

**Characterizing neutrino and neutron fluxes
from the Spallation Neutron Source for the
COHERENT experiment**

by

Rebecca L. Rapp

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
at
Carnegie Mellon University
Department of Physics
Pittsburgh, Pennsylvania

Advised by Professor Diana Parno

August 16, 2022

Abstract

Coherent elastic neutrino-nucleus scattering (CEvNS), a process by which a neutrino scatters from the whole nucleus, has been observed by the COHERENT collaboration in multiple detectors using pulsed neutrinos produced by the Spallation Neutron Source (SNS) at Oak Ridge National Laboratory. This work contributes to the improved understanding of neutron and neutrino fluxes, which all COHERENT subsystems use to interpret their observations, as COHERENT transitions from discovery to precision measurements of CEvNS.

In the first half of this work, we walk through the prediction of the neutrino flux and spectra at the SNS using Geant4 simulation, and establish that there is an associated 10% uncertainty on these calculations due to a lack of pion-production data for SNS operating conditions. This large systematic is now the dominant uncertainty limiting COHERENT's precision physics goals. To contribute to the reduction of this uncertainty, we also perform design studies for a new deployed D₂O detector to experimentally normalize our neutrino flux.

The only observable signature of CEvNS is a low-energy nuclear recoil, which means neutrons cause a similar signal to neutrinos in our sensitive detectors. In the second half of this work, we transition to the characterization of the beam-related neutron flux from the SNS. We monitor the neutron flux near COHERENT detectors using a dedicated subsystem: the Multiplicity and Recoil Spectrometer (MARS). In this work, we simulate this mobile, Gd-doped plastic scintillator and begin to develop a neutron response matrix that can be used to unfold the incident neutron spectrum. We also perform a rate analysis to study the incoming flux at multiple locations near our CEvNS detectors.

Acknowledgements

Science is not done in a bubble, and this work is no exception. In these pages which mark the completion of my most difficult task to-date, I want to express my gratitude to the village behind the curtain, each of whom helped me navigate my transition to scientific adulthood.

Let me focus first on the personal development I brought into my study of physics. My grandparents — Tom (Sr.) and Helen Rapp and the late Jack and Eileen O’Toole — built a rock-solid foundation for our entire family, on which each success could blossom and be celebrated, and any setback was met with a safety net to regroup and try again. My parents — Tom (Jr.) and Amy Jo Rapp — nourished my work ethic and love of knowledge from this foundation, and provided every imaginable resource so that I could pursue anything of interest¹ as long as I made every effort to be present in my pursuits and *learn*. My siblings — Jessica, Sara, and Thomas — were privy to my phase of stampeding through obstacles like a bull in a china shop, especially when *somebody* decided to move “surprising” new obstacles into a previously clear path; growing up figuring out how to balance our collective needs and interests shaped who I’ve become. My extended family — the O’Toole (and Buza, Gladstone, Bernero, etc.) and Rapp (and Leroux, Ianno, Supernovich, etc.) clans alike — are an eclectic mix (in the best way) who have always lived by example to demonstrate what it means to be a member of a community; they offer welcome perspectives, new opportunities, and endless reserves of support between exhausting adventures.

I had the pleasure of being trained to think from a young age by some truly incredible people. My math and science teachers — especially Janet Dykstra, Christy Galilei, Eric Weimerskirch, and Tim Wesolowski — laid the foundation for my day-to-day analysis toolkit and cultivated my sense of wonder at the world around me. My English teachers — especially Tara Wemyss and Jennifer Spiegel — counseled quick (but effective) writing and trained me both to interpret works through multiple different lenses and to look for the lenses that others apply in their own analyses. My marching band directors — namely

¹Well — except for the one time 8-year-old me was convinced I could be a lawyer. Good call, Mom.

Dave Cornelius, Justin Turpin, and Jon Stolarz — ingrained the immense value of individual effort and accuracy within a more global picture. I also formed lasting friendships with my peers — most especially Robbie Johnson and Alex Poremba — which taught me how to learn in a group and find fun even when long days of hard work, whether assigned or chosen, consume any free time. These skills have served me exceptionally well in my undergraduate and graduate education, and will continue to benefit my experiences as both a researcher and educator.

The connections I formed with the wonderful community at Washington & Jefferson College put me on a path that fit my personal and professional goals. I found the very best friends I could ever hope to have — Aric Brown, Mike Chepanoske, Ashley Knoerdel, and Lottie Murray — who have each inspired shenanigans which became cherished memories, and who were equally motivated to spend time together while focused on our individual pursuits. My undergraduate advisor — Mike McCracken — opened the door into a field of people doing science that I had never dreamed possible and convinced me that I was capable of joining them². I was mentored by many phenomenal faculty and staff across the college, and my interactions with them ultimately sparked the desire to become a professor myself.

The rigorous, collaborative nature of CMU’s physics graduate program was immensely improved by my cohort — especially Nick Bower, Caitlin Carnahan, Sam Foley, Matt Ho, Emma Oxford, Alejandro Sanchez, Nishchay Suri, and Daniel Zile — with whom I share a friendship forged in the fires of some of our most trying times; I am lucky to have gone through the program alongside their brilliance³. The effort all our faculty and staff pour into improving the graduate student experience does not go unnoticed, and I am especially indebted to Steve Garoff, Manfred Paulini, Fred Gilman, and Markus Deserno for their guidance and support throughout the program. Amanda Bodnar is one of the true forces behind the department, and I am thankful for her mastery of all things organizational.

I joined the Parno group in my first semester of research and found an exceptional working relationship with the COHERENT collaboration. The students and postdocs in my working groups — especially Igor Bernardi, Conan Bock, Matthew Heath, Max Hughes, Karla Tellez-Giron-Flores, and Justin Raybern — make work enjoyable and exciting, and I am immensely grateful for their insights and expertise, as well as for our shared experiences in the lab and at collaboration meetings. The senior-member mentors in my working groups (who aren’t separately acknowledged later on) — Belkis Cabrera-Palmer,

²Fun fact: he was right. Don’t tell him I said that — he’s already won so many of our debates.

³I am also grateful for their forgiveness when I ignore the group chat and associated event invitations...

Yuri Efremenko, Jon Link, Jing Liu, Diane Markoff, and Kate Scholberg — have been a constant source of encouragement and always lead me to believe the work I contribute is important to our physics goals; I have definitely grown under their expert tutelage and consider it an enormous privilege to have my work scrutinized by their accomplished minds.

My fellow graduate students in the Parno group — Byron Daniel, Gen Li, Kyungmin Park, Ana Paula Vizcaya Hernandez, and Larisa Thorne — are exceptionally friendly and supportive humans boasting a shared fascination with neutrinos. I have learned so much from our discussions during weekly journal clubs, and I think our *skribbl.io* game-days might be one of the biggest highlights during the pandemic years. I have also gained a great deal from interacting with the undergraduate students rotating through our group — especially Alex Arteaga, Duy Hoang, M Mahoney, Aria Salyapongse, Cole Skuse, Samuel Vasquez, and Shuaixiang Zhang — and they certainly improved my work. I absorbed most of what I know about life as a physicist from my office buddies — Husni Almoubayyed, Zach Baldwin, Sayan Mandal, Will McGinley, and Naomi Jarvis — and our shared lunches, chats, and breaks kept me sane throughout the program.

My thesis committee — Curtis Meyer, Jason Newby, Riccardo Penco, and Brian Quinn — have made thoughtful inquiries throughout these long years to guide my focus and planning, immeasurably improving this work. I am especially appreciative of their efforts in helping me develop the confidence to speak with increasing levels of scientific authority.

Finally, and arguably most importantly, my advisor — Diana Parno — has been the most incredible role model in my transition to scientific adulthood. Her patience, support of non-research interests, and insistence on work-life balance have made me a better researcher. She believed in my abilities and foresaw the contributions I could make to our research before I even knew how to compile a C++ program, and her optimism and creativity in the face of the challenges we encountered have made the work both possible and enjoyable. I still have so much I want to learn from her, but I am grateful for the work we have done so far and look forward to our future collaborations.

Contents

1	Introduction	2
1.1	Historical Overview	3
1.2	Neutrinos and the Standard Model	8
2	Coherent Elastic Neutrino-Nucleus Scattering	13
2.1	Neutrino Scattering	14
2.2	Standard Model Cross Section	16
2.3	Physics Sensitivities	17
2.4	Experimental Requirements	22
3	COHERENT	25
3.1	The Spallation Neutron Source	25
3.2	Subsystem Overview	29
4	Simulating Particle Production at the Spallation Neutron Source	36
4.1	Introduction	36
4.2	Validation of Simulation Physics	37
4.3	Modeling the Spallation Neutron Source in Geant4	47
4.4	Neutrino Flux for the First Target Station	51
4.5	Neutrinos at the Second Target Station	58
4.6	Light Dark Matter Production at the SNS	59
4.7	Ongoing Efforts	61
5	Design Studies for a Flux Normalization Detector at the SNS	63
5.1	D ₂ O Detector Overview	63
5.2	Detector Geometries	69
5.3	Simulation Physics	75

5.4	Expected Response to Signal and Background	77
5.5	Fiducialization Analyses	81
5.6	Optical Measurements	86
5.7	Conclusion	95
6	Neutrons in Neutrino Alley	97
6.1	Neutrons as a CEvNS Background	97
6.2	Strategies to Measure BRN Backgrounds <i>In Situ</i>	100
7	Developing a Simulation of MARS to Unfold the BRN Spectrum	111
7.1	Geant4 Simulation of MARS	111
7.2	Characterizing the MARS Response	120
7.3	Efforts towards Unfolding Neutron Spectra	136
8	Monitoring BRN Rates with MARS	139
8.1	Data Quality	139
8.2	Analysis Strategy	141
8.3	Production Run Results	150
9	Conclusion	161
A	Daily Data Collected During MARS Production Runs	163
A.1	Production Run 3	164
A.2	Production Run 4	193
A.3	Production Run 5	196

List of Tables

3.1	Systematic uncertainty breakdown for the COH-CsI and COH-Ar-10 CEvNS measurements.	33
4.1	Geometric volume breakdown of π^+ production at the SNS.	49
4.2	Production process breakdown for ν created at the SNS.	52
4.3	Fit parameters for the dependence of the SNS neutrino luminosity on the incident proton energy illustrated in Figure 4.14.	55
5.1	Simulated signal and background rate estimates for the D ₂ O detector.	81
8.1	A comparison of the BRN flux throughout Neutrino Alley.	159
A.1	MARS Data: January 2019	164
A.2	MARS Data: February 2019	165
A.3	MARS Data: March 2019	166
A.4	MARS Data: April 2019	167
A.5	MARS Data: May 2019	168
A.6	MARS Data: June 2019	169
A.7	MARS Data: July 2019	170
A.8	MARS Data: August 2019	171
A.9	MARS Data: September 2019	172
A.10	MARS Data: October 2019	173
A.11	MARS Data: November 2019	174
A.12	MARS Data: December 2019	175
A.13	MARS Data: January 2020	176
A.14	MARS Data: February 2020	177
A.15	MARS Data: March 2020	178
A.16	MARS Data: April 2020	179
A.17	MARS Data: May 2020	180

A.18	MARS Data: June 2020	181
A.19	MARS Data: July 2020	182
A.20	MARS Data: August 2020	183
A.21	MARS Data: September 2020	184
A.22	MARS Data: October 2020	185
A.23	MARS Data: November 2020	186
A.24	MARS Data: December 2020	187
A.25	MARS Data: January 2021	188
A.26	MARS Data: February 2021	189
A.27	MARS Data: March 2021	190
A.28	MARS Data: April 2021	191
A.29	MARS Data: May 2021	192
A.30	MARS Data: July 2021	194
A.31	MARS Data: August 2021	195
A.32	MARS Data: September 2021	196
A.33	MARS Data: October 2021	197
A.34	MARS Data: November 2021	198
A.35	MARS Data: December 2021	199
A.36	MARS Data: January 2022	200
A.37	MARS Data: February 2022	201
A.38	MARS Data: March 2022	202
A.39	MARS Data: April 2022	203

List of Figures

1.1	A diagram of the Standard Model of Particle Physics.	3
1.2	Left: Observed spectrum of products from alpha decay, reproduced from [2]. Right: Observed spectrum of products from beta decay, reproduced from [3].	4
1.3	Feynman diagrams for β -decay (left), inverse beta decay (center), and neutral-current neutrino-electron scattering (right). These diagrams can be interpreted as time flowing from left to right.	12
2.1	Neutrino-electron scattering cross section on a broad energy scale.	14
2.2	Insights into the weak mixing angle and the effective neutrino charge radius from CEvNS measurements.	18
2.3	Constraints on non-standard interactions from CEvNS measurements.	20
2.4	A plot of the “neutrino fog” important for WIMP dark matter searches.	21
3.1	Predicted timing and energy distributions for SNS neutrinos.	27
3.2	Diagrams illustrating the location of COHERENT detector deployments.	29
3.3	Plots of the CEvNS cross section and expected nuclear recoil energy for COHERENT detectors.	30
3.4	Predicted neutron-dependence of the CEvNS cross section.	31
4.1	Validation of Geant4 physics models against the Norbury-Townsend parameterization of total π^+ production.	39
4.2	Double-differential comparisons of Geant4 physics models against HARP/HARP-CDP π^+ -production cross section data.	41
4.3	Comparisons of Geant4 physics models against HARP/HARP-CDP π^+ -production cross section data integrated over either production angle or momentum.	43

4.4	Comparisons of Geant4 simulation predictions for HARP and HARP-CDP angular regions.	44
4.5	Comparisons of Geant4 physics models against HARP/HARP-CDP π^+ -production cross section data integrated over both production angle and momentum.	45
4.6	Comparisons of double-differential cross sections of π^+ production to Abaev <i>et al.</i> measurements.	45
4.7	Comparisons of differential cross sections of π^+ production to Abaev <i>et al.</i> measurements.	46
4.8	ORNL technical drawings illustrating components included in our Geant4 model of the SNS neutrino production.	48
4.9	Illustrations of the Geant4 model of the SNS.	48
4.10	A comparison of the different proton beam window designs of the SNS.	50
4.11	Simulated timing and energy spectra for neutrinos produced at the SNS.	51
4.12	Distributions of neutrino energy broken down by creation process for each neutrino flavor produced at the SNS.	52
4.13	The SNS beam energy from January 2013 – August 2021.	53
4.14	Best fits for the proton energy dependence of the SNS neutrino flux.	54
4.15	Simulated production positions and angles of neutrinos produced at the SNS.	55
4.16	A comparison of the simulated neutrino production positions for varied incident energy and SNS configuration.	56
4.17	A plot illustrating the anisotropy of the neutrino production at the SNS.	58
4.18	Geant4 implementation of the Second Target Station target and moderators.	59
4.19	π^0 distributions before and after the PPU.	60
4.20	π^+ energy loss in the SNS target	61
5.1	Predicted cross sections and interaction rates for neutrino interactions in the D ₂ O detector.	66
5.2	Simulated distribution of photoelectrons observed from isotropic electron events in the D ₂ O detector.	68
5.3	Simulated energy reconstruction and resolution for the Tank geometry.	69
5.4	Geant4 screenshots of possible D ₂ O geometries.	70
5.5	Simulated signal comparison between the Box and Axle geometries.	73
5.6	Comparisons of the energy reconstruction and resolution between all possible geometries for the D ₂ O detector.	73

5.7	An engineering diagram of the Tank geometry.	74
5.8	The quantum efficiency of simulated PMTs for D ₂ O.	77
5.9	Simulated energy response from varied signal and background sources for the Tank and No-D ₂ O geometries.	78
5.10	An estimate of the expected signal and statistical precision for the Tank geometry.	79
5.11	Distribution of the simulated event positions meeting specified energy cuts in D ₂ O simulation.	82
5.12	A plot illustrating the impact of energy cuts on the fiducial volume in the D ₂ O simulation.	84
5.13	Quantifying the fiducial volume of the Tank and No-D ₂ O geometries in the radial and vertical dimensions.	86
5.14	A comparison of the simulated response of the D ₂ O detector using different values for the total reflectivity.	87
5.15	Visualization of the apparatus for D ₂ O optical reflectivity measurements.	88
5.16	A diagram of the UNIFIED model in Geant4.	91
5.17	An illustration of the difference between specular lobe and specular spike reflection.	92
5.18	An example of the spectral intensity data for D ₂ O specularity studies.	93
5.19	A comparison of the specularity of different samples in the D ₂ O reflectivity studies.	94
5.20	A comparison of the light yield for 30 MeV electrons (left) and fractional energy resolution (right) uniformly distributed throughout the Tank geometry, implemented with varied reflectance types. The Diffuse, Spec-Spike, and SpecLobe labels indicate the relevant coefficient is 1 with others set to zero. The Combo label is a simulation in which $C_{ss} = C_{sl} = 0.33$ and $C_{dl} = 0.34$. Our baseline simulations implement the Diffuse curves.	95
6.1	Diagrams illustrating the location of Neutrino Alley relative to the SNS target and the concrete/gravel shielding.	99
6.2	Diagram illustrating the locations of BRN detectors in Neutrino Alley.	101
6.3	Visualization of the NSC and SciBath BRN detectors.	102
6.4	Plots from the NSC background campaign in the SNS target hall.	103
6.5	Visualization of the Timing Cart and the SNS beamline locations relative to Neutrino Alley.	105
6.6	A summary of the Timing Cart BRN flux measurements.	106

