign) or not (wrong-sign). Events have been weighted to somewhat better approximate events that would appear in the DUNE FD.

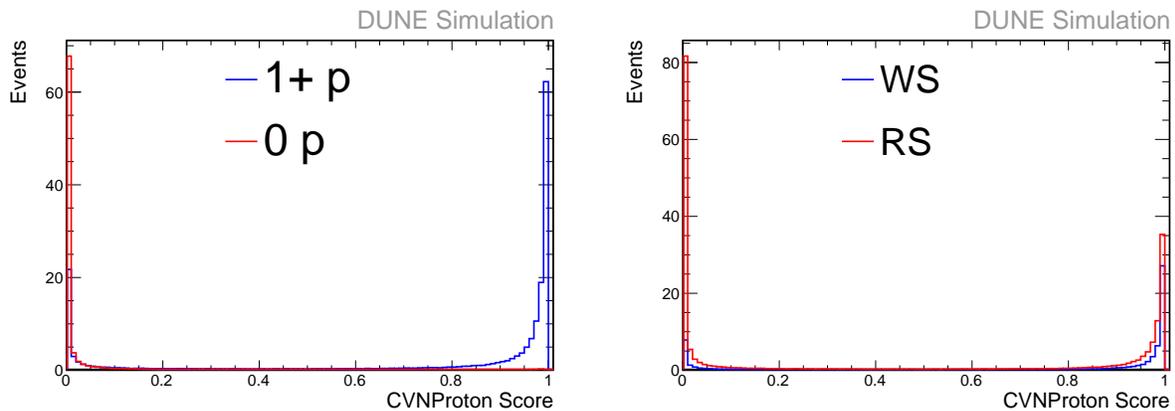


Figure 9.12: Spectrum of a summed CVN proton score for events selected with a cut on having protons or not (left) or being true antineutrino, right-sign, or not, wrong-sign (right). Events have been weighted to somewhat better approximate events that would appear in the DUNE FD.

Examples of the  $\bar{\nu}_e$  appearance predictions in the antineutrino beam with two forms of wrong-sign splitting applied are shown in Figure 9.13. These are formed by adding another equivalent set of bins shifted by 10 units to the right, and then recording the spectra like this with wrong-sign-like events put in this set of bins. Note that these do a good job of creating a wrong-sign-enhanced

sample while leaving a large sample of antineutrino appearance events in an right-sign-enhanced sample. This is especially true with the CVN anti-neutrino score version. Furthermore, note that the contribution from higher oscillation maxima can be seen as well: the second oscillation maximum occurs a little under 1GeV (recall Figure 1.2). Comparing with the right-hand side of Figure 9.10 shows that this shape is more easily seen with these split samples than in the full selection.

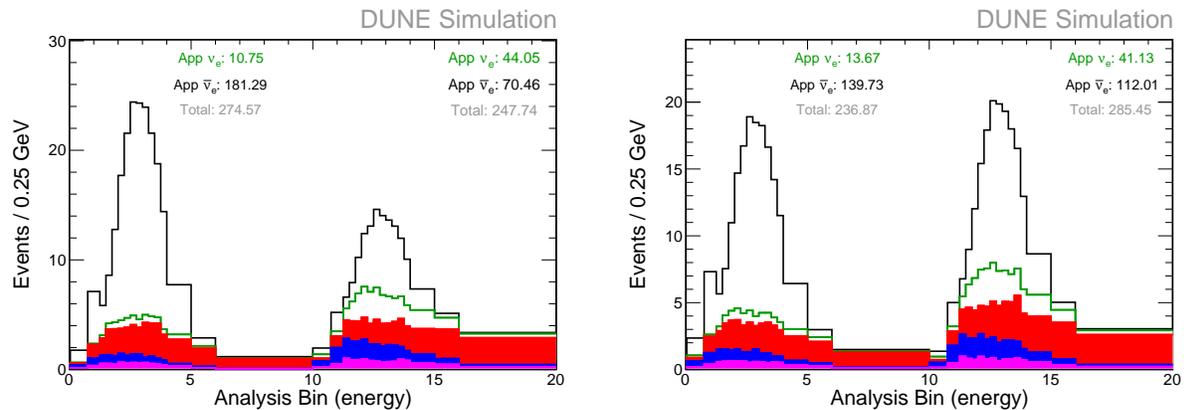


Figure 9.13: Unextrapolated FD predictions for the antineutrino beam for DUNE, with various methods of splitting into right-sign enhanced (bins below 10) and wrong-sign enhanced samples (bins above 10): based on the CVN antineutrino score being above or below 0.5 (left) and based on the CVN score for containing protons above or below 0.5 (right).

### 9.3 Oscillation analysis in DUNE and wrong-sign separation

For DUNE oscillation analysis, the exposure is more complicated to describe. The detector and beamline will both change over time, with one or two modules running at the start of beam and the full 40kt being completed after the beam has started. Furthermore, the beam will be upgraded from  $\sim 1\text{MW}$  power to  $\sim 2\text{MW}$  power over the course of DUNE’s lifetime. Therefore, the most precise definition of exposure for DUNE is one which combines the variables for detector fiducial volume, beam power, and time, giving the unit  $\text{kt}\cdot\text{MW}\cdot\text{yr}$ . With these in mind, a “staged” approach to oscillation analysis is adopted, with the following conditions:

Years	0.04	0.21	0.42	1.18	1.73	3.09	5.17
Exposure (kt·MW·yr)	1	5	10	30	50	100	200
Years	7.0	8.19	10.0	11.83	15.0	17.04	19.12
Exposure (kt·MW·yr)	336	450	624	800	1104	1300	1500

Since the NOvA study showed the qualitative differences in  $\sin^2 \theta_{23}$  values, only  $\sin^2 \theta_{23} = 0.441$  is used here, as this value was chosen as a plot point for externally presented DUNE materials [85].

The method of running fits is very similar to that for the preceding NOvA study, with a loop over parameters and conducting tests at many values of  $\delta_{CP}$  to examine slices and CP conservation rejection, both with and without wrong-sign separation. For exploratory purposes and to more directly mirror the study presented above, the fit is made very similar to the one described earlier in this chapter. Since DUNE will generally have an easier time telling the hierarchies apart, only the NH fits are performed and shown. Relatedly, the CP violation rejection plots check only the NH curves for the sensitivity at  $\delta_{CP} = 0, \pi$  (rather than both NH and IH).

Figure 9.14 shows the resulting slices for the 4 most interesting  $\delta_{CP}$  points with 7 years of running (split into half of the exposure in the neutrino beam and half of the exposure in the antineutrino beam). The version here has the wrong-sign separation using the two CVN antineutrino score values, both in longer dashed lines; the short dashed lines indicate perfect wrong-sign separation<sup>3</sup>. Unlike the method tested in NOvA, the chosen wrong-sign separation here does show some marked improvement. Furthermore, perfect wrong-sign separation shows more room for potential improvement as well. Figure 9.15 shows the resulting CP conservation rejection plot, which again shows that wrong-sign separation in DUNE seems to show even more potential than that in NOvA<sup>4</sup>. Taking the values to reject CP conservation at all tested values of exposure, Figure 9.16 shows the rejection power at  $\delta_{CP} = \pi/2$  and  $3\pi/2$  and potential improvements gained by considering wrong-sign separation in both the perfect case and currently-available methods. Further plots are shown in Appendix D, including one with the CVN Proton separation technique<sup>5</sup>. Appendix D further considers the case where only antineutrino beam is run and the same study performed.

With a larger and more advanced detector, longer baseline, and enhanced statistics, DUNE will

---

<sup>3</sup>The minimum in the case of  $\delta_{CP} = 3\pi/2$  appears to slightly miss the exact true value; it is unclear exactly why this is happening, but since the main goal is to compare analysis techniques, this is deemed a minor issue.

<sup>4</sup>These plots are fairly similar to externally shown DUNE plots (for example the one in Reference [85]), though there may be some shape difference and the sensitivity seems lower in my figure. This could be due to different assumptions or procedures between my study and what went into such plots, where for preliminary and exploratory purposes, I am mostly repeating the methodology of the preceding NOvA study.

<sup>5</sup>These do show some similar discontinuities in the fit results to those seen in NOvA, especially in cases where the sensitivity reaches high levels, even in the same hierarchy.

usher in a precision era in accelerator-based long-baseline neutrino physics. This chapter explored potential enhancements to the oscillation analysis. Specifically, this chapter demonstrated that effective sign separation of the  $\bar{\nu}_e$  appearance samples can provide improved sensitivity to CP violation above current expectations. Continued improvement to sign separation tools could further enhance the sensitivity, as evidenced by the “perfect” separation test. Experiments may see benefit from including atmospheric neutrinos, from analysis upgrades, and from other improvements yet to be considered, though this was only briefly noted in this chapter. The discovery of CP violation in the neutrino sector would be a major milestone in the history of physics, now nearly 64 years after the discovery of the neutrino.

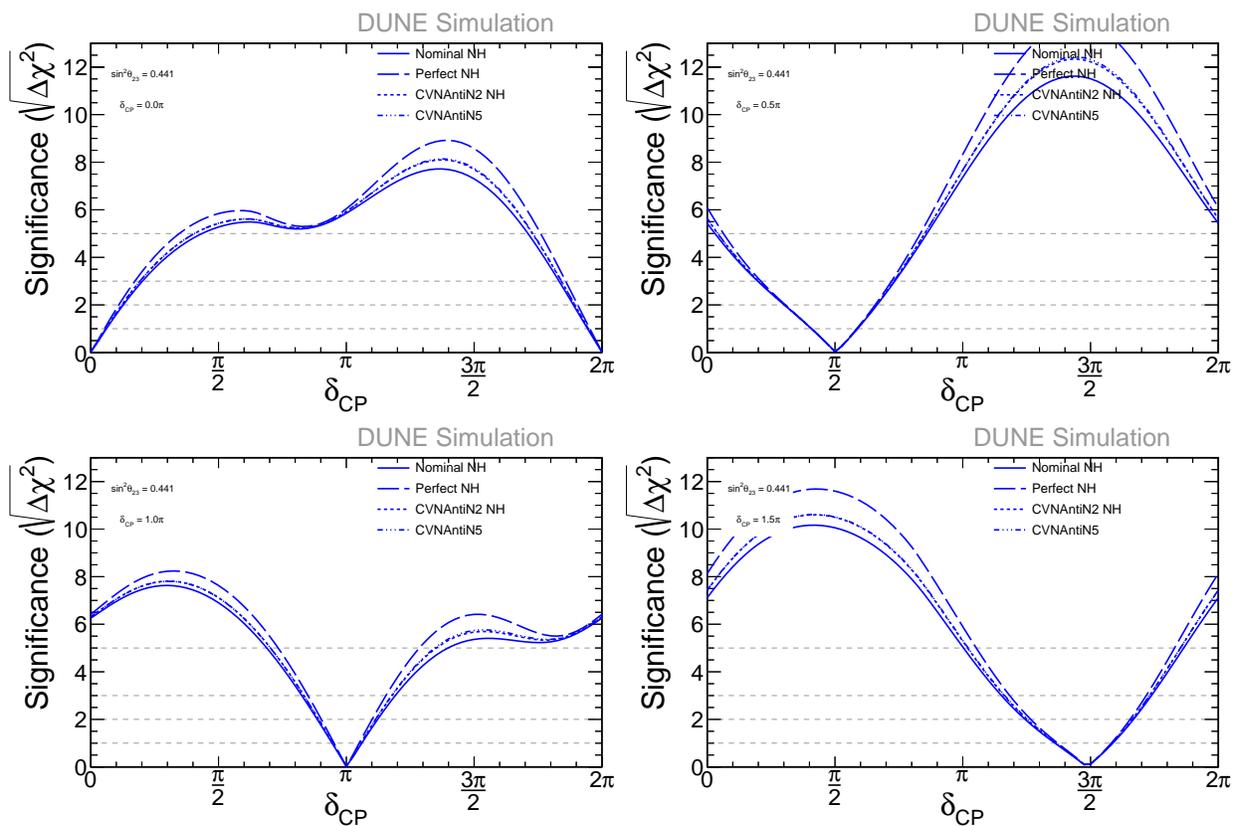


Figure 9.14: Slices for true  $\delta_{CP} = 0$  (top left),  $\pi/2$  (top right),  $\pi$  (bottom left),  $3\pi/2$  (bottom right) with 7 years of exposure for DUNE, split equally between neutrino beam and antineutrino beam. The wrong-sign separation techniques here use CVN antineutrino scores.

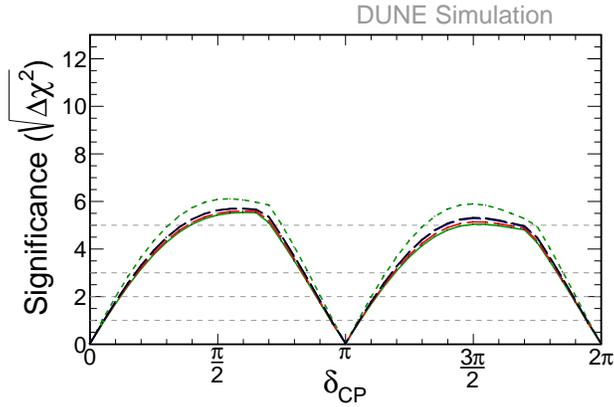


Figure 9.15: CP conservation rejection at 7 years exposure for DUNE (green solid). CVN Proton separation (red dash) shows some improvement, though the CVN antineutrino score (blue and black dash) shows more. Perfect sign separation (green dash) reveals further room for improvement.

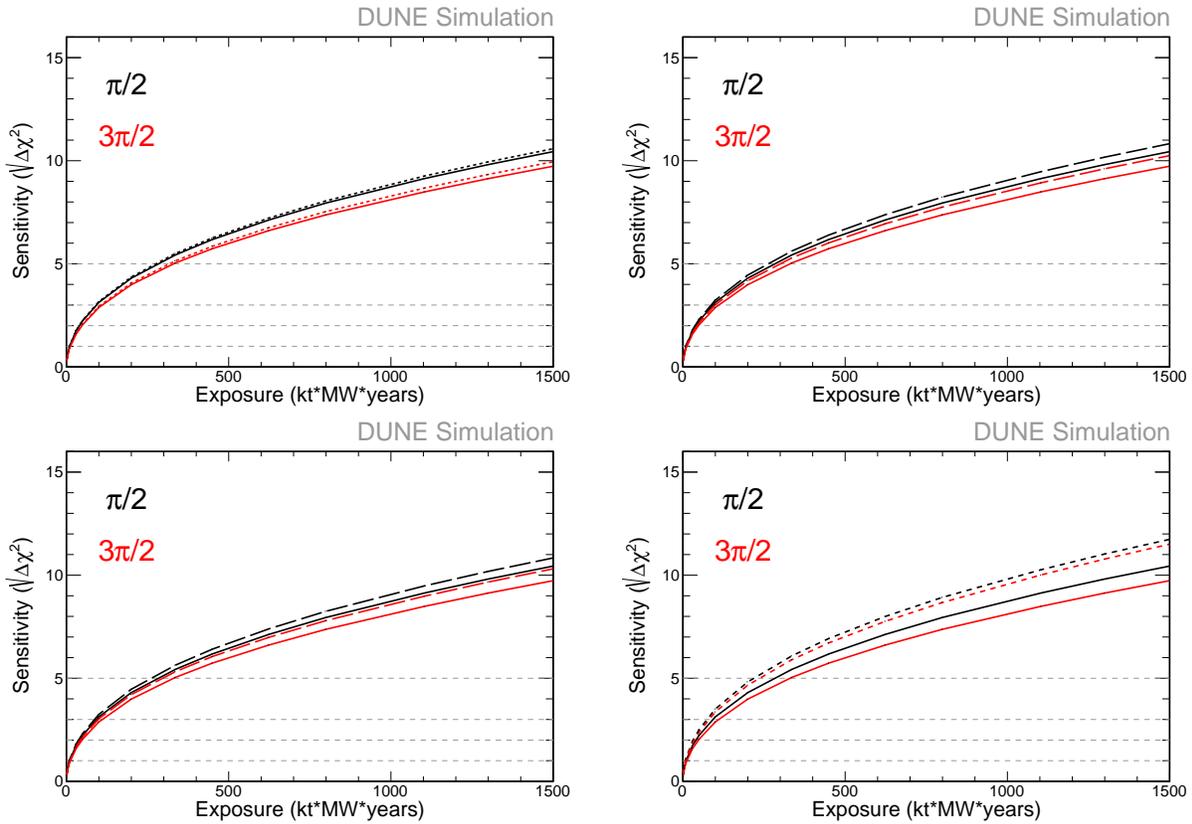


Figure 9.16: CP conservation rejection at  $\delta_{CP} = \pi/2$  (black) and  $3\pi/2$  (red) vs. beam exposure for DUNE with no sign separation (solid) and with sign separation (dashed): CVN Proton (top left), CVN antineutrino at 0.2 (top right) and at 0.5 (bottom left), and perfect separation (bottom right).

## APPENDIX A

### Studies on Drift in the NOvA Energy Scale Calibration

To understand the neutrino energy in signal events, the energy of the byproducts of the neutrino interactions provides the input. Ultimately, one wishes to find the correlation between energy deposited and signal read out, for example photoelectrons/MeV. In NOvA, energy scale calibration takes place in two steps. Scintillation produced by charged particles in the cells are converted and transported to the readout by fibers, and attenuation will affect the amount of light which reaches the readout. To account for this effect, through-going cosmic ray muons are studied for response as a function of position along a cell. On top of this attenuation, the effects of shadowing caused by energy depositions in other parts of the large detector and thresholding are also accounted for. Thresholding is an artificial increase in mean signal at the far end of a cell due to the bias caused by lower energy deposits not creating enough light to be seen. This is called the relative calibration, and correcting for these effects converts the raw photoelectrons (PE) digitized by the electronics to a fixed scale ( $PE_{\text{Corr}}$  or  $PECorr$ ) [42]. The second aspect of calibration is to then set the energy scale ( $PECorr/\text{MeV}$ ) by comparing the corrected signal from minimum ionizing stopping muon hits to the expected energy deposition.

This second energy calibration step is the focus of this appendix. As of the recent analysis of neutrino beam and antineutrino beam data discussed in this thesis, the methodology is generally unchanged relative to the previous analysis with just neutrino beam data [123]. Since then, mostly checks and exploratory studies have been undertaken. Hits are taken from stopping muons using the tricell selection wherein hits are required in both adjacent cells in the same plane, hence tricell [124]. This selection ensures that the muon has gone through a cell in a well-defined way, allowing the path length of the muon in the cell to be reasonably well reconstructed. An illustration of a tricell hit is shown in Figure A.1 [42]. Hits 1-2m from the determined muon stopping point, in the most minimally-ionizing portion of the energy deposition, are used such that small variations in the

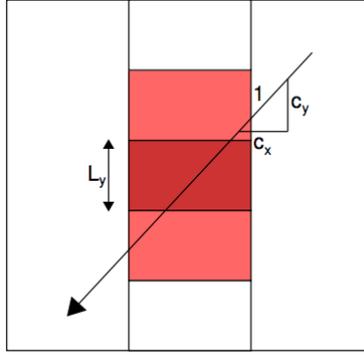


Figure A.1: Illustration of a tricell hit [42], as shown in Chapter 3.

expected energy deposited by moving the stopping point will not impact the results. Additionally, for this most recent iteration [125], hits are selected only from regions of the cells where the relative calibration has left a relatively flat distribution of mean PECorr/cm, to minimize potential biases related to the relative calibration.

### A.1 The drift in calibration

For expediency, the collaboration reused the calibration from the previous oscillation analysis, as described above, for the 2018 oscillation analysis<sup>1</sup>. That is, the calibration for data in Period 5 (the most recent data-taking period used in said analysis) would be reused for data-taking Periods 6 and 7, with the anticipation that the calibration should not change much from period to period. The energy scale [MeV/PECorr] in the Far Detector (FD) high gain periods (3-5) only varied between  $4.432 \times 10^{-2}$  and  $4.436 \times 10^{-2}$  [125]. This constitutes a change of less than .1%. The Near Detector (ND) experienced somewhat more variation when including Period 2, which for the ND is the same gain setting, but between Periods 3 and 5 had only a variation between  $4.298 \times 10^{-2}$  and  $4.301 \times 10^{-2}$  MeV/PECorr [125].

However, data from more recent run periods were checked to ensure consistency with the previous results. The MEU is the average of mean x-view and y-view visible signals per unit track length (PECorr/cm). For the calibration used in Periods 3 through 5 at the FD, the MEU ranged from 40.24 PECorr/cm to 40.27 PECorr/cm [125]. Unexpectedly, the data from Epoch 6a was found to clearly

<sup>1</sup>This was the first look at the neutrino and antineutrino beam, which then led to the recent top-up with more antineutrino beam presented.

shift toward lower response, with the resulting MEU of 40.05 PECorr/cm. Characterizing this by the difference over the average, this is a change of 0.57% from Epoch 5a.

Data from Epochs 7a and 7b, two further data-taking periods after Period 6, were in hand with enough time to perform further calibration checks on these epochs. First, a small sample of Epoch 7a FD files were produced to perform a quick test. When, later, a fuller set of data from Epoch 7a was processed, the manner in which the definitions were handled may have possibly allowed an estimated  $\sim 150$  files out of  $\sim 10,000$  to be present twice [126]. This could skew numbers from Epoch 7a for the FD a little bit, but they are presented in the plots and tables here for illustrative purposes regardless. No conclusions will be drawn specifically from Epoch 7a FD values herein.

Upon checking the calibration constants for Epochs 7a and 7b, the MEU was found to continue drifting downward with time, as is summarized in Table A.1 for both the ND and FD. This downward drift is potentially indicative of a degradation, and potential culprits will be discussed in Section A.2.5. Typically, the calibration for a given period had come from one epoch or a few epochs within the data-taking period, so the Period 5 calibration comes from Epoch 5a. Likewise, Period 4 uses calibration data from an epoch in Period 4, etc. The MEU was found to drift from the value found in Epoch 5a by 1.58% in the FD and 0.90% in the ND by Epoch 7b and further to 1.62% in the ND by Epoch 8b. This scale of drift from the calibration constants used (Period 5) was fine for amount of time passed for the original analysis, using data up to Period 7b in the FD. However, for data after Period 7b in the recent top-up analysis, more recent calibration constants were used. Further, the nature and cause of this drift is a matter of study.

Assuming a constant drift, the data-weighted midpoint of each epoch would set the timescales over which the drift occurs from epoch to epoch. Taking roughly the calendar midpoint of each epoch, we can approximate the drift. The midpoints of Epoch 5a and Epoch 7b are  $\approx 12$  Dec 2016 and  $\approx 9$  Feb 2018, respectively, roughly a separation of 1.16 years. Taking this per cent scale linearly, the ND would drift at  $0.9\% / 1.16 \text{ years} = 0.78\%/\text{year}$ , and the FD would drift at  $1.58\% / 1.16 \text{ years} = 1.36\%/\text{year}$ . This scale would predict a shift of 1.28% if one were to project out to Epoch 7d, while in reality this was found to be 1.18%. Such predictions are therefore rough, but in

the ballpark.

In fact, prior studies with older versions of calibration had actually indicated such a drift, with fits to PECorr/cm showing a drift of 1-2% per year [127]. Though these studies would have been conducted with different selection cuts and such, they are largely consistent. The application of the relative calibration pinning the response to a fixed value had the effect of masking this drift.

	FD MEU PECorr/cm	Date range of epoch	% Difference 100x(diff/avg)	ND MEU PECorr/cm	% Difference 100x(diff/avg)
Epoch 5a	11/2016 - 1/2017	40.28	–	41.71	–
Epoch 6a	2/2017 - 7/2017	40.05	0.57%	41.54	0.42%
Epoch 7a	11/2017 - 12/2017	39.73*	1.37%	41.40	0.76%
Epoch 7b	1/2018 - 3/2018	39.65	1.58%	41.34	0.90%
Epoch 7d	4/2018 - 7/2018	–	–	41.22	1.18%
Epoch 8b	10/2018 - 2/2019	–	–	41.04	1.62%

Table A.1: The MEU [PECorr/MeV] found for recent epochs, comparing to that from Epoch 5a, the last used for calibration in the previous analysis. The second column in each set uses the difference of the given epoch and Epoch 5a (and average), to demonstrate the percentage change between the epochs. The Epoch 7d and 8b ND values were calculated by an expert-in-training and double-checked by me. (\*) As noted in the text, a slight file issue may be present in the Epoch 7a FD sample, but the numbers are presented here for illustrative purposes.

To understand this drift with time, the MEU calculation was done with much finer granularity and plotted versus run number in groups of 10, sometimes with the FD broken into its 14 diblocks separately. Breaking up the FD data in various ways such as this allows for more insight into the nature of this drift. For example, is it uniform throughout the detector? In one version of the plot, the relative calibration was pinned to the values from Period 3 to remove the effect masking the drift in Periods 4 and 5. This is shown in Figure A.2, with arrows drawn on from a presentation by the calibration conveners [128]. Clear drifts can be seen in the PECorr/cm over time, and one can note three further things:

1. There are occasional upward bumps in the MEU value followed by the continuing downward trends

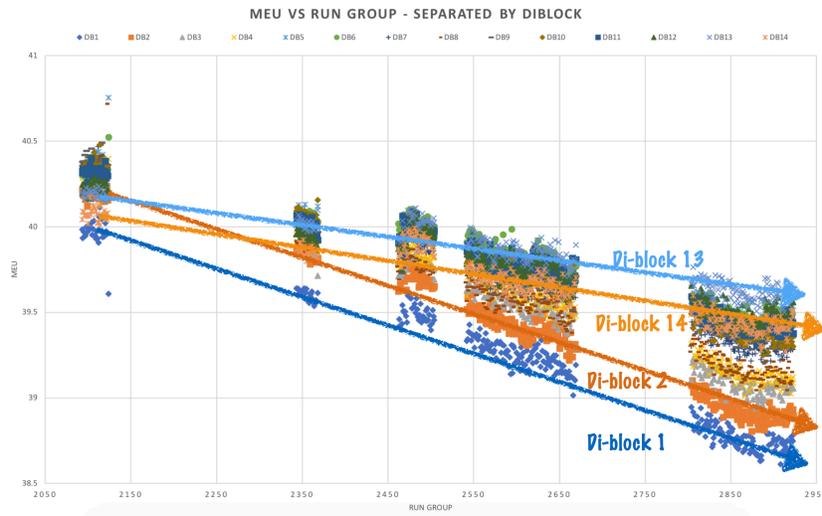


Figure A.2: Plot of MEU vs. groups of 10 run numbers where each FD diblock is shown separately. The calibration was forced to use Period 3 for the relative calibration step (PECorr). This shows that the drift was occurring in earlier periods and over time. Note both the spread among diblocks in this figure and the fact that the diblocks seem to be evolving with different slopes [128].

2. The drift seems to occur with different slopes in different diblocks. Perhaps a unique calibration constant for each diblock would be beneficial: this is explored herein.
3. The drift shows even before it appears in the numbers in Table A.1. This is because the relative calibration will reset the average to this fixed scale for each period.