6.7	A diagram of the MARS detector and an example event.	107
6.8	A diagram of the MARS data acquisition.	109
6.9	The recorded waveforms for a single MARS event.	110
7.1	A screenshot of the Geant4 description of the MARS detector.	112
7.2	Comparisons of the simulated response of the MARS detector using different physics models.	113
7.3	An illustration of the event-processing in MARS.	115
7.4	A comparison of the quenching model parameterization we apply in simulation to data.	116
7.5	A comparison of the quenching model for different particles in the MARS simulation.	117
7.6	Measurements of the signal acceptance fraction at multiple locations in time for MARS.	120
7.7	Plots illustrating the discrepancy between the front and back light collection in MARS.	121
7.8	A visualization of the setup for our cobalt-60 calibrations of the MARS detector.	123
7.9	Smearing parameter fit results for an example position in the cobalt-60 grid scan.	125
7.10	A comparison of the measured smearing parameters in MARS over a few years.	126
7.11	A comparison of the extrapolated smearing parameter maps for the front and back cobalt-60 scans performed in May 2022.	128
7.12	Simulations of the MARS response to Michel electrons compared to data.	129
7.13	Visualization of the backing- α pixelation for the MARS DT calibration.	131
7.14	Comparisons between simulation and data for the capture time constant in MARS.	133
7.15	Comparisons between simulation and data for the spectral response of MARS to incident 14.1 MeV neutrons.	134
7.16	A comparison of the Michel electron feature in MARS data during different beam-operating conditions.	136
7.17	A plot of the BRN spectrum observed by MARS and the preliminary energy-dependent detection efficiency from simulation.	137
8.1	A plot of the daily trigger rate of the MARS detector for Run 3 (2019 - 2021).	141

8.2	An illustration of a single event-pair in the MARS analysis.	142
8.3	A comparison of the observed energy spectrum from event-pairs in MARS during a week of beam-on and beam-off data.	143
8.4	A comparison of example t and E_2 daily distributions for MARS Runs 3, 4, and 5.	144
8.5	Plots illustrating the identification of the Michel electron and muon-induced neutron event populations in MARS.	146
8.6	An example week of muon-induced neutron data illustrating the derivation of variable E_2 cuts in MARS.	147
8.7	A plot illustrating the correlation between observed excesses using constant E_2 cuts in the Michel electron and muon-induced neutron populations in MARS.	148
8.8	Plots showing the derivation of variable E_2 cuts for MARS Runs 3, 4, and 5.	150
8.9	Monitoring of the BRN rates observed by MARS throughout the Run 2. . .	152
8.10	Monitoring of the BRN rates observed by MARS throughout the Run 3. . .	153
8.11	Distributions of BRN arrival time t and spectrum E_1 observed during MARS Run 3.	154
8.12	Monitoring of the BRN rates observed by MARS throughout the Run 4. . .	155
8.13	Distributions of BRN arrival time t and spectrum E_1 observed during MARS Run 4.	156
8.14	Monitoring of the BRN rates observed by MARS throughout the Run 5. . .	157
8.15	Distributions of BRN arrival time t and spectrum E_1 observed during MARS Run 5.	158
8.16	A comparison of all BRN flux measurements throughout Neutrino Alley. .	160

Glossary

BRN Beam-Related Neutron – a neutron produced by the proton beam spill at the SNS which could be in Neutrino Alley in a time window coincident with the neutrino signal.

CC-D Charged current neutrino interaction with a deuteron: $\nu_e + d \rightarrow p + p + e^-$.

CC-O Charged current neutrino interaction with oxygen: $\nu_e + {}^{16}\text{O} \rightarrow e^- + {}^{16}\text{F}^*$.

CEvNS Coherent elastic neutrino-nucleus scattering; pronounced “sevens”. This acronym might be found elsewhere as CE ν NS or CENNS.

COHERENT A collaboration formed in 2013 to unambiguously observe CEvNS by deploying a suite of detectors at the SNS.

CRY A Monte-Carlo toolkit external to but compatible with Geant4 used to generate cosmic ray events for a user-defined date, geographic location, and overburden.

FIR Finite Impulse Response.

HOG Hot off gas – a pipe in Neutrino Alley which carries radioactive waste from the SNS target, which produces a steady-state 511-keV γ background for all COHERENT detectors.

IBD Inverse Beta Decay.

LS Liquid Scintillator.

MARS Multiplicity and Recoil Spectrometer.

Neutrino Alley A basement hallway in the SNS target building with 8 meters water equivalent of overburden. The physical location of COHERENT detectors.

NLL Negative Log Likelihood.

NSC Neutron Scatter Camera.

ORNL Oak Ridge National Laboratory.

PDF Probability distribution function.

PE Photoelectron.

PMT Photomultiplier tube.

POT Proton on target.

QE Quantum efficiency.

SNS Spallation Neutron Source.

Chapter 1

Introduction

The classification of observables is a powerful tool for scientists when describing new phenomena. The explicit name we apply to something holds immense value in communicating unique features, properties, or behaviors. For example, a mammal has key traits which are distinct from those of birds – in identifying a new creature, we would start by comparing observations with traits of known classes to determine the kinds of questions we should ask to understand the creature. In particle physics, we similarly group like-behaviors and characteristics together in our investigations in order to develop strategies to further our exploration into the behaviors and interactions of fundamental particles.

Yet, with the goal to rigidly classify all things, there will undoubtedly be outliers. Imagine a fish; it must have gills and have some method of controlled propulsion through water. In order to accurately describe a seahorse, however, one must account for male pregnancy – it is still a fish, but it’s an atypical fish whose classification must explain the additional observations. We thus introduce subdivisions to our established framework so that we may appropriately discuss these interesting creatures.

Such is the history of particle physics! As we make new observations, we add the appropriate labels to our dictionary of particles. Sometimes these even overhaul our entire classification scheme to account for a shift in perspective – consider the atom or the proton and what the term “fundamental particle” could have meant before their discovery – but our most recent overhaul occurred in the 1960s and 70s. We now have a classification of all known fundamental particles and of the interactions which govern their observed behavior, and we use this to predict and understand the known universe with significant accuracy. This is called the Standard Model of Particle Physics, and a diagram of the classification scheme physicists have developed is shown in Figure 1.1.

A lepton is a fundamental particle in the Standard Model with half-integer spin that

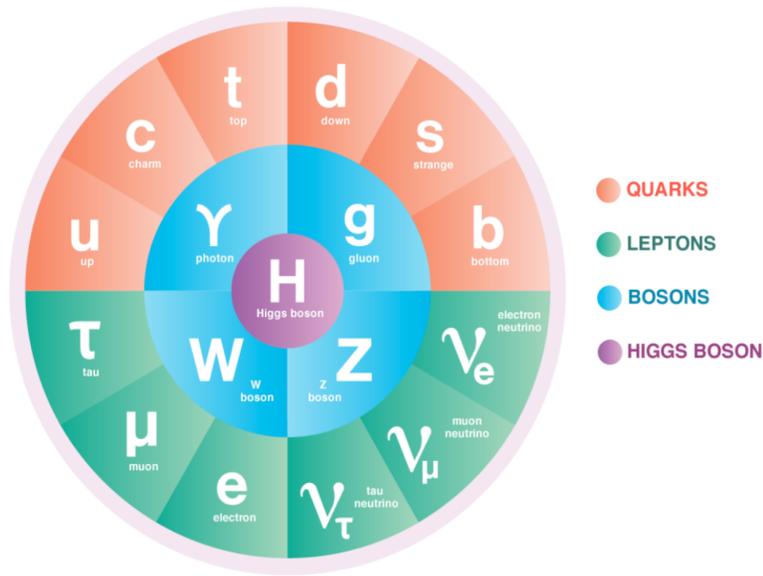


Figure 1.1: A diagram of the Standard Model of Particle Physics, as presented by Symmetry magazine and the Department of Energy’s Office of Science [1]

does not interact via the strong force — all the leptons have a green fill in Figure 1.1. The most well-known lepton is the electron — its fishy analogue would be a carp or trout. Some leptons, however, feature oddities that merit closer study. In this chapter, I discuss the historical observations, unique characteristics, and typical interactions of the neutrino — the particle physics analogue of the seahorse.

1.1 Historical Overview

Neutrino physics is closely tied to the weak force, which governs radioactive beta decay, and can also be associated with nuclear processes like fusion or fission given their role in producing vast quantities of neutrinos. The history of this subfield is very modern, by physics standards, having begun in the early 20th century. By this time, scientists had determined that nuclei existed and could undergo three distinct types of radioactive nuclear decay; these were classified using observed differences in the product of each decay as α (positively charged product), β (negatively charged product¹), and γ (charge-neutral product). Physicists developed the technology to resolve the energy spectrum for

¹Technically this is β^- decay. There is also β^+ decay, which is the process by which an up quark becomes a down quark. For the purposes of this text, $\beta = \beta^-$.

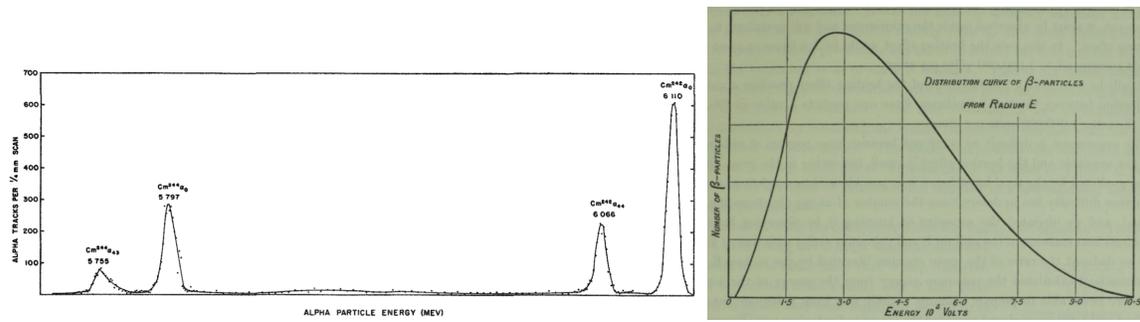


Figure 1.2: Left: Observed spectrum of products from alpha decay, reproduced from [2]. Right: Observed spectrum of products from beta decay, reproduced from [3].

each type of decay, and it was the spectrum of betas emitted in nuclear decay that inspired the idea that undetected neutrinos must exist.

1.1.1 Nuclear Beta Decay

Radioactive decay is both the underpinning of neutrino physics and the first investigation into the weak force. These studies were well-motivated in developing an understanding of the nuclear structure, as there was a measurable change to the nucleus when a decay took place. Three distinct types of radiation were quickly identified in a variety of “Radium” nuclides²:

- α : Positively charged, easily stopped particles. Ernest Rutherford established that an α particle is the nucleus of a helium atom (which we now know contains two protons and two neutrons). A sample spectrum of emitted alpha particles for the decays of ^{242}Cm and ^{244}Cm is shown in the left panel of Figure 1.2.
- β : Negatively charged, moderately penetrating particles. Henri Becquerel correctly suggested in 1900 that the emitted β particles are electrons, due to their identical charge-to-mass ratios. The spectrum of emitted beta particles from Radium E is shown in the right panel of Figure 1.2; this result confirmed Chadwick’s initial observation of a continuous spectrum [4].
- γ : Neutral, highly penetrating particles. Paul Villard correctly suspected in 1900 that γ -rays were electromagnetic radiation (photons) with shorter wavelengths than

²At this point in history, isotopes of different nuclei were known to exist, but several decades of study would pass before modern notation would be adopted. Radium E, for example, is now known to be ^{210}Bi .

X-rays. A sample observation of the γ -ray spectrum from a given isotope is not reproduced in this work, but we extensively use the well-known energies of γ -rays from nuclear decay to calibrate our MARS detector in Chapter 7.

The beta spectrum was observed to be continuous in 1914. This was ground-breaking and *confounding* because conservation of energy and momentum dictate that an emitted particle from a two-body decay process cannot have a continuous spectrum — it should have monoenergetic resonances, just as observed for α decays (see left panel of Figure 1.2).³ Imagine the first observation that male seahorses give birth — this is analogous to what happened in physics. Several scientists worked to re-observe the phenomenon and confirmed that the observation is valid, while many others worked to explain the observation within or outside of the established framework of physics. The initial observation of a continuous spectrum was later rigorously investigated and proven, with a confirmation measurement shown in Figure 1.2.

1.1.2 Proposition of a New Particle

The knowledge that beta decay did not produce a monoenergetic electron sent theorists into a flurry as they tried to develop a model which could predict the data. Some theories suggested that energy conservation must not be a universal law — that the foundational axiom respected by all physicists must certainly be incorrect. In 1930, Wolfgang Pauli resolved this new knowledge with the proposition of a new particle. In a letter to colleagues⁴, he established that by extending beta decay to have *three* products instead of only two, the continuous spectrum would be consistent with known physics. The nuclear decay, under the approximation that the nuclear recoil is negligible, would then emit two particles whose energies sum to the mass-difference of the nucleus before and after the decay, but the proposed new particle would be undetectable with current technology.

It was Enrico Fermi in 1934, however, who built the first theory of beta decay including Pauli’s proposed particle, which Fermi dubbed the “neutrino” [5]. He understood that the particle would have to be very light⁵ and have no electric charge (hence “neutrino”,

³Due to the importance the claim that a two-body decay should result in monoenergetic products will hold for COHERENT, a more rigorous and mathematical explanation is given as Equation 3.3.

⁴Fun fact: Pauli did not formally publish this idea until 1934, by which point Enrico Fermi had developed a (mostly) complete theory of β -decay based on Pauli’s proposal.

⁵Fermi’s theory also describes how precision measurements of the beta decay spectrum can help physicists determine the mass of the neutrino — this strategy is still in use in modern cutting-edge research to kinematically determine the neutrino mass (see Section 1.2.1).

meaning “little neutral one”) in order to have avoided detection, as well as having a spin of $1/2$ in order to conserve angular momentum. After two decades, Fermi had developed a theory which could predict the experimental observations.

1.1.3 First Detection

Soon after Fermi’s model was published, Hans Bethe and Rudolf Peierls posed that if nuclei can undergo β -decay to generate a neutrino, it must also be possible for the nucleus to annihilate a neutrino – this process would come to be known as inverse beta decay (IBD). They calculated a cross section⁶ of $\sigma < 10^{-44}$ cm², corresponding to a penetrating power of more than 10^{16} kilometers of “solid matter”, and concluded that “there is no practically possible way of observing the neutrino” [6]. In some sense, they were right – it would take more than two decades of technological advancement and planning to enable the detection of these “ghostly” particles.

During these decades, physicists would need to develop a detector to search for neutrinos and explore possible sources to identify possible deployment locations for the detector. Neutrinos were known to be born in β -decay processes, so the idea was to identify a large source of radioactive isotopes – like a nuclear reactor⁷! The IBD process ($\bar{\nu}_e + p \rightarrow n + e^+$) results in two products – a neutron and a positron; observing the coincidence⁸ of both particles would give a very clean signature that an IBD process has taken place.

In 1953, C.L. Cowan and F. Reines published a tentative observation of the neutrino [7] by deploying a detector to search for both the neutron and the positron near the Hanford nuclear reactor. Their detector used a Cd-doped scintillator⁹, which enabled a detection of IBD on hydrogen nuclei by searching for γ s produced when the positron annihilated with electrons in the scintillator, and a coincident observation of γ s following a neutron-

⁶In particle physics, the cross section of a process is an estimate of the effective area presented to the beam for causing the process of interest. We’ll discuss this further in Chapter 2.

⁷Note the importance of this timing in history – the Manhattan project (1942–1946) occurred during the development of neutrino detection efforts (proposed 1934, observed 1953), with Frederick Reines as a participating member of the project’s Theoretical Division. As a result, using a nuclear reactor as the neutrino source was actually a secondary plan. Initially, physicists at Los Alamos National Laboratory proposed to use an atomic explosion as a source for their nearby neutrino detector.

⁸In particle physics, coincidence means that two detector events occurring at a predictable time, position, direction, etc. are likely correlated and together produce a more identifiable signature of a process than either event alone. For the IBD observation, the time-coincidence of the e^+ and n reduces backgrounds from single-particle events.

⁹Scintillating materials generally contain stable hydro-carbons and emit detectable light when charged particles move through the material (ionizing radiation).

capture on Cd¹⁰. By observing an excess of counts in their detector when the reactor was on compared to when the reactor was not operating, they had tentative evidence of a neutrino detection. They confirmed their work by positioning their detector at the Savannah River plant to reduce cosmic-ray backgrounds, and published the first confirmed detection of neutrinos in 1956 [8].

1.1.4 Observation of Multiple Flavors

While the detection of the neutrino confirmed Pauli's proposition and Fermi's model of beta-decay, it did not answer every question. Neutrinos were observed to be produced in radioactive processes, such as the beta decays which provide fuel for the nuclear fusion processes which generate energy inside every star. John Bahcall made a prediction using the standard solar model [9] to estimate the rate of neutrinos with high enough energy to convert chlorine to argon¹¹ and be observed in Ray Davis's detector [10]. With a first observation of solar neutrinos in-hand, characterization of the rate of neutrinos produced by the Sun would give a new level of understanding to our model of the solar core, as the "ghostly" neutrinos were able to escape from the Sun without interacting. However, the experiment only observed 1/3 of the neutrinos predicted — the "solar neutrino problem" [11].

Meanwhile, a 1962 experiment at Brookhaven National Laboratory by Lederman, Schwartz, and Steinberger established that there must be more than one "flavor" of neutrino [12]. These were distinguished as "electron" and "muon" neutrino flavors for consistency with the classification of the multiple generations of charged leptons. This scheme suggested a third generation of neutrino associated with the "tau" lepton should exist, but it was not detected until the DONUT experiment in 2000 [13]. These three flavors are the only known neutrinos which can interact via the weak force¹² [14].

In 1957, Bruno Pontecorvo proposed that neutral particles like neutrinos could undergo mixing, or transition between matter and antimatter states similar to kaons [15]. The detection of multiple neutrino flavors reignited the neutrino mixing theory, and within

¹⁰This is a similar strategy to the way we monitor neutron rates in Neutrino Alley with MARS! More in Chapter 8.