## A.2 Supporting evidence: digging into the drift

Section A.1 establishes that there is indeed a drift to the calibration, that it seems to have been present but masked by relative calibration for a long period of time, and that it is present in both detectors, though seemingly somewhat different in magnitude. This section will go into more detail on digging that has been done since the original (re)discovery of this drift to try to understand its effects and potential causes.

### A.2.1 Digging into the drift structure

It is interesting to note that that the drift broken into run ranges as in Figure A.2 shows interesting structure. The drift is generally always negative within fluctuations, yet there are occasional spikes or odd structures with increased response. In particular, there is a spike in MEU in Epoch 6a (fourth grouping of points) somewhat visible in Figure A.2, roughly 2/3 of the way between 2550

and 2650 on the x-axis. After this spike, the general downward trend continues, but seemingly from this elevated level. An effort was made to attempt to see if these elevations correlate with physical conditions of the detector, though it is important to note that these may be second order effects on top of whatever drift is happening: such an explanation need not also explain the drift.

One spike of note is the structure in Period 5 (third grouping of points) in Figure A.2. The large shift in calibration levels during Period 5 occur near group 2465x (that is 24650-24659). Indeed, this run group correlates with the “Sprinkler Incident” at the FD, in which the detector was down for several days while the computing systems recovered from the sprinkler having been set off in the computing room.

For less significant downtimes and more regular maintenance periods, one can turn to the NOvA electronic logbook (ECL), which is used to keep notes and records during shift-taking and work on detectors, for more information. The spike in Period 6 correlated with some load-shed testing (a procedure by which the computing systems drop unnecessary processes to mitigate unwanted increases in rack temperature) though this would not seem important to the environmental conditions of the detector. Further, there was a dryer filter replacement around a similar time (ECL 105540), but an expert suggested these were probably red herrings [129]. Said expert also looked through a few monitoring plots and found no long-term trends [129]. However, the spike in Period 7 seems to correlate with a power trip at the FD on 9 February 2017 in which APDs were warmed.

Conversely, one could pick out particular maintenance periods and try to check for any spikes in the trend. For example, it was noted that there was a maintenance day in run group 2862x. This data point is hard to distinguish as different from a fluctuation by eye. With further looks at plots of each diblock over time, perhaps one could perform a deeper study of maintenance periods in the detectors. Though it seems as if there could be some correlation of spikes to changes in detector conditions, not every case could be explained as such. Further detector hardware/environment questions will require more thought/digging into.

### *A.2.2 Digging into the diblock breakdowns*

An interesting note from Figure A.2 is that the calibration of the FD energy scales is not doing the best job of calibrating out differences between diblocks. In the figure, one can see that even when the relative calibration is run on a given period, a few diblocks are systematically lower in response than the rest of the main group of diblocks. Further, when one examines the evolution of response including the drift, it also seems to be the case that the diblocks drift differently, with the slope of the drift smaller for diblock 13 and 14 than for diblock 1 and 2. The scale of the drift was presented in a more averaged sense in Table A.1.

The spread in energy scale calibration among diblocks can be understood by following the same procedures and cuts used to perform the monolithic calibration used in the analyses but instead calculating separate numbers for each diblock. The studies show that the diblocks in the Near Detector MEU [PECorr/cm] and the simulations in both the FD and ND generally agree better than that in the FD data, which shows a more appreciable and growing difference between diblocks. Table A.2 below shows the split between the lowest MEU and highest MEU diblocks and also the average PE/cm (with no relative calibration corrections), both under the same selection cuts and averaged between X and Y views as they are in the main calibration. While the relative calibration narrows the spread between diblocks from  $\sim 20\%$  to  $\sim 1\%$ , some spread is still present and growing with time due to the drift. Therefore, it seems like more could be done to account for these differences, as will be discussed in Section A.3. Interestingly, in the PE/cm calculations, the FD simulated samples (both low gain and high gain) showed separations of  $\sim 13\%$  among the diblocks that was seemingly fixed via the relative calibration, as the spread was not seen in MEU. Finally, note that the PE/cm before any relative calibration is applied is also decreasing with time. This further speaks to the downward drift in detector signal.

One further point of note is that the MEU values from fiber brightness bins also seem to differ as a function of diblock. That is, the MEUs from even the dimmest fiber bin in the brightest diblocks (on average) are greater than that of the brightest fiber bin in the dimmest diblocks (on average). This can be seen in Figure A.3, where the x-axis is diblock number, the y-axis is MEU, and the dots

	Lo MEU (Diblock)	Hi MEU (Diblock)	Diff 1-[Lo/Hi]	Lo PE/cm (Diblock)	Hi PE/cm (Diblock)	Diff 1-[Lo/Hi]
Period 2	39.91 (1)	40.26 (5)	0.0087	22.30 (3)	28.07 (9)	0.2056
Period 3	39.98 (1)	40.39 (9)	0.0101	22.14 (3)	28.21 (9)	0.2150
Period 4	39.94 (1)	40.36 (9)	0.0102	21.95 (3)	27.99 (9)	0.2158
Period 5	39.98 (1)	40.39 (9)	0.0102	21.91 (3)	27.93 (9)	0.2158
Period 6	39.73 (1)	40.18 (6)	0.0113	21.77 (3)	27.77 (9)	0.2161
Epoch 7a	39.28 (1)	39.89 (5)	0.0152	21.54 (3)	27.57 (9)	0.2188
Epoch 7b	39.19 (1)	39.83 (5)	0.0160	21.49 (3)	27.49 (9)	0.2184

Table A.2: Spread in MEU [PECorr/cm] and PE/cm among FD diblocks. Recall that the values in Epoch 7a may be slightly skewed due to a small file issue, but the general trend is apparent. This table was also presented in Chapter 3.

are the 9 different fiber brightness levels. Note that the simulation is much more uniform. This may indicate that the initial conditions of the fibers themselves are not fully what causes the spread in brightness among diblocks.

### A.2.3 Digging into the distributions

In the course of examination, different subsets of variables plotted for each period provide useful comparisons for understanding drift effects. The cosmic muon distributions, hit distributions, PE, PE/cm, etc. can be broken down into different dependences such as fiber brightness, position along cell, plane number, etc. Each of these provide a slightly different view of the phase space of this drift. Some of the more interesting plots are shown below. A description of what is plotted is given in the figure captions. Implications on potential causes will be discussed later.

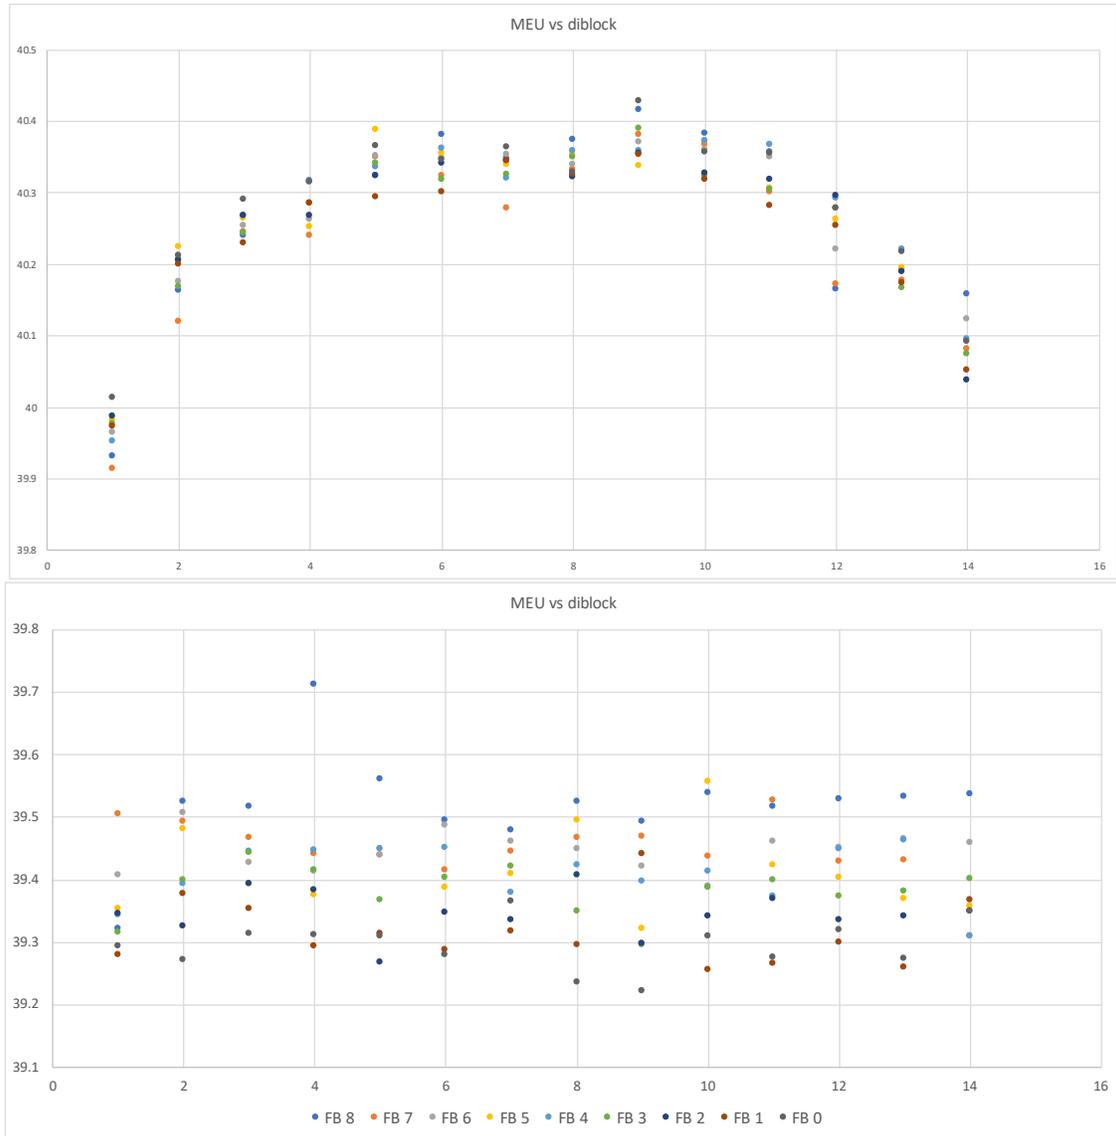


Figure A.3: Breakdown of the PECorr/cm scales (y-axis) seen in the diblocks of the FD (x-axis) for the various fiber brightness bins (each dot) for the Period 3 data (top) and high gain simulation (bottom). The legend in the bottom plot may not correlate to the same colors in the top plot, but the trends described in the text are clear.

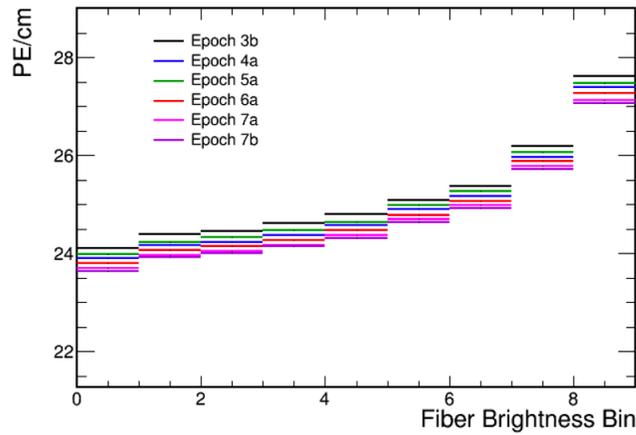


Figure A.4: A look at the response in PE/cm for the FD X-view vs. fiber brightness bin. The degradation in response is visible across cells with fibers of all brightness ranges.

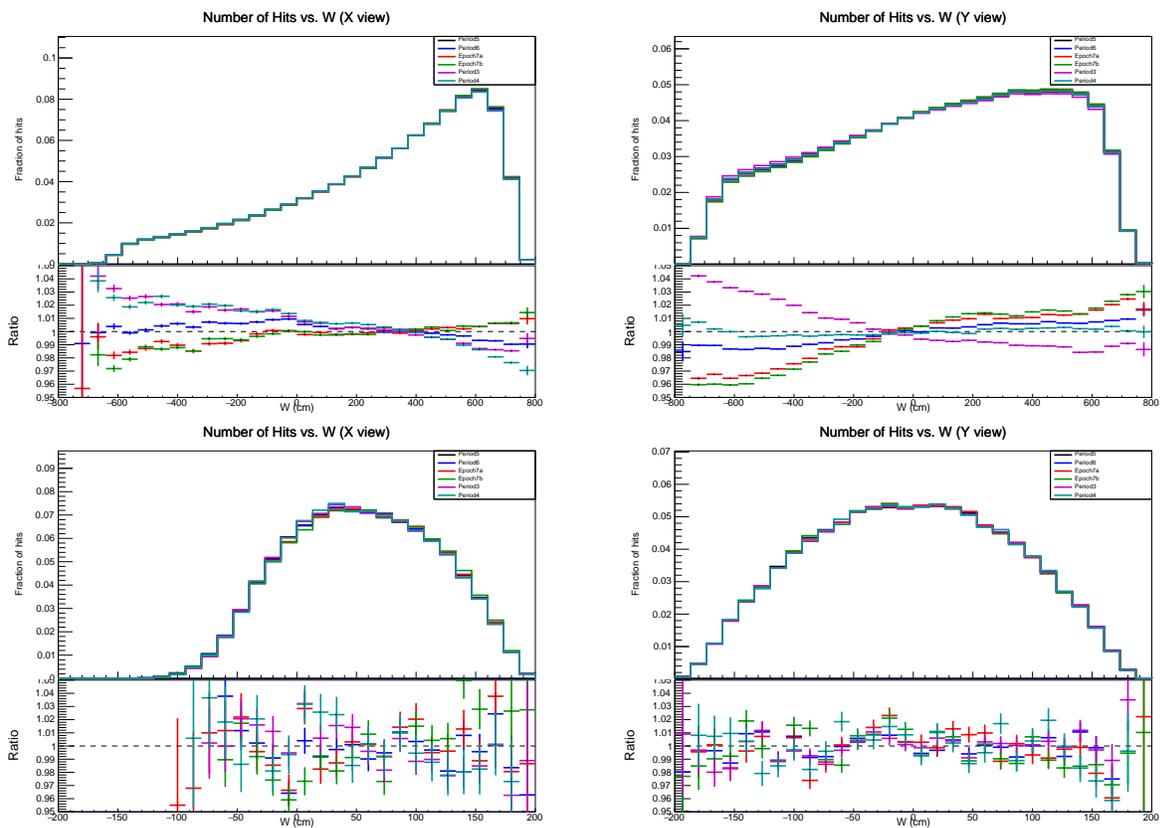


Figure A.5: Fraction of hits in the sample vs. position along cell. The x-view is on the left, with the y-view on the right. Black is Period 5. Blue, red, and green are more recent periods, and purple and teal are older periods. (top) FD: this spectrum clearly tilts over time, indicating fractionally fewer hits at positions in the cell further from the readout (closer to negative end), a potential thresholding. (bottom) ND: little such thresholding is visible here.

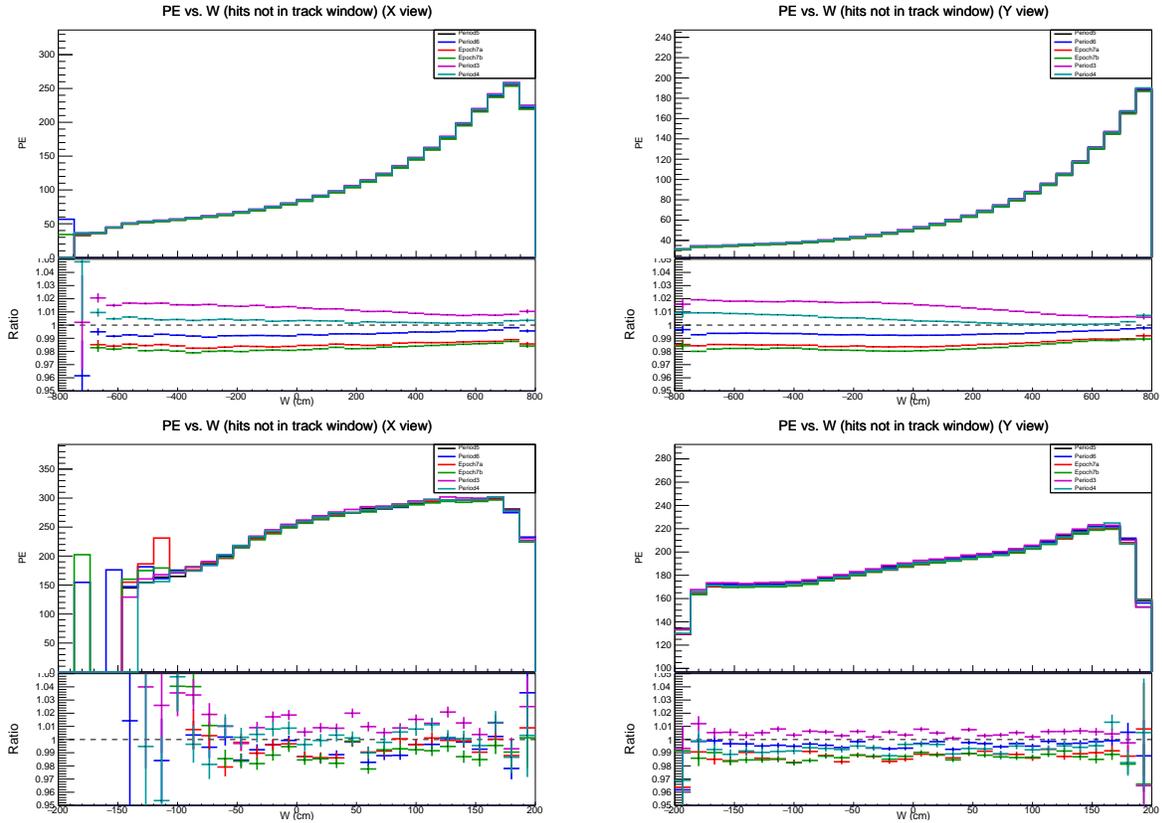


Figure A.6: Average response in PE vs. position along cell. The x-view is on the left, with the y-view on the right. Black is Period 5. Blue, red, and green are more recent periods, and purple and teal are older periods. Both the FD (top) and ND (bottom) show clear signs of PE dropping over time across the entire length of the cell. Interestingly, especially in the FD, it looks like the degradation is slightly less at the readout end (closer to positive end).

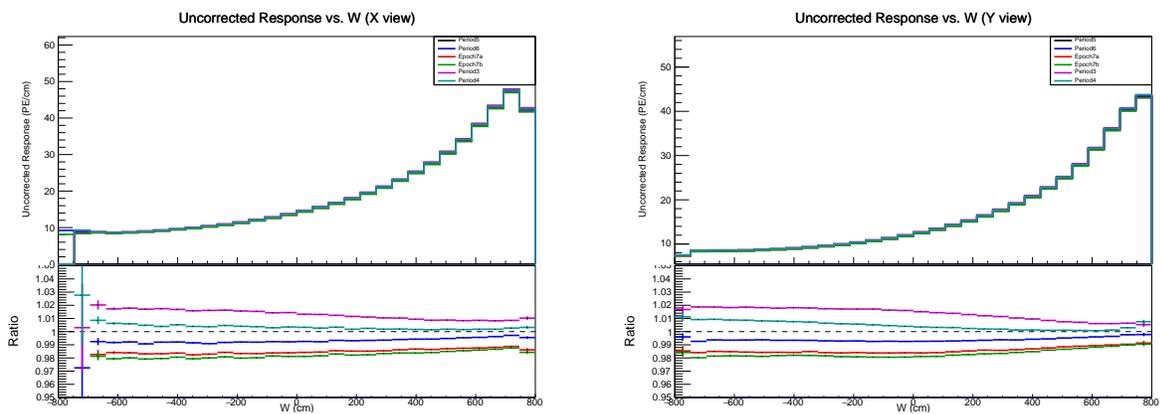


Figure A.7: Average response in PE/cm vs. position along cell in the FD. The x-view is on the left, with the y-view on the right. Black is Period 5. Blue, red, and green are more recent periods, and purple and teal are older periods. Again, the detector shows clear signs of response degrading over time across the entire length of the cell. As before, it looks like the degradation is slightly less at the readout end (closer to positive end). This figure is also shown in Chapter 3.

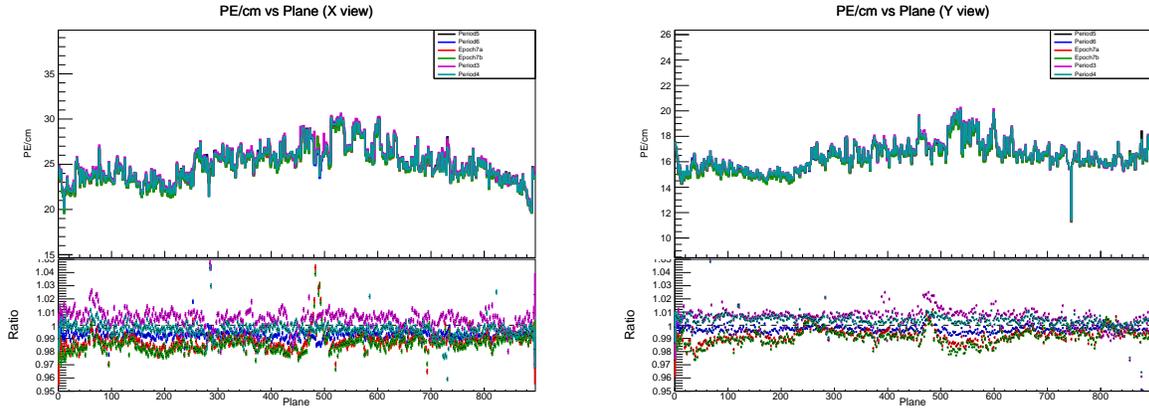


Figure A.8: Average response in PE/cm vs. plane in the FD. The x-view is on the left, with the y-view on the right. Black is Period 5. Blue, red, and green are more recent periods, and purple and teal are older periods. Both show clear signs of PE/cm dropping over time across the detector.

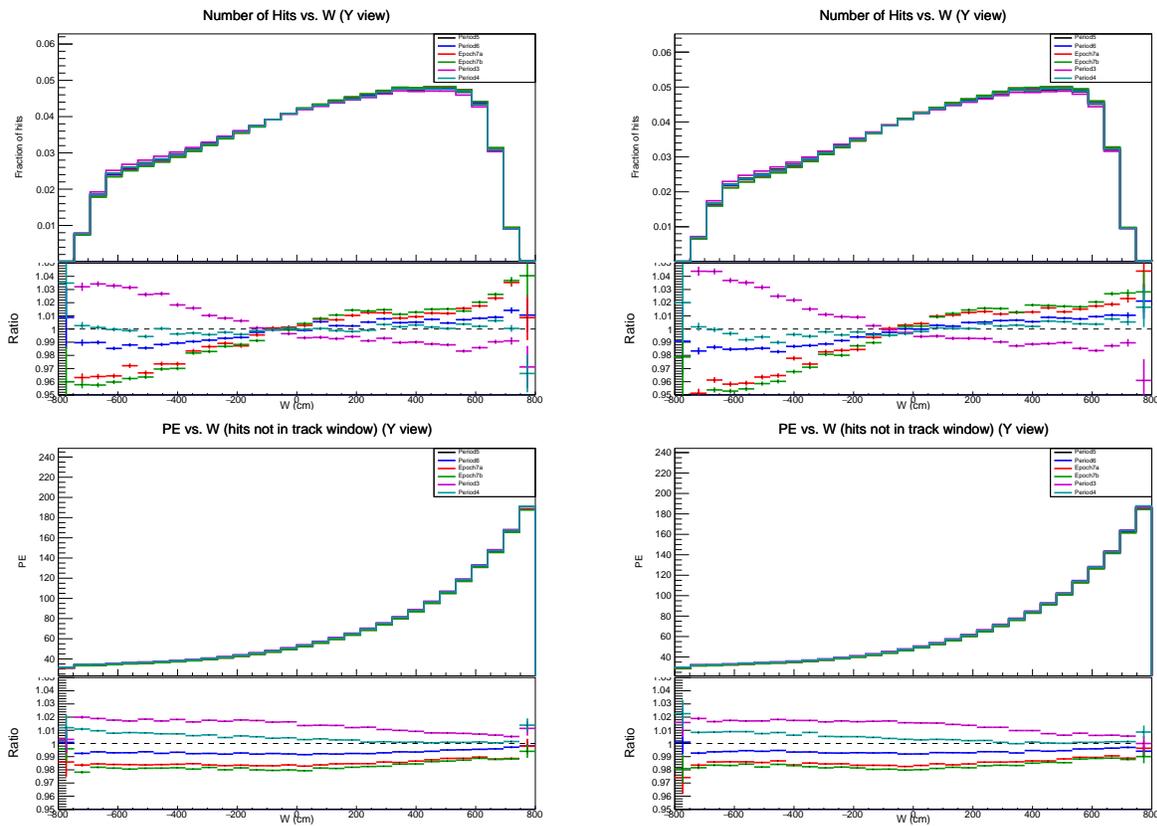


Figure A.9: (top) Fraction of hits in the sample vs. position along cell for the FD y-view, in the 3rd brightest fiber bin (left) and 3rd dimmest fiber bin (right). (bottom) Average response in PE vs. position along cell for the FD y-view, in the 3rd brightest fiber bin (left) and 3rd dimmest fiber bin (right). Note that while the effect of thresholding going from brighter fibers to dimmer fibers appears to perhaps increase by some amount, the change in response is similar.

#### A.2.4 Digging into the drift over time

Furthermore, the combination/reformulation of parts of the energy scale calibration code in the Framework for Enhancing and Noting the NOvA Energy Calibration (FENNEC) package allowed for additional useful plotting. Major additions were profile histograms (ROOT TProfile) of response (in PECorr/cm, PE/cm, and PE) versus the actual time of the data, in various splittings of the relevant data. This produces time-series versions of plots such as the one in Figure A.2 that was made based on run numbers, making the plots more intuitive and able to be fit. These plots are shown for the FD in Figure A.10 and for the ND in Figure A.11. To save space, only the y-view is given. For the ND, the difference in structure at the earliest times was traced to a change in the trigger occurring in Period 2. Fits are shown in Table A.3 and Figure A.12 and reveal the drift. Slight changes were made to the time stamp usage in FENNEC, which may have changed the timing bins in which some hits fall. However, re-doing the fits for all diblocks combined at the FD after changing the timing, the slopes agreed to at least 3 significant digits and are therefore considered fine. The fits will be somewhat further changed by the usage of TProfiles as opposed to profiles of a two-dimensional histogram as originally used. Fits of PE/cm degradation in each diblock are shown for the FD in Table A.4, and figures A.16-A.20 at the end of this appendix show these. Given Figure A.2, one would expect the slope to be higher in diblocks 1 and 2 than 13 and 14: indeed this is the case.

	X-view		Y-view		Average Drift (PE/cm/year)
	PE/cm/yr	$\chi^2$ /DOF	PE/cm/yr	$\chi^2$ /DOF	
ND	$-0.344 \pm 0.010$	1883/1112	$-0.407 \pm 0.004$	3412/1112	-0.376
FD	$-0.284 \pm 0.001$	6296/640	$-0.221 \pm 5 \times 10^{-4}$	12222/640	-0.253

Table A.3: Results of a preliminary linear fit to the PE/cm time series.

It is interesting to consider if there is some clear correlation of the effects seen as compared with the fiber brightnesses. As discussed previously, Figure A.4 shows that the degradation is happening regardless of fiber brightness. However, that figure does not further break down the

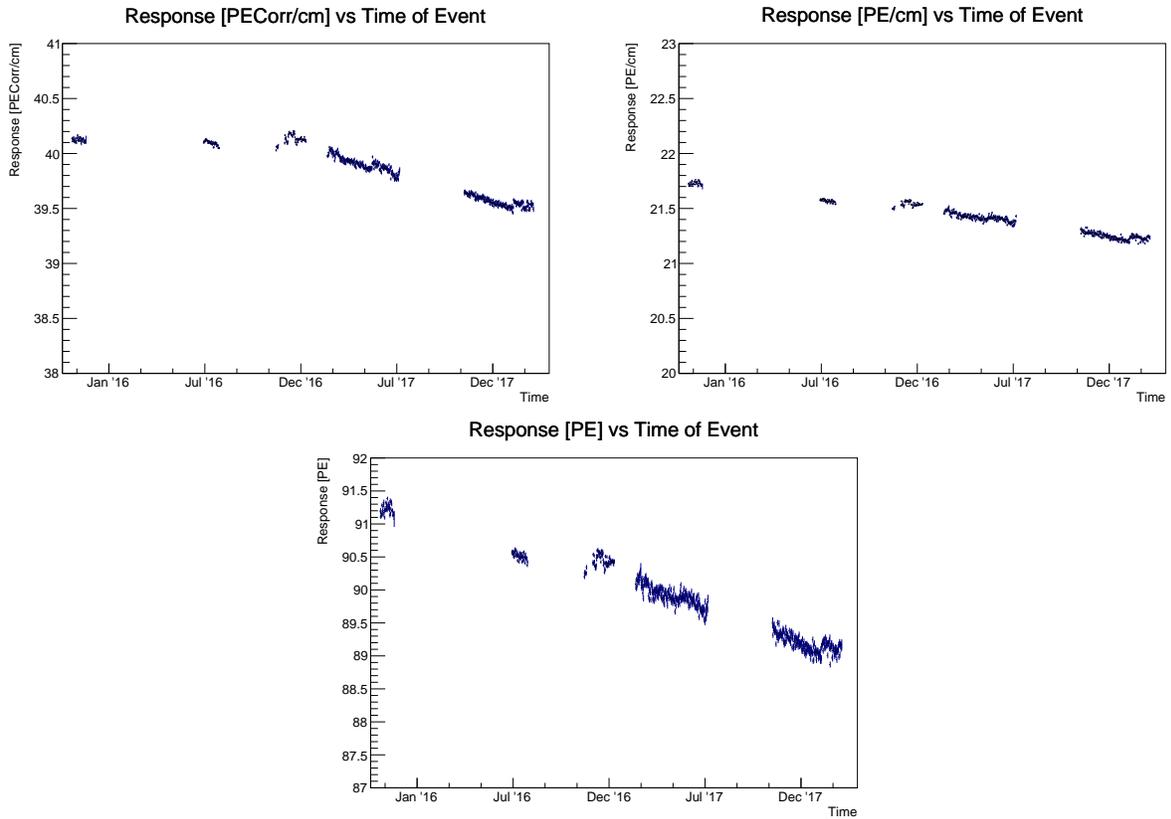


Figure A.10: Profile of response for selected hits over time in the FD since the start of high gain running. Gaps occur where calibration data was not processed and/or where the beam was off. For some data-taking periods, especially earlier ones, only one sub-period (epoch) of FD data was used for calibration.

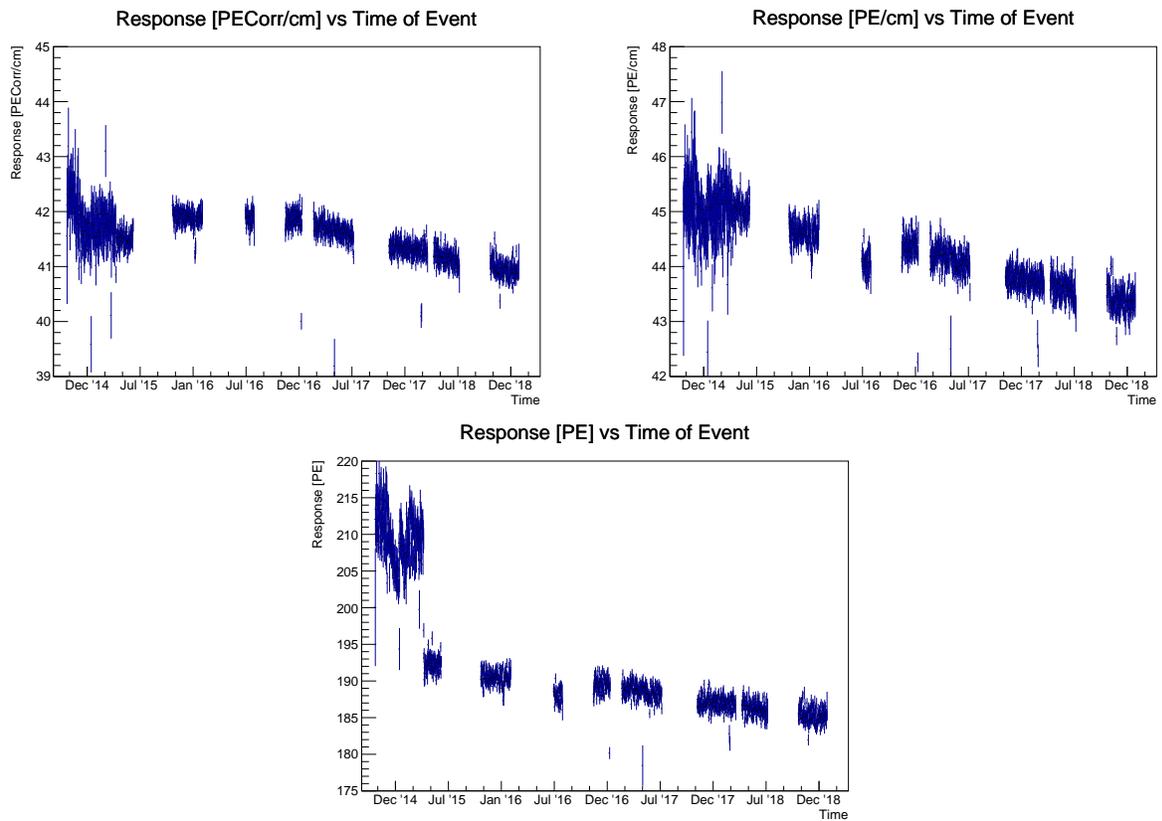


Figure A.11: Profile of response for selected hits over time in the ND since Period 2. Differences in structure at the earliest times was traced to changes in the trigger during Period 2.

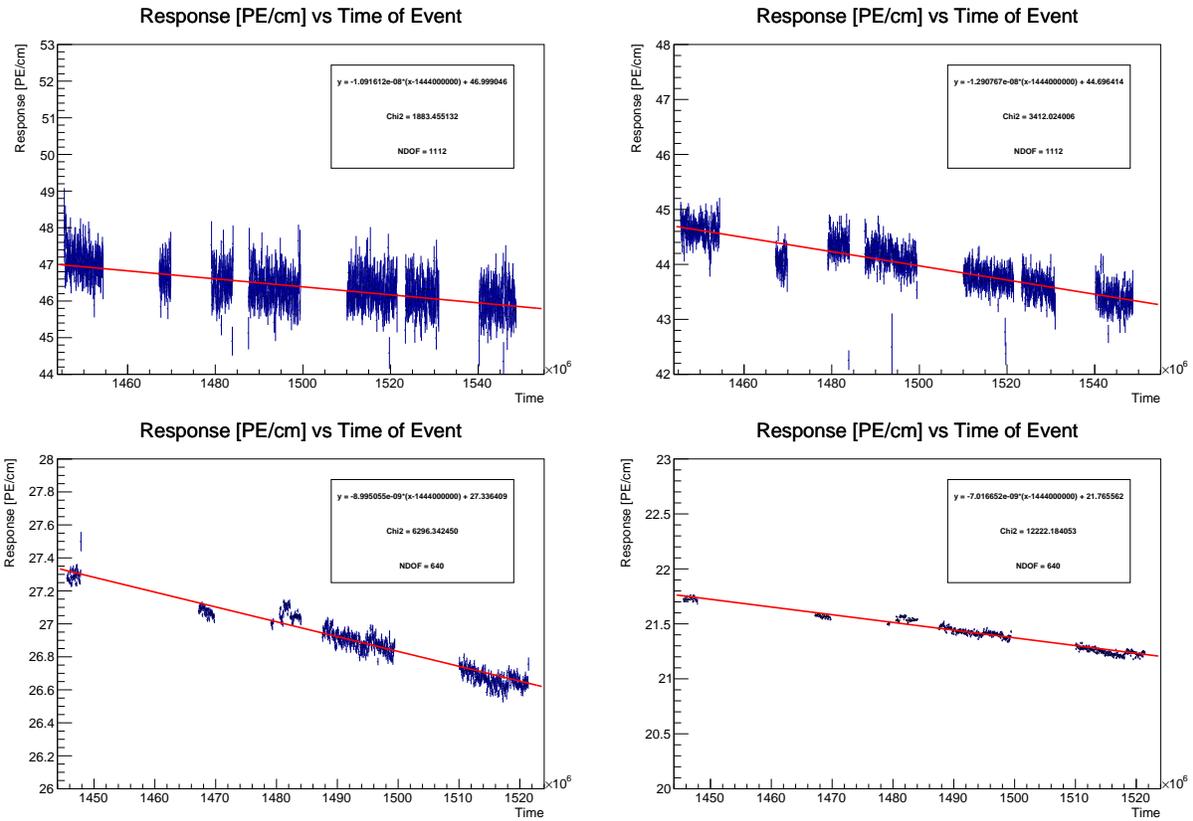


Figure A.12: Linear fits to the PE/cm distributions, for the ND x-view (top left), ND y-view (top right), FD x-view (bottom left), and FD y-view (bottom right). The box in the plot shows the fit (in PE/cm/s),  $\chi^2$ , and number of degrees of freedom.

Diblock	X-view PE/cm/yr	$\chi^2$ /DOF	Y-view PE/cm/yr	$\chi^2$ /DOF	Average Drift (PE/cm/year)
1	$-0.360 \pm 0.004$	2471/640	$-0.278 \pm 0.002$	5848/640	-0.319
2	$-0.341 \pm 0.004$	1677/639	$-0.325 \pm 0.002$	1927/639	-0.333
3	$-0.336 \pm 0.004$	1723/639	$-0.246 \pm 0.002$	1877/639	-0.291
4	$-0.303 \pm 0.004$	1687/639	$-0.247 \pm 0.002$	1878/639	-0.275
5	$-0.276 \pm 0.004$	1329/639	$-0.199 \pm 0.002$	1508/639	-0.237
6	$-0.258 \pm 0.004$	1571/639	$-0.191 \pm 0.002$	2446/639	-0.224
7	$-0.344 \pm 0.004$	1486/639	$-0.226 \pm 0.002$	1761/639	-0.285
8	$-0.376 \pm 0.004$	1902/639	$-0.286 \pm 0.002$	1889/639	-0.331
9	$-0.319 \pm 0.004$	1318/639	$-0.311 \pm 0.002$	2229/639	-0.315
10	$-0.294 \pm 0.004$	1377/639	$-0.259 \pm 0.002$	1883/639	-0.277
11	$-0.247 \pm 0.004$	1618/639	$-0.237 \pm 0.002$	2128/639	-0.242
12	$-0.226 \pm 0.004$	1850/639	$-0.144 \pm 0.002$	2236/639	-0.185
13	$-0.211 \pm 0.004$	1487/639	$-0.090 \pm 0.002$	2173/639	-0.151
14	$-0.167 \pm 0.004$	1631/639	$-0.135 \pm 0.002$	3851/639	-0.151

Table A.4: Results of a preliminary linear fit to the PE/cm time series for each FD diblock.

data into regions of the detector. To this end, it is interesting to note from Table A.4 that diblocks 8 and 9 have a similar PE/cm/year rate of degradation as diblocks 1 and 2, even though the starting response for diblocks 1 and 2 make it seem like a higher fractional change. Meanwhile, diblocks 12, 13, and 14 have the lowest average drift. A simple check is looking at the fiber brightness maps at the FD for relative brightness levels. By eye, it appears as if diblocks 8 and 9 have a significantly higher fraction of high bin fibers compared to diblocks 1 and 2. Furthermore, by eye, it looks like diblocks 12-14 may also have a higher fraction of high bin fibers than diblocks 1 and 2. This would seem to imply no such correlation.

#### A.2.5 Investigating causes

Since the process of producing and detecting scintillation light from stopping muons requires many sequential steps, it can be hard to disentangle effects from each other to separate out the

cause of the drift: it could be any of a number of different causes from the broader categories of scintillator, fiber, photodiode/electronics hardware. Therefore, it is important to consider the distributions and points raised above alongside other external checks which have been run in the past, to narrow in on a potential cause.

First, consider the possibility that the drift is related to the detector readout/hardware. It was already noted that it could be a second order effect, causing some of the spikes in response that are seen, with some correlation between such spikes and detector status noted. A way that the hardware could cause drift is if, for example, the gain of the APDs were dropping over time while we believe it to remain constant. However, the detector threshold settings have remained fairly constant over time. Further, it was noted that if the issue were with APDs or FEBs, one would expect the entire detector to drift together, instead of the different slopes seen among diblocks [130]. Some questions related to other parts of the detector hardware (voltage readout, chiller piping, etc.) remain to be investigated as of the conclusion of the work described here.

Beyond this, there are two potential sources listed above meriting discussion. Consider the fibers. A collaborator previously tested fibers for amount of lost core: besides one peculiar sample, the conclusions suggest that the degradation of the core should not be noticeable on the timescales of interest here [131, 132]. However, the slight deviation in the degradation of PE and PE/cm versus  $W$  near the readout end could point to a different fiber effect: the attenuation lengths. If the visibility of the fibers were degrading, one would expect less signal to reach the APDs and one might further expect that this degradation would have some shape related to the distance at which light was picked up in the cell. Looking at the relative calibration fits from Period 3 to Period 5 in both detectors, the fit attenuation length is a little shorter in Period 4 than Period 3, but is longer (about the same) in the FD (ND) in Period 5 than Period 4. The values of the ratios of Period 4 to Period 3 and Period 5 to Period 3 for each cell are shown in Figure A.13 as taken from the calibration constants used in recent analysis. The mean from the profile of these are given in the caption. There is no clear sign that this is the culprit based on the attenuation fit parameter in the relative calibration, given that the ratio seems to recover in the FD.

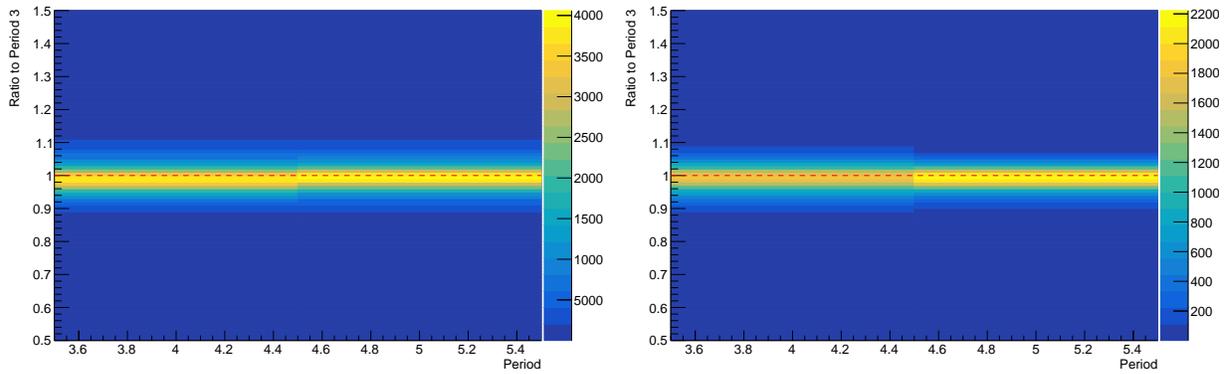


Figure A.13: The ratios of attenuation length for each cell from the relative calibration fit CSVs for the given period with respect to the corresponding attenuation length in Period 3, for the FD (left) and ND (right). The tails are clearly quite large, but the bulk of cells fall slightly below 1 (marked by the red dashed line). The distribution remains about the same in the ND for Period 5 as for Period 4, but for the FD it is a little closer to 1 in Period 5 than in Period 4. The means from a profile of this plot rounded to 3 significant figures are: Period 4/3: 0.993 (FD) and 0.992 (ND); Period 5/3: 0.997 (FD) and 0.992 (ND).

However, it is hard to fully judge this based on the fits of the relative calibration, since only one exponential is used, and it is mainly aimed to pick up bulk attenuation effects in the central regions of cell, with roll-offs in the profiles at the near and far end and extra bits added to attempt to fit these [42]. One may be interested in what happens as one changes the attenuation length only. A toy model was constructed which uses a single attenuation and has both long- and short-path light. In this plot, the readout end is at 0 and the far end is at -1600cm for the FD and -400cm for the ND. Figure A.14 shows what happens if one changes the attenuation length here. This does not match the shape of PE or PE/cm degradation from Figures A.6 and A.7, so any potential change is not accounted for by this term alone. This seems especially true for the FD.

However, since the point here is to consider degraded attenuation along the fibers, a more realistic treatment of the fibers would be ideal. Specifically, a collaborator has worked out new attenuation functions from fiber stringing data [133, 134], which uses a double attenuation length formula separately fit to the ND and FD. I modified these equations to add in “mirror” image terms for the long-path light. This shifts the exact value that one will get at  $x=0$  (the readout end of the cell), but since the point is to determine change from one period to the next, a floating constant

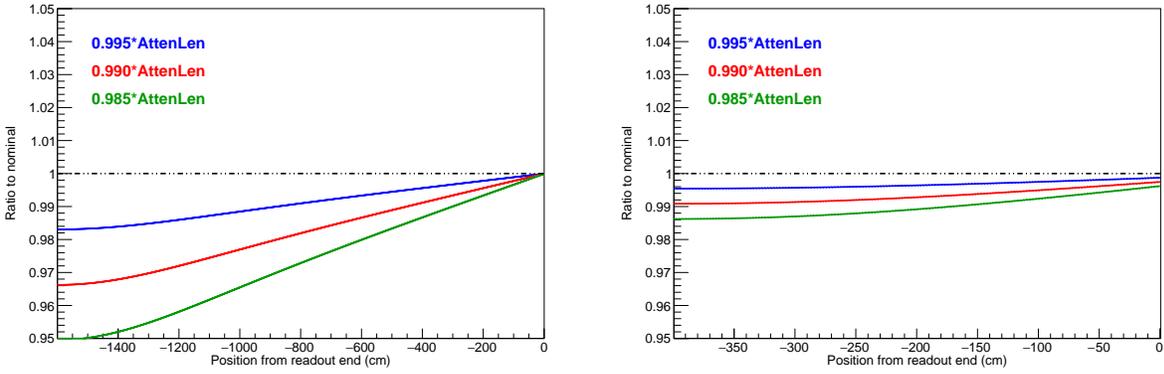


Figure A.14: Taking a naive, single-exponential model and changing the attenuation length on the distance scales of the FD (left) and ND (right). The FD clearly does not match the degradation patterns seen in data. The starting attenuation lengths chosen here roughly correspond to a truncated mean value in the relative calibration fit from Period 3 and then fractional shifts to that number.

multiplicative factor will not change the shape. Furthermore, to account for the extra length of fiber between the end of the cell and the APD, I add in 1m extra offset on each detector as an approximation by making the mirror term for long-path light be 2m more offset and adjusting  $x$  to always measure 1m offset from the input  $x$ . This is why 3400cm appears in the equation for the FD and not 3200cm even though I am assuming the cells are 16m long.

1. ND:

$$0.4624e^{x/175.1\text{cm}} + (1 - 0.4624)e^{x/687.2\text{cm}} +$$

$$0.4624e^{(-1000\text{cm}-x)/175.1\text{cm}} + (1 - 0.4624)e^{(-1000\text{cm}-x)/687.2\text{cm}}$$

evaluated at  $x = w - 100\text{cm}$  for  $w \in [-399\text{cm}, 0\text{cm}]$  as described in text

2. FD:

$$0.6245e^{x/322.3\text{cm}} + (1 - 0.6245)e^{x/902.3\text{cm}} +$$

$$0.6245e^{(-3400\text{cm}-x)/322.3\text{cm}} + (1 - 0.6245)e^{(-3400\text{cm}-x)/902.3\text{cm}}$$

evaluated at  $x = w - 100\text{cm}$  for  $w \in [-1599\text{cm}, 0\text{cm}]$  as described in text

If one plots how this function changes as one reduces the short attenuation length only, the shape

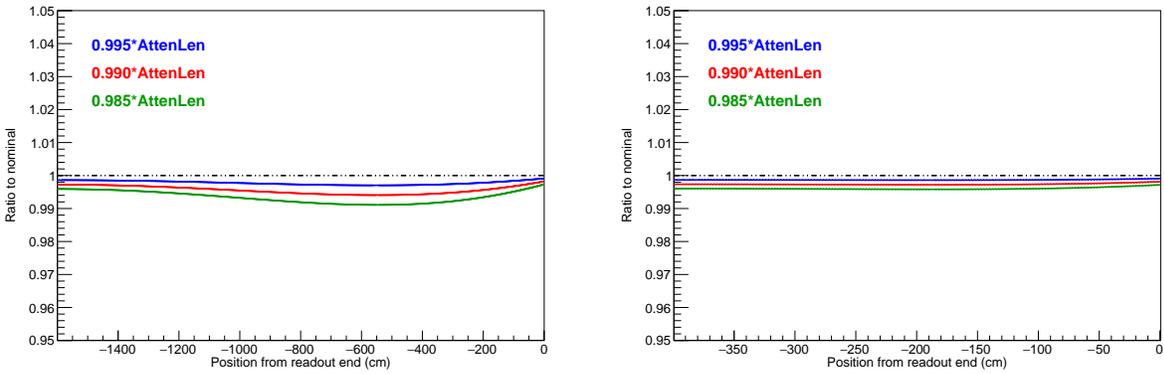


Figure A.15: A somewhat more realistic treatment of fiber attenuations was determined separately for the FD and ND [133, 134]. The functions were adapted to roughly approximate signal as a function of position in cell from point 0 (the readout end) to the most negative point (far end of a cell). This adapted function approximates 1m of extra fiber between the cell and the APD/front-end electronics, and includes a “mirror” term for the long-path light. When one alters only the short attenuation length in the function, one finds a somewhat similar shape to the degradation seen in the detectors.

one finds is actually somewhat similar to the degradation seen in PE/cm and PE. Versions of this are shown in Figure A.15 and should be compared to the data plots of PE and PE/cm in the ND and FD. Note especially the shape in the FD model plot, where the characteristic pinching is seen, and the shape in the ND model plot, which seems a bit flatter. In this way, the toy model has some similarity to the observation. Physically testing fibers after being situated in liquid scintillator is an example of a small test or experiment which could benefit this line of study in the future.

Next, consider the scintillator. It is interesting to note that the scintillator was produced with two variations of fractional composition [40] and the somewhat dimmer scintillator was put mostly into the first diblock (and a little of the second diblock) in the FD. Additionally, there was some discussion that the higher numbered diblocks tended towards slightly lower fill percentages (one block in particular appears less filled in the documentation [135]). A few mechanisms for scintillator aging have been thought about:

1. An idea that has come up is that if the scintillator survived the original fill without being oxygen-saturated, it is possible that over time it could absorb oxygen and result in reduced yield. In this scenario, one would expect the degradation to plateau as the scintillator reached

saturation.

2. Another idea that has come up is whether it is possible that pseudocumene could evaporate out of the scintillator over time. This would be expected to reduce yield.
3. One could imagine a scenario where the scintillator “yellowed” and therefore has increasingly less visibility. In discussions with collaborators (for example during meetings), this may be a less-favored explanation.

It is hard to distinguish given the current data in Figures A.12 if a linear fit is truly the best fit function that could be improved by allowing it to account for the spikes and other periods of greater stability or if another fit function would be more appropriate.

Further, it is hard to say that the drift is slowing down, but the current evidence seems inconclusive. Moreover, it is hard to distinguish if the slightly less rapidly changing PE/cm nearest the readout end in the FD is related to fibers or not. A simple look at attenuation functions shows an intriguing similarity in shapes between the PE/cm degradation and a scaling of the short attenuation length in a two-attenuation model at the FD, though it is not a perfect match. If not for this effect, one might have otherwise thought that the evidence is more in favor of the scintillator being the culprit. Scintillator degradation under some conditions is an effect that has been observed, for example in Borexino [136]. Interestingly, Daya Bay has also noticed the detector response degrading at the level of  $\sim 1\%$  per year [26]. The degradation of PE/cm seems fairly consistent across fiber brightness bins, meaning that if this an attenuation effect, fibers are likely undergoing the same process regardless of their initial conditions. Further study should look into means of distinguishing effects between fiber attenuation and scintillator aging.

### **A.3 Discussion**

This appendix lays out clear evidence for a degradation in detector response over time in both the ND and FD and examines the effects of this drift on both the raw PE signal and the energy scale calibration. Many distributions were presented which show clear signs of this drift, which is at the order of 1% per year in the calibration constants or  $\sim 1/3$  of a PE/cm per year. Section A.2.5 looked

at the available evidence from this line of work and other studies to discuss possible causes.

There is more that can be done moving forward and more studies which can be done to isolate the culprit. While it was noted that detector state/environmental conditions seem plausible as a second order effect (e.g. to explain the bumps in time series plots), there are remaining questions about voltage readout, chiller piping, etc.

Further, while there is some substructure in the plot of PE/cm vs.  $W$  close to the readout end that could be interpreted as a fiber effect, this effect is not a huge change and appears more flat further from the readout end. An attempt was made to use a toy model of attenuation in fibers and does perhaps point to some similarities in the shape of degradation versus position in a cell. This line of study may benefit from a suggestion to run some sets of simulations, one with different fiber attenuation lengths and one with different scintillator brightnesses. Then, producing plots of the same variables as presented above, one could see which set of shifts looked more like the drift effects seen in the data. Such a study is well-motivated to provide further evidence for or against each potential culprit discussed, but it is outside the scope of this work. Work on understanding potential thresholding of cell hits due to this effect are also warranted.