¹¹The "high"-energy neutrinos which could be observed were predicted to occur from a small fraction of Boron decays.

¹²These are called the "active" flavors, and are the only neutrinos that could be lighter than the Z^0 boson and interact via the weak force. There are theories which propose the existence of "sterile" neutrinos that are defined by their lack of interaction via the weak force.

the decade theories of neutrino oscillation between the known flavor states were developed¹³ [16, 17]. It would take a few more decades, though, for evidence from the Super-Kamiokande experiment in 1998 [18] and final confirmation from the SNO experiment in 2001 [19] to observe that neutrinos do indeed “oscillate” or change flavor states. This resolved the solar neutrino problem — it wasn’t that only 1/3 of the neutrinos were produced inside the Sun, it was that 2/3 of the electron-flavored neutrinos that were produced in the core of the Sun had oscillated to a different flavor that could not be observed by Ray Davis’s detector.

1.2 Neutrinos and the Standard Model

In the development of the Standard Model, which unified the electromagnetic and weak forces on this fundamental scale, several assumptions about the properties of neutrinos enable predictions of their interactions with other fundamental particles. In this section, we’ll take a look at the assumptions we make, and the general extension to the standard model which enables us to make predictions about neutrino interactions and oscillations.

1.2.1 Properties

There are three types, or flavors, of neutrinos that each form a doublet with one of the charged leptons. As such, we have electron neutrinos (ν_e), muon neutrinos (ν_μ), and tau neutrinos (ν_τ), and their associated anti-particles ($\bar{\nu}_e, \bar{\nu}_\mu, \bar{\nu}_\tau$). The flavors are important for determining which interactions will take place, as lepton number is believed to be conserved for each flavor in weak interactions — it must be the same for the initial and final states. Beta decay ($d \rightarrow u + \bar{\nu}_e + e^-$) and IBD ($u + \bar{\nu}_e \rightarrow d + e^+$), for example, are processes that only $\bar{\nu}_e$ will undergo, as they have lepton number -1 to balance the electron’s lepton number of $+1$ during β -decay and the positron’s lepton number of -1 during IBD.

As we previously mentioned, neutrinos possess no electric charge. All leptons, being spin-1/2 particles, are fermions. Neutrinos also have a non-zero but very small mass. The three active flavors are superpositions of three unique mass states (ν_1, ν_2 , and ν_3). The absolute mass scale of these mass states has not yet been determined ($m_\nu < 0.8$ eV at 90% confidence level [20]). However, because the absolute neutrino mass scale is sig-

¹³This mixing is analogous to but *not* the same as kaon mixing. The matter and antimatter neutrino states mix separately, such that $\nu_X \rightarrow \nu_Y$ and $\bar{\nu}_X \rightarrow \bar{\nu}_Y$ are allowed, but $\nu_X \rightarrow \bar{\nu}_Y$ is not, even for $X = Y$.

nificantly smaller than that of any of the other Standard Model particles and the nature (Majorana or Dirac) is still unknown [21], neutrino mass is purposefully excluded from the Standard Model because a renormalizable mass term cannot be constructed [22]. Cross-section calculations treat neutrinos as massless with minimal loss of accuracy, but massless neutrinos cannot describe the observed oscillations. Instead, the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix is the commonly accepted framework for making predictions about neutrino oscillations [22]. This extension to the Standard Model describes the probabilities of each flavor of neutrino oscillating into the others as a function of time, and offers a robust formalism that can predict the observables of interest.

Neutrino oscillations arise because, as we previously stated, the flavor eigenstates of neutrinos are superpositions of neutrino mass eigenstates. Let's consider a case with two neutrino flavor states ν_α and ν_β that are superpositions of mass states ν_a and ν_b . We introduce a "mixing angle" θ , such that the superposition of states can be explicitly written:

$$\begin{aligned} |\nu_\alpha\rangle &= \cos\theta |\nu_a\rangle + \sin\theta |\nu_b\rangle, \\ |\nu_\beta\rangle &= -\sin\theta |\nu_a\rangle + \cos\theta |\nu_b\rangle. \end{aligned}$$

Let's assume that a particle ν_α is created through some interaction, and we want to determine the probability that some time t later the neutrino will be detected as the other flavor ν_β . The time evolution of a particle is dependent on eigenstates of the Hamiltonian H ; for neutrinos, these are the mass states, with eigenvalues $E_{(a,b)}$ that are described by their mass $m_{(a,b)}$ and momenta $\vec{p}_{(a,b)}$. With the assumption of natural units so that $\hbar = c = 1$, we can write the time evolution of this particle as

$$|\nu(t)\rangle = e^{-iHt} |\nu_\alpha\rangle = e^{-iE_a t} \cos\theta |\nu_a\rangle + e^{-iE_b t} \sin\theta |\nu_b\rangle,$$

such that the probability P of finding the neutrino in a particular state after some time t is

$$\begin{aligned} P(\nu_\alpha \rightarrow \nu_\beta) &= |\langle \nu_\beta | \nu(t) \rangle|^2 \\ &= |(-\sin\theta \langle \nu_a | + \cos\theta \langle \nu_b |) (e^{-iE_a t} \cos\theta |\nu_a\rangle + e^{-iE_b t} \sin\theta |\nu_b\rangle)|^2 \\ &= |e^{-iE_b t} \sin\theta \cos\theta - e^{-iE_a t} \sin\theta \cos\theta|^2 \\ &= (\sin\theta \cos\theta)^2 (2 - 2\cos((E_a - E_b)t)). \end{aligned}$$

The eigenstate energies $E_i = \sqrt{m_i^2 + |\vec{p}_i|^2}$ can be expressed using the particle energy E for the ultra-relativistic limit, such that $|\vec{p}_i| \approx |\vec{p}| \approx E$ since $|\vec{p}_i| \gg m_i$. It then follows

that $E_i \approx E + \frac{m_i^2}{2E}$. The time t can be expressed as the distance the neutrino has travelled in that time, such that $t \approx L$. With these modifications, we find

$$\begin{aligned} P(\nu_\alpha \rightarrow \nu_\beta) &= \frac{\sin^2 2\theta}{4} \left(2 - 2 \cos \left(\frac{m_a^2 - m_b^2}{2E} L \right) \right) \\ &= \sin^2 2\theta \sin^2 \left(\frac{m_a^2 - m_b^2}{2E} L \right). \end{aligned}$$

This oscillation probability depends only on the mass ‘‘splitting’’ $\Delta m_{ab}^2 = m_a^2 - m_b^2$, the mixing angle θ , and the ratio L/E .

The probabilities for the full three-neutrino mixing are calculated using the PMNS matrix, which can be parameterized as

$$U = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_{\text{CP}}} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta_{\text{CP}}} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta_{\text{CP}}} & s_{23}c_{23} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta_{\text{CP}}} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta_{\text{CP}}} & c_{23}c_{13} \end{pmatrix}, \quad (1.1)$$

where $s_{ij} \equiv \sin \theta_{ij}$ and $c_{ij} \equiv \cos \theta_{ij}$. The rows are organized by flavor-state (ν_e, ν_μ, ν_τ) while the columns indicate the mass-state (ν_1, ν_2, ν_3)¹⁴. This matrix (and thus, the oscillation probabilities) then depends on four parameters: the three mixing angles θ_{12}, θ_{13} , and θ_{23} and the CP-violating phase δ_{CP} ¹⁵. Mixing-angle measurements are a significant target of interest for a variety of oscillation experiments, while δ_{CP} is primarily constrained by global fits to experimental observations¹⁶ [22].

Similar to the two-neutrino example, the probabilities will depend on the mixing angles and the mass splittings, or the differences between the neutrino mass states ($\Delta m_{12}^2, \Delta m_{13}^2, \Delta m_{23}^2$, where $\Delta m_{ij}^2 \equiv m_j^2 - m_i^2$). To date, experiments have measured the mass splittings Δm_{12}^2 and $|\Delta m_{13}^2| \approx |\Delta m_{23}^2|$. The sign of Δm_{13}^2 has not been measured, leading to two distinct possibilities for the neutrino-mass ordering: the ‘‘normal ordering’’ $m_1 < m_2 < m_3$ and the ‘‘inverted ordering’’ $m_3 < m_1 < m_2$. Thus, neutrino physicists are very determined to unambiguously determine the mass ordering with future oscillation experiments like DUNE [23] and HyperK [24], while current and future neutrino-mass experiments like KATRIN [25] and Project-8 [26] seek to set the absolute scale for the neutrino mass.

¹⁴To connect the matrix representation to the two-neutrino notation, note that the muon-flavored neutrino state can be written as $|\nu_\mu\rangle = U_{\mu 1} |\nu_1\rangle + U_{\mu 2} |\nu_2\rangle + U_{\mu 3} |\nu_3\rangle$ and $U_{\mu 1} = -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta_{\text{CP}}}$.

¹⁵This is technically only true for Dirac neutrinos. If the neutrino is found to be a Majorana particle, there are two additional complex phases.

¹⁶Current-generation experiments are not well tuned for direct δ_{CP} measurements.

The anti-particles of all fermions are defined by charge conjugation from their normal matter counterparts. Since neutrinos are charge-neutral, this reduces the condition to a difference in helicity, the projection of a particle’s spin vector onto its momentum. More simply put, helicity is a measure of whether a particle’s spin is aligned with its direction of travel (right-handed) or antialigned (left-handed). All neutrinos are left-handed, and all anti-neutrinos are right-handed. Since this difference in helicity and the sign of the lepton numbers are the only differences between neutrinos and anti-neutrinos, it is possible for neutrinos to be their own antiparticles if lepton-number symmetry is violated (consider a Lorentz boost into a frame where the momentum is moving in the opposite direction). This is another important open question in modern neutrino physics – we try to ascertain whether neutrinos are Majorana or Dirac in nature by searching for a very rare decay process, neutrinoless double beta decay, that can only occur if neutrinos are their own antiparticles (Majorana) [27].

1.2.2 Interactions

As leptons, neutrinos do not interact via the strong force, and thus do not couple to gluons. As charge-neutral particles, neutrinos also do not interact via the electromagnetic force, and thus do not couple to photons. Their mass is very small, thus gravitational interactions are negligible. Thus, the study of neutrino processes is primarily concerned with the weak force. However, the weak force has a very short range (much less than the radius of a proton) and small coupling, resulting in very low interaction rates for all neutrino processes.

With the historical context we established, it’s useful to discuss beta decay and inverse beta decay in our modern language as a segue into physics using the Standard Model. At the fundamental level in this theory, β^- decay is a weak process by which a down quark becomes an up quark, thus changing a neutron (valence quark content: up, down, down) to a proton (valence quark content: up, up, down). The left panel of Figure 1.3 is the Feynman diagram for this process. The initial down quark undergoes a decay, and becomes an up quark by emitting a W boson. The W boson is far off its mass shell, and decays very quickly (considered instantaneous for most purposes) into an electron and an anti-electron neutrino. This is the representation of a short-ranged interaction in particle physics; the initial and final states of the system we are observing can be distinguished from left to right in the diagram, and the interaction is mediated by the exchange of a boson. The IBD process is quite similar.

While the W^\pm bosons are responsible for all charged-current (CC) weak interactions,

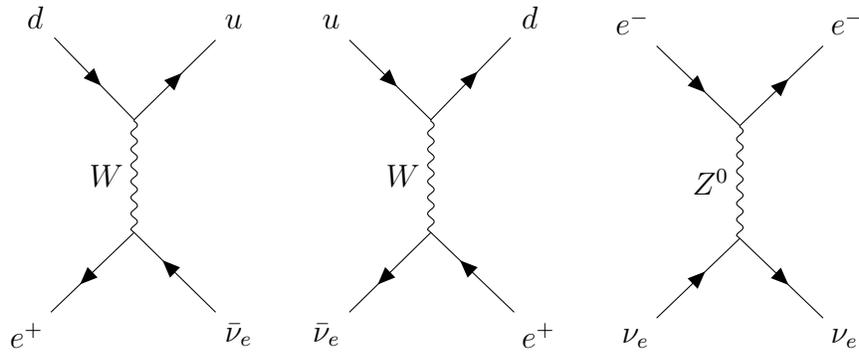


Figure 1.3: Feynman diagrams for β -decay (left), inverse beta decay (center), and neutral-current neutrino-electron scattering (right). These diagrams can be interpreted as time flowing from left to right.

there are also neutral-current (NC) interactions which are mediated by the Z^0 boson. A neutral-current scattering process is shown in the right-most panel of Figure 1.3. My work primarily focuses on investigations of a particular neutral-current interaction: coherent elastic neutrino-nucleus scattering.

Chapter 2

Coherent Elastic Neutrino-Nucleus Scattering

Neutrinos are created across a broad energy scale from a variety of terrestrial and astrophysical sources, as shown in Figure 2.1. Some interactions can take place across all incident neutrino energies, such as the neutrino-electron scattering shown in Figure 2.1, but other interactions can have upper or lower energy bounds. To describe these constraints, we often relate the momentum transfer q to the de Broglie wavelength λ of the mediating force-carrier, such that $q \propto \lambda^{-1}$. Deep-inelastic scattering, for example, is a process by which an incident neutrino scatters off a single quark inside a nucleon and causes a hadronic shower ($\nu + q \rightarrow \nu + \text{h.s.}$). This interaction has a lower threshold requiring that the mediator for this interaction must have a wavelength small enough to resolve individual quarks¹.

The opposite is also true – there are some interactions which have a ceiling, such that the momentum transfer cannot exceed a certain amount. Coherent elastic neutrino-nucleus scattering (CEvNS) is the process by which an incident neutrino scatters off a nucleus without changing the nuclear state or interacting with any individual nucleon². The mediator for this process must have $\lambda > R$, where R is the size of the specific nucleus,

¹The “resolution” we describe here is analagous to microscopic resolution, or the ability to distinguish two closely positioned features. If the resolution of an optical microscope is too low, features may blur together and be inobservable. In particle physics, the resolution of an interaction is similar – if the wavelength of an incident particle is too large (low momentum = low resolution), the “features” of the target particle will be ignored.

²We use the acronym “CEvNS” in this work, pronounced “sevens”. In the literature, CEvNS may also appear as “CE ν NS”, “CNS”, or “CENNS”, though the latter acronyms can also refer to neutron-nucleus scattering.

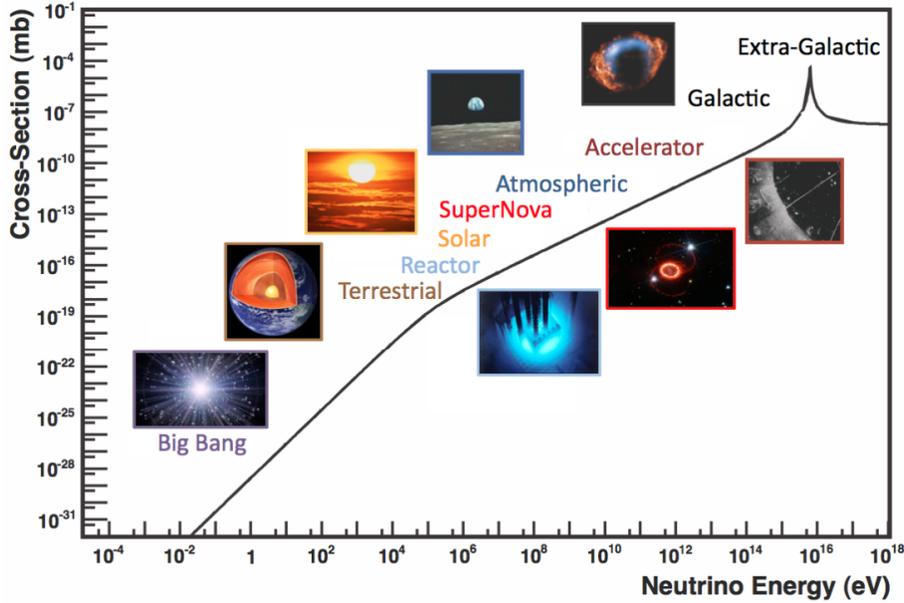


Figure 2.1: An illustration of the broad range of neutrino energies. Sources of neutrinos with a given energy are shown in comparison to the neutrino-electron scattering cross section for $\bar{\nu}_e + e^- \rightarrow \bar{\nu}_e + e^-$. Reproduced from [28].

such that there is a strict constraint ($qR < 1$). As a result, this process is associated with only low-energy incident neutrinos, up to energies in the tens-of-MeV for a moderately-sized nuclei.

In this chapter, we'll examine CEvNS in detail. We lay a foundation for the particle physics formalism we'll use throughout this work in Section 2.1 before describing the Standard Model prediction in Section 2.2. Section 2.3 describes some of the interesting physics that can be explored with observations of CEvNS, and we detail varied experimental approaches in Sections 2.4 and 2.4.2.

2.1 Neutrino Scattering

Neutrino *scattering* is a type of interaction that is easily thought of as an interaction involving a fundamental particle that appears in both the initial and final states³. Consider the scattering process plotted across a broad range of neutrino energy scales in Figure 2.1, $\bar{\nu}_e + e^- \rightarrow \bar{\nu}_e + e^-$, in which an incident neutrino scatters off an electron. We can further

³Generally, most interactions can be thought of as scattering processes; the description I use here is intended to relate the particle-physics regime to the classical picture of scattering.

specify that this is also an *elastic* scattering process; the initial and final states are completely identical. Compare this with the deep *inelastic* scattering we discussed previously $\nu_e + q \rightarrow \nu_e + \text{h.s.}$, in which the hadronic shower did not exist in the initial state and is created as a result of the interaction.