In addition to re-discovering and learning more about the drift in the energy scale in the detectors, this line of study has noted a number of potential deficiencies in the way that calibration constants are determined for analysis periods. For one, Section A.2.2 noted that while the relative calibration removes most of the differences among diblocks, there is still  $\sim 1\%$  spread between high- and low-response diblocks remaining after its application at the FD. Further, the spread is seen to increase over time as the diblocks drift with different slopes (see Figure A.2). Also, calibrating once per period, as had been the standard method in the past, is likely insufficient to account for all the drift which occurs within the given period.

Performing calibration, especially this energy scale calibration, in somewhat shorter periods is warranted. Additionally, the calibration constants are stored in the tagged release as CSV files with one entry for each diblock over some range of runs (say Period 4, which covers run 23420 to 24613 in the FD). Rather than calculating the constants monolithically, it should therefore be straight-

forward to store a different constant for each diblock. FENNEC, which built up the calibration code from what existed to contain more features, allowed one to ask for a given number of runs per “era” and allows one to ask it to calculate a separate number for each diblock. An easier solution is to separate the files into finer groups and calculate calibration constants for each. Furthermore, rather than using the mean of a (slightly) truncated PECorr/cm distribution and dE/dx distribution to make the calculation, one could use other characterizations (median, Landau function, etc.). Finally, NOvA effort in the past had developed a drift term which could be applied to the calibration process (as in Reference [137], for example) but never was. However, by just shortening the time periods used for each calibration, it may be possible to mitigate much of the effects of this drift on the constants.

Prior to the recent analysis, the calibration tag had the most recent relative calibration and absolute calibration pinned to Period 5. This means that more than a year passed since the last relative calibration before the initial  $\nu + \bar{\nu}$  analysis. Since the drift effects start to become appreciable on a timescale of a few years, the absolute calibration constants were updated for Periods 7 and 8 in the ND using the constants in Table A.1 for Epochs 7d and 8b. More importantly, FD data after the initial  $\nu + \bar{\nu}$  analysis was given a new constant from Epoch 7b for the top-up (as in Table A.1). Only the data added into the top-up was given this new calibration: no data that had been used in the initial analysis was re-calibrated. This should effectively “reset” the drift and alleviate systematics concerns for the recent analysis, which is the focus of Chapters 4-6. In the longer term, more work and further study is needed.

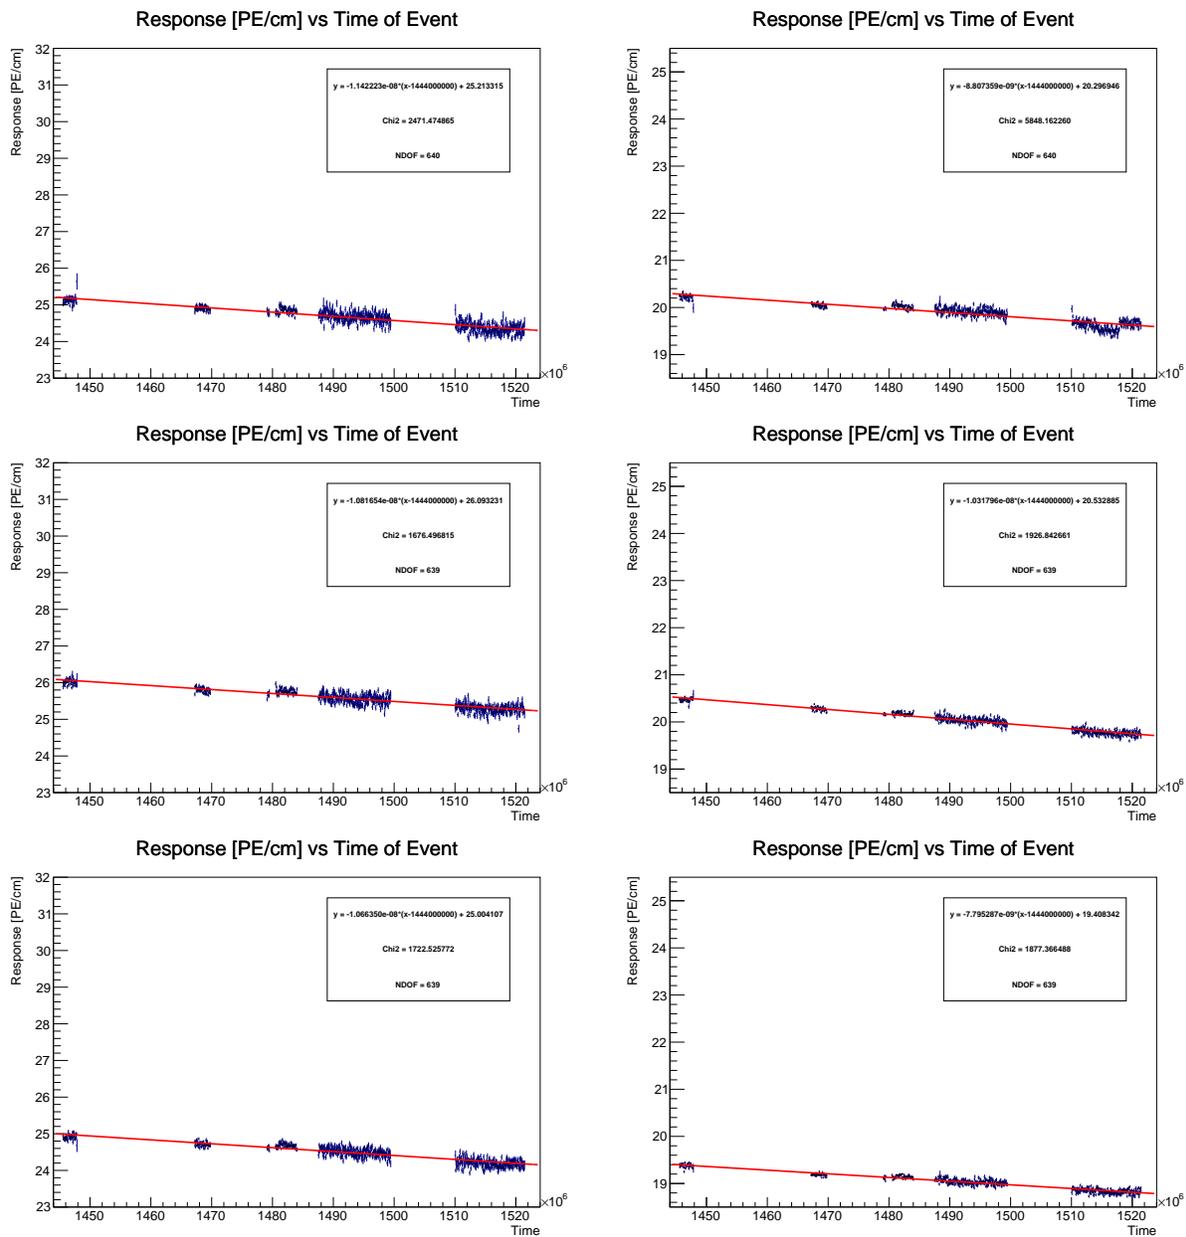


Figure A.16: Linear fits to the PE/cm distributions for diblocks 1-3 at the FD. The box in the plot shows the fit (in PE/cm/s),  $\chi^2$ , and number of degrees of freedom. The left column is x-view, and the right column is y-view. The top row is diblock 1, and each successive row is the next diblock, such that the last row is diblock 3.

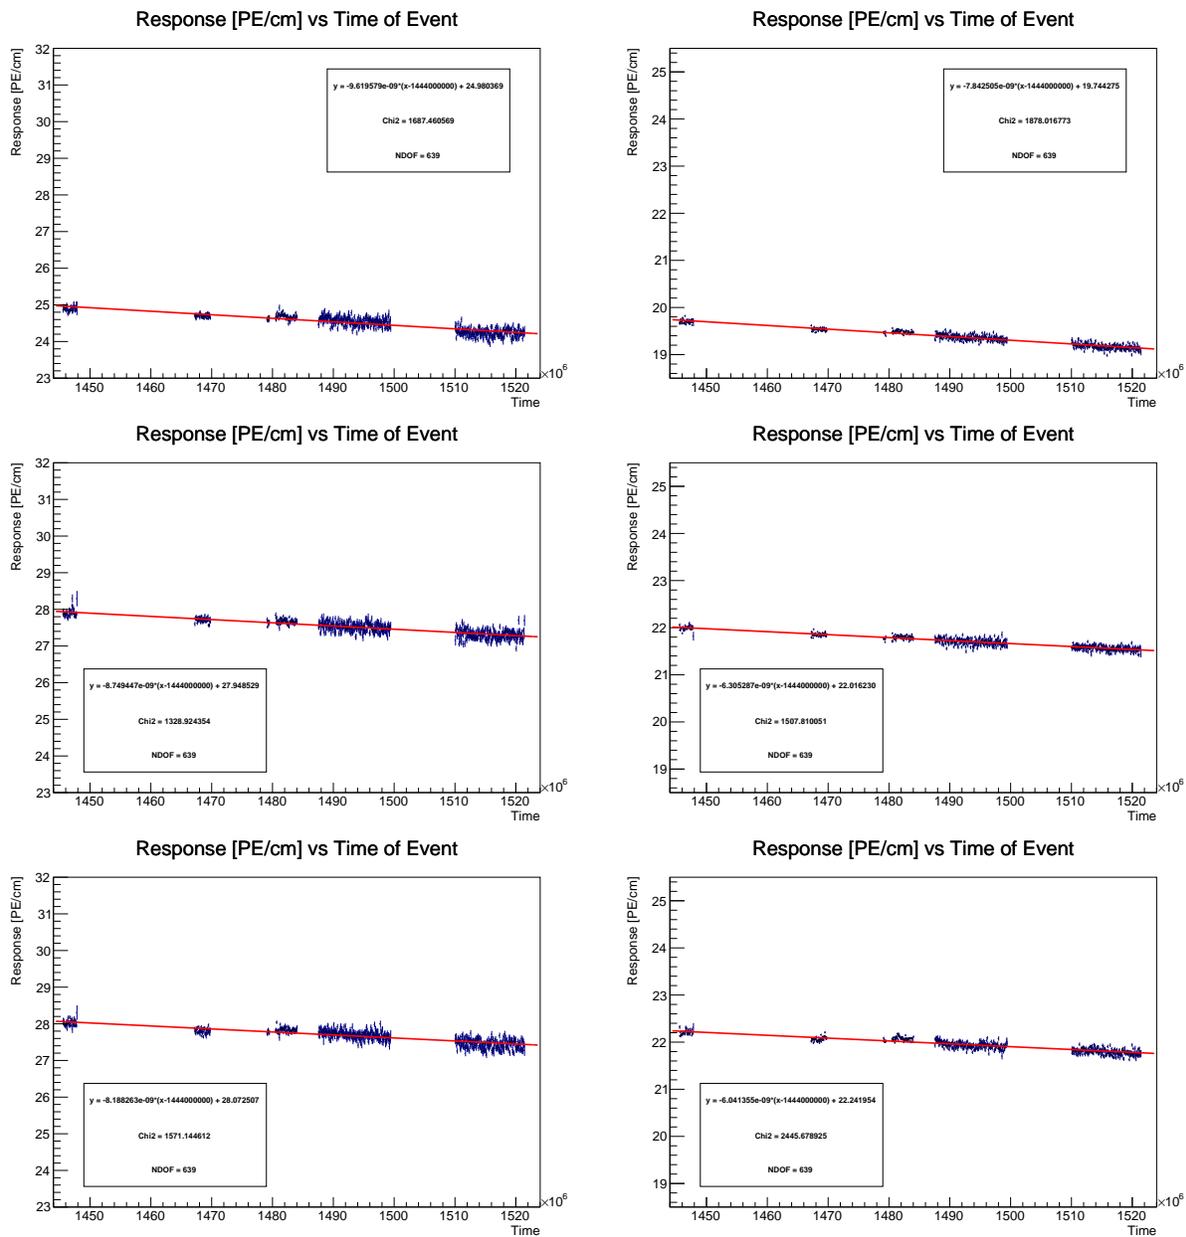


Figure A.17: Linear fits to the PE/cm distributions for diblocks 4-6 at the FD. The box in the plot shows the fit (in PE/cm/s),  $\chi^2$ , and number of degrees of freedom. The left column is x-view, and the right column is y-view. The top row is diblock 4, and each successive row is the next diblock, such that the last row is diblock 6.

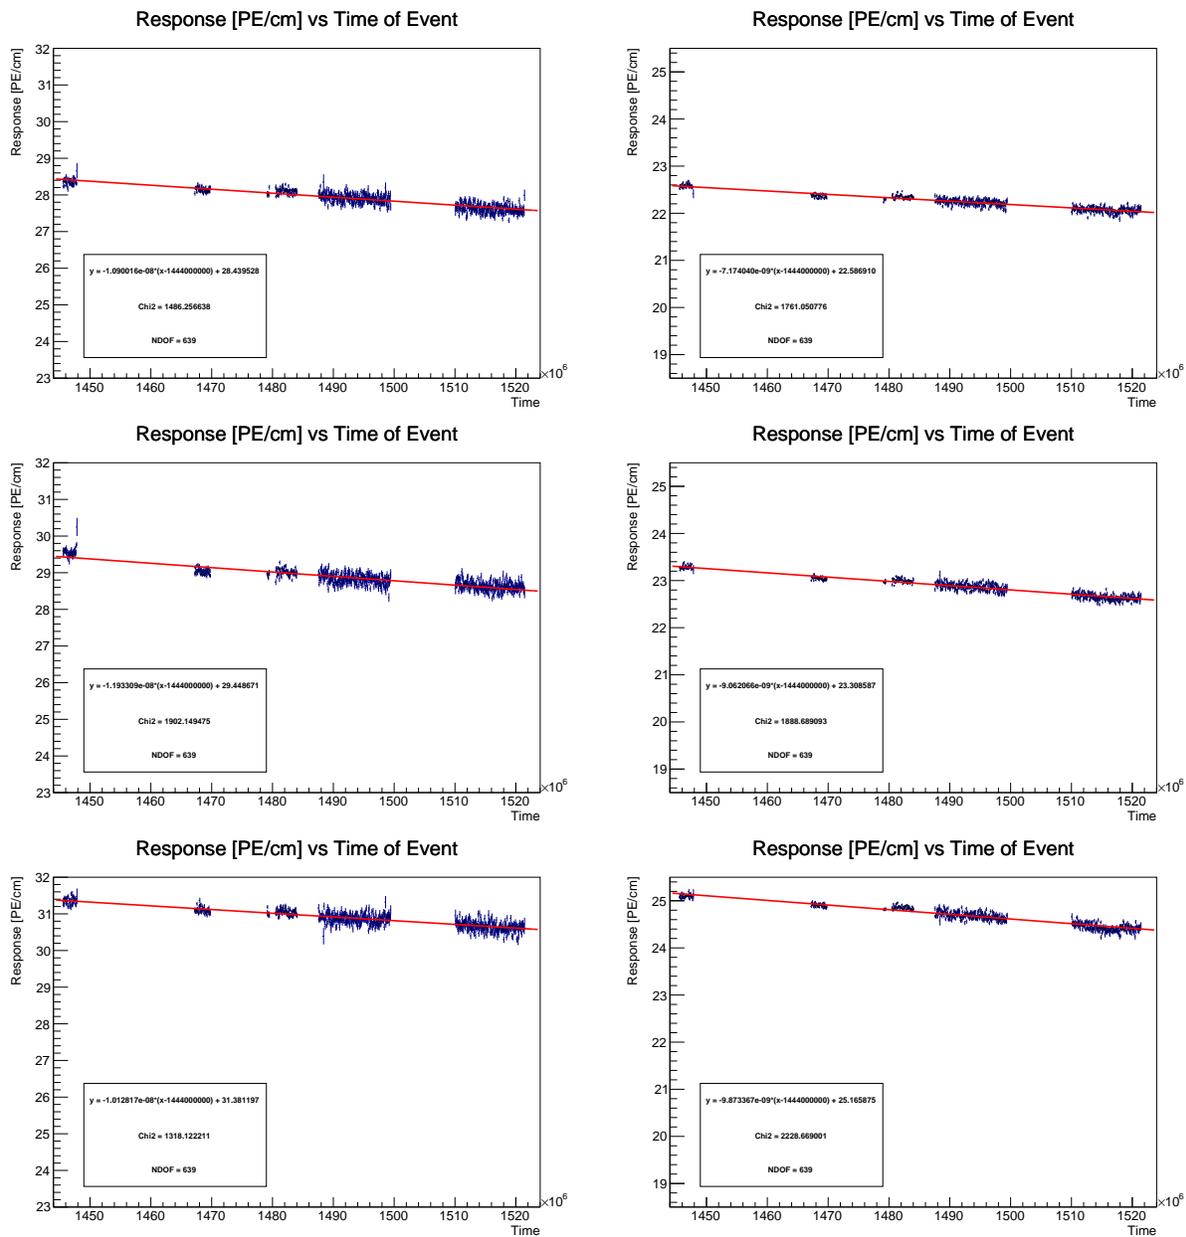


Figure A.18: Linear fits to the PE/cm distributions for diblocks 7-9 at the FD. The box in the plot shows the fit (in PE/cm/s),  $\chi^2$ , and number of degrees of freedom. The left column is x-view, and the right column is y-view. The top row is diblock 7, and each successive row is the next diblock, such that the last row is diblock 9.

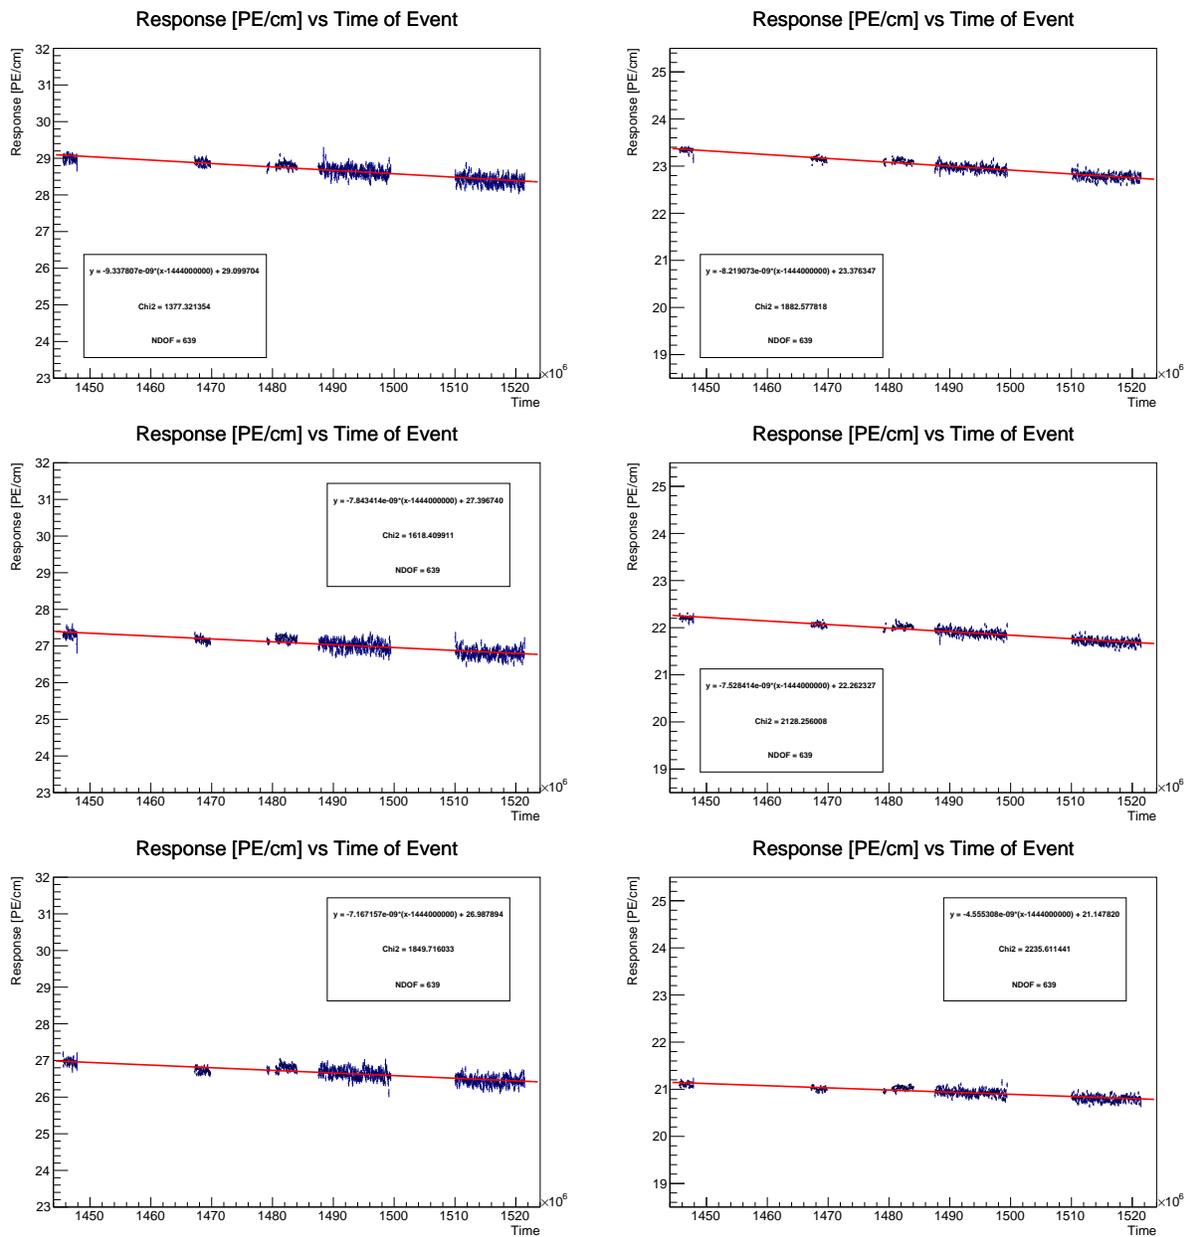


Figure A.19: Linear fits to the PE/cm distributions for diblocks 10-12 at the FD. The box in the plot shows the fit (in PE/cm/s),  $\chi^2$ , and number of degrees of freedom. The left column is x-view, and the right column is y-view. The top row is diblock 10, and each successive row is the next diblock, such that the last row is diblock 12.

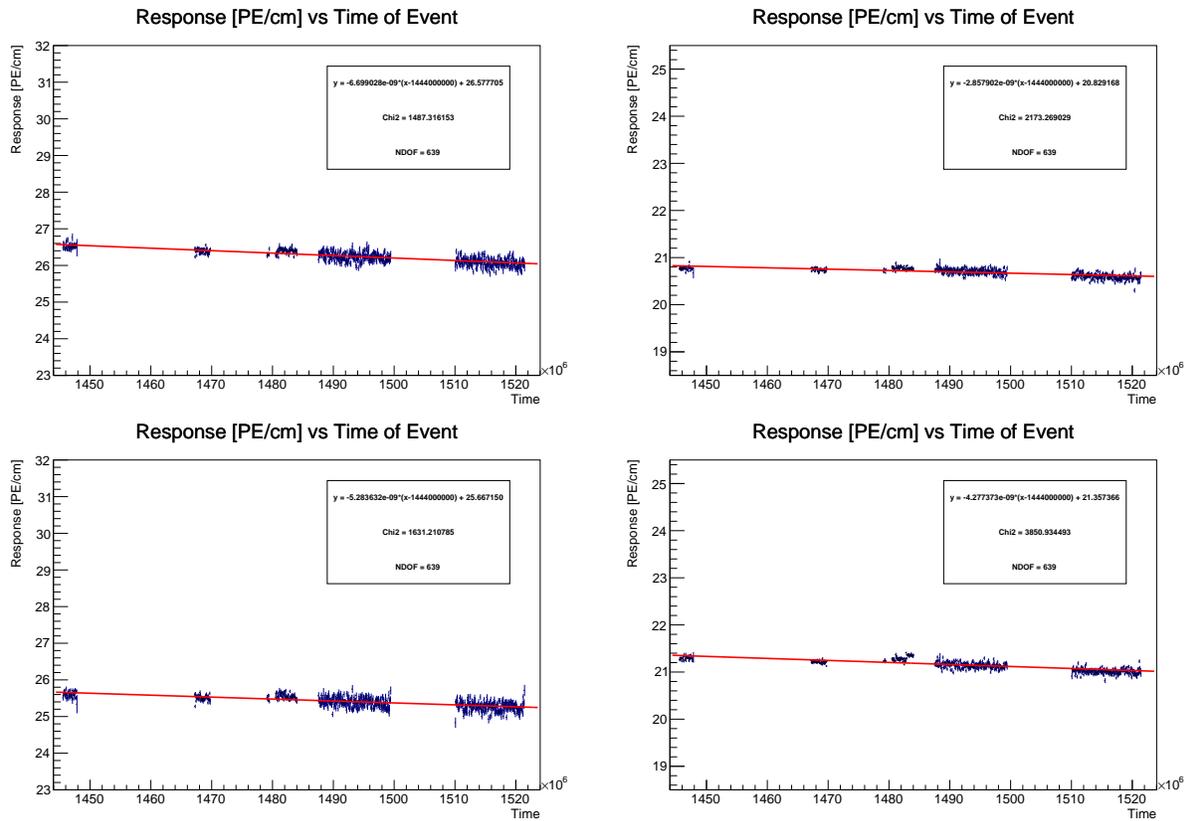


Figure A.20: Linear fits to the PE/cm distributions for diblocks 13-14 at the FD. The box in the plot shows the fit (in PE/cm/s),  $\chi^2$ , and number of degrees of freedom. The left column is x-view, and the right column is y-view. The top row is diblock 13, and each successive row is the next diblock, such that the last row is diblock 14.

## APPENDIX B

### Studies on Timing Resolution of SiPM Readout Electronics

A major purpose of the photon detector system for the Deep Underground Neutrino Experiment (DUNE) is to provide precise timing for events occurring in the volume of Liquid Argon (LAr). While a time stamp is given by the beamline monitoring for neutrinos incident from the Fermilab beam, the reconstruction of events in the time-projection chamber (TPC) can be improved by having more precise timing than is given by the drifting electrons. Additionally, DUNE will be sensitive to non-beam events (e.g. supernova neutrinos, atmospheric neutrinos, and proton decay events), wherein no beamline timing is available. Here, one needs precise timing of events; a timing resolution of at least approximately a microsecond is needed [138]. Scintillation photons from interactions in LAr can be utilized to provide such timing, and work is being undertaken to equip DUNE with a photon detection system.

Using an LED and fibers to send signals to prototype photon detection/detector (PD) modules submerged in cryogenic liquids provides a method of assessing the timing resolution of detector components. In this study, two signals separated temporally by a known amount provides access to the timing characteristics of the readout system. Offline processing and analysis of the collected waveforms is used to determine the timing resolution of the readout.

#### B.1 Experimental Setup

The Indiana University dewar provides a useful test stand for DUNE PD prototyping work. The dewar was used with a prototype PD module and an Argonne Silicon photo-multiplier Signal Processor (SSP) to digitize waveforms to test the timing resolution of the PD readout. Furthermore, the TallBo facility, housed in the Proton Assembly Building (PAB) at Fermilab and consisting of an 84-inch tall dewar, provides a larger volume test stand with very pure liquid argon that allows one to use cosmic rays to test PD prototypes [139] (see also related Reference [140]). One can use methods just like those used at the local dewar to study the timing resolution of the readout system

at this larger test stand, as well. Several prototypes have been tested at the TallBo facility, and in one such instance this setup provided a set of timing tests.

To perform studies using an LED, a driver was used to output a short signal from an LED. The driver takes as inputs an external voltage supply and a function generator. The function generator acts as a trigger for the driver. In order to extract the timing resolution of the PD readout, the LED must be pulsed quickly enough that two pulses appear in each digitized waveform from the SSP. To achieve this, a RIGOL function generator output a burst-mode square pulse function, so that two triggers were separated by  $5\mu\text{s}$  and repeated at a set interval. A picture of the LED driver is shown in Figure B.1 [141], along with a schematic of the function.

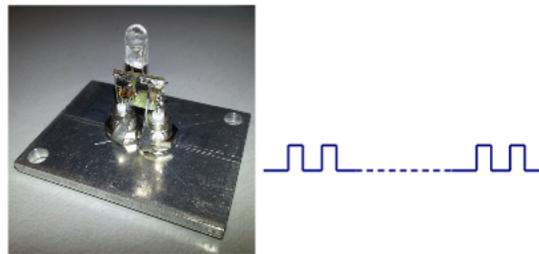


Figure B.1: Photo of LED driver with LED installed and a mock-up of the pulse structure created by the RIGOL generator [141]. The rising edges of the double pulse structure are separated by  $5\mu\text{s}$  and the first rising edge of each burst of pulses is separated by a larger amount of time.

The 370nm driven LED was placed adjacent to the end of an optical fiber. The main fiber splits and connects to feed-throughs slightly above the top of the dewar. The dewar itself contains a cryogenic liquid. Local tests were conducted with liquid nitrogen (LN2) to keep the electronics cold without introducing scintillation signals, and TallBo testing was conducted with LAr, as reading scintillation signals were a key point in that experiment. The fibers inside the dewar carry the LED signals to a piece of plastic placed horizontally along the PD paddle(s). This piece serves the purpose of diffusing the LED signals and illuminating the bars covered with wavelength shifter (WLS). The WLS types used shift ultraviolet wavelengths nearly to the visible spectrum. A picture of the PD paddle with fibers and diffusing plastic is shown in Figure B.2 [141].

As seen in Figure B.2, the PD paddle is a prototype PD module, holding 4 bars and 12 SiPMs (SensL B-Series and C-Series have been used), with each bar having 3 corresponding silicon photo-

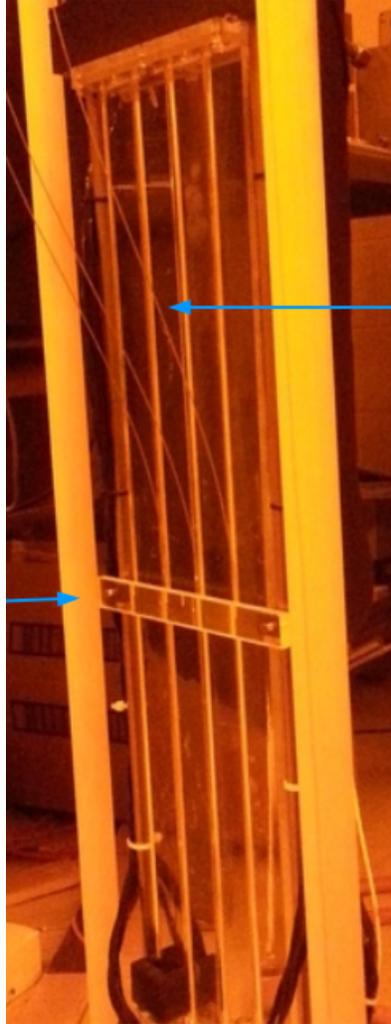


Figure B.2: PD paddle used in the Indiana University dewar [141]. The arrow on the left points to the diffusing plastic and the arrow on the right points to one of the feed-through wires that transmits the LED signal to the diffusing plastic.

multipliers (SiPMs). The photons from the LED that are converted in the bar and propagated down its length can then be detected by the SiPMs. The SiPM signals are read by the SSP, which can handle 12 input channels. In this study, the SiPMs are numbered such that each takes up a channel, and channels 0-2 are on one bar, channels 3-5 are on another bar, channels 6-8 are on a third bar, and channels 9-11 are on a fourth bar.

Though the bars in the PD module vary in crafting technology and/or WLS used, this is not expected to be an issue. The signals read by an SiPM and compared in the analysis are propagating through the same bar from a similar point along the bar, meaning that the effects of a given bar having a different composition, technology, or small defects should not affect the timing results. Then, the spread in the distribution of timing differences between the two LED signals on a given channel is reflective of the timing resolution of the readout system. The SSP had a working goal of better than 30ns timing resolution [107].

Using the SSP along with SiPMs provides excellent readout capabilities [139] (see also related References [140, 142]). For example, this system provides excellent photoelectron (PE) resolution. The SSP records the sum of 20 samples in the peak region. Dividing this sum by 20 then yields the average ADC counts per bin in the peak region, providing a feel for the PE peak resolution in this study. This is plotted below, in Figure B.3, as an example. Though the pulsed LED signal on the SiPM covers a wide range of PE, the resolution makes each PE peak clearly visible. Looking at the plot, there are 31 PE peaks in the shown range of about 380 (averaged) ADC counts, so there are  $\sim 12$  (averaged) ADC counts per PE peak for this example. With this information, the leftmost peak (at about 60 (averaged) ADC counts) corresponds to 5 PE. Counting peaks from there, the two highest peaks then correspond to 17 or 18 being the most likely PE count.

Simply, one could also note that there are about 4 PE peaks for every 50 (averaged) ADC counts in Figure B.3. For example, there are four peaks between 50 and 100 (averaged) ADC counts. This again yields  $\sim 12$ -13 counts per PE peak. Assuming 12 counts per PE peak, 17 PE would correspond to  $\sim 204$  (averaged) ADC counts, and 18 PE would correspond to  $\sim 216$  (averaged) ADC counts. Indeed, the two PE peaks with the most events fall close to these numbers. Thus, this method also

corresponds to 17 or 18 being the most likely number of PE. One could similarly find the most likely PE for other SiPMs and in other data sets as well. Though there will be some variation in (averaged) ADC counts per PE peak from channel to channel, the most likely (averaged) ADC counts from each SiPM channel will be reported in the rest of the document to provide a general idea of the correlations between signal strength and timing resolution.

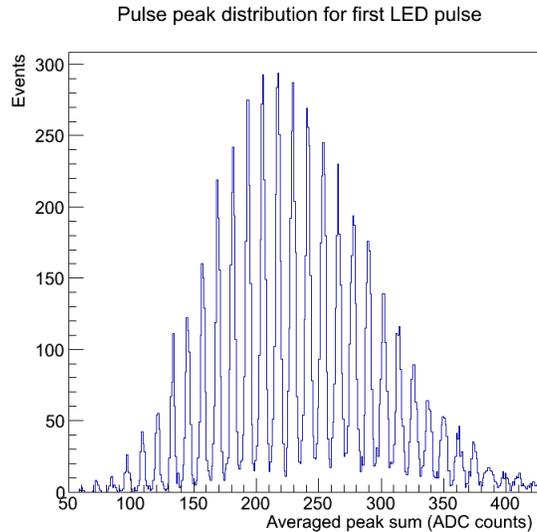


Figure B.3: Signal distribution showing the ADC counts recorded (divided by 20 samples to yield an average) in the region of the peak corresponding to the first LED pulse over a large number of events. There are 31 PE peaks in the approximately 380 (averaged) ADC counts shown, or about 12 averaged counts per PE in this case. This means that 17-18 is the most likely number of PE in an event. Using the alternative method provided in the text yields the same result.

The RIGOL function generator also outputs a logic pulse for every burst, and this pulse is used to trigger the SSP to take data, utilizing its 150MHz processing capability to sample approximately every 6.67ns. Using this external trigger reduces the influence of dark noise or scintillation signals. Offline processing and data analysis are then used to extract the desired quantities, and these are the subject of the following section.

## B.2 Data Analysis and Results

Data were taken using the SSP and stored in ROOT files. Offline processing to subtract the waveform from the baseline creates positive-going waveforms as shown in Figure B.4.

The two LED pulses read by the SiPM are readily visible in Figure B.4 and are separated by

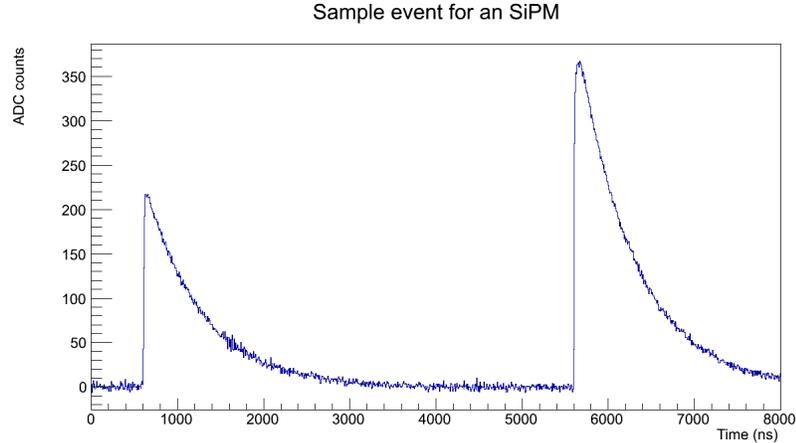


Figure B.4: A sample waveform adjusted to take into account the baseline background. The SiPM response shows clearly the signals corresponding to the two LED pulses, separated by approximately  $5\mu\text{s}$ .

approximately  $5\mu\text{s}$ . A baseline subtracted waveform is then analyzed. First, the given waveform is searched for two peaks. A peak here is defined to be a sample with ADC counts greater than or equal to the preceding and following 10 samples.

The external voltage supply was set to a value high enough that almost all events to be analyzed have peaks above 35ADC. To avoid calling small background fluctuations peaks, a candidate for a peak must have ADC counts greater than 35. This selection criterion also acts to cut events with especially small peaks, avoiding any oddities that may be caused by using this method to try to describe the leading edges for such small peaks. Additionally, to avoid calling small background that sometimes occurs between the LED signals the second peak, a further selection cut is made to search for the SiPM signal from the second LED pulse within a region from approximately 4746.67 to 5253.33ns (712 to 788 samples) after the first pulse peak.

The analysis focuses on the middle of the pulse's rising edge to determine its characteristics. The reason for doing this is that one can approximate the middle of the rising edge by a line, without worrying about the rising edge's curved beginning and end, and extrapolate to the point at which half the peak was reached; this method was explored by GlueX as a method to obtain timing

resolutions [143]. The linear approximation then gives

$$B'_i = mt_i + b \quad (\text{B.1})$$

for the signal in the  $i^{\text{th}}$  sample, where  $m$  and  $b$  are constants to be determined and  $t_i$  is the time of the  $i^{\text{th}}$  sample.

The sample occurring directly before and the sample at or directly after the waveform reaches half-maximum intensity (indicated by  $B_{\text{bef}}$  and  $B_{\text{aft}}$ , respectively) are then used to obtain the values of the constants [143]. By using the equations for  $B'$  before and after, one finds

$$m = \frac{B_{\text{aft}} - B_{\text{bef}}}{t_{\text{sample}}}, \quad (\text{B.2})$$

where  $t_{\text{sample}}$  is the time between samples and

$$b = B_{\text{aft}} - mt_{\text{aft}}. \quad (\text{B.3})$$

This then gives the result that

$$t_{\text{char}} = \frac{\frac{B_{\text{peak}}}{2} - b}{m}, \quad (\text{B.4})$$

where  $t_{\text{char}}$  is the characteristic time of the pulse and  $B_{\text{peak}}$  the signal at the peak.

Then  $\tau = (t_2 - t_1) - 5000$ , where  $t_2$  and  $t_1$  refer to the characteristic times for the second and first pulses respectively, represents the difference (in ns) in pulse start times relative to the expected value. For each output channel, the results for  $\tau$  over many events are plotted in a histogram ranging from -20ns to +20ns with bin size 0.4ns. The width of the distributions, as given by ROOT, reveal the time resolution of the PD readout system. As the distribution is for two separate timings, the timing resolution for a single signal is  $\frac{\text{RMS}}{\sqrt{2}}$ .

The first set of data analyzed was taken on one day. Due to some SiPMs not functioning properly, only three SiPMs have been analyzed for this data set. More than 20,000 double pulse events were analyzed for each SiPM. On a second day, data were taken with the same settings as before

and provide a second set of data for analyzing the timing resolution, this time for 11 SiPMs, as only one did not work properly. The testing at TallBo included in this study used 3 full PD modules (denoted as paddles A, B, and C) of 4 bars and 12 SiPMs each, and one module (denoted as paddle D) holding a bar and a fiber-based technology, thus using 5 SiPMs. Of the 41 SiPMs in use, some SiPMs did not function properly or had an unexpected PE distribution. Other SiPMs saw only faint signals or were receiving signals from a bar with no WLS and have thus also been cut. Therefore, the results of analysis of 25 SiPMs have been included from testing at TallBo. The results presented from TallBo are based on at least 12,000 double pulse events passing the selection cuts for each SiPM. While additional troubleshooting might have allowed for the use of additional SiPMs, the reduced number of SiPMs still give adequate results.

The data sets are summarized in Table B.1 below. In the following analyses, the numbers 0-11 will be used for the channels in data sets 1 and 2, as they used the same paddle of 4 bars and 12 SiPMs. This numbering scheme will allow for comparison of the channels between data sets. For the analysis of data set 3, a similar numbering scheme is used, but with a preceding letter to denote the paddle containing that channel. For example, channel A0 refers to the first SiPM on paddle A.

<b>Data Set</b>	<b>Facility</b>	<b>SiPMs analyzed</b>
1	Indiana University	3
2	Indiana University	11
3	TallBo (Fermilab)	25

Table B.1: Data sets and the important characteristics of each set.

The distribution of characteristic times for the first and second pulses for one SiPM is shown in Figure B.5. As the SSP samples approximately every 6.67ns, the feature in the plot on the left of Figure B.5 at approximately 607ns is very close to a boundary between samples.

The characteristics of the timing distributions are given by ROOT and appear below for each of the three data sets in Tables B.2, B.3, and B.4. Figure B.6 shows three such timing distributions. Though not perfect, the shapes of the timing distributions are reasonably Gaussian. The value given by  $\frac{\text{RMS}}{\sqrt{2}}$  specifies the timing resolution and is thus the value quoted in the tables to two significant digits. Further inspection revealed that there were some events in the underflow/overflow bins of

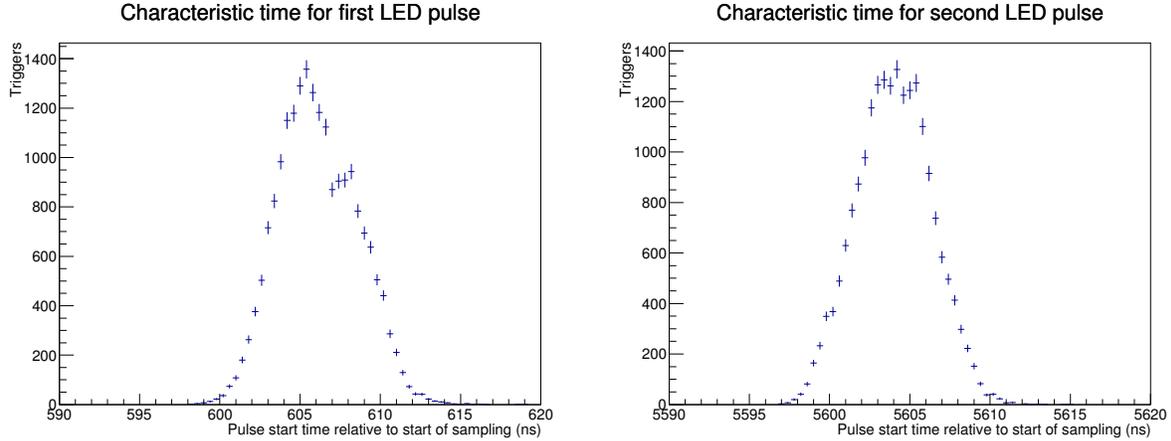


Figure B.5: Characteristic times (extrapolated time of half-maximum) for the first (left) and second (right) LED pulses for a channel from data collected in the Indiana dewar. The feature at approximately 607ns is close to a division between samples.

the histograms, though clearly the bulk of events were contained in the distributions. The SiPMs from Paddle C in a different subset of the same TallBo test were used to understand the scintillation time constants in a published paper [92] and it is therefore expected that these SiPMs should be well-behaved. Using a wider window ( $\pm 100\text{ns}$ ) would only change the time resolution reported of 2 of the 12 channels and only by 0.1ns. While there were a few events (an exceedingly small fraction) even outside this window, these events were found to pick up on stray features in the waveform or coincident pulses not from the LED signal, accounting for the large differences. This is one advantage to placing a cut of  $\pm 20\text{ns}$  on the histogram: it cuts away undesired events. For all these reasons, it was determined  $\pm 20\text{ns}$  is a fine time window for the purposes of this study.

The analysis of the data from the Indiana dewar suggests the resolution lies between 0.85 and 1.8ns, and the analysis of the TallBo data suggests the resolution lies between 0.48 and 2.1ns. As seen in the plots in Figure B.6, the means of the distributions are not the expected value of 0ns. This shift is indicative of a systematic effect, perhaps in some instrumental response. Also included in the tables are the most likely (averaged) ADC counts from the first LED pulse. For the example given in Figure B.3, this would correspond to the location of the peak. Though there will be variation from channel to channel in terms of number of (averaged) ADC counts per PE, this estimates the signal strength. For cases where two equal peaks occur, the average of the two is

reported (to the nearest integer) to characterize its signal strength. There appears to be a correlation between the width of the timing distribution and the strength of the signal due to photon statistics, hence the range of RMS values.

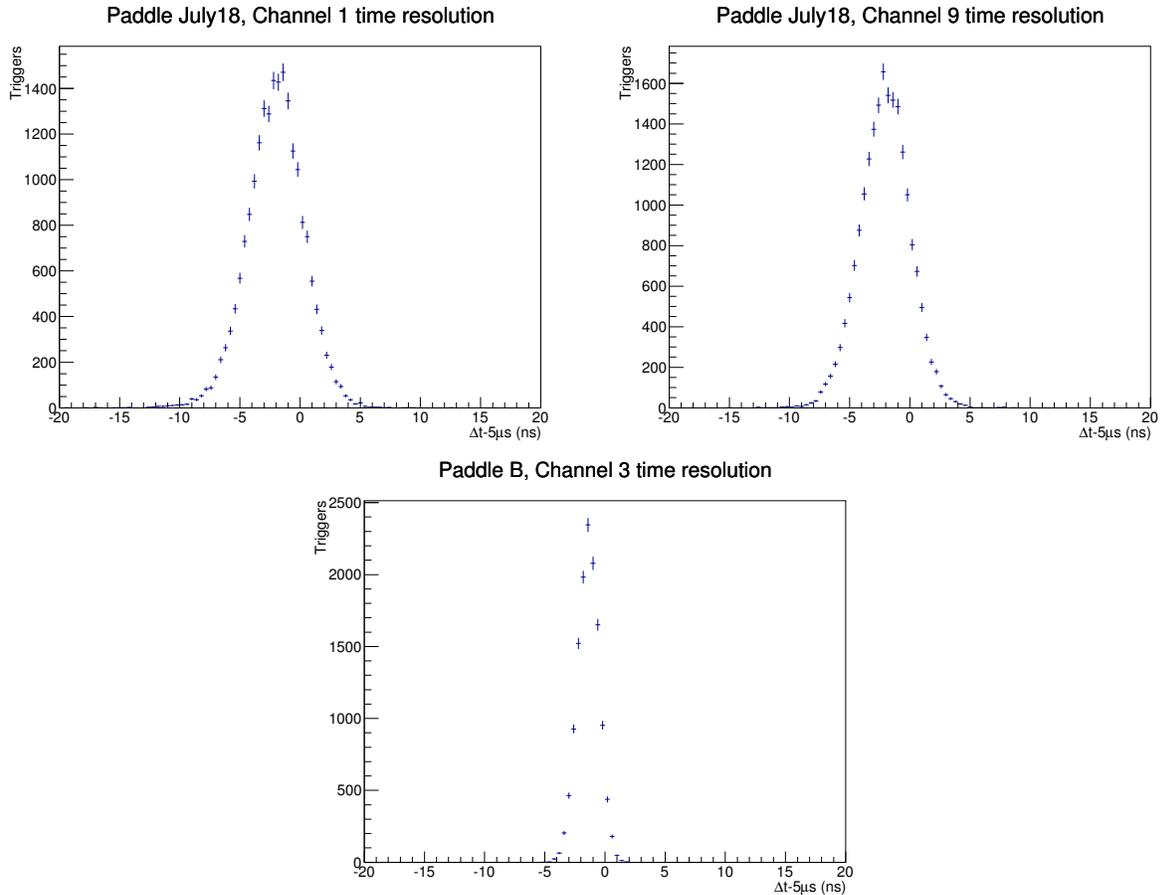


Figure B.6: The timing distributions in  $\tau$  for 3 channels using this analysis method. The distribution from paddle B corresponds to a very bright signal (see Table B.4). Errors shown correspond to statistical counting uncertainties. That the mean is shifted negative from the expected 0ns is likely due to a systematic effect.

<b>SiPM channel</b>	<b>RMS/<math>\sqrt{2}</math> (ns)</b>	<b>Most likely signal strength (ADC)</b>
1	1.7	126
2	1.7	157
9	1.5	217

Table B.2: Results for the RMS for the timing distribution of the readout and approximate signal strength characterization for data taken in the first Indiana data set. The break separates the channels by bar. The most likely signal strength is defined as the peak in the channel's signal distribution and is thus given in (averaged) ADC counts.

<b>SiPM channel</b>	<b>RMS/<math>\sqrt{2}</math> (ns)</b>	<b>Most likely signal strength (ADC)</b>
1	1.8	177
2	1.8	173
3	.94	603
4	.93	624
5	.92	579
6	.87	776
7	.85	878
8	.85	790
9	1.8	186
10	1.8	187
11	1.8	155

Table B.3: Results for the RMS for the timing distribution of the readout and approximate signal strength characterization for data taken in the second Indiana data set. Channels 1, 2, and 9 refer to the same SiPMs as in the first Indiana data set. Breaks separate the channels by bar. The most likely signal strength is defined as the peak in the channel's signal distribution and is thus given in (averaged) ADC counts.

SiPM channel	RMS/ $\sqrt{2}$ (ns)	Most likely signal strength (ADC)
A0	1.7	240
A1	1.6	299
A4	1.9	222
A6	.94	639
A7	1.0	591
B0	1.0	955
B1	1.0	843
B2	1.0	916
B3	.62	1840
B4	.63	1765
B6	.93	729
B7	.91	915
B8	.93	850
C0	2.1	117
C1	2.0	138
C2	2.1	117
C3	.49	1470
C4	.48	1460
C5	.51	1315
C6	1.2	287
C7	.94	450
C8	.94	558
C9	1.5	231
C10	1.4	276
C11	1.6	204

Table B.4: Results for the RMS for the timing distribution of the readout and approximate signal strength characterization for data taken at TallBo. Breaks separate the channels by bar. The most likely signal strength is defined as the peak in the channel's signal distribution and is thus given in (averaged) ADC counts.

## APPENDIX C

### Studies on multiplexing SiPM readout

Chapter 8 raised the concept of ganged (multiplexed) silicon photomultipliers (SiPMs) to facilitate adding SiPMs and/or potentially reducing the number of readout channels necessary to implement the Photon Detection (PD) system for a DUNE single-phase detector. In the main text, it was noted that passive ganging - coupling SiPMs in parallel - was perhaps the easiest solution, but leads to altered waveforms (with the SensL SiPMs). Active ganging presents multiple paths forward and can overcome some of the drawbacks of passive ganging, but has issues of its own and is more complicated.

This appendix will briefly discuss early testing of SiPMs, especially comparing individual SiPM signals to those of a gang of SensL SiPMs. In short, degradation of the peak pulse height and a lengthening of the recovery time is seen. A gang of three SensL SiPMs was found to produce distinguishable signals even at the 1 photoelectron (PE) level, but a gang of six was less favorable for 1 PE signals. The major lightguide bar-based PD instrumentation used in the protoDUNE experiment chose to gang the SiPMs by 3 at a time. There are some indications that these properties are dependent on the type of SiPM used, and/or one might reach a different conclusion if changing components or conditions (e.g. deciding 1PE signals are unnecessary). Therefore the best option may vary for future PD systems.

Because the dark rate of each SiPM is on the order of 10Hz at LAr temperatures, running an experiment for several minutes of live time is a sufficient method to acquire thousands of waveforms for analysis of low level signal pulses (1 or 2 PE). A small open dewar was used to house a single board of SiPMs with 3 individual SiPMs, a passive gang of 3 SiPMs, and a passive gang of 6 SiPMs. The SiPM board was suspended in LN<sub>2</sub> in the dewar using thin screw rods. A picture of a different electronics board in the same setup is shown in Figure C.1 [144]. The setup was further used to test an active ganging board, signals over varying cable lengths, and warm versus cold cables [144].



Figure C.1: Photograph of an electronics board used to test active ganging properties in the setup used for ganging tests [144]. The board has SensL SiPMs but is not the board used for the tests detailed here. Bubbling LN2 can be seen in the dewar.

Only the passive ganging board over a short readout cable (10m) is discussed here.

For this study, the header information saved by the SSP was used in part to analyze the data. A variable stores the sum of several samples around the peak of the waveform: a histogram of this variable then shows very clearly a set of individual peaks corresponding to discretized levels of PE. The peak bin was found from the first few PE peaks and fit with a Gaussian in the region  $\text{peak}-2\text{ADC}$  to  $\text{peak}+2\text{ADC}$ , such that the mean and width of the PE peak could be determined. Waveforms within 1 sigma deviation from this peak are then averaged to compute the average waveform for a given channel at a given number of PE. In this way, the waveforms corresponding to 1 SiPM, a gang of 3 SiPMs, and a gang of 6 SiPMs could be compared. A PE distribution for individual SiPMs is given in Figure C.2, with a Gaussian fit to the 1PE peak overlaid.