In scattering processes involving nuclei, there are several possibilities, listed here in order of increasing momentum transfer:

- Low-energy: The incident particle interacts with the whole nucleus in phase — this is called *coherent* scattering, and generally does not change the nucleus.
- Moderate-energy: The incident particle interacts with individual nucleons — this is called quasi-elastic or incoherent scattering, and can knock a nucleon out of the nucleus or destabilize it by changing a nucleon ($p \leftrightarrow n$).
- High-energy: The incident particle interacts with the partons within a nucleon — this is called deep inelastic scattering, and generally destroys the nucleus.

Notably, all low-momentum coherent scattering processes feature an increase to the interaction probability because the incident particle interacts in phase with the whole nucleus. In this way, there are constructively interfering terms, where incoherent interactions can have destructively interfering terms⁴.

Generally, particle physicists discuss scattering rates (and other processes) in terms of the **cross section** of a process. This term is rooted in classical scattering: the cross-sectional area of an object is directly proportional to the likelihood of striking it. The cross section for a particular reaction, denoted σ with units of area⁵, is a measure of the probability for a particular process to occur given the properties and quantum numbers of all particles involved in the initial and final states. We often calculate differential cross sections to examine the probability given a particular condition (e.g., the probability of an interaction which produces a final-state particle with energy E can be written as the differential cross section $d\sigma/dE$). Detector event rates can then be calculated from the cross sections by convolving the interaction probabilities with the incident flux, or rate of particles per second per unit area, and the detector-specific properties. For the rest of this work, we'll focus on measurements of the highest-probability interaction for low-energy neutrinos and moderately sized nuclei, CEvNS.

⁴This is similar to how ocean waves can combine to make one large wave (constructive) or “cancel” each other out (destructive).

⁵Particle physicists often use units of “barns” (1 b = 10^{-28} cm²), a measure derived from the famous phrase “can’t hit the broad side of a barn.”

2.2 Standard Model Cross Section

There are both charged-current and neutral-current interactions within the unified theory of electroweak physics established by Glashow, Weinberg, and Salam in 1968 [29–31]. Their proposal of a neutral current, confirmed by the Gargamelle experiment [32], inspired Daniel Freedman to posit the existence of CEvNS in 1973 [33]. Since CEvNS is an elastic scattering interaction without any other lepton in the interaction, it is flavor independent – any neutrino or antineutrino flavor has the same interaction probability. In the Standard Model, this cross section can be expressed as [34]

$$\frac{d\sigma}{dT} = \frac{G_F^2 M}{2\pi} \left[(G_V + G_A)^2 + (G_V - G_A)^2 \left(1 - \frac{T}{E_\nu}\right)^2 - (G_V^2 - G_A^2) \frac{MT}{E_\nu^2} \right], \quad (2.1)$$

where G_F is the Fermi constant, M is the mass of the nucleus, E_ν is the neutrino energy, and T is the nuclear recoil energy. The expressions G_V and G_A are given by

$$G_V = (g_V^p Z + g_V^n N) F_{\text{nucl}}^V(q^2), \quad (2.2)$$

$$G_A = (g_A^p (Z_+ - Z_-) + g_A^n (N_+ - N_-)) F_{\text{nucl}}^A(q^2), \quad (2.3)$$

where Z_\pm and N_\pm are the number of spin up (+) or down (-) protons (Z) and neutrons (N). The couplings $g_{V,A}^{p,n}$ refer to the vector and axial couplings to protons and neutrons, and the $F_{\text{nucl}}^{V,A}(q^2)$ represent the nuclear form factors, or (non-relativistically) the Fourier transforms of the density distribution of nucleons within the nucleus, as a function of the momentum transfer q . The axial contribution G_A depends on the difference between spin-up and spin-down nucleons. Often, experimentalists choose detectors with even-even⁶ nuclei, such that the axial contributions drop out of the cross section entirely. The vector coupling to the proton $g_V^p \approx 1 - 4 \sin^2 \theta_W$ is suppressed because $\sin^2 \theta_W \approx 0.23$ [22]. As such, $G_V \propto N$, so the cross section is proportional to N^2 . The recoil energy $T = \frac{q^2}{2M}$ of an interaction can take values from zero to $T_{\text{max}} = \frac{2E_\nu^2}{M+2E_\nu}$. Thus, experiments must balance the growth of the cross section with M and N^2 against the ability to observe the small nuclear recoils.

By integrating over all possible recoil energies, T , it can be shown that the total cross section $\sigma \propto E_\nu^2$. Thus, the likelihood of a CEvNS interaction *increases* with the neutrino energy, though this must be balanced against the possibility for the momentum transfer to be too large and break the coherence condition $qR < 1$. Typically, neutrinos with $E_\nu < 50$ MeV will scatter via CEvNS for moderately-sized nuclei.

⁶The term “even-even” here means that there are equal numbers of spin up and down nucleons, such that $Z_+ - Z_- = N_+ - N_- = 0$.

2.3 Physics Sensitivities

Measurements of the CEvNS cross section can make a broad impact, ranging from constraining neutrino properties, to testing the Standard Model, to enabling new particle searches. In this section, we overview a few areas in which measurements of the CEvNS cross section can make a substantial contribution.

2.3.1 Electroweak Parameters and Neutrino Properties

In the cross section (Equation 2.1), there is direct sensitivity to the weak mixing angle through the coupling to the proton. By comparing experimental results to the predicted cross section, we can provide a test of the weak mixing angle with unique systematics compared to other measurements. Measurements that are inconsistent with the Standard Model expectation of the event rate could point to new physics in the low-energy regime, such as additional neutral currents [35]. Presently, measurements of the weak mixing angle through CEvNS [36, 37] are not competitive with other methods (e.g., [38–40]) as shown in Figure 2.2 (left), but they do provide a complementary approach that is specific to neutrino-quark interactions (more in Section 2.3.3).

The neutrino, though neutral, couples to charged force carriers (W^\pm) and can be interpreted as if there was a charged field surrounding it. In this way, it can be interpreted to have an effective charge radius [41]

$$\langle r_{\nu_X}^2 \rangle = \frac{G_F}{4\pi^2\sqrt{2}} \left[3 - \ln \left(\frac{m_X^4}{M_W^4} \right) \right], \quad (2.4)$$

where G_F is the Fermi constant, X represents the flavor of the neutrino, and m_X is the mass of the charged lepton associated with the flavor. This introduces a small flavor-dependent effect to the CEvNS cross section with a modification to the weak mixing angle $\sin^2 \theta_W \rightarrow \sin^2 \theta_W \left(1 - \frac{2}{3} M_W^2 \langle r_{\nu_X}^2 \rangle \right)$. By examining the flavor-dependent contribution, the effective neutrino charge radius can be constrained as shown in Figure 2.2 (right) [37, 42–44].

The cross section can also be written in terms of the neutrino magnetic moment for a spin-zero nucleus as [45]

$$\frac{d\sigma}{dT} = \frac{\pi\alpha^2\mu_\nu^2 Z^2}{m_e^2} \left(\frac{1 - T/E_\nu}{T} + \frac{T}{4E_\nu^2} \right), \quad (2.5)$$

where α is the fine structure constant, m_e is the mass of the electron, and μ_ν is the neutrino magnetic moment in units of Bohr magnetons. In the Standard Model, the neutrino

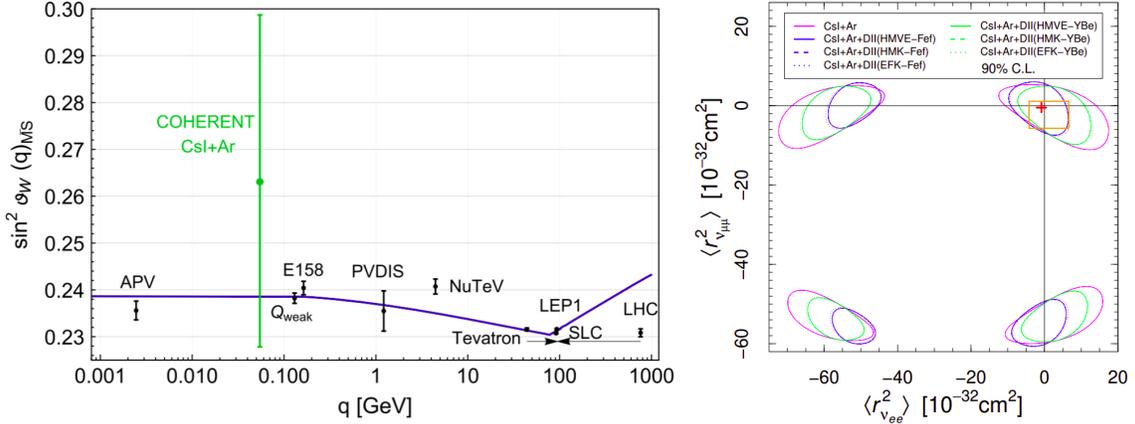


Figure 2.2: Left: A measurement of the weak mixing angle using combined CsI and Ar CEvNS data from the COHERENT collaboration. Reproduced from [43]. Right: Constraints on the effective neutrino charge radius, jointly fit with CEvNS measurements from COHERENT and constraints from the Dresden-II reactor experiment. This figure is reproduced from [44], which also studies non-standard interactions that allow the incident and final-state neutrino to be different flavors; this plot, however, constrains that model so that the neutrino flavor cannot change in the interaction, such that their $\langle r_{\nu_{xx}}^2 \rangle$ is equivalent to our $\langle r_{\nu_x}^2 \rangle$ from Equation 2.4.

magnetic moment is extremely small ($\mu_\nu = 3.2 \times 10^{-19} \mu_B \left(\frac{m_\nu}{1\text{eV}}\right)$ [22]), and any observation that would be consistent with a larger value is an indication of new physics. Current constraints from CEvNS measurements are detailed in [46].

2.3.2 Nuclear Form Factor

Similar to the use of parity-violating electron scattering in PREX [47] and CREX [48], CEvNS is an interesting channel to use as a probe of the nuclear neutron density. Non-relativistically, the weak nuclear form factor $F_{\text{nucl}}^V(q^2)$ is the Fourier transform of the density distribution of protons and neutrons within the nucleus. Given the suppression of the proton contribution to the coupling, CEvNS is mainly sensitive to the neutron distribution within the nucleus.

Because the proton density distribution (with radius R_p) is typically easier to measure and is predicted to be smaller than the neutron distribution, measurements of the nuclear form factor through CEvNS can give insight into the radius of the neutron density distribution R_n and thickness of the neutron “skin” ($R_{\text{skin}} = R_n - R_p$), or how much the

neutron distribution extends past the proton distribution. In particular, R_{skin} is important to understand the behavior of neutron stars because it is strongly correlated to the symmetry energy, a quantity related to the equation of state of neutron stars that describes the increase in the energy of a system when there are not equal numbers of protons and neutrons. Measuring the neutron-skin thickness of nuclei, then, gives important insights into the properties of neutron-rich matter [49].

Measurements of the CEvNS cross section by COHERENT (see Chapter 3) and constraints on the cross section from the Dresden-II reactor experiment have produced constraints on the neutron-skin thickness for Ar, Ge, Cs, and I [44, 50].

2.3.3 Beyond the Standard Model

As mentioned in Section 2.3.1, new measurements that are inconsistent with the Standard Model predictions can point to new physics. Measurements of the CEvNS cross section, in particular, can constrain the strength of non-standard interactions (NSI), or the likelihood for neutrinos to interact with quarks in ways that we don't yet know about, by extending our theoretical model.

Generally, NSI would arise from a modification to the vector couplings $g_V^{p,n}$ by allowing for new couplings $\epsilon_{\alpha\beta}^{qV}$, where $q = u, d$ is the quark participating in the interaction⁷ and α, β correspond to the incident and outgoing neutrino flavor, respectively. These, like modifications to the cross section associated with the effective neutrino charge radius, introduce a neutrino-flavor dependence to the predicted flavor-blind cross section of CEvNS. The cross section from Equation 2.1 can then be rewritten as [34]

$$\begin{aligned} \left(\frac{d\sigma}{dT}\right)_{\nu\alpha} &= \frac{G_F^2 M}{\pi} (F_{\text{nucl}}^V(2MT))^2 \left[1 - \frac{MT}{2E_\nu^2}\right] \times \\ &\times \left\{ \left[Z(g_V^p + 2\epsilon_{\alpha\alpha}^{uV} + \epsilon_{\alpha\alpha}^{dV}) + N(g_V^n + \epsilon_{\alpha\alpha}^{uV} + 2\epsilon_{\alpha\alpha}^{dV}) \right]^2 \right. \\ &\left. + \sum_{\alpha \neq \beta} \left[Z(2\epsilon_{\alpha\beta}^{uV} + \epsilon_{\alpha\beta}^{dV}) + N(\epsilon_{\alpha\beta}^{uV} + 2\epsilon_{\alpha\beta}^{dV}) \right]^2 \right\}, \end{aligned} \quad (2.6)$$

where α is the flavor of the incoming neutrino and β is the outgoing neutrino flavor. Experiments can then measure the CEvNS cross section, and determine what values of the couplings $\epsilon_{\alpha\beta}^{qV}$ could be consistent with the observations. An example of these kinds of constraints is shown in Figure 2.3.

⁷We do not allow for other quark flavors in this model since only up and down quarks can appear as valence quarks in the protons and neutrons comprising the nucleus, and will dominate the interaction probabilities.

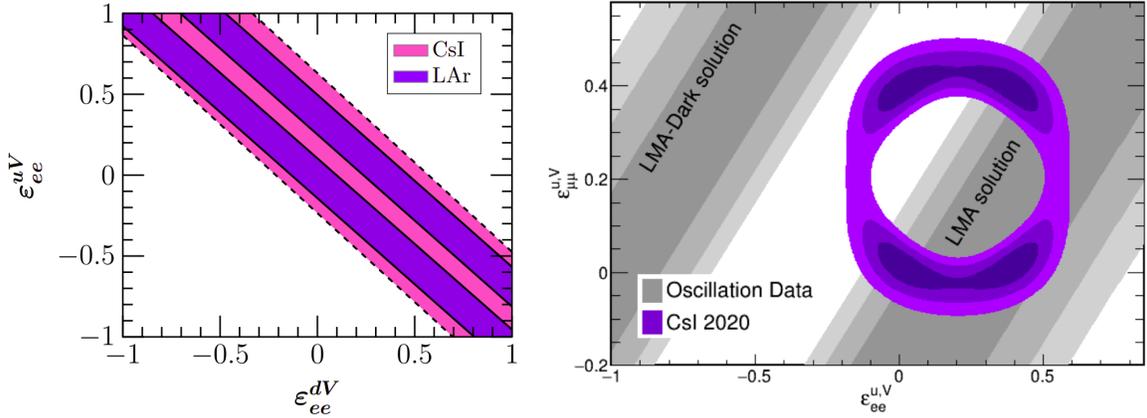


Figure 2.3: Left: Constraints on NSI parameters from CEvNS measurements made by the COHERENT collaboration, reproduced from [37]. Right: Constraints on NSI parameters from CEvNS measurements made by the COHERENT collaboration compared to the allowed values from neutrino oscillation data, reproduced from [54].

The values of each ϵ can be either positive or negative, leading to respective enhancement or suppression of the CEvNS cross section; for some combinations of NSI parameters, a specific combination of Z and N values in the nucleus might even result in the same Standard Model cross section, so measurements of CEvNS on multiple targets help to break any accidental degeneracies [51].

Models of NSI parameters, however, also impact other areas of neutrino physics. In particular, the existence of NSI allows a degeneracy in neutrino oscillation measurements; the so-called “LMA-Dark” degeneracy, for example, affects the interpretation of measurements of θ_{12} [52] and measurements of the neutrino mass ordering [53]. However, CEvNS measurements help to break this degeneracy, as shown in Figure 2.3 (right).

2.3.4 Model Inputs

With its (relatively) high cross section, measurements of the CEvNS cross section are useful as input into models of physical systems which depend on the way neutrinos will interact with matter. For example, CEvNS is expected to play a role in the dynamics of core-collapse supernova. The immense flux of neutrinos with energies in the tens-of-MeV produced during the supernova, and the high cross section of CEvNS, provide a way to reignite a stalled shock wave [55].

Measurements of CEvNS also inform searches for other processes with an observable nuclear recoil as the only signature, such as searches for WIMP dark matter. Solar and

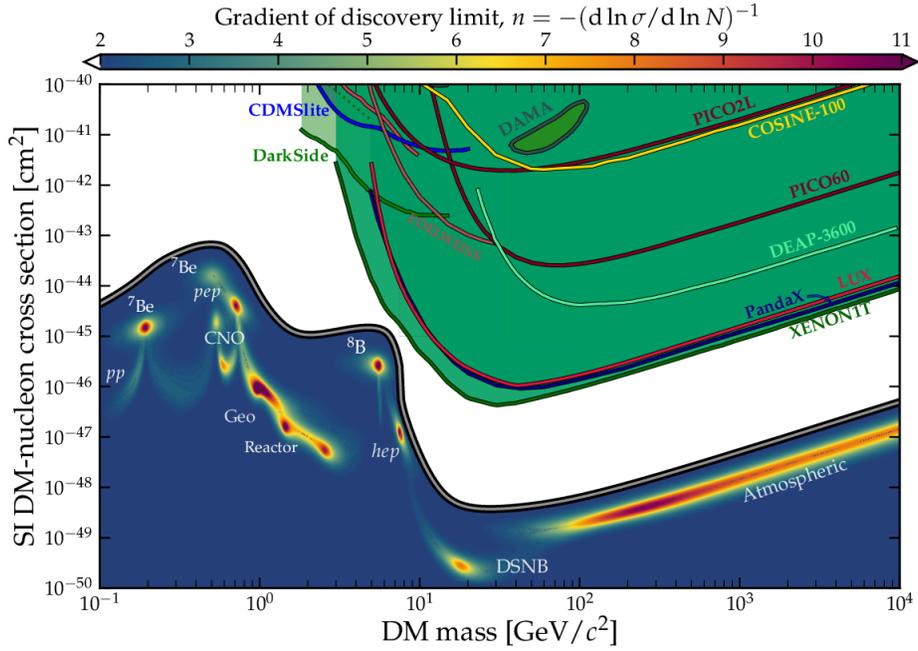


Figure 2.4: A plot reproduced from [57] illustrating the status and projections of xenon-based direct-detection WIMP experiments as exclusion curves (the lines on this plot exclude towards the upper right corner, with the shaded-green region indicating the current best limit) on the spin-independent (SI in the figure) cross section approach the “neutrino fog”. CEvNS interactions from neutrinos of varied sources (labeled in figure) will contribute background nuclear-recoil events to WIMP experiments and limit the statistical precision of any possible signal. The discovery limit, n , measures the impact of background events using the significance, $\sigma \propto N^{-1/n}$; the feasibility of observing a WIMP excess N decreases as the discovery limit increases. In the regular Poissonian regime, $n = 2$, but any possible measurement loses significance as statistical fluctuations in the large and irreducible neutrino background cloud the signal ($n > 2$).

atmospheric neutrinos, along with neutrinos from the diffuse supernova neutrino background, undergoing CEvNS interactions will produce an irreducible background in WIMP detectors, historically referred to as the “neutrino floor”, illustrated in Figure 2.4 [56, 57]. Precision measurements of the CEvNS cross section, however, transition this background from an opaque “floor” to a translucent (but dense) “fog”, as this irreducible contribution to the nuclear recoil spectrum observed by WIMP detectors can be better understood and modeled so that next-generation experiments can delve into this parameter space [58].