The triggering sample is placed  $\sim 150$  samples into the waveform ( $1\mu\text{s}$ , recall a sample is  $\sim 6.67\text{ns}$ ), such that the baseline RMS can be determined by calculating the RMS of the first 100 samples. This RMS is  $\sim 2.5\text{ADC}$ , depending on the cable lengths. The triggering threshold for the SSP's self-trigger functionality is ideally at or about  $1/2$  PE, such that it will trigger on essentially all real signals. However, one wants to be careful to not introduce a lot of baseline noise. With

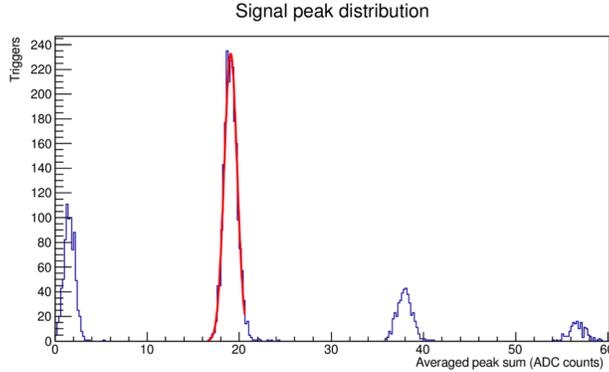


Figure C.2: Histogram of peak pulse heights for individual SiPMs, from Reference [145]. The Gaussian fit to the 1PE peak is overlaid.

the RMS at approximately 2.5ADC, lowering the threshold much below 7ADC quickly introduces dominating triggers from baseline noise. This is demonstrated in Figure C.3. The leftmost peak in each plot are triggers from baseline noise, and the second is the 1PE peak. Note for a threshold of 7ADC (bottom figure), the two are similar in height, but for a threshold of 5ADC (top left), the baseline triggers peak is  $\sim 100$  times more prevalent than the 1PE peak. The result for a threshold of 6ADC (top right) is intermediate.

Since the signals are SiPM dark noise, there is less 2PE than 1PE, less 3PE than 2PE, etc. This is seen in Figure C.3 and is true because to get 2 or more PE from a dark noise signal, one essentially needs afterpulsing or cross-talk. For the gang of six SiPMs, it was found that the 1PE peak was too close to the trigger threshold to reliably trigger, and detected 1PE and 2PE signals were similar in rate. For this reason, the 1PE peak for the gang of six was not used, but the 2PE peak was instead used. For the individual SiPMs and the gang of 3, both the 1PE and 2PE peaks were used. An average waveform for 2PE signals was made in the manner described.

The maximum of a running of average of 5 samples gives the ADC of the peak in the average waveform, and the ADC/PE was taken from the average of PE peaks for the first two consecutive PE analyzed. For the gang of 6, this is 2PE and 3PE. A straight line is determined from two bins in the rising edge. The rise time was characterized by the time taken for such line to traverse 10% to 90% of the pulse height. Errors were determined by using the standard deviation from lines corresponding to variations of the data in the waveform using the min-to-max and max-to-min of

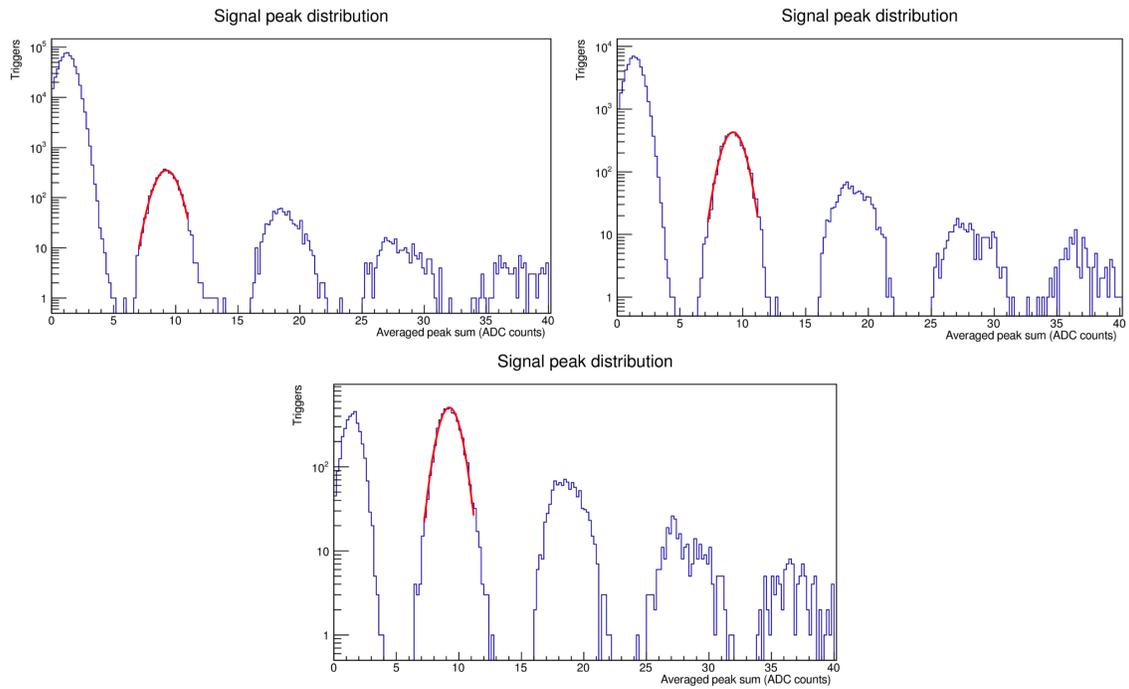


Figure C.3: Histogram of peak pulse heights in three minutes of live time with the same configurations but for a different trigger threshold, from Reference [145]. The leftmost peak in each plot are triggers from baseline noise, and the second is the 1PE peak. Note for a threshold of 7ADC (bottom figure), the two are similar in height, but for a threshold of 5ADC (top left), the baseline triggers peak is  $\sim 100$  times more prevalent than the 1PE peak. The result for a threshold of 6ADC (top right) is intermediate.

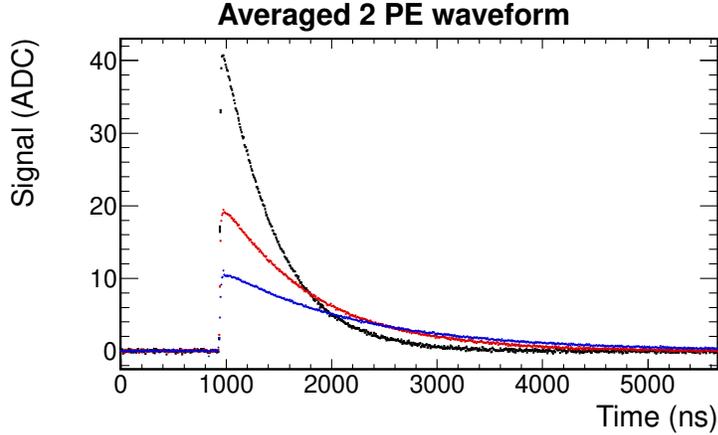


Figure C.4: Averaged 2PE waveforms for tests of individual SiPMs (black), a passive gang of 3 SiPMs (red), and a passive gang of 6 SiPMs (blue). Note that the peak ADC/PE trends downward with added SiPMs. Correspondingly, the recovery time gets longer with additional SiPMs.

the error ranges of the two bins. Finally, the recovery time is determined by a decay time constant, that is a single exponential fit. A fit for the time region after the peak where the five-bin running average is between 70% and 10% of the maximum was found to be sufficiently explanatory.

The results can be seen qualitatively by looking at the averaged waveforms presented in Figure C.4, and quantities are presented in Table C.1 as in [144]. The error of RMS comes from the RMS of the histogram of baseline RMS. There is clear reduction in the peak ADC/PE with an increase in number of SiPMs ganged, and a clear increase in the recovery time.

SiPM	ADC/PE	1PE peak		2PE peak		Baseline RMS (ADC)
		Rise Time (ns)	Recovery Time (ns)	Rise Time (ns)	Recovery Time (ns)	
Individual	20.3	$13.4 \pm 0.5$	473	$12.8 \pm 0.7$	468	$2.45 \pm 0.30$
Gang of 3	9.5	$19.0 \pm 0.6$	892	$16.2 \pm 0.6$	863	$2.51 \pm 0.18$
Gang of 6	5.3	–	–	$19.9 \pm 1.0$	1317	$2.53 \pm 0.18$

Table C.1: Properties of SensL SiPMs, individually and in passive gangs of 3 and 6, as in [144].

## APPENDIX D

### Further NOvA and DUNE oscillation sensitivities with wrong-sign sign separation

#### D.1 NOvA

The study detailed in the main text in Chapter 9 took place at a number of exposures, and the following set of Figures D.1-D.4 show how the one-dimensional profiles evolve as exposure is increased. Further, it was noted that the study was conducted for 3 values of  $\sin^2 \theta_{23}$ , but for expediency most plots only showed the results for the 2019 analysis best-fit value. Figure D.5 shows an example for  $\delta_{CP} = 3\pi/2$  at  $36 \times 10^{20}$  POT in both beams, with all three  $\sin^2 \theta_{23}$  values.

While none of these plots show much improvement, one can check that the idea should work in principle. Figure D.6 shows what the slice would look like if the experiment ran only antineutrino beam after the 2019 analysis and managed to collect  $3.6 \times 10^{22}$  POT. Indeed, the addition of higher antineutrino beam stats with lower neutrino beam stats does show more noticeable results for the current wrong-sign separation techniques.

Moving on from one-dimensional sensitivities to fully explore the fits, contours explore two-dimensional variable space under the same types of constraints detailed for the one-dimensional case. For exploratory purposes, Figure D.7 shows likelihood contours at 68% (approximated as  $1\sigma$ ),  $2\sigma$ , and  $3\sigma$  for  $\sin^2 \theta_{23}$  vs  $\delta_{CP}$  with the default oscillation parameters ( $\delta_{CP} = 0$ ,  $\sin^2 \theta_{23} = 0.5$ ). These use the ultimate planned exposure for NOvA,  $36 \times 10^{20}$  POT in both beams. The plots show the small gain in sensitivity from the application of the wrong-sign  $\nu_e$  BDT.

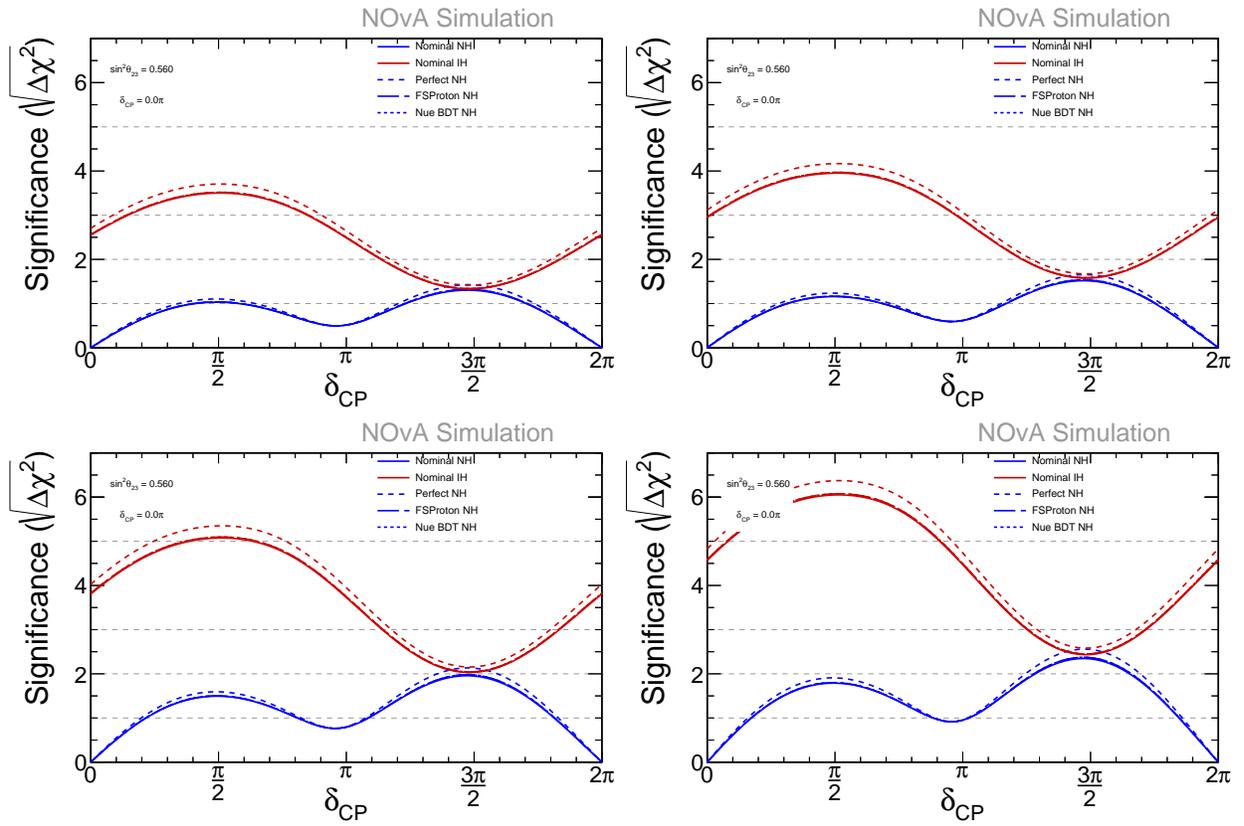


Figure D.1: Slices for true  $\delta_{CP} = 0$ ,  $\sin^2 \theta_{23} = 0.56$  at the 2019 analysis POT (top left),  $15 \times 10^{20}$  POT in both beam modes (top right),  $25 \times 10^{20}$  POT in both beam modes (bottom left), and a full  $36 \times 10^{20}$  POT in both beam modes (bottom right).

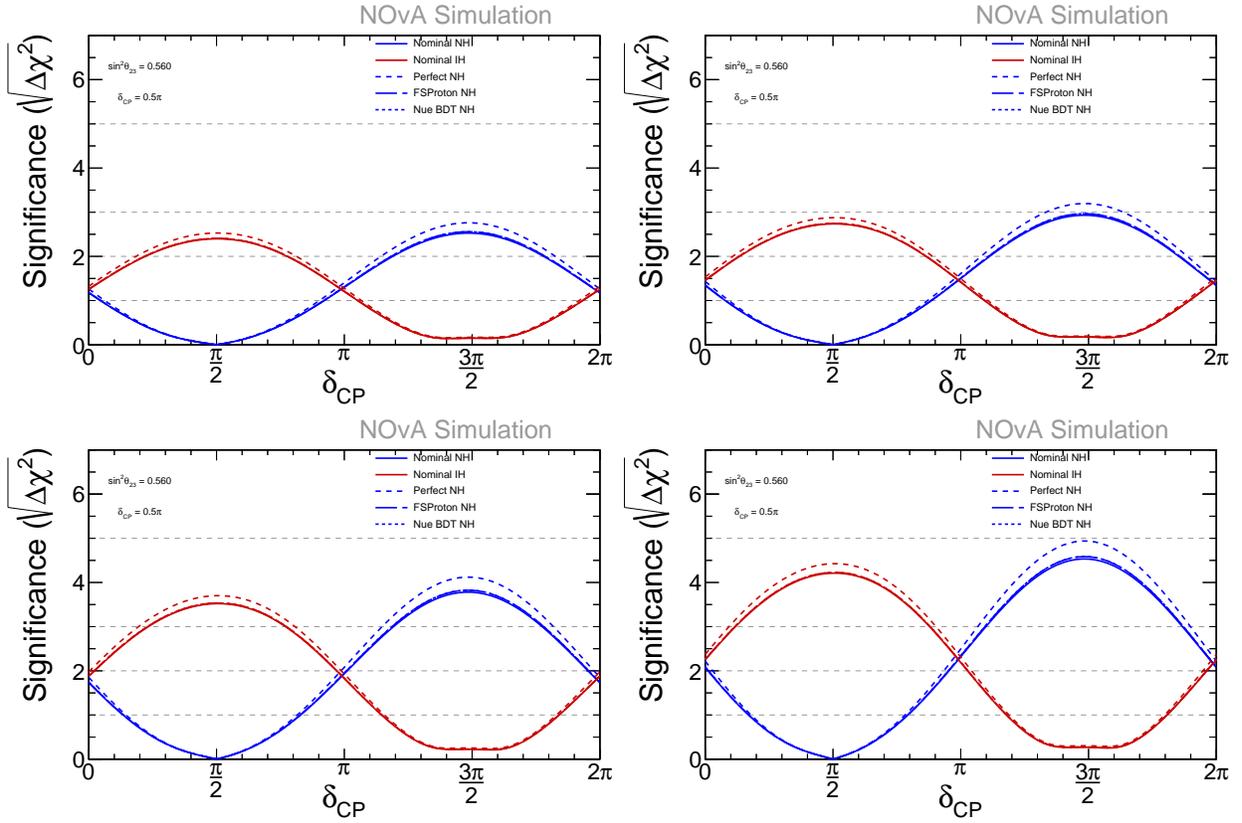


Figure D.2: Slices for true  $\delta_{CP} = \pi/2$ ,  $\sin^2 \theta_{23} = 0.56$  at the 2019 analysis POT (top left),  $15 \times 10^{20}$  POT in both beam modes (top right),  $25 \times 10^{20}$  POT in both beam modes (bottom left), and a full  $36 \times 10^{20}$  POT in both beam modes (bottom right).

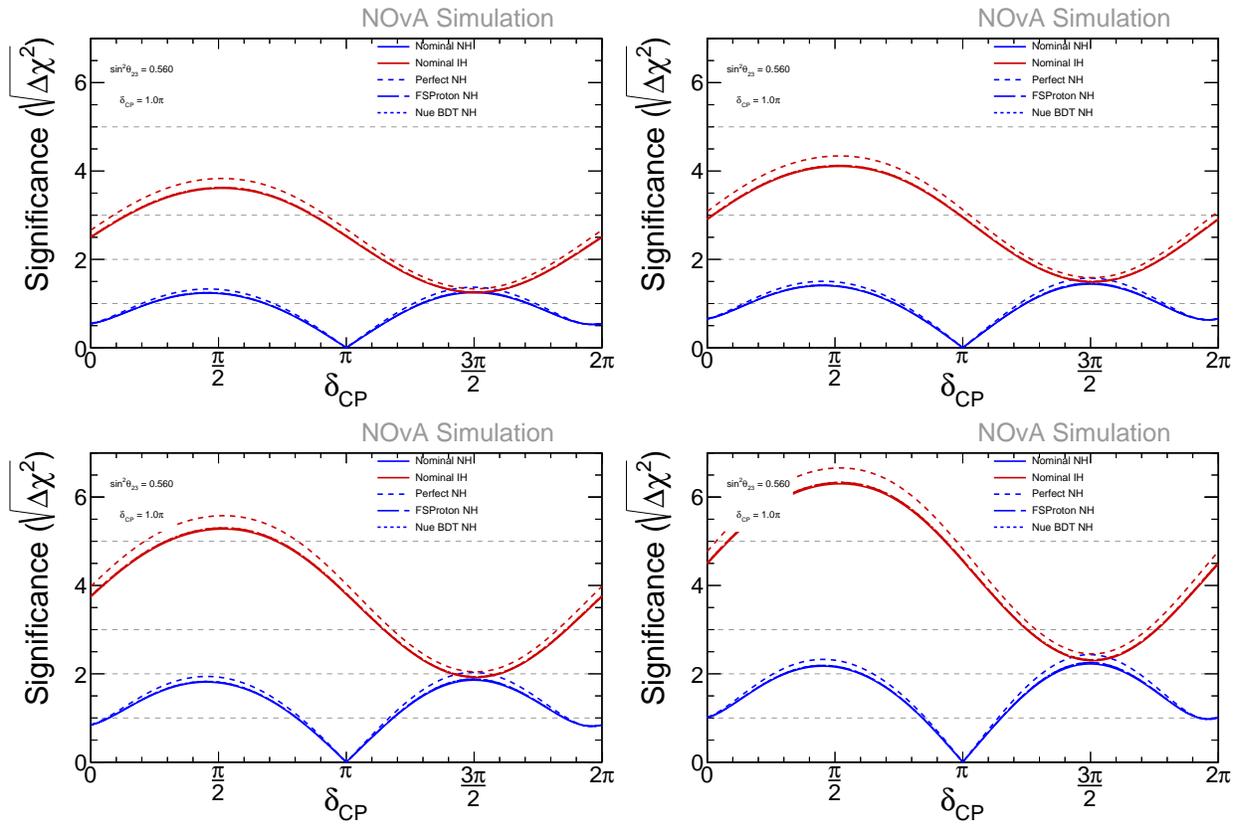


Figure D.3: Slices for true  $\delta_{CP} = \pi$ ,  $\sin^2 \theta_{23} = 0.56$  at the 2019 analysis POT (top left),  $15 \times 10^{20}$  POT in both beam modes (top right),  $25 \times 10^{20}$  POT in both beam modes (bottom left), and a full  $36 \times 10^{20}$  POT in both beam modes (bottom right).

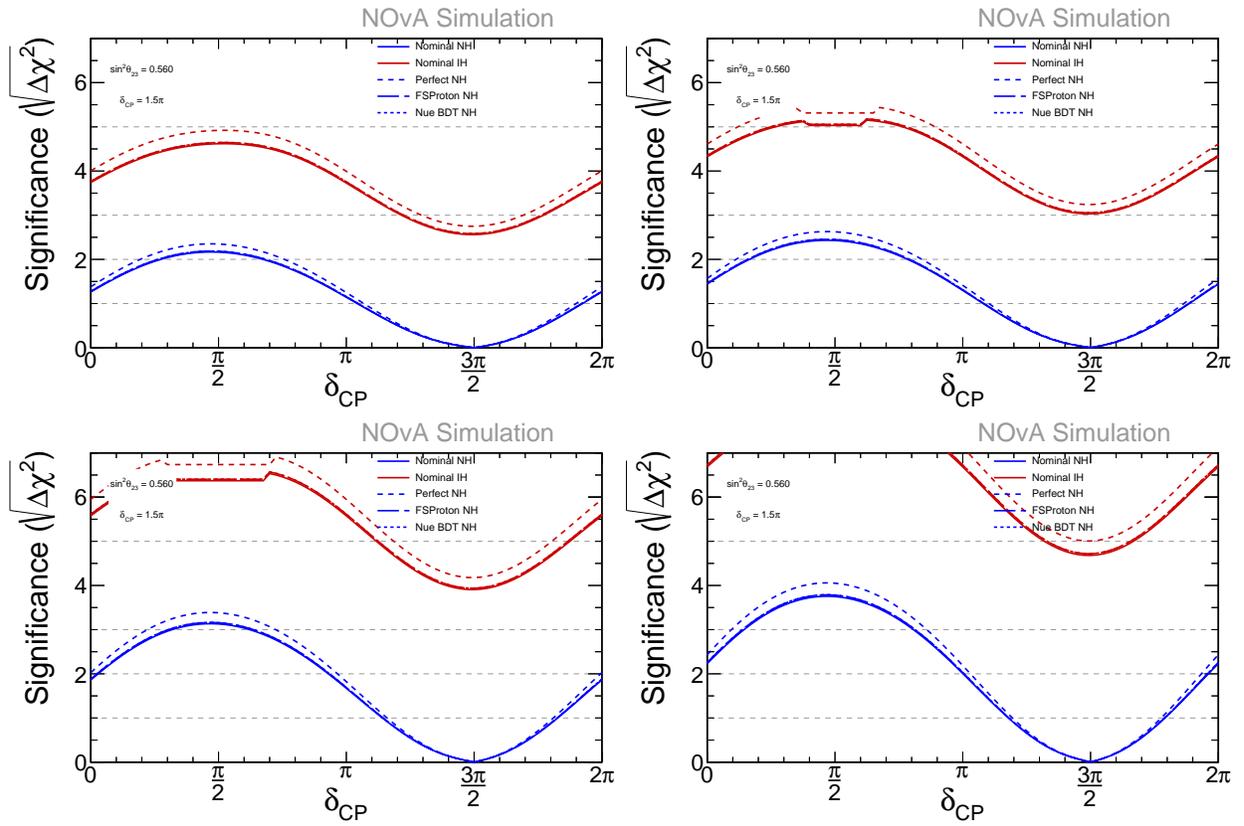


Figure D.4: Slices for true  $\delta_{CP} = 3\pi/2$ ,  $\sin^2\theta_{23} = 0.56$  at the 2019 analysis POT (top left),  $15 \times 10^{20}$  POT in both beam modes (top right),  $25 \times 10^{20}$  POT in both beam modes (bottom left), and a full  $36 \times 10^{20}$  POT in both beam modes (bottom right).

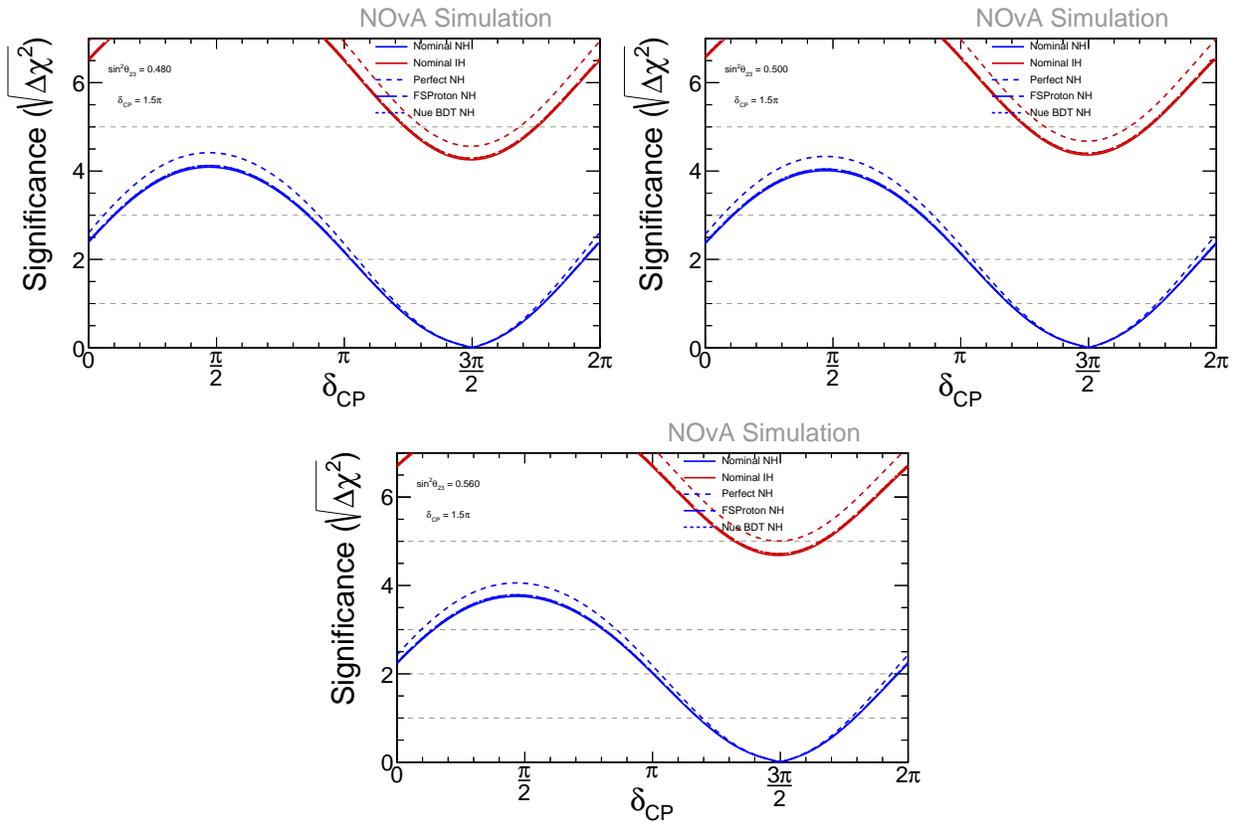


Figure D.5: Slices for true  $\delta_{CP} = 3\pi/2$  with a full  $36 \times 10^{20}$  POT in both beam modes with true  $\sin^2 \theta_{23} = 0.48$  in the lower octant (top left), true  $\sin^2 \theta_{23} = 0.5$  for maximal mixing (top right), and true  $\sin^2 \theta_{23} = 0.56$  in the upper octant (bottom).