2.3.5 Future Detection Strategies

CEvNS measurements also enable future physics searches. With precise knowledge of the CEvNS cross section, this process can be used as a signature to search for new physics, like sterile neutrinos. Since CEvNS is a flavor-independent process, a sterile search using low-energy neutrinos can look for the absence of all active flavors, rather than searching in individual flavor channels like inverse beta decay experiments do [59].

The high cross section of CEvNS also provides a new avenue to design and deploy relatively compact detectors for non-proliferation efforts [60] to monitor nuclear reactors [61] or spent nuclear fuel [62]. These efforts generally contribute to the NuTools exploration of the role of neutrino physics in nuclear energy and security [63].

2.4 Experimental Requirements

As mentioned in its initial proposition, CEvNS detection opens the doors to numerous physical insights, but there are incredible challenges surrounding the measurements. Freedman declared it an “act of hubris” [33] to try to observe this process due to the low interaction rate of neutrinos and the incredibly difficult-to-detect observable. In this section, we overview the strategies and technological advancements that made cross-section measurements of CEvNS possible.

2.4.1 Detector Sensitivity

A large energy response is a necessity to ensure that a detector can search for small nuclear recoils⁸. This was the biggest hurdle preventing CEvNS detection; it took decades of technological advancement (most notably associated with the search for WIMP dark matter) to be able to achieve the sensitivity to observe nuclear recoils in the tens-of-keV.

The quenching factor of a detector, or the conversion from nuclear recoil energy (e.g., keVnr) to observed energy (e.g., keVee, meaning “keV electron equivalent”) is important to correctly interpret measurements. Dedicated efforts are underway at the Triangle Universities Nuclear Laboratory (TUNL) to support CEvNS searches (among other physics endeavors) by making precise measurements of the quenching factors for a variety of nuclear targets and detector technologies (e.g., [64–67]).

⁸For scintillating detectors, this means a high light yield — a large amount of scintillation. For semiconductors, this means a large amount of ionization.

2.4.2 Knowledge of Incident Fluxes

Given a good understanding of the detector response, the remaining requirements for any experimental cross-section measurement are to understand the incident flux of particles that could initiate the interaction of interest and to understand any incident flux of particles that could mimic the measured observable. With extremely sensitive detectors, there are many backgrounds to consider. Some, like cosmic-ray or ambient γ backgrounds, can be mitigated with choices of shielding and overburden that suit the choice of detector. Neutron backgrounds, however, are especially important to understand because elastic neutron scatters can cause the same nuclear-recoil observable as CEvNS. It is impossible to distinguish between neutron- and neutrino-induced recoils, which necessitates precise knowledge of the neutrino flux (signal) and the neutron flux (background) through the detector⁹. There are several known sources which can provide a neutrino flux in the right energy range for CEvNS experiments, and we overview a few of these here.

Astrophysically, the production of low-energy neutrinos is primarily associated with solar neutrinos or supernova neutrinos. Supernovae, however, are notoriously unpredictable (the most recent supernova observed through neutrino detection was in 1987) and consequently not the best choice for a dedicated CEvNS experiment¹⁰. Solar neutrinos generally have energies with hundreds of keV, but decays of ${}^8\text{B}$ can produce neutrinos up to ~ 14 MeV; the XENON1T [69] and PandaX [70] collaborations recently performed the first searches for CEvNS on xenon using ${}^8\text{B}$ solar neutrinos, with sensitive detectors deployed deep underground to limit backgrounds.

Artificial neutrino sources, however, generally have larger measurable fluxes because detectors can be placed significantly closer to the sources. Nuclear reactors, for example, have an extremely high flux of $\bar{\nu}_e$ neutrinos with $E_\nu < 1$ MeV. Several experiments are aiming to measure CEvNS at reactors, but measurements at reactors are quite difficult due to the high neutron background rates and lower-energy recoils [71–80].

Particle accelerators, however, provide the easiest path forward for CEvNS experiments (e.g., [81–84]). In particular, pion decay-at-rest sources feature the controlled production of neutrinos with $E_\nu < 50$ MeV. This is typically done by colliding energetic protons with massive nuclei, creating unstable mesons that quickly decay into neutrinos. The biggest benefit, though, is that the production of neutrinos using accelerator pulses results in a distinctive arrival-time distribution at deployed detector locations. This is es-

⁹Hence: the importance of this work for the COHERENT collaboration.

¹⁰Measurements of the CEvNS cross section can, however, validate a new channel for observing supernova neutrinos and further motivate CEvNS experiments (e.g., [68]).

pecially beneficial for constraining parameters in beyond Standard Model interactions, as the flavor-dependence can be isolated in the distinguishable arrival times of muon and electron neutrinos. The rest of this work will focus on the efforts of the COHERENT collaboration to measure CEvNS and other interesting cross sections at the Spallation Neutron Source (SNS) at Oak Ridge National Laboratory. We'll provide an overview of the experiment in Chapter 3, then closely examine the expected neutrino (Chapters 4 and 5) and neutron (Chapters 6, 7, and 8) fluxes at specified COHERENT detector locations in the SNS target hall.

Chapter 3

COHERENT

The COHERENT experiment is a collaboration of 21 institutions in the US, Canada, Russia, and South Korea, formed in 2013 with the goal to unambiguously observe CEvNS. With $\mathcal{O}(100)$ active members, the collaboration found an institutional home at Oak Ridge National Laboratory (ORNL) and made the first discovery of CEvNS in 2017 [85]. While the only measurements of the CEvNS cross section are from COHERENT’s CsI and Ar detectors [54, 86], experimental programs around the globe are using different detectors and neutrino sources to search for CEvNS [87]. In this chapter, I overview the experimental efforts of the COHERENT collaboration in past, present, and future, and highlight the specific tie-ins of my own work to our collaborative goals.

3.1 The Spallation Neutron Source

The Spallation Neutron Source (SNS) at ORNL is currently the most intense neutron production facility in the world. Featuring a proton accelerator and a liquid mercury target, the SNS currently operates at a nominal 1.4 MW, with a 60 Hz, 1 GeV proton beam. The incident protons spall the Hg nuclei, producing large numbers of neutrons and, as a byproduct, large numbers of charged pions. While the majority of π^- will capture on the positively-charged nuclei within the dense Hg target, the π^+ will quickly ionize and come to rest.

Pions are unstable mesons which undergo weak decay. The neutrino production chain

$$\pi^+ \rightarrow \mu^+ + \nu_\mu \tag{3.1}$$

$$\mu^+ \rightarrow e^+ + \bar{\nu}_\mu + \nu_e \tag{3.2}$$

will generally occur after the pions have been brought to rest (more in Chapter 4). The kinematics of these interactions then create a well-defined rate of neutrinos dependent on the rate of pions produced. In the case of the two-body π^+ decay, we can directly obtain the energies of the emitted particles from the foundational conservation of the four-momentum $p_i = p_f$. Recall that the energy-momentum relation $E^2 = |\vec{p}|^2 c^2 + m^2 c^4$ dictates that $p^\lambda p_\lambda = E^2 - |\vec{p}|^2 c^2 = m^2 c^4$. In the following derivation, we use natural units such that $c = 1$ and adopt the pion frame as our reference. With this assumption, we know that $\vec{p}_\pi = 0$, and $\vec{p}_\mu = -\vec{p}_\nu$, such that the following holds:

$$\begin{aligned}
p_i^2 &= p_f^2 \\
p_\pi^\lambda p_{\pi\lambda} &= (p_\mu + p_\nu)^\lambda (p_\mu + p_\nu)_\lambda \\
m_\pi^2 &= m_\mu^2 + m_\nu^2 + 2(E_\mu E_\nu - \vec{p}_\mu \cdot \vec{p}_\nu) \\
m_\pi^2 - m_\mu^2 - m_\nu^2 &= 2E_\mu E_\nu + 2|\vec{p}_\nu|^2 \\
&= 2E_\mu E_\nu + 2(E_\nu^2 - m_\nu^2) \\
m_\pi^2 - m_\mu^2 + m_\nu^2 &= 2E_\nu(E_\nu + E_\mu) \\
&= 2E_\nu m_\pi \\
\therefore E_\nu &= \frac{m_\pi^2 - m_\mu^2 + m_\nu^2}{2m_\pi}. \tag{3.3}
\end{aligned}$$

Treating $m_\nu \approx 0$, and using the known masses of the pion $m_\pi \approx 140$ MeV and muon $m_\mu \approx 105$ MeV, the two-body decay of a pion at rest will produce neutrino with the monoenergetic spectrum $E_\nu \approx 30$ MeV. The derivation of the three-body decay spectra do not follow as simply, but can be determined using Fermi's Golden Rule to calculate the differential decay rates.

The flavor-dependent SNS neutrino fluxes $\phi_\nu(E)$ are estimated analytically (assuming the μ^+ decays at rest) in the kinematically allowed region $E_\nu \in \left[0, \frac{m_\mu^2 - m_e^2}{2m_\mu}\right]$ for a given pion production rate R_{π^+} and distance L from the target location as [88]

$$\phi_{\nu_\mu}(E_{\nu_\mu}) = \frac{R_{\pi^+}}{4\pi L^2} \frac{2m_\pi}{m_\pi^2 - m_\mu^2} \delta\left(1 - \frac{2E_{\nu_\mu} m_\pi}{m_\pi^2 - m_\mu^2}\right), \tag{3.4}$$

$$\phi_{\bar{\nu}_\mu}(E_{\bar{\nu}_\mu}) = \frac{R_{\pi^+}}{4\pi L^2} \frac{64}{m_\mu} \left(\frac{E_{\bar{\nu}_\mu}}{m_\mu}\right)^2 \left(\frac{3}{4} - \frac{E_{\bar{\nu}_\mu}}{m_\mu}\right), \tag{3.5}$$

$$\phi_{\nu_e}(E_{\nu_e}) = \frac{R_{\pi^+}}{4\pi L^2} \frac{192}{m_\mu} \left(\frac{E_{\nu_e}}{m_\mu}\right)^2 \left(\frac{1}{2} - \frac{E_{\nu_e}}{m_\mu}\right). \tag{3.6}$$

The spectra for $\bar{\nu}_\mu$ and ν_e have similar, but not identical shapes due to the time-ordering of the muon decay; the $\bar{\nu}_\mu$ must always be created first, and thus has a larger phase space

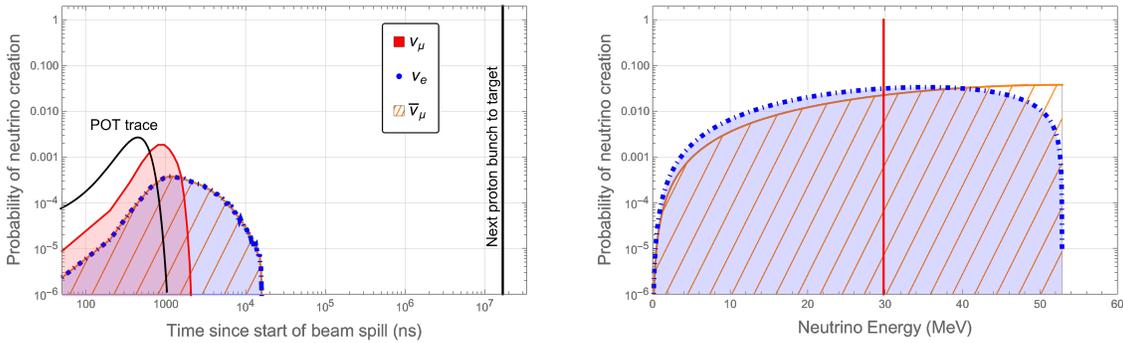


Figure 3.1: Expected distributions of neutrinos in time (left) and energy (right) for an ideal pion decay-at-rest source with the SNS proton-on-target (POT) trace illustrating the beam timing.

available at higher energies than the ν_e . These distributions are plotted in Figure 3.1 (right) treating the normalization as unity $\left(\frac{R_{\pi^+}}{4\pi L^2} = 1\right)$.

Perhaps the biggest advantage of the SNS, however, is its duty cycle. In Figure 3.1 (left), we show the timing distribution of neutrinos per incident pulse; the majority of neutrinos are expected to be produced within 10 μ s after a nuclear spallation. With a beam spill of 350 ns FWHM (the proton-on-target POT trace in Figure 3.1) sent to target 60 times per second (every 16(ish) ms), and short pion (26 ns) and muon (2.2 μ s) lifetimes, there is ample opportunity to analyze and subtract the steady-state backgrounds every beam spill. As we discussed in Chapter 2, this is one of the biggest benefits of CEvNS experiments at decay-at-rest neutrino sources — reactor experiments are more difficult because they do not have the background rejection of accelerators.

Within the next year, the SNS will begin a Proton Power Upgrade (PPU) [89] to increase the beam energy and power of the SNS accelerator. This upgrade will ultimately increase the neutrino production at the SNS, and, as we'll examine in detail in Chapter 4, the new 1.3 GeV beam energy is not high enough to initiate new production mechanisms (kaon production, high fraction of decay-in-flight neutrinos, etc.); the timing and energy distributions of our neutrinos will stay the same. The PPU will increase the nominal SNS operating parameters from 1 GeV/1.4 MW operations to 1.3 GeV/2.0 MW by 2024.

This process is, however, only the first stage of the fully upgraded SNS; within the next decade, ORNL is expected to construct a Second Target Station (STS) [90]. This facility will complement the operations of the current First Target Station (FTS) and feature new neutron beamlines optimized for cold neutrons, a rotating tungsten target assembly, and

the opportunity to create a neutrino hall with the ability to support 10-ton-scale detectors closer to the target with dense, high-Z shielding to reduce neutron backgrounds. This facility will use the same accelerator system as the FTS, and operate at 1.3 GeV. However, the 60 Hz accelerator will now split pulses between the FTS and the STS; 3 of every 4 pulses will go to the FTS (averaged 45 Hz, 2.0 MW operations), and the final pulse will be directed to the STS (15 Hz, 0.8 MW operations).

At the SNS, Neutrino Alley has become the experimental home to the full suite of COHERENT detectors. This basement hallway provides deployment locations between 19 and 28 m from the target center. The neutrino flux (signal) in the alley is of order $10^7 \nu \text{ cm}^{-2} \text{ s}^{-1}$. With the intense, pulsed flux of the SNS and the large CEvNS cross section, our neutrino detectors can be made significantly smaller than those searching for cosmogenic signatures or situated at long-baseline locations.

The dominant backgrounds in Neutrino Alley are cosmic rays, as our overburden is a modest 8 meters-water-equivalent, and beam-related neutrons which cannot be characterized in time regions anticoincident with a beam spill. We observe orders-of-magnitude reduction in the beam-related neutron flux (background) compared to other locations in the SNS target hall [85, 91]. A background of 511-keV γ s is present in Neutrino Alley during beam operations due to a hot-off-gas (HOG) pipe which extends throughout the alley carrying radioactive waste (notably β^+ emitters) away from the target.

Perhaps the biggest challenge of this location, however, is that it was not designed to be an experimental hall. It is a basement hallway whose initial purpose was for fire safety and radiation monitoring, so we must abide by strict codes to ensure we do not interfere with the general safety of our colleagues at the SNS. We must not interfere with the 3-ft walkway in Neutrino Alley, essentially limiting the overall size of our detectors (and any shielding surrounding them) to fit (depth-wise) on a standard 3-ft industrial pallet. The COHERENT detectors are generally not the enormous detectors one might envision when discussing neutrino experiments, with CsI even earning notoriety as “the world’s smallest neutrino detector” [92].

Two illustrations of Neutrino Alley relative to the target are shown in Figure 3.2; the left is a snapshot of collaboration subsystems from 2018, and the right is a projection of what our deployments might look like in the near future. In both cases, there is a mix of detectors dedicated to CEvNS detection, and detectors with the purpose of improving our understanding of systematics. In the next section, I first overview the detectors designed for studying CEvNS, then discuss the subsystems which support these efforts.

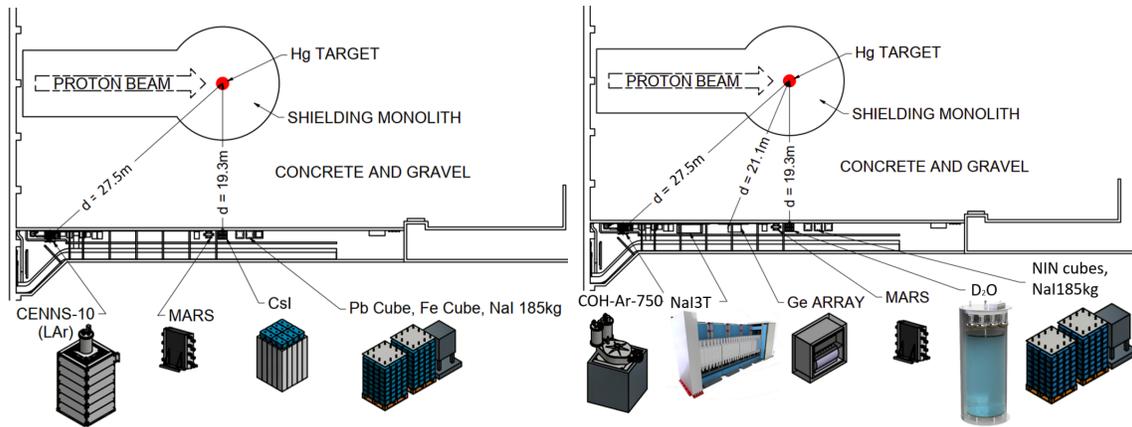


Figure 3.2: Deployment locations for COHERENT detectors in 2018 (left) and in the near-term future (right).

3.2 Subsystem Overview

The initial goal of the COHERENT experiment was to detect CEvNS for the first time, and our multi-faceted approach to the measurement has lent itself well to growing a diverse program of low-energy neutrino cross-section measurements. The cross-sections of particular interest are plotted in Figure 3.3, where curves in blue are CEvNS cross sections, curves in green are charged-current interactions of interest to the neutrino community, and curves in black are specific cross sections that are a background for COHERENT and many other neutrino experiments which use similar shielding materials. Since our CEvNS detectors are sensitive to small nuclear recoils, we also can search for nuclear recoils from accelerator-produced dark matter, for which CEvNS becomes an important background to understand.

As we discussed in Chapter 2, the Standard Model prediction for the CEvNS cross section predicts an N^2 dependence, where N is the number of neutrons in the target nucleus. At present, the COHERENT observations are consistent with this dependence, with measurements on a small nucleus (Ar) and a heavy nucleus (Cs/I), and we have plans to make CEvNS measurements on Na and Ge as well. Figure 3.4 shows the SNS flux-averaged cross section prediction and the COHERENT measurements across the range of target nuclei we deploy. We'll spend some time here discussing the specific technologies we employ for each of these target nuclei, in a roughly chronological order by deployment date. More details for each subsystem can be found in [81, 91].

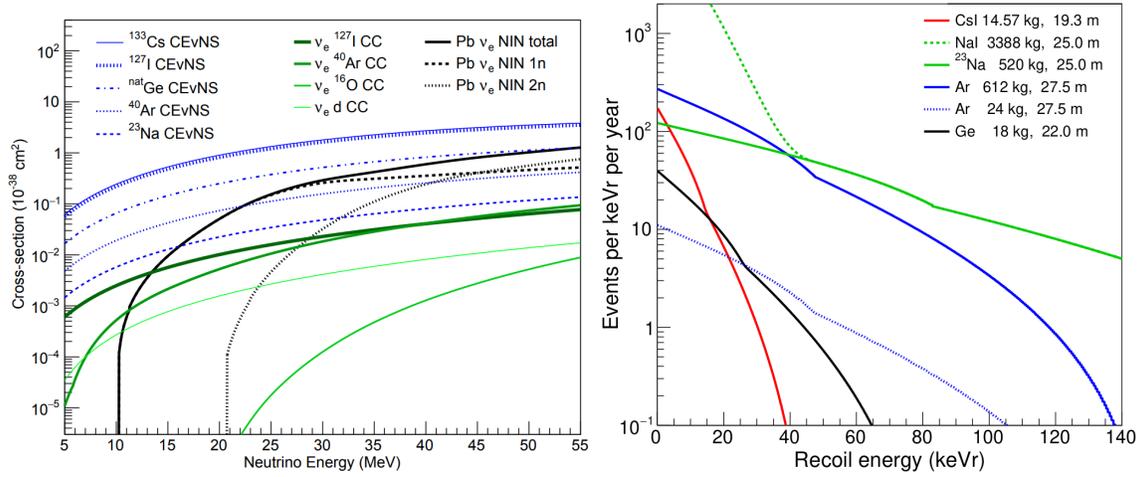


Figure 3.3: Left: Energy-dependent cross sections for many processes of interest to the COHERENT collaboration. CEvNS cross-sections on varied target nuclei are shown in blue, general charged-current cross sections are shown in green, and neutrino-induced-neutron cross sections on lead are shown in black. Reproduced from [81]. Right: Expected nuclear recoil energies from CEvNS interactions in COHERENT subsystems, updated from [91]. Note that this does not include any detector effects; this is the predicted recoil energy spectrum for a given detector, not the predicted observation of the nuclear recoils within that detector.

3.2.1 CEvNS Detectors

3.2.1.1 CsI

Cesium and iodine are nuclei of similar size and neutron number, and a sodium-doped crystal (CsI[Na]) produces an observable light yield with acceptable afterglow from the scintillation associated with small nuclear recoils. A 14.6-kg CsI[Na] detector with a single photomultiplier tube (PMT), in combination with substantial shielding to protect from backgrounds in the alley, was deployed in 2013. The first measurement of CEvNS was reported in 2017 at 6.7σ significance over the null hypothesis [85]. A final result including all data along with an improved measurement of the CsI quenching factor [94] was reported in late 2021, rejecting a no-CEvNS prediction at 11.6σ [54].

The CsI detector was decommissioned in summer 2019, having reached the desired statistical precision for CEvNS study. Analyses with flavor-dependent cross section measurements and searches for accelerator-produced dark matter are also possible [95–97]. Future efforts to improve the precision of the CEvNS measurement on CsI will delve into

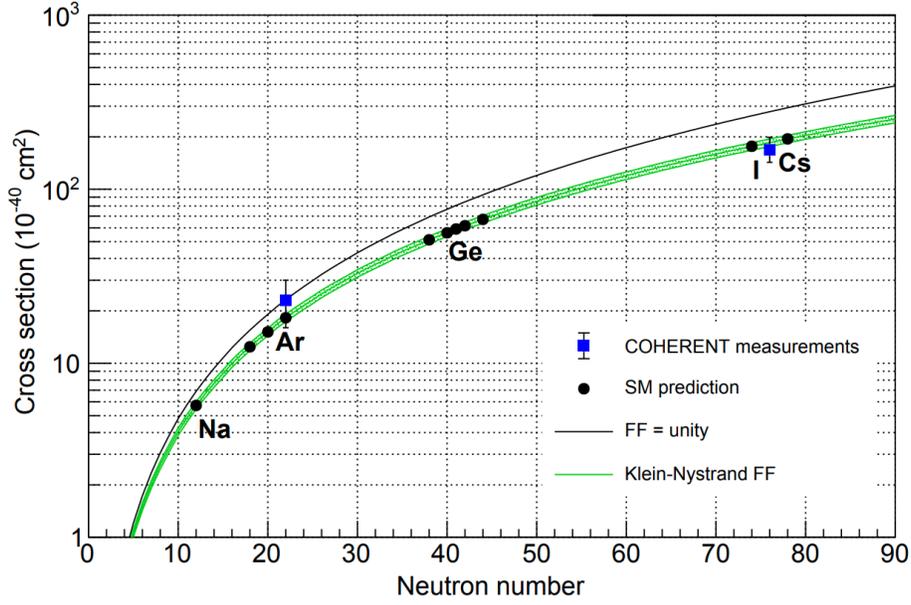


Figure 3.4: The Standard Model prediction for the neutron-dependence of the CEvNS cross section, flux-averaged over the SNS neutrino energy distributions and implementing the Klein-Nystrand form-factor (FF) [93] and associated uncertainties. Reproduced from [81].

the realm of undoped cryogenic scintillators using silicon-photomultipliers (SiPMs) instead of PMTs to remove the limiting background of Cherenkov radiation within the quartz window of a PMT [98].

3.2.1.2 LAr

Liquid argon is one of the most widely used detector technologies in neutrino physics. Its pulse-shape discrimination abilities, along with multi-phased approaches to create time-projection chambers, make it a useful target nucleus for a variety of experiments. The DUNE collaboration, in particular, is set to build a 68-kiloton LAr far detector for neutrino oscillation studies to determine the neutrino mass hierarchy and the CP-violating phase δ_{CP} [99]. Given the enormous volume of the detector, DUNE will also search for supernova neutrinos with energies in the tens of MeV, and CEvNS will be a viable signature for detection given the high cross section. Low-energy charged-current interactions on Ar are not well known, and the deployed LAr system at the SNS provides an opportunity to go after these interesting processes as well.

A 10-kg detector with two PMTs and a wavelength-shifting coating is currently deployed [86], with 750-kg and 10-ton installments planned for the future. A LAr-TPC is

also projected for deployment at the Second Target Station in the next decade [81].

3.2.1.3 Ge

Most of the sensitive detector technologies COHERENT use to detect scintillation light from the recoiling nucleus from a CEvNS interaction, but a semiconductor is also an effective strategy. High purity P-Type Point Contact (PPC) germanium detectors, in particular, have an excellent energy resolution and a low threshold, and this technology will further reduce beam-related backgrounds. Germanium is a nucleus of moderate size and will further test the N^2 dependence in the nucleon range between Ar and CsI. Furthermore, this array will enable research-and-development (R&D) for a potential larger array to pursue new physics [81].

The collaboration is in the process of modularly deploying 2.2-kg detectors in a shielded cryostat with room for 8 detectors (total estimated 17.3 kg active mass) as they become available from the manufacturer.

3.2.1.4 NaI

As the lightest nucleus COHERENT considers, Na will have both the largest CEvNS recoil signature and the lowest cross section. In preparation for a multi-ton array of NaI(Tl) crystals (NaI ν E ν Te), a 185-kg prototype detector (NaI ν E) was deployed to perform R&D work and begin testing with dual-gain PMT bases to simultaneously search for CEvNS and CC-interactions on ^{127}I .

However, the collaboration's interest in Na is not only due to its small nuclear size; while most of the nuclei we consider are even-even, and thus are not affected by the axial terms in the cross-section calculation, Na has an unpaired proton, and the CEvNS cross section may be sensitive to those axial contributions (see Eqn. 2.3). This provides a unique approach to studying the axial coupling g_A , which has high importance for $0\nu\beta\beta$ experiments [100].

At present, the deployment of a multi-ton array (NaI ν E ν Te) is underway in Neutrino Alley with dual-gain bases to search for both the low-energy signature of CEvNS *and* the high-energy signature of charged-current interactions on ^{127}I : $\nu_e + ^{127}\text{I} \rightarrow e^- + ^{127}\text{Xe}^*$. The installation of this array is ongoing; a total of seven modules will be deployed, providing a total active mass of 3388 kg.

COH-CsI-1 Systematics Budget		COH-Ar-10 Systematics Budget		
Contribution	Rate Uncertainty	Contribution	Analysis 1	Analysis 2
Steady-state background	3.0%	Energy calibration	0.8%	4.6%
Beam-related neutron	0.9%	Pulse shape calibration	7.8%	3.3%
Neutrino-induced neutron	0.5%	Beam-related neutron	6.7%	10.7%
Form factor	3.4% (th.), 0.6% (exp.)	Form factor	2%	2%
Neutrino flux	10%	Neutrino flux	10%	10%
Quenching factor	3.8% (25% in 2017)	Quenching factor	1.0%	1.0%
Efficiency	4%	Efficiency	6.4%	7.2%

Table 3.1: Left: A table of the systematic uncertainties in the final CsI results [54]. Right: A table of the systematic uncertainties in the LAr results, shown for both of the independent analysis groups [86].

3.2.2 Supporting Efforts

For the CEvNS detectors with reported measurements, CsI and LAr, we show the reported error budgets in Table 3.1. The dominant uncertainty is the neutrino flux, and the neutron background uncertainties in LAr also contribute heavily to the total systematic uncertainty on the measurement. Investigations into these uncertainties are supported by dedicated subsystems deployed to Neutrino Alley. We overview the active systems here, and will discuss some additional historical efforts in Chapter 6.

3.2.2.1 MARS

Beam-related neutrons (BRN) are a background for each CEvNS detector, and unlike the neutrino flux, the observed BRN flux is notably different throughout the alley due to varied neutron paths from the target to Neutrino Alley. The Multiplicity and Recoil Spectrometer (MARS) is a mobile neutron spectrometer used to monitor these neutron rates at multiple detector locations in Neutrino Alley *in situ*. We provide more details about this detector in Chapter 6, detail our characterization and calibration in Chapter 7, and overview our rate monitoring efforts in Chapter 8.

A small NaI crystal is affixed unshielded to the top of the MARS detector with a single PMT. By observing the spectral response of this system over long periods of time, we can independently monitor the ambient backgrounds (e.g., a nearby hot-off-gas pipe). This detector supports the MARS measurements, as MARS cannot be shielded from the γ s without affecting the neutron flux we aim to monitor.

3.2.2.2 Timing Cart

The “Timing Cart” is a small and extremely portable cart that houses four liquid-scintillator (LS) cells used to identify the arrival times of BRNs throughout Neutrino Alley within roughly 100 ns. While MARS is a large-volume spectrometer, the timing cart can quickly estimate differences in neutron rates as we develop a model of the paths neutrons take to reach Neutrino Alley. We’ll discuss some of the results from these detectors alongside the MARS analyses in Chapter 8.

3.2.2.3 NUBEs

The “Neutrino Cubes”, or NUBEs, are detectors designed to measure the cross section for an incident neutrino to produce a neutron. The neutrino-induced-neutron (NIN) cross sections in common shielding materials such as lead or iron have never been measured, and these detectors serve both to inform an estimate on the expected NIN rates for the well-shielded CEvNS detectors and to measure the cross section of this process for the larger neutrino community. In particular, the Helium and Lead Observatory (HALO) uses NINs from Pb as its primary detection mechanism to search for ν_e from supernovae [101].

The collaboration deployed two NUBEs in Neutrino Alley: a Pb detector in 2016 and a Fe detector in 2017. Both systems feature the designated material surrounding six liquid scintillator cells, each with its own PMT, to measure the neutron rate. Both detectors are shielded with water bricks to mitigate the prompt BRN background in Neutrino Alley. The production of NINs can occur in both charged-current ($\nu_e + A_N^Z \rightarrow e^- + A_{N-1}^{*Z+1}$) and neutral-current ($\nu_X + A_N^Z \rightarrow \nu_X + A_{N-1}^{*Z}$) processes, where A represents an arbitrary nucleus, Z is the number of protons in the nucleus, N is the number of neutrons, and the asterisk indicates a nucleus that will emit neutrons (and γ s) that can be detected with the LS cells [91].

Very recently, this effort has concluded. The NUBEs were decommissioned in May 2022, having completed their mission by measuring σ_{NIN} for Pb, uncovering a 4σ disagreement with current predictions using the Monte Carlo event generator MARLEY [102–104], and finding that the additional deployment time needed to improve the result was incompatible with the physics goals of other subsystems in Neutrino Alley. The analysis for the Fe NUBE is still to be completed.

3.2.2.4 D₂O

For all our CEνNS detectors, and most notably for the public results shown in Table 3.1, the largest systematic is our knowledge of the neutrino flux, because there is no data on pion production from proton-Hg interactions at 1 GeV. My work, described in Chapter 4, is an investigation into this systematic and a discussion of our understanding of the neutrino flux from the SNS through simulation, but without new data we cannot reduce the uncertainty. The D₂O detector will directly measure the flux of ν_e in Neutrino Alley and reduce this uncertainty to the few-percent level. A detailed description of the design efforts which contributed to the recent deployment of a commissioning geometry can be found in Chapter 5.

Chapter 4

Simulating Particle Production at the Spallation Neutron Source

This chapter pulls largely from our soon-to-be published work [105], with supplemental details and plots provided to support more recent findings. A group from the University of Florida developed the baseline simulation of the SNS, and I adapted their work to provide information specific to COHERENT and complete all included analyses for the neutrino production at the SNS. This involved significant collaboration input and conversation, and the advice and efforts of Yuri Efremenko, Diane Markoff, Diana Parno, and Daniel Salvat in particular were crucial to the development of the simulation. Aria Salyapongse contributed extensively to the simulation of thin-target geometries for the validation of our physics model in Section 4.2, and Shuaixiang Zhang built the geometries for the proton beam window in Section 4.4 and the Second Target Station geometry in Section 4.5. The simulation code we developed for this chapter is publicly available at [106].

4.1 Introduction

Precise knowledge of the SNS neutrino flux is essential to unlocking the full physics potential of the COHERENT cross-section measurements. The uncertainty on the overall normalization of the neutrino flux is the dominant systematic uncertainty in the Ar results [107] and the second-largest systematic in the initial CsI results [85]. Thanks to updated measurements of the quenching factor in CsI [94], the neutrino flux is the dominant systematic in the final CsI results [108].

In this chapter, we describe a detailed model of the SNS using the Geant4 Monte Carlo

framework [109, 110] to characterize the neutrino flux for the COHERENT detectors. In addition to the geometry, the simulation accuracy relies on the underlying implementation of pion production in the Geant4 physics model. Section 4.2 describes our validation efforts using four standard physics lists against the available world π^+ -production data. World data, however, are imperfect — Hg-target data are not available at low proton energies, data sets at proton energies near 1 GeV are very limited, and most pion-production cross sections are measured using thin targets that do not replicate the half-meter of dense material the protons at the SNS encounter¹. Although the existing data are insufficient for a precise validation, we estimate the uncertainty of our simulated flux with the QGSP_BERT physics list to be about 10%. Section 4.3 describes our simulation of the SNS, along with our tools for studying the characteristics of the resulting neutrinos. We also discuss the effect of changes to SNS operating conditions; for example, the incident proton kinetic energy has ranged from 0.83 – 1.011 GeV during COHERENT’s lifetime in Neutrino Alley. Section 4.4 summarizes the properties of our simulated neutrino flux using the selected physics list.