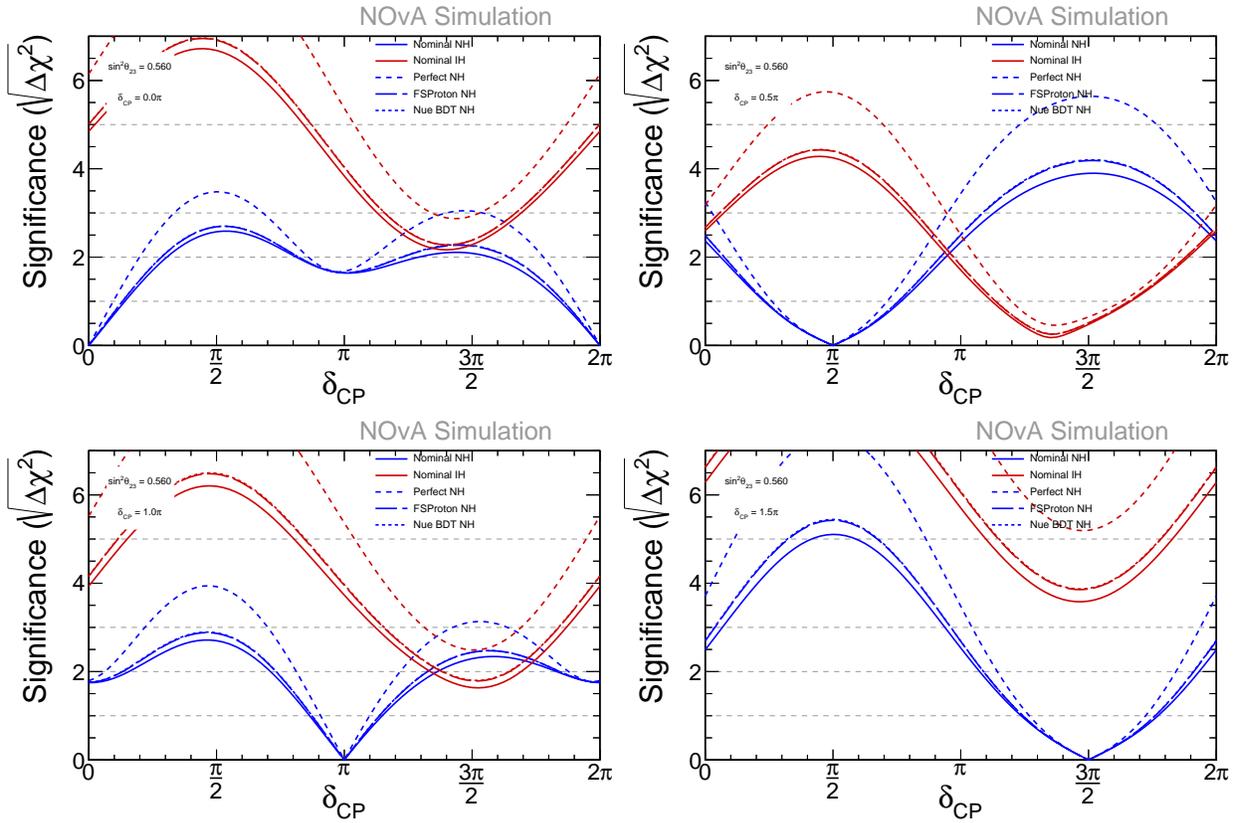


Figure D.6: Slices for true  $\delta_{CP} = 0$  (top left),  $\pi/2$  (top right),  $\pi$  (bottom left),  $3\pi/2$  (bottom right) at  $\sin^2 \theta_{23} = 0.56$  with the 2019 analysis POT in neutrino beam and a hypothetical  $360 \times 10^{20}$  in the antineutrino beam. The current wrong-sign separation techniques can be seen to have some effect here, which shows that with a large amount of additional stats one could actually see some benefit from this technique. Of course, NOvA will not collect this many POT, but this figure is meant to be instructive.

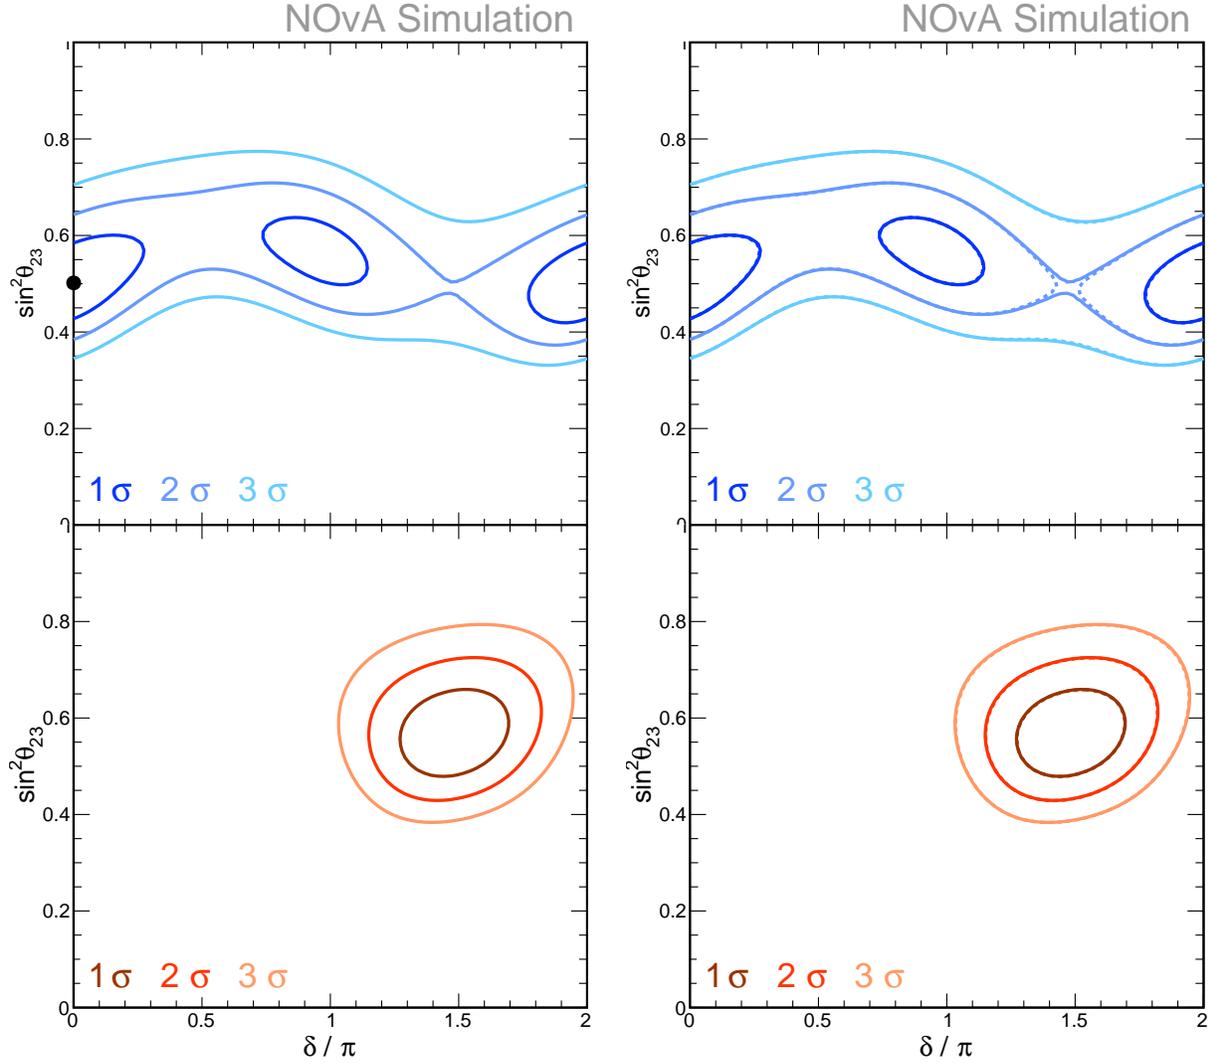


Figure D.7: Contours for  $\sin^2 \theta_{23}$  vs.  $\delta_{CP}$  at the projected ultimate exposure for NOvA,  $36 \times 10^{20}$  POT in both beams. The left plot shows the nominal case with the best-fit lining up well with the chosen true parameters. The right plot shows the nominal again in solid lines and the version found when applying the wrong-sign  $\nu_e$  BDT as a separation technique in dashed lines. The 1 sigma line is an approximation, as this used the 68% Gaussian surface. Only small improvements are found when using this BDT.

## D.2 DUNE

The main text showed the results plots for 7 years of data-taking (336 kt\*MW\*yr). The text noted that as exposure increased, some of the profiling issues arise in regions where the sensitivities are very high. Figure D.8 shows the curves with 15 years of exposure for DUNE (1104 kt\*MW\*yr), split equally between neutrino beam and antineutrino beam. Though there are issues with the profiling, especially at high sensitivities, it appears to be okay especially at the lowest significance CP conserving point. The CP conservation rejection curves are drawn in Figure D.9.

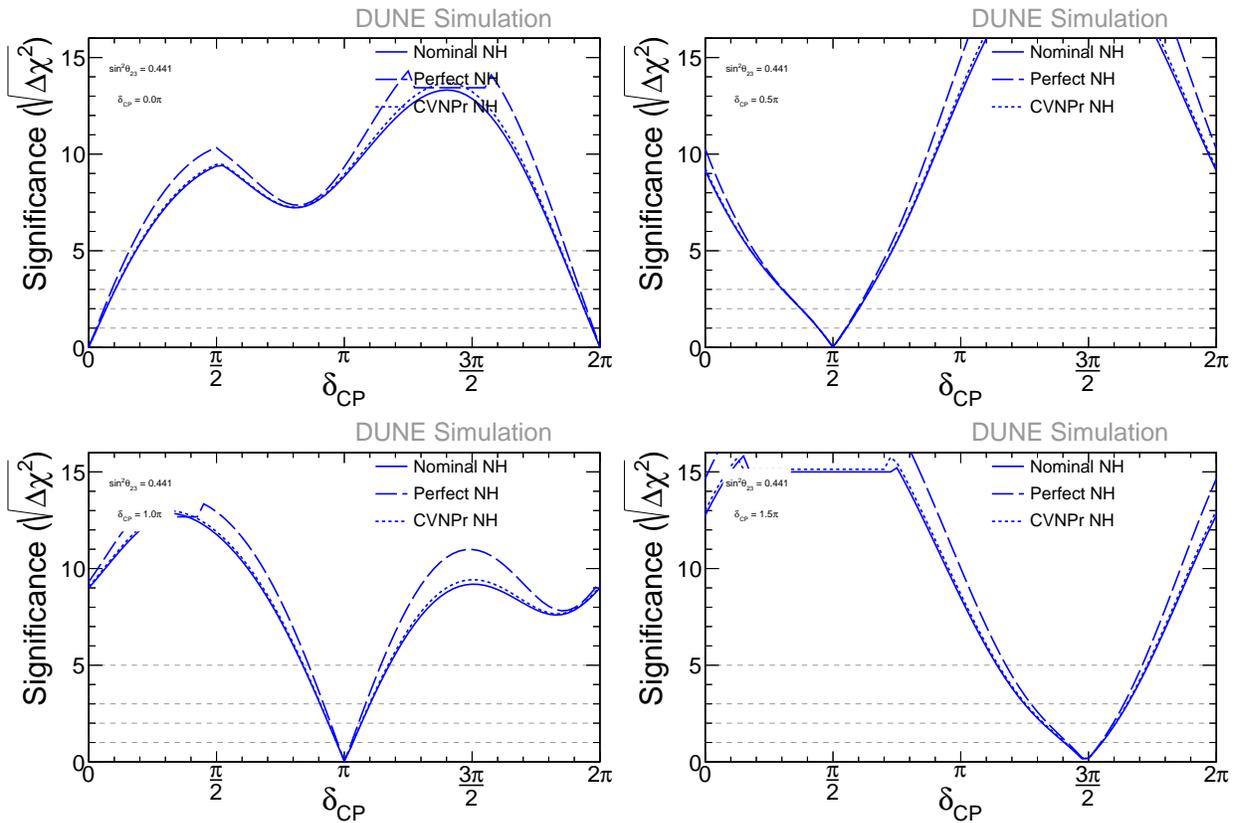


Figure D.8: Slices for true  $\delta_{CP} = 0$ (top left),  $\pi/2$ (top right),  $\pi$ (bottom left),  $3\pi/2$ (bottom right) with 15 years of exposure for DUNE (1104 kt\*MW\*yr), split equally between neutrino beam and antineutrino beam. The wrong-sign separation techniques plotted for this slice use cut-offs in the CVN antineutrino score of 0.2 and 0.5.

One further study conducted was to test the situation if *only* antineutrino beam is run. In this case, the significance gained from studying  $\nu_e$  appearance is gone, and in a nominal sense one would have a line on the biprobability plot rather than a point. Then, one might expect to gain back some

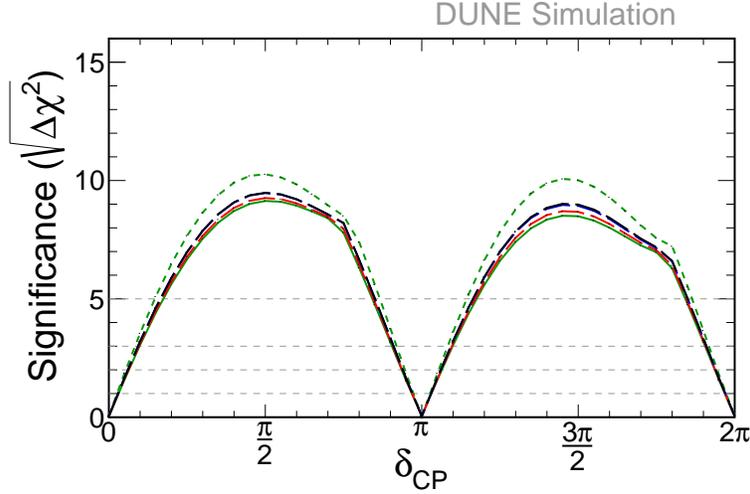


Figure D.9: CP conservation rejection curves for true  $\sin^2 \theta_{23} = 0.441$  at 15 years exposure for DUNE, split equally between neutrino beam and antineutrino beam. The red dashed curve is for the CVN Proton score method and provides only modest improvement over the nominal case. The CVN antineutrino score appears to provide considerably improved significance, though there is still room for potential improvement. A perfect WS separation is shown by the green dashed line.

sensitivity to both neutrino and antineutrino oscillation from performing wrong-sign separation. This is somewhat like the idea of performing sign separation in an atmospheric study, where the flux does not change but having a neutrino-enhanced and antineutrino-enhanced sample allows for higher sensitivity to parameters such as CP violation. The perfect wrong-sign separation case would give both a strictly  $\nu$  and strictly  $\bar{\nu}$  sample for analysis. Analysis with the wrong-sign separation techniques would then be expected to be somewhere between the two cases. Indeed, even the CVN Proton separation produces reasonable improvement. Figures D.10-D.13 show results of this study. Recall that AntiN in the legends refer to separation based on the CVN Antineutrino score, and CVNPr refers to separation using the CVN Proton score. Interestingly, the case of perfect separation still yields potential  $5\sigma$  sensitivity to CP violation within a reasonable amount of time, even if only the antineutrino beam is run. This will certainly depend on the precision with which backgrounds are understood, but is an interesting result nonetheless.

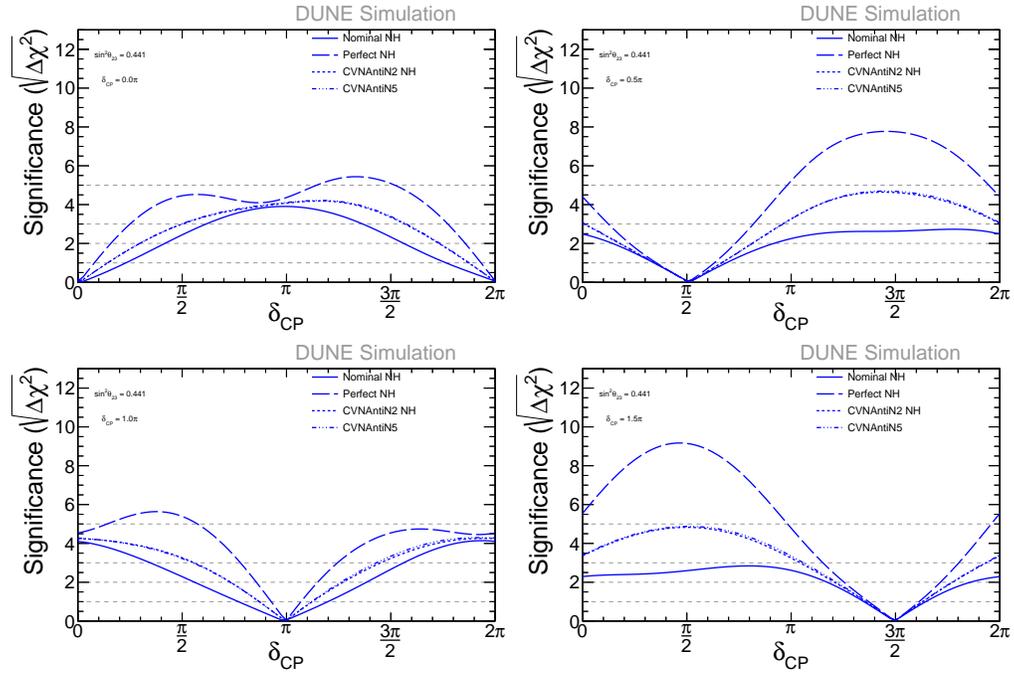


Figure D.10: Slices for true  $\delta_{CP} = 0$  (top left),  $\pi/2$  (top right),  $\pi$  (bottom left),  $3\pi/2$  (bottom right) with 7 years of exposure for DUNE in only the antineutrino beam.

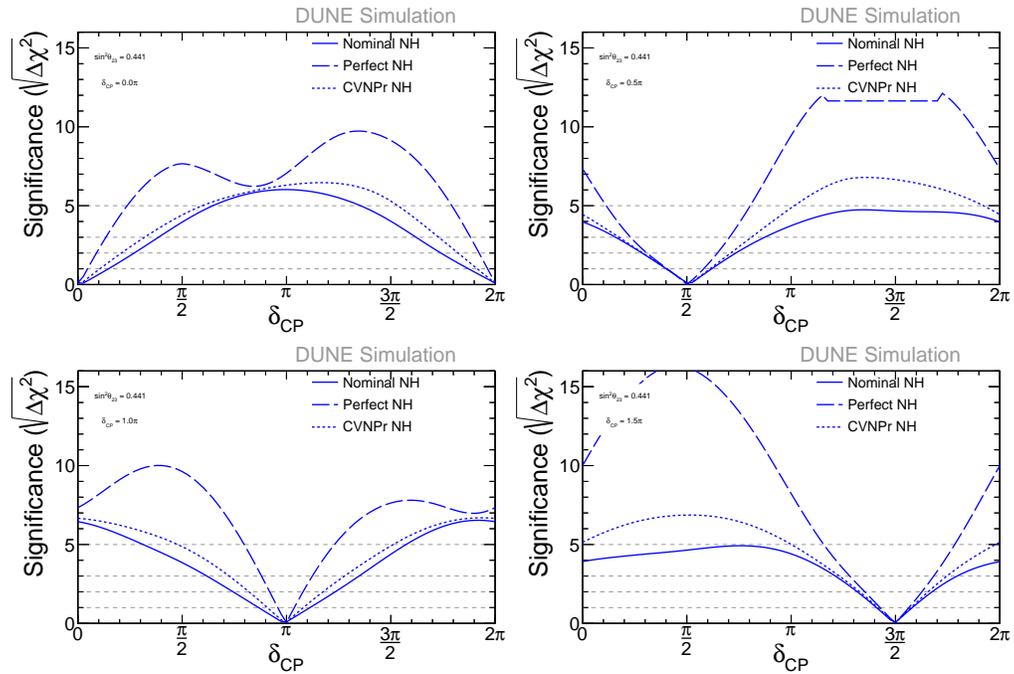


Figure D.11: Slices for true  $\delta_{CP} = 0$  (top left),  $\pi/2$  (top right),  $\pi$  (bottom left),  $3\pi/2$  (bottom right) with 15 years of exposure for DUNE (1104 kt\*MW\*yr) in only the antineutrino beam.

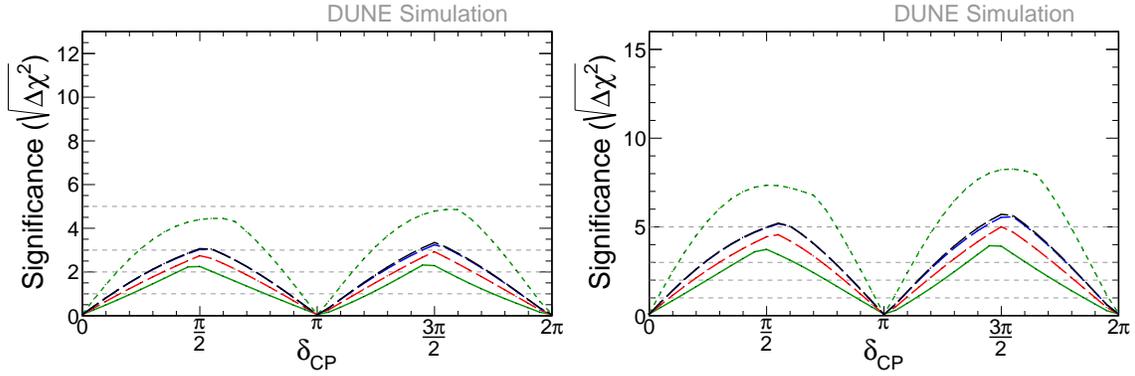


Figure D.12: CP conservation rejection curves for true  $\sin^2 \theta_{23} = 0.441$  at 7 years (left) and 15 years (right) exposure for DUNE, in only the antineutrino beam. The red dashed curve is for the CVN Proton score method, and the CVN antineutrino score cases are in black and blue dashes. All three methods provide improved significance. A perfect WS separation is shown by the green dashed line.

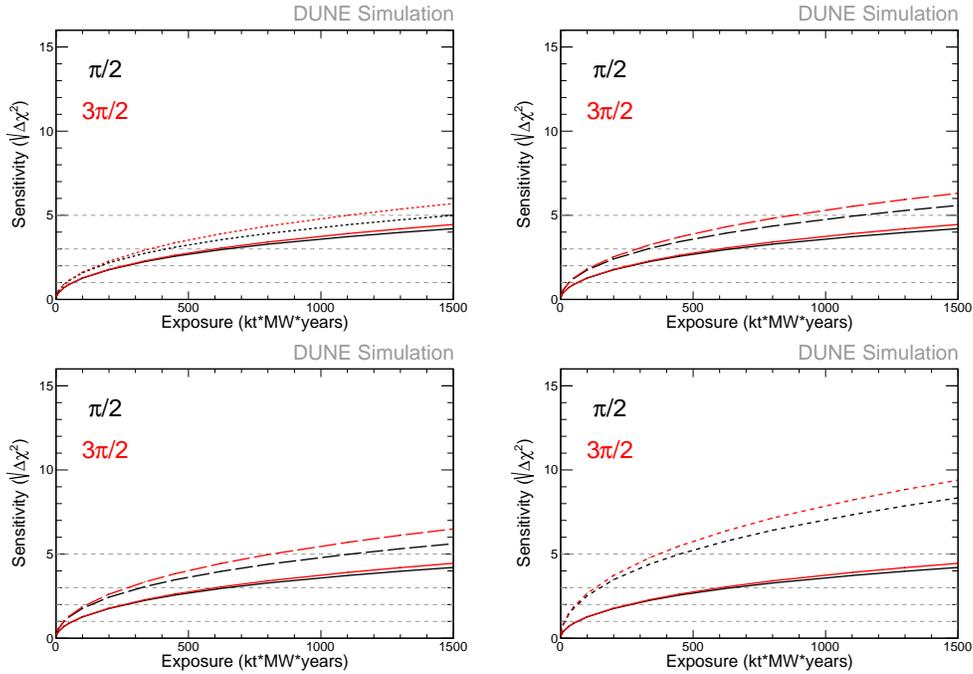


Figure D.13: CP conservation rejection at  $\delta_{CP} = \pi/2$  (black) and  $3\pi/2$  red vs. beam exposure for DUNE in the case of running antineutrino beam only. Solid lines are the procedures with no wrong-sign separation and the dashed lines in each figure correspond to the result with a wrong-sign separation technique: CVN Proton (top left), CVN antineutrino at 0.2 (top right), CVN antineutrino at 0.5 (bottom left), and perfect wrong-sign separation (bottom right). The conversion from  $\text{kt}\cdot\text{MW}\cdot\text{yr}$  to running years is given in the text.

## Bibliography

- [1] C. D. Ellis and W. A. Wooster. The average energy of disintegration of radium e. *Proceedings of the Royal Society A*, 117:109–123, 1927.
- [2] F. Reines and C. L. Cowan, Jun. The neutrino. *Nature*, 178(4531):446–449, 1956.
- [3] W. Pauli. Open letter. Tr. Kurt Riesselmann, MicroBooNE-doc-953: <http://microboone-docdb.fnal.gov/cgi-bin/RetrieveFile?docid=953;filename=pauli%20letter1930.pdf>, 1930.
- [4] M. Tanabashi et al. (Particle Data Group). *Phys. Rev. D*, 98:030001, 2018.
- [5] G. Danby et al. Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos. *Phys. Rev. Lett.*, 9:36–44, 1962.
- [6] K. Kodama et al. Observation of tau neutrino interactions. *Physics Letters B*, 504(3):218–224, 2001.
- [7] The ALEPH Collaboration, The DELPHI Collaboration, The L3 Collaboration, The OPAL Collaboration, The SLD Collaboration, The LEP Electroweak Working Group, The SLD Electroweak and Heavy Flavour Groups. Precision electroweak measurements on the z resonance. *Physics Reports*, 427(5):257–454, 2006.
- [8] A.B. McDonald. Evidence for neutrino oscillations I: Solar and reactor neutrinos. *Nuclear Physics A*, 751:53–66, 2005.
- [9] H. Nunokawa, S. Parke, J. Valle. CP violation and neutrino oscillations. *Progress in Particle and Nuclear Physics*, 60(2):338–402, 2008.
- [10] K. Abe et al. Measurement of muon antineutrino oscillations with an accelerator-produced off-axis beam. *Phys. Rev. Lett.*, 116:181801, 2016.