Our simulation has applications to additional nuclear and particle physics experiments proposed at the SNS. In Section 4.5, we present a neutrino-flux simulation based on preliminary design work for a proposed Second Target Station (STS) with a tungsten target; our results suggest that the STS could be a very productive site for next-generation neutrino experiments. Section 4.6 describes the use of our simulation to study π^0 and π^- production at the SNS, relevant to accelerator-based searches for light dark matter. We then discuss several future avenues for reducing uncertainties related to the SNS neutrino flux in Section 4.7.

4.2 Validation of Simulation Physics

We investigated four standard physics models (or “physics lists”) as implemented in a standard installation of Geant4.10.06.p01: FTFP_BERT, QGSP_BERT, QGSP_BIC, and QGSP_INCLXX. With all SNS protons well below 10 GeV, the differences in the underlying string models of FTFP_BERT and QGSP_BERT were found to be negligible; in this work we focus only on QGSP_BERT. We note here that the plots within this section use natural units, such that $c = 1$.

Each candidate for our physics list models nuclear structure in a specific way. With an

¹There are significant proton-energy losses within the thick target at the SNS that are not accounted for by thin-target data — see Figure 4.20.

implementation of the classical Bertini Cascade model [111] for incident hadrons below 3 GeV, QGSP_BERT is a favored model for the production of hadrons (and subsequently, neutrinos) with its treatment of the nucleus as a Fermi gas of nucleons that can be solved on average using the Boltzmann equation for a projectile moving through the gas [112]. The QGSP_BIC physics list differs only for protons and neutrons, for which it implements a Binary Cascade and models the nucleus as an isotropic sphere. In this model, the nucleons are placed at specific positions so that projectiles can interact with them individually, and each nucleon carries a random momentum, uniformly sampled between zero and the Fermi momentum [113]. Finally, QGSP_INCLXX extends the Liège Intranuclear Cascade model [114] benchmarked against spallation studies below 3 GeV [115] by modeling the nucleus in a very similar manner to QGSP_BIC, but adding the possibility to emit nucleon clusters that can cause secondary reactions after a projectile interacts with the nucleus. Both QGSP_BIC and QGSP_INCLXX require increased computation time (compared to QGSP_BERT) to model the interactions of projectiles with more massive nuclei [116].

In prior estimations the COHERENT collaboration has used the QGSP_BERT physics list with an assigned 10% uncertainty on any flux predictions coming from simulation efforts. This estimate was informed by prior studies using an implementation of the Bertini model in the LAHET Monte Carlo framework [117] to make predictions for the LSND and KARMEN experiments [118–120]. World data at the time of their investigation did not agree with LAHET predictions, and LAHET predictions that were renormalized to match available data were lower than Geant4 predictions [118, 121, 122]. The 10% systematic was assigned to our neutrino-flux calculations to conservatively account for this discrepancy [85].

Since the lack of pion-production data from 1 GeV proton-mercury interactions prevents a direct comparison, our choice of physics model must be validated via other targets, usually at higher energies. In Section 4.2.1 we compare the total π^+ -production cross section to the Norbury-Townsend parameterization developed to match data from proton-nucleus and nucleus-nucleus collisions [123]. Not included in the development of the Norbury-Townsend parameterization, however, are newer results focusing on double-differential measurements, such as those from the thin-target HARP experiment [124]. We detail our validations against the HARP measurements in Section 4.2.2. Older experiments also collected double-differential pion-production data at energies closer to the SNS, such as Abaev *et al.* in 1989 [125], but their data have a very limited angular coverage. We use these data sets to check the model behavior at lower proton energies (Section 4.2.3) since they cannot constrain our total neutrino flux. We discuss the effects of modeling the thick

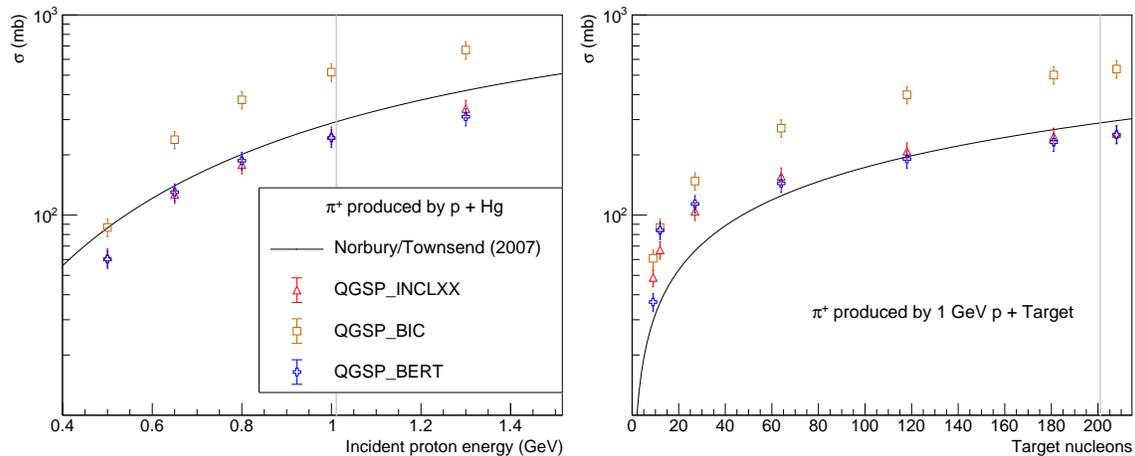


Figure 4.1: Comparisons of the Norbury-Townsend parameterization and Geant4 model predictions of total pion-production cross section. Left: Dependence of total cross section on incident proton energy for a mercury target. The vertical line indicates the current SNS operating energy of 1.011 GeV, but COHERENT still sees π^+ production at energies below this value due to proton energy loss within the thick target (see Fig. 4.20). Right: Dependence of total cross section on target nucleus for a proton energy of 1 GeV. The vertical line represents a mercury target.

target of the SNS in Section 4.2.4 and interpret all of our validation work to estimate a neutrino flux systematic for COHERENT in Section 4.2.5.

4.2.1 Norbury-Townsend Parameterization

The Norbury-Townsend parameterization is an empirical function developed to parameterize pion-production data from proton-nucleus and nucleus-nucleus interactions measured by Nagamiya *et al.* [126]. While developed in the right energy range for SNS operations at ~ 1 GeV, π^+ production data were only taken for subsets of Ne + NaF, Ne + Cu, Ne + Pb, C + C, C + Pb, Ar + KCl, Ar + Pb for 0.4, 0.8, and 2.1 GeV per incident nucleon — only π^- production data were available from the proton-nucleus studies [123]. Although our focus is π^+ production in this work, we note that future effort to check the candidate physics models against the parameterizations for π^- and π^0 production will be useful to validate the flux predictions for dark-matter-producing particles at the SNS that we present in Section 4.6.

The Norbury-Townsend parameterization of the π^+ production cross section (σ_{π^+} in mb) from incident protons is shown in Equation 4.1, where A_t is the number of target

nucleons, and E_i (in GeV) is the energy per incident nucleon:

$$\sigma_{\pi^+} = \frac{A_t^{2.2/3}}{0.00717 + 0.0652 \frac{\log(E_i)}{E_i} + \frac{0.162}{E_i^2}}. \quad (4.1)$$

This parameterization overpredicts the π^+ -production from the Ar + Pb and Ne + Pb data used in the development by between 15 and 30%, though we do note that this demonstrates a 1σ consistency due to the large errors of the datasets from the early 1980s.

Using a thin simulated target to minimize proton energy loss ($5 \times 5 \times 0.5 \text{ cm}^3$) and specifying the isotope, molar mass, and density for each target, we counted the total number of pions produced. We then scaled this event rate by our simulated number of target nuclei and incident flux of protons to convert to a total cross-section prediction. Figure 4.1 shows comparisons of these results to the parameterization across incident energies (top) and target nucleus (bottom), with a 10% uncertainty applied to the cross sections from each potential physics list.

The left panel of Figure 4.1 demonstrates that for p + Hg at 1 GeV, QGSP_BERT and QGSP_INCLXX agree with the parameterization at a $\sim 15\%$ level, and generally have better agreement at lower energies ($\sim 10\%$ at 0.8 GeV) than at higher energies ($\sim 25\%$ at 1.3 GeV). The QGSP_BIC model, however, consistently overpredicts the p + Hg parameterization by more than 50% from 0.8 - 1.3 GeV. The right panel of Figure 4.1 illustrates that for 1 GeV p + A, QGSP_BERT and QGSP_INCLXX agree with the parameterization at the $\sim 15\%$ level for nuclear targets near Hg. Due to the high uncertainties in the parameterization and its underlying datasets, we use these results to demonstrate a reasonable consistency at an $\sim 15\%$ level for the predictions of both QGSP_BERT and QGSP_INCLXX, and to identify the overall normalization problem of QGSP_BIC.

4.2.2 HARP and HARP-CDP

HARP, the Hadron Production Experiment (PS214), operated at CERN's Proton Synchrotron from 2000 to 2002. With a nearly 4π acceptance and incident proton momentum range from 1.5 GeV/c to 15 GeV/c, HARP measured 7 different solid targets (Be, C, Al, Cu, Sn, Ta, Pb) as well as 4 cryogenic liquid targets (H_2 , D_2 , N_2 , O_2). The HARP collaboration disagreed on their TPC calibrations causing a subgroup, HARP-CDP, to promote different calculations of pion momenta and identification of protons and pions [127]. Both of these differences impact the final analysis, such that HARP-CDP reports 16-34% lower integrated cross sections than the HARP analysis, depending on the target. In this paper, both sets of cross-section results were checked against our Monte Carlo simulations.

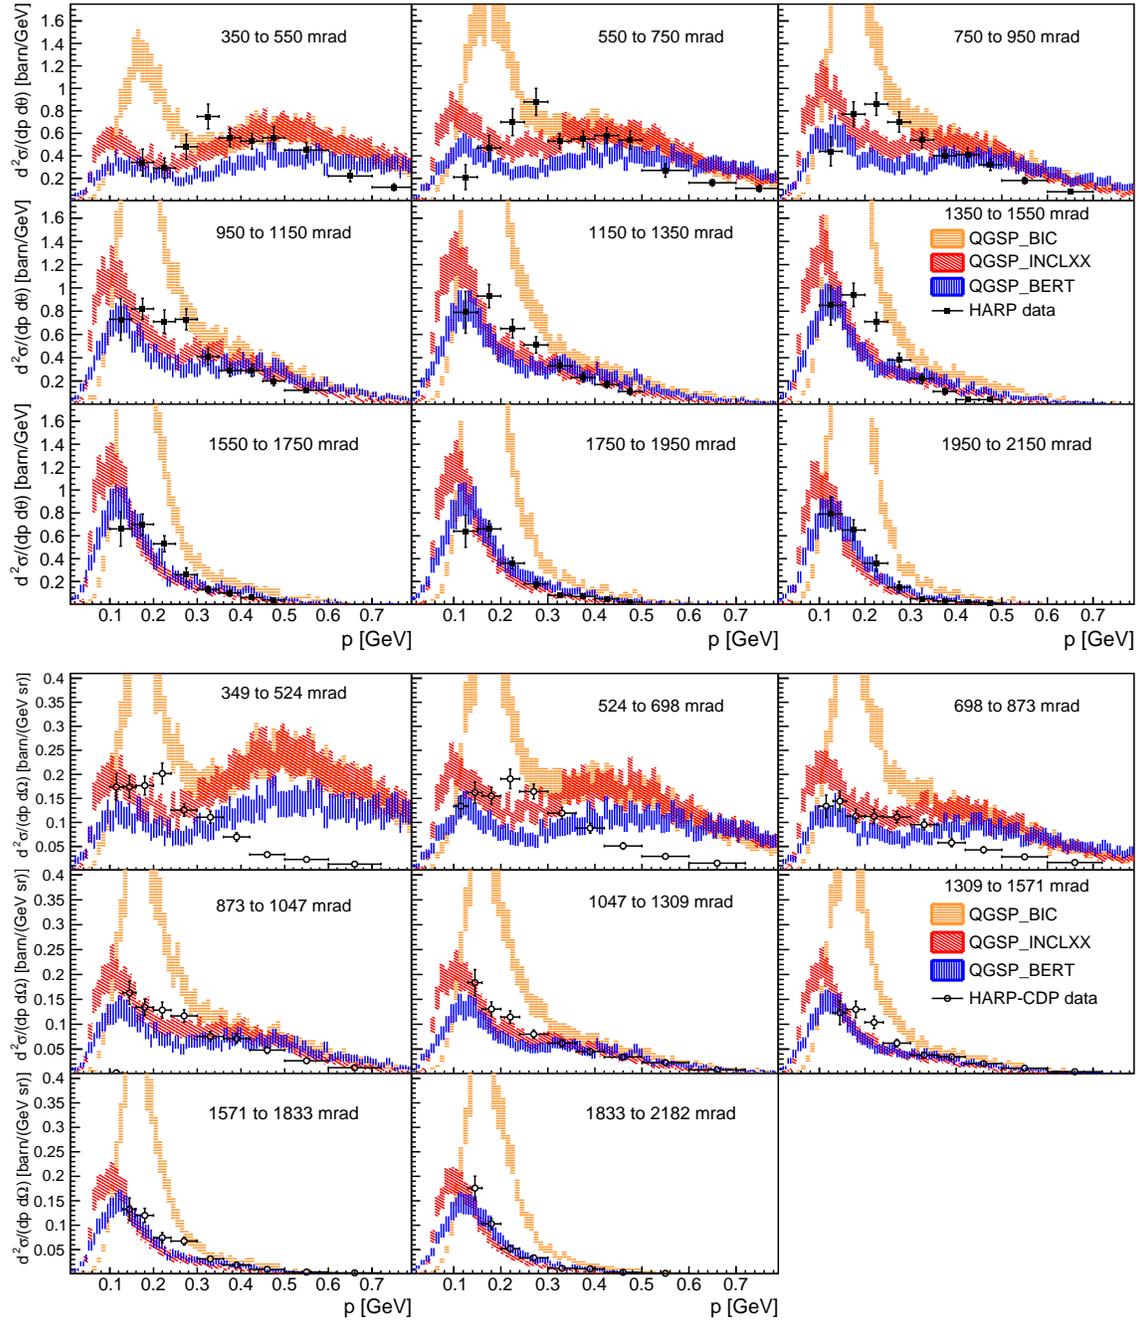


Figure 4.2: Comparisons of double-differential cross sections of π^+ production from 3 GeV/c $p+^{208}\text{Pb}$ as predicted by the different Geant4 physics lists to the measurements from HARP (top) and HARP-CDP (bottom). The error band shown for each physics list include a 10% systematic uncertainty and the statistical uncertainties from simulation. We highlight ^{208}Pb because, among HARP targets, this isotope is closest in mass to the SNS mercury target.

Data were not collected for incident protons at 1 GeV; therefore we compare to the HARP and HARP-CDP analyses of 3 GeV/c data on a large range of nucleon numbers: Be [128], C [129, 130], Al [128, 131], Cu [129], Sn [129, 132], Ta [133, 134], and Pb [128, 135]. We follow a similar procedure to our Norbury-Townsend comparisons and simulate monoenergetic protons with 2.205 GeV of kinetic energy (calculated from the 3 GeV/c beam momentum) incident on a thin target ($5 \times 5 \times 0.5 \text{ cm}^3$), though here counting pions produced per pion momentum and production angle rather than total number of pions. We then scale the stored event rates by our simulated target details to convert to a doubly differential cross-section prediction from each simulation model. Figure 4.2 illustrates the direct comparison of our simulation to the HARP and HARP-CDP results for 3 GeV/c $p + {}^{208}\text{Pb}$. The simulation error bands combine statistical uncertainty with the estimated 10% systematic uncertainty on the simulation prediction. The HARP and HARP-CDP data consistently disagree with the simulation predictions by more than 10% in the low-angle bins, but generally start to demonstrate consistency with QGSP_BERT and QGSP_INCLXX in the higher-angle bins. We also note that the simulation generally better-predicts the data in the higher-momentum bins, and the normalization problem of QGSP_BIC is evident in the low-momentum bins.

Since the Geant4 models predict the HARP and HARP-CDP data better in some bins than others and the SNS neutrino flux is known to arise from π^+ decay at rest (see Table 4.2), we integrate away the angular or momentum dependence to compare singly differential cross sections. Since there is The comparisons shown in the left panel of Fig. 4.3 integrate our simulation prediction over the angular region of the HARP analysis; there is less than a 1% difference from the simulation prediction integrated over the HARP angular region and over the HARP-CDP angular region (see Figure 4.4), so we show HARP-CDP data on the same axes.

The SNS mercury target is thick and dense enough to stop the majority of the pions regardless of production angle or momentum. Therefore, to a very good approximation, we require only the total cross section to simulate the neutrino production. We integrate away the dependence on both production angle (350 - 2150 mrad), again ignoring any difference between HARP and HARP-CDP angular regions, and momentum (0.1 - 0.8 GeV) and show the total cross-section comparisons to both HARP and HARP-CDP data in Fig. 4.5. The ratio of the Geant4 model prediction to the HARP or HARP-CDP result determines how well we predict the data and is therefore considered in this work as the most convincing estimate of our systematic uncertainty (more in Section 4.2.5).