- [11] L. Wolfenstein. Neutrino oscillations in matter. *Phys. Rev. D*, 17:2369–2374, 1978.
- [12] S. P. Mikheyev and A. Y. Smirnov. Resonant amplification of  $\nu$  oscillations in matter and solar-neutrino spectroscopy. *Il Nuovo Cimento C*, 9(1):17–26, 1986.
- [13] M. Maltoni and A. Y. Smirnov. Solar neutrinos and neutrino physics. *The European Physical Journal A*, 52(4):87, 2016.
- [14] A. Himmel. First Oscillation Results with Neutrino and Antineutrino Beams in NOvA, June 2018. Joint Theoretical-Experimental Physics Seminar (Fermilab).
- [15] Y. Fukuda et al. Atmospheric  $\nu_\mu/\nu_e$  ratio in the multi-gev energy range. *Physics Letters B*, 335(2):237–245, 1994.
- [16] G. Feldman. History of long-baseline accelerator neutrino experiments. 2019. arxiv:1901.09431 [hep-ex].
- [17] T. J. Haines et al. Calculation of atmospheric neutrino-induced backgrounds in a nucleon-decay search. *Phys. Rev. Lett.*, 57:1986–1989, 1986.
- [18] R. M. Bionta et al. Contained neutrino interactions in an underground water detector. *Phys. Rev. D*, 38:768–775, 1988.
- [19] R. Becker-Szendy et al. Search for muon neutrino oscillations with the Irvine-Michigan-Brookhaven detector. *Phys. Rev. Lett.*, 69:1010–1013, 1992.
- [20] Q. R. Ahmad et al. Direct evidence for neutrino flavor transformation from neutral-current interactions in the Sudbury Neutrino Observatory. *Phys. Rev. Lett.*, 89:011301, 2002.
- [21] W.W.M. Allison et al. The atmospheric neutrino flavor ratio from a 3.9 fiducial kiloton-year exposure of soudan 2. *Physics Letters B*, 449(1):137–144, 1999.
- [22] M Ambrosio et al. Measurement of the atmospheric neutrino-induced upgoing muon flux using macro. *Physics Letters B*, 434(3):451–457, 1998.

- [23] P.F. de Salas et al. Status of neutrino oscillations 2018:  $3\sigma$  hint for normal mass ordering and improved CP sensitivity. *Physics Letters B*, 782:633–640, 2018.
- [24] I. Esteban et al. Global analysis of three-flavour neutrino oscillations: synergies and tensions in the determination of  $\theta_{23}$ ,  $\delta_{CP}$ , and the mass ordering. *Journal of High Energy Physics*, 2019(1):106, 2019.
- [25] (The DUNE Collaboration) R. Acciarri and others. LBNF and DUNE Conceptual Design Report, Volume 2: The physics program for DUNE at LBNF, 2016. arXiv:1512.06148 [ins-det].
- [26] D. Adey et al. Measurement of the electron antineutrino oscillation with 1958 days of operation at Daya Bay. *Phys. Rev. Lett.*, 121:241805, 2018.
- [27] Y. Fukuda et al. Evidence for oscillation of atmospheric neutrinos. *Phys. Rev. Lett.*, 81:1562–1567, 1998.
- [28] G. Barr et al. Uncertainties in atmospheric neutrino fluxes. *Phys. Rev. D*, 74:094009, 2006.
- [29] M. Messier. *Evidence for Neutrino Mass From Observations of Atmospheric Neutrinos with Super-Kamiokande*. PhD thesis, Boston University, 1999.
- [30] K. Abe et al. Atmospheric neutrino oscillation analysis with external constraints in Super-Kamiokande I-IV. *Phys. Rev. D*, 97:072001, 2018.
- [31] P. Adamson et al. The numi neutrino beam. *NIM A*, 806:279–306, 2016. arxiv:1507.06690.
- [32] NuMI/NOvA target. Webpage, November 2017. [https://targets.fnal.gov/NuMI\\_target.html](https://targets.fnal.gov/NuMI_target.html) (Accessed 8 July 2019).
- [33] J. Hylen. NUMI Status, October 2016. NOvA-doc-16362 (Internal).
- [34] J. Hylen. NUMI Status, September 2018. NOvA-doc-32745 (Internal).

- [35] P. Adamson et al. Search for sterile neutrinos in MINOS and MINOS+ using a two-detector fit. *Phys. Rev. Lett.*, 122:091803, 2019.
- [36] M. A. Acero et al. First measurement of neutrino oscillation parameters using neutrinos and antineutrinos by NOvA. 2019. arXiv:1906.04907 [hep-ex].
- [37] K. Abe et al. Search for CP violation in neutrino and antineutrino oscillations by the T2K experiment with  $2.2 \times 10^{21}$  protons on target. *Phys. Rev. Lett.*, 121:171802, 2018.
- [38] M. Wascko. T2K Status, Results, and Plans. Talk at XXVIII International Conference on Neutrino Physics and Astrophysics, 4-9 June 2018, Heidelberg, Germany, DOI: 10.5281/zenodo.1286752, URL: <https://doi.org/10.5281/zenodo.1286752>.
- [39] K. Abe et al. Hyper-Kamiokande Design Report. 2018. arXiv:1805.04163 [ins-det].
- [40] S. Mufson et al. Liquid scintillator production for the NOvA experiment. *NIM A*, 799:1–9, 2015.
- [41] A. Himmel. New Oscillation Measurements from NOvA, January 2018. CERN Particle Physics Seminar.
- [42] C. Backhouse et al. The attenuation and threshold calibration of the NOvA detectors. NOvA-doc-13579 (Internal), 2017.
- [43] T. Alion. Third analysis calorimetric energy scale in the NOvA detectors. NOvA-doc-23372 (Internal), 2017.
- [44] L. Aliaga. NuMI beam prediction for the NOvA 2017 analyses. NOvA-doc-23441 (Internal), 2017.
- [45] L. Aliaga. *Neutrino Flux Prediction for the NuMI Beamline*. PhD thesis, The College of William and Mary, 2016. FERMILAB-THESIS-2016-03.
- [46] L. Aliaga. 2017-2018 beam plots, May 2018. NOvA-doc-20843 (Internal).

- [47] R. Gran. Model uncertainties for Valencia RPA effect for MINERvA, 2017. FERMILAB-FN-1030-ND, arXiv:1705.02932 [hep-ex].
- [48] T. Katori. Meson exchange current (MEC) models in neutrino interaction generators. *AIP Conference Proceedings*, 1663(1):030001, 2015. proceedings of the 8th International Workshop on Neutrino-Nucleus Interactions in the Few-GeV Region (NuInt12). arXiv:1304.6014 [nucl-th].
- [49] A. Aurisano et al. A convolutional neural network neutrino event classifier. *JINST*, 11(09):P09001, 2016.
- [50] F. Psihas. *Measurement of Long Baseline Neutrino Oscillations and Improvements from Deep Learning*. PhD thesis, Indiana University, 2018. FERMILAB-THESIS-2018-07.
- [51] D. Doyle. nue 2018 ND data MC decomp, June 2018. NOvA-doc-29633 (Internal).
- [52] C. Backhouse and A. Himmel for the NOvA  $\nu_e$  group. Overview of the 2018 NOvA  $\nu_e + \bar{\nu}_e$  appearance analysis. NOvA-doc-26699 (Internal), 2018.
- [53] E. Catano-Mur. Blessing package - Ana2019 - nue selection 2d plots, May 2019. NOvA-doc-37644 (Internal).
- [54] F. Psihas et al. Context-enriched identification of particles with a convolutional network for neutrino events, 2019. arXiv:1906.00713 [ins-det].
- [55] S. Yu and Z. Djurcic. Energy reconstruction for 2018  $\nu_e$  appearance analysis. NOvA-doc-26696 (Internal), 2018.
- [56] E. Catano-Mur. *Constraints on neutrino oscillation parameters with the NOvA experiment*. PhD thesis, Iowa State University, 2018. FERMILAB-THESIS-2018-29.
- [57] D. Pershey. *A Measurement of  $\nu_e$  Appearance and  $\nu_\mu$  Disappearance Neutrino Oscillations with the NOvA Experiment*. PhD thesis, California Institute of Technology, 2018. FERMILAB-THESIS-2018-17.

- [58] S. Yu. Blessing package: BEN decomp Ana2018, April 2019. Original version 2018, NOvA-doc-29618 (Internal).
- [59] T. Nosek. Blessing package: nue MichelDecomp 2018, May 2018. NOvA-doc-29702 (Internal).
- [60] S. Yu. Blessing package: Extrapolation cartoon for Ana2018, June 2018. NOvA-doc-29612 (Internal).
- [61] M. Groh and N. Nayak. Ana2019 nue data spectra, June 2019. NOvA-doc-37584 (Internal).
- [62] D. Doyle. Accelerating Feldman-Cousins for NOvA using NERSC Supercomputers, June 2018. New Perspectives 2018 (Fermilab).
- [63] D. Pershey and E. Catano-Mur. Technical note for fitting and Feldman-Cousins procedure for the 2017 analysis. NOvA-doc-23126 (Internal), 2017.
- [64] G. Feldman and R. Cousins. Unified approach to the classical statistical analysis of small signals. *Phys. Rev. D*, 57:3873–3889, 1998.
- [65] N. Nayak. Far detector exposure accounting for the 2017 analysis. NOvA-doc-23147 (Internal), 2017.
- [66] N. Agofonova and others. Final results of the OPERA experiment on  $\nu_\tau$  appearance in the CNGS beam. *Phys. Rev. Lett.*, 120:211801, 2018. arXiv:1804.04912 [hep-ex].
- [67] M. A. Acero et al. Measurement of neutrino-induced neutral-current coherent  $\pi^0$  production in the NOvA near detector. 2019. arXiv:1902.00558 [hep-ex].
- [68] A. Mislivec. Private communication: April 23, 2018.
- [69] J. Wolcott. Neutrons – plots for blessing, June 2018. NOvA-doc-29973 (Internal).
- [70] G. Davies et al. CVN 2018. NOvA-doc-27467 (Internal), 2018.

- [71] B. Howard, N. Nayak, E. Niner. Wronsign identification techniques for  $\nu_e$  appearance analysis. NOvA-doc-27248 (Internal), 2018.
- [72] J. Urheim. Covariance matrix for system of linear equations. NOvA-doc-38505 (Internal), 2018.
- [73] D. Pershey. 2-prong  $\nu_\mu$ CC events, December 2017. NOvA doc-25584 (Internal).
- [74] M. Strait. Using neutron captures to measure  $\nu_\mu$  contamination in the RHC beam. NOvA-doc-22955 (Internal), 2019.
- [75] M. Groh. Ana2019 nue FHC event displays, May 2019. NOvA-doc-37593 (Internal).
- [76] M. Groh. Ana2019 nue RHC event displays, May 2019. NOvA-doc-37596 (Internal).
- [77] A. Booth. Ana2019 numu spectra blessing package, May 2019. NOvA-doc-37806 (Internal).
- [78] N. Nayak. Blessing package : Ana2019 RHC bi-event and mono-event plots, June 2019. NOvA-doc-37500 (Internal).
- [79] S. Calvez and L. Kolupaeva. Blessing package : Joint fit Ana2019 results, June 2019. NOvA-doc-37794 (Internal).
- [80] E. Catano-Mur. Blessing package - Ana2019 - projected significances “reach” - joint nue + numu, June 2019. NOvA-doc-37791 (Internal).
- [81] S. Yu. Blessed package for ana2018: Future sensitivity 36+36, September 2018. NOvA-doc-29958 (Internal).
- [82] A. Bernstein et al. Report on the depth requirements for a massive detector at Homestake. arXiv:0907.4183 [hep-ex], 2009.
- [83] (The DUNE Collaboration) B. Abi and others. The DUNE Far Detector Interim Design Report, Volume 2: Single-phase module, 2018. arXiv:1807.10327 [ins-det].

- [84] (The DUNE Collaboration) B. Abi and others. The DUNE Far Detector Interim Design Report, Volume 3: Dual-phase module, 2018. arXiv:1807.10340 [ins-det].
- [85] E. Worcester. DUNE: Status and Science. Talk at XXVIII International Conference on Neutrino Physics and Astrophysics, 4-9 June 2018, Heidelberg, Germany, DOI: 10.5281/zenodo.1286763, URL: <https://doi.org/10.5281/zenodo.1286763>.
- [86] (The DUNE Collaboration) B. Abi and others. The DUNE Far Detector Interim Design Report, Volume 1: Physics, technology & strategies, 2018. arXiv:1807.10334 [ins-det].
- [87] B. Jones. Introduction to scintillation light in liquid argon, July 2014. LArTPC14: <https://indico.fnal.gov/event/8381/session/4/contribution/12/material/slides/1.pdf>.
- [88] R. S. Mulliken. Potential curves of diatomic rare-gas molecules and their ions, with particular reference to Xe<sub>2</sub>. *Journal of Chemical Physics*, 52(10):5170–5180, 1970.
- [89] T. Doke et al. Absolute scintillation yields in liquid argon and xenon for various particles. *Japanese Journal of Applied Physics*, 41(Part 1, No. 3A):1538–1545, 2002.
- [90] T. Heindl et al. The scintillation of liquid argon. *EPL*, 91(6):62002, 2010.
- [91] S. Kubota, M. Hishida, J. Raun. Evidence for a triplet state of the self-trapped exciton states in liquid argon, krypton, and xenon. *Journal of Physics C: Solid State Physics*, 11(12):2645–2651, 1978.
- [92] D. Whittington, S. Mufson, B. Howard. Scintillation light from cosmic-ray muons in liquid argon. *JINST*, 11(05):P05016, 2016.
- [93] R. Acciarri et al. Effect of nitrogen contamination in liquid argon. *JINST*, 5(06):P06003, 2010.
- [94] R. Acciarri et al. Oxygen contamination in liquid argon: combined effects on ionization electron charge and scintillation light. *JINST*, 5(05):P05003, 2010.

- [95] A. Hitachi et al. Effect of ionization density on the time dependence of luminescence from liquid argon and xenon. *Phys. Rev. B*, 27:5279–5285, 1983.
- [96] P. Peiffer et al. Pulse shape analysis of scintillation signals from pure and xenon-doped liquid argon for radioactive background identification. *JINST*, 3(08):P08007, 2008.
- [97] P. Cennini et al. Detection of scintillation light in coincidence with ionizing tracks in a liquid argon time projection chamber. *NIM A*, 432(2):240–248, 1999.
- [98] T. Doke et al. Let dependence of scintillation yields in liquid argon. *NIM A*, 269(1):291–296, 1988.
- [99] T. Doke, K. Masuda, E. Shibamura. Estimation of absolute photon yields in liquid argon and xenon for relativistic (1 MeV) electrons. *NIM A*, 291(3):617–620, 1990.
- [100] A. Himmel. Photon Detectors Simulation Status for SN/LowE/Radiopurity, May 2016. DUNE collaboration meeting (Internal):  
<https://indico.fnal.gov/event/10612/session/19/contribution/80/material/slides/0.pdf>.
- [101] K. Mavrokoridis et al. Argon purification studies and a novel liquid argon recirculation system. *JINST*, 6(08):P08003, 2011.
- [102] B. Howard et al. A novel use of light guides and wavelength shifting plates for the detection of scintillation photons in large liquid argon detectors. *NIM A*, 907:9–21, 2018.
- [103] V. Gehman et al. Fluorescence efficiency and visible re-emission spectrum of tetraphenyl butadiene films at extreme ultraviolet wavelengths. *NIM A*, 654(1):116–121, 2011.
- [104] SensL. An introduction to the silicon photomultiplier. Technical Note. Avail. [sensl.com](http://sensl.com), 2011.
- [105] (The DUNE Collaboration) B. Abi and others. The single-phase ProtoDUNE Technical Design Report, 2017. arXiv:1706.07081 [ins-det].

- [106] J. T. Anderson et al. SiPM Signal Processor User Manual, July 2016. DUNE-doc-1571 (Internal).
- [107] G. Drake. Photon Detector Electronics, May 2015. LBNE-doc-10923 (Internal).
- [108] K. Spurgeon. Photon Detection System Simulations Utilizing Geant4 and GATE, July 2018. DUNE Photon Detector Simulation and Physics working group meeting: <https://indico.fnal.gov/event/17750/contribution/1/material/slides/0.pdf>.
- [109] K. Spurgeon and D. Whittington. Private communication: August 8, 2018.
- [110] S. Coutu et al. Searching for TeV cosmic electrons with the CREST experiment. *Nuclear Physics B - Proceedings Supplements*, 215(1):250–254, 2011. Proceedings of the 12th Topical Seminar on Innovative Particle and Radiation Detectors (IPRD10).
- [111] M. Antonello et al. Analysis of liquid argon scintillation light signals with the ICARUS T600 detector, 2006. Technical Report ICARUS-TM/06-03.
- [112] A. Neumeier et al. Attenuation measurements of vacuum ultraviolet light in liquid argon revisited. *NIM A*, 800:70–81, 2015.
- [113] N. Ishida et al. Attenuation length measurements of scintillation light in liquid rare gases and their mixtures using an improved reflection suppresser. *NIM A*, 384(2):380–386, 1997.
- [114] A. Neumeier et al. Attenuation of vacuum ultraviolet light in liquid argon. *The European Physical Journal C*, 72(10):2190, 2012.
- [115] L. Gallego et al. Modeling crosstalk in silicon photomultipliers. *JINST*, 8(05):P05010, 2013.
- [116] E. Catano-Mur and N. Nayak. 2018 bi-event plots – blessing package, May 2018. NOvA-doc-29663 (Internal).

- [117] C. Patrignani et al. (Particle Data Group). *Chin. Phys. C*, 40:100001, 2016 and 2017 update.
- [118] E. Niner. Wrong-Sign sensitivity with CVN prongs, January 2018. NOvA doc-25899 (Internal).
- [119] B. Howard, N. Nayak, and E. Niner. Sensitivity studies using wrong-sign ID methods, February 2018. NOvA doc-26364 (Internal).
- [120] E. Catano-Mur. Joint oscillation analysis with CAFAna (A bird's-eye view), June 2019. NOvA doc-37926 (Internal).
- [121] L. Kolupaeva. NOvA joint fit tutorial, June 2019. NOvA doc-37929 (Internal).
- [122] S. Monsalve and L. Whitehead. DUNE CVN Update, September 2018. DUNE collaboration meeting (Internal):  
<https://indico.fnal.gov/event/16526/session/30/contribution/94/material/slides/0.pdf>.
- [123] M. A. Acero et al. New constraints on oscillation parameters from  $\nu_e$  appearance and  $\nu_\mu$  disappearance in the NOvA experiment. *Phys. Rev. D*, 98:032012, 2018.
- [124] L. Vinton. Calorimetric energy scale calibration of the NOvA detectors. NOvA-doc-13579 (Internal), 2015.
- [125] T. Alion. Third analysis calorimetric energy scale in the NOvA detectors. NOvA-doc-23372 (Internal), 2017.
- [126] K. Warburton. Private communication: April, 2018.
- [127] T. Blackburn and L. Vinton. Drift Calibration, November 2015. NOvA doc-14320 (Internal).
- [128] L. W. Koerner and R. J. Nichol. Calibration Plenary, May 2018. NOvA doc-29537 (Internal).

- [129] L. Mualem. Private communication: April, 2018.
- [130] L. Mualem. Private communication: November, 2018.
- [131] M. Strait. Fiber Aging WBS 2.2, 2.3, June 2007. NOvA doc-2014 (Internal).
- [132] M. Strait, D. Cronin-Hennessy, K. Arms. Interaction rate of polystyrene fiber optics with pseudocumene-based liquid scintillator. *NIM A*, 615(1):33–36, 2010.
- [133] J. Blair. Center Brightness, December 2018. NOvA doc-34702 (Internal).
- [134] J. Blair. Fiber brightness technote. NOvA-doc-35674 (Internal), 2019.
- [135] R.J. Tesarek. Filling the NOvA Far Detector with liquid scintillator. NOvA-doc-13705 (Internal), 2015.
- [136] J. Benziger et al. A scintillator purification system for the borexino solar neutrino detector. *NIM A*, 587(2):277–291, 2008.
- [137] B. Behera, J. Paley, A. Giri. Drift Calibration, February 2017. NOvA doc-16984 (Internal).
- [138] LBNE Collaboration. The Long-Baseline Neutrino Experiment: Exploring Fundamental Symmetries of the Universe, 2014. arXiv:1307.7335v3.
- [139] R. Acciarri et al. Summary of the second workshop on liquid argon time projection chamber research and development in the United States. *JINST*, 10(07):T07006, 2015.
- [140] D. Whittington. Design testing for a large-area photon detection system, July 2014. LArTPC14: <https://indico.fnal.gov/getFile.py/access?contribId=17&sessionId=4&resId=0&materialId=slides&confId=8381>.
- [141] D. Whittington. Recent Results from Indiana University, July 2014. LBNE-doc-9466 (Internal).

- [142] R. Wasserman. LBNE Photon Detection System Design and Prototypes, July 2014.  
LArTPC14: <https://indico.fnal.gov/getFile.py/access?contribId=16&sessionId=4&resId=0&materialId=slides&confId=8381>.
- [143] J.V. Bennett et al. Precision timing measurement of phototube pulses using a flash analog-to-digital converter. *NIM A*, 622(1):225–230, 2010.
- [144] C. Macias and B. Howard. Passive Ganging Studies, July 2016. DUNE-doc-1464 (Internal).
- [145] B. Howard and C. Macias. SiPM Passive and Active Ganging with Cable Variation, May 2016. DUNE Photon Detection Workshop at CSU:  
<https://indico.fnal.gov/event/12089/session/4/contribution/13/material/slides/0.pdf>.

## CURRICULUM VITAE

### Education

---

PhD, Physics — 2019

Indiana University, Bloomington

*Thesis Title:* Toward a precision era of neutrino oscillation physics: liquid argon scintillation detector development for DUNE and neutrino oscillation studies with NOvA

*Advisor:* Dr. Jon Urheim

MS, Physics — 2015

Indiana University, Bloomington

BS with Honors, Physics and Astronomy — 2014

University of Pittsburgh

*Summa cum laude*

*Minor in mathematics*

### Research Activities

---

Collaborator on NuMI Off-axis  $\nu_e$  Appearance (NOvA) experiment

Work on electron neutrino appearance analysis, developing understanding of effects and systematics for the first antineutrino beam analysis and continuing analysis, mostly in efforts relating to the understanding of the wrong-sign component of the sample, but also as relates to other effects and systematics, including interaction modes and neutrons. Worked on developing calibration efforts assessing detector response. In conducting this research, also interfaced with grid computing and file management.

*Research Direction:* Dr. Jon Urheim

Collaborator on Deep Underground Neutrino Experiment (DUNE), formerly LBNE

Work on photon detection system for scintillation light in liquid argon (LAr) time-projection chambers. Some of this work has involved studies of individual components, most notably hands-on work in studies using silicon photomultipliers (SiPMs) and a waveform digitizer. Further work has involved an integrated prototype test in a liquid argon test stand at Fermilab and work on simulations. By comparing a simulation combining component studies and vendor information with data from an integrated prototype module, information on the scintillation

yield of LAr due to minimum ionizing cosmic-rays could be determined. I have also begun looking into dE/dx studies for proton decay in DUNE FD, and I have taken protoDUNE shifts.

*Research Direction:* Dr. Stuart Mufson and Dr. Jon Urheim

---

## Honors

---

2019 – Outstanding Graduate Student in Research Award

IU Physics Dept.

---

## Summer schools, symposia

---

March 2018 – Excellence in Detector and Instrumentation Technologies (EDIT)

Fermilab

---

## Publications (*papers listed first*)

---

M. A. Acero et al, *First measurement of neutrino oscillation parameters using neutrinos and antineutrinos by NOvA*. arXiv:1906.04907 [hep-ex]

B. Howard, S. Mufson, D. Whittington, et al, *A Novel Use of Light Guides and Wavelength Shifting Plates for the Detection of Scintillation Photons in Large Liquid Argon Detectors*. NIM A (2018), vol 907. doi:10.1016/j.nima.2018.06.050, arXiv: 1710.11233

*Part of special issue: Advances in Instrumentation and Experimental Methods (Special Issue in Honour of Kai Siegbahn)*

M. A. Acero et al, *New constraints on oscillation parameters from  $\nu_{\mu e}$  appearance and  $\nu_{\mu \mu}$  disappearance in the NOvA experiment*. Phys.Rev.D (2018), **98**, 032012. arXiv: 1806.00096

D. Whittington, S. Mufson, and B. Howard. *Scintillation light from cosmic-ray muons in liquid argon*. JINST **11** P05016

B. Howard, on behalf of the DUNE Collaboration. *Liquid argon scintillation detection utilizing wavelength-shifting plates and light guides*. JINST **13** C02006 (LIDINE2017)

B. Howard, on behalf of the DUNE Collaboration. *Developing Detectors for Scintillation Light in Liquid Argon for DUNE*. PoS (ICHEP2016) 1058

---

## Invited talks

---

October 2016 – CPAD Instrumentation Frontier Meeting 2016

“Plate and Light-Guide Based Photon Detection System Design for DUNE”

June 2016 – New Perspectives 2016

“DUNE in 10 Minutes”

### **Other Talks and Posters**

---

June 2018 – New Perspectives 2018 (talk)

“New electron (anti-)neutrino appearance analysis from NOvA”

*URA Travel Funding – \$400 – for New Perspectives and User Meeting attendance*

September 2017 – Light Detection in Noble Elements (LIDINE) 2017 (talk)

“Liquid argon scintillation detection utilizing wavelength-shifting plates and light guides”

August 2016 – 38<sup>th</sup> International Conference on High Energy Physics (ICHEP) (poster)

“Developing Detectors for Scintillation Light in Liquid Argon for DUNE”

June 2016 – Fermilab Users Meeting (poster)

“Developing a Photon Detection System for DUNE”

January 2016 – Neutrino Detector R&D Facilities Workshop (Fermilab) (talk)

“TallBo User Experience”

April 2015 – APS April Meeting (talk)

“Scintillation photon detection in liquid argon at the Long-Baseline Neutrino Facility”

### **Selected teaching, organizational, and outreach-related activities**

---

June 2018 - June 2019

President, Young NOvA organization

Plan and lead meetings related to Young NOvA business

Plan and led day-long tutorial session (June 2018)

Work on NOvA internal web site upgrades: e.g. writing page on use of Slack on NOvA

Spring 2016 - Spring 2017

Co-coordinator of seminars for Bridge Program

Plan student-organized seminars aimed at students in Indiana University APS Bridge Program, typically with additional broader appeal to advanced undergrad students and new grad students.

Fall 2014

Associate Instructor, Indiana University

Instructor for P201 (algebra-based physics I) lab sections

Maintained office hours in department’s Physics Forum

## **Professional and Honor Society Memberships**

---

American Physical Society (APS)

Sigma Pi Sigma (Honor Society)

Phi Beta Kappa (Honor Society)

## **Selected Relevant Skills**

---

Experience with C++ and the analysis software ROOT

Some experience with Python, Mathematica

Experience with Linux systems and shell-scripting

Trained to use machining tools – IU student machine shop