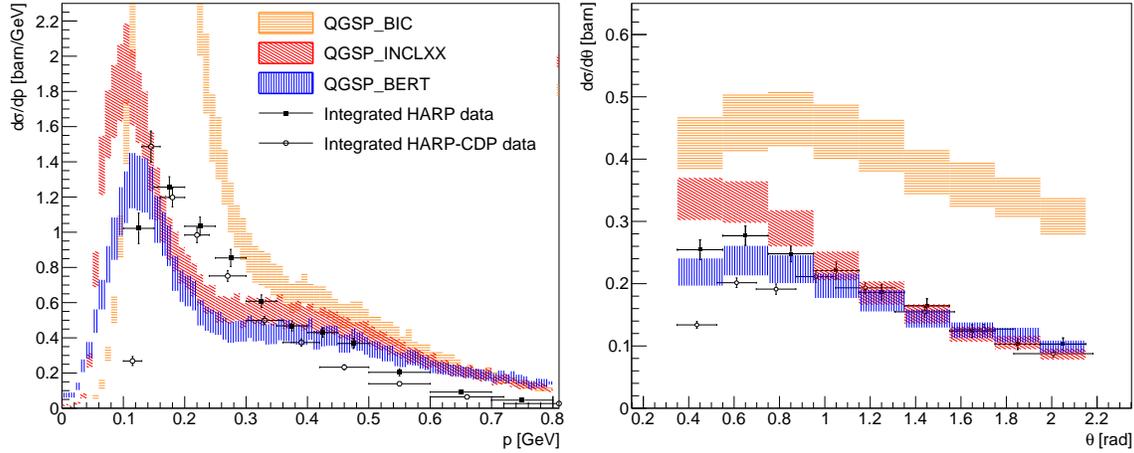


Figure 4.3: Comparison of measured differential cross sections of π^+ production from 3 GeV/c $p+^{208}\text{Pb}$ to Geant4 physics lists. Left: HARP and HARP-CDP data were integrated over their respective angular regions and compared to simulation integrated from 350 to 2150 mrad in production angle (see Figure 4.4). Right: HARP and HARP-CDP data were integrated from 0.1 to 0.8 GeV/c in momentum and compared to simulation integrated on the same region.

4.2.3 Low-energy Pion-production Data

Using the proton synchrotron at the Leningrad Nuclear Physics Institute (Gatchina, Russia) with a beam kinetic energy of 997 ± 5 MeV, Abaev *et al.* measured pion production on 16 different targets (isotopes of H, B, C, O, Mg, Al, Cu, Sn, Ta, W, and Pb) at 0° and 57.8° with 0.01 steradian angular acceptance [125]. We compare to a range of nucleon numbers, but we exclude comparisons to different isotopes of the same nucleus in this work as no significant difference between different isotopes was found in the cross sections from data or simulation. The double-differential comparisons of the Geant4 models to Abaev *et al.* are shown in Fig. 4.6, and the momentum-integrated comparisons are shown in Fig. 4.7.

The right panel of Fig. 4.6 at 57.8° (~ 1009 mrad) is similar to the center-left panels of the HARP (950 - 1150 mrad) and HARP-CDP (873 - 1047 mrad) results shown in Fig. 4.2, albeit with different incident proton energy and angular acceptance. We observe a similar disagreement with the double-differential comparisons, but a slightly better agreement with the momentum-integrated comparison in Fig. 4.7. The points at 208 nucleons in the momentum-integrated comparisons of Fig. 4.7 are similar to the behavior of a single bin we showed in the right panel of Fig. 4.3, albeit with larger angular acceptance. We observe

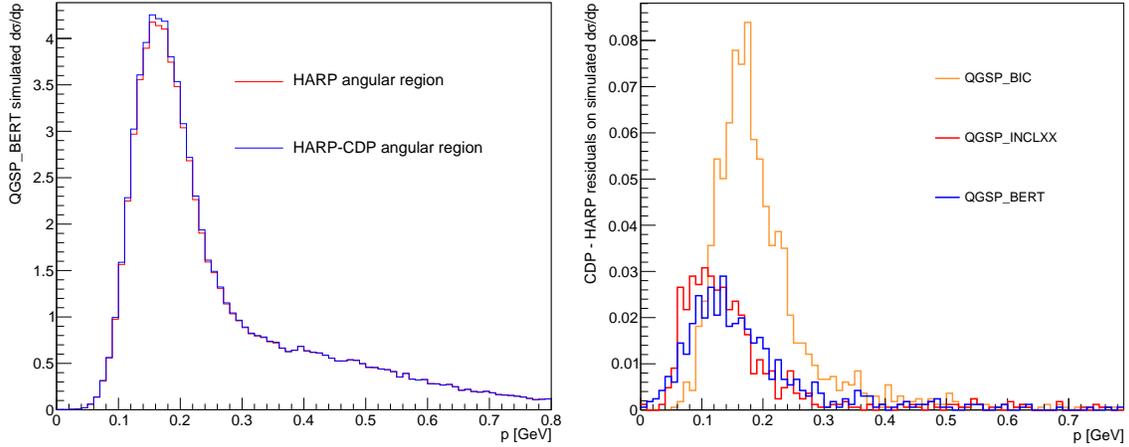


Figure 4.4: Left: Simulated π^+ production cross section using QGSP_BERT, integrated over the different angular regions of HARP and HARP-CDP. Right: Residual distribution of the simulated cross sections integrated over HARP and HARP-CDP angular regions for all Geant4 models investigated in this work. We demonstrate that the impact of this difference in angular regions between the analyses is of order 1% (note that the right panel show the direct residuals, not the relative residuals), and neglect this difference going forward.

that single angular bins do disagree with the HARP and HARP-CDP data by greater than 10%, which suggests that we cannot draw conclusions about the uncertainty on the total pion production from the limited angle coverage of the Abaev data. Focusing specifically on the comparisons to ^{208}Pb and noting that the 0° prediction of QGSP_BERT is consistent with the Abaev data at the 10% level, we infer that the 10% uncertainty consistent with HARP predictions is a reasonable estimate at lower proton energy.

4.2.4 Secondary Particle Interactions

The pion-production model of QGSP_BERT has the best agreement for thin-target data, but we must also model the proton energy loss and the interactions of any secondary particles that are produced. For example, π^+ scattering will affect our predictions on how many π^+ decay at rest, and pion interactions such as charge exchange or absorption will impact the number of π^+ that decay into neutrinos. We do not perform any specific validation of these processes in this work, but we offer some comments here about the models we implement.

Proton energy-loss models are important in Geant4, as they have broad applications

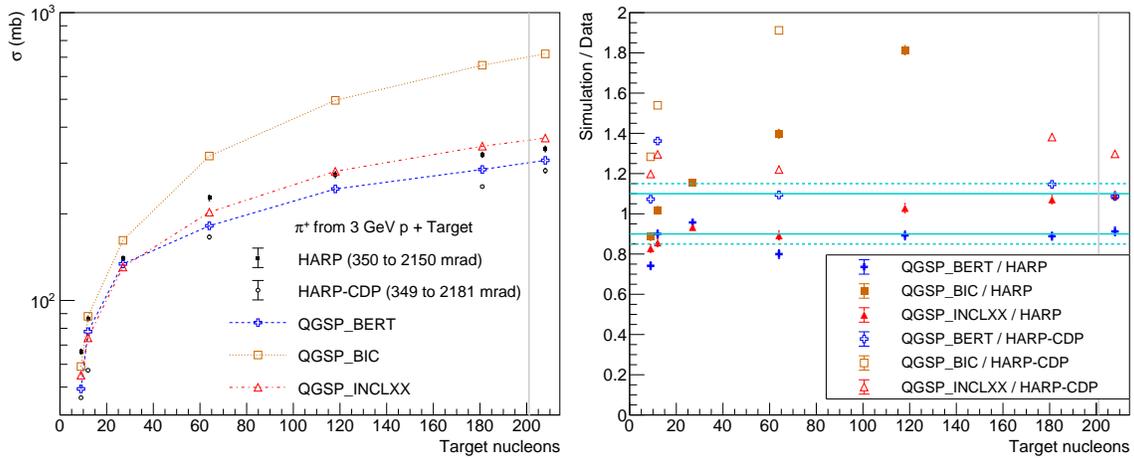


Figure 4.5: Left: The HARP data and Geant4 model predictions of the pion-production cross section integrated over 350 - 2150 mrad and 0.1 - 0.8 GeV. The HARP-CDP data are also shown but are integrated over 349 - 2181 mrad (determined by the data binning) and 0.1 - 0.8 GeV. Right: Ratio of the Geant4 simulated predictions to the central values of the data, plotted with an uncertainty on all three simulations shown as data error / central value (the HARP-CDP error bars are small enough to be hidden by the points themselves). The horizontal solid cyan lines mark a $\pm 10\%$ 1σ uncertainty band, and the dashed cyan lines mark a $\pm 15\%$ 1σ uncertainty band. The vertical gray line on each plot represents a mercury target.

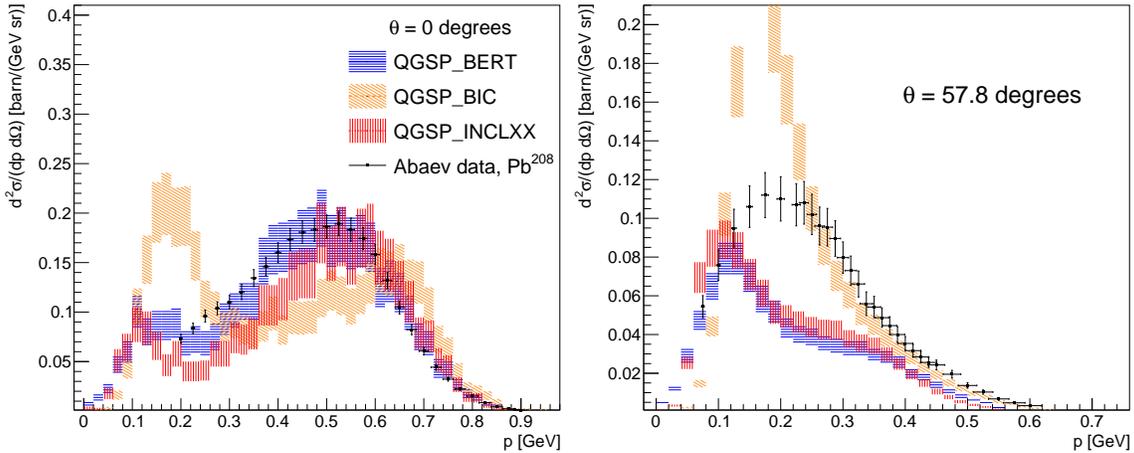


Figure 4.6: Comparisons of double-differential cross sections of π^+ production from 1 GeV $p+^{208}\text{Pb}$ at 0° and 57.8° as predicted by the different Geant4 physics lists to the measurements from Abaev *et al.* [125].

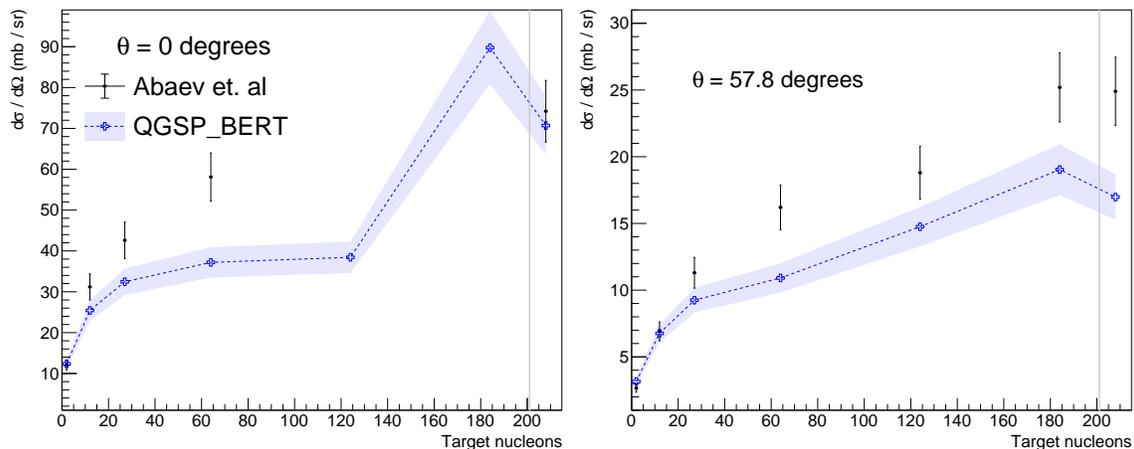


Figure 4.7: Comparisons of differential cross sections at 0° and 57.8° as predicted by QGSP_BERT to measurements by Abaev *et al.* [125]. The vertical gray lines represent a mercury target.

within both the particle physics and medical communities. As a result, these models have undergone extensive validation efforts [136, 137]. The energy-loss profile of protons is generally accurate to within 2% above 1 MeV. We do not perform any further validation of the energy losses within the SNS target, and we neglect any uncertainty from the proton energy-loss profile within the thick SNS target in this work.

Pions and other secondary hadrons created at the SNS are well below 10 GeV and are simulated using the default cross-section tables implemented in the Bertini Cascade model (primarily the Baranshenkov and Glauber-Gribov parameterizations) [111]. In our simulations of the SNS, $\sim 25\%$ of the π^+ tracks that are produced end in non-decay processes (labeled “pi+Inelastic” in Geant4). There is less than a 1% difference between this fraction of non-decay destruction of pion tracks between QGSP_BERT and QGSP_INCLXX, though QGSP_BIC simulations estimate that nearer to 30% of π^+ tracks end in non-decay processes. A small fraction (less than 15%) of these “pi+Inelastic” track deaths can still result in neutrino production due to the production of additional protons or pions. We do not perform any specific validation of these processes in this work, and we neglect the uncertainty here given the difference between the QGSP_BERT and QGSP_INCLXX predictions and the previously identified normalization problem of the QGSP_BIC predictions.

4.2.5 Discussion

We are not aware of any data from p+Hg and very few data sets exist at these energies, so this work is intended as a cross-check of prior estimates rather than as a derivation of our neutrino-flux systematic while new experimental efforts are developed (see Sec. 4.7). We find that QGSP_BERT is the only model which agrees at the 10% level with the cross section measurements of both HARP and HARP-CDP for the ^{208}Pb dataset (the closest to Hg); the other lists overpredict the HARP-CDP results. For lighter nuclei like W (see Sec. 4.5), these comparisons suggest that a 15% uncertainty on the QGSP_BERT prediction would be more appropriate. The comparison against the Norbury-Townsend parameterization shows reasonable consistency at the 10% level with QGSP_BERT and QGSP_INCLXX, and further demonstrates an overall normalization problem with QGSP_BIC, despite noteworthy agreement in the tails of the 57.8° Abaev measurement. While QGSP_INCLXX is acceptable, QGSP_BERT has better agreement with the data and the added bonus of being more computationally efficient since QGSP_INCLXX produces energy conservation errors in some of the SNS volumes. We note that while the momentum-integrated Abaev data may disagree with QGSP_BERT predictions at more than the 10% level, similar disagreement is shown in the right panel of Fig. 4.3 for single points of the momentum-integrated HARP and HARP-CDP comparisons; it is only after an additional integration over angle that good agreement is achieved. Ultimately, the limited angular coverage of the Abaev data limits our ability to investigate this effect.

In light of these studies and prior work using the Bertini cascade for neutrino flux calculations [118–120, 138, 139], we choose to use QGSP_BERT with a 10% uncertainty on the flux predictions that come from our Geant4 simulations. This systematic uncertainty needs to be reduced for the future precision program of COHERENT, and we describe future avenues for experimentally reducing this uncertainty in Section 4.7.

4.3 Modeling the Spallation Neutron Source in Geant4

The design of the SNS target and moderator suite was optimized for neutron production and related science [141]. We define simplified components of the SNS target monolith that are expected to contribute to pion production or to the stopping of pions and muons. The simplification process is demonstrated in the left panel of Fig. 4.8, where the technicalities of the target vessel are reduced to the mercury-containing region shaded red. The right panel of Fig. 4.8 highlights the target monolith and Neutrino Alley to illustrate

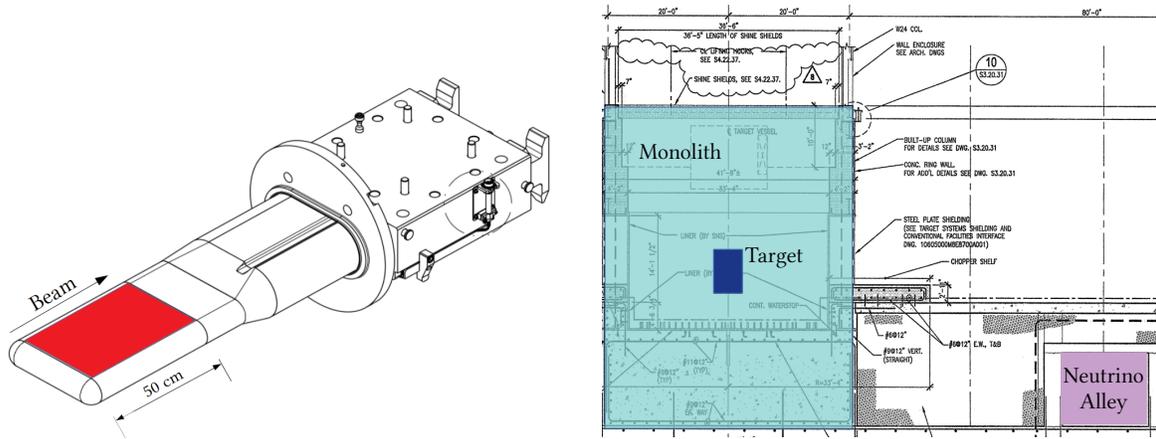


Figure 4.8: Left: ORNL technical drawing of the target vessel. The red section highlights the main Hg target as implemented in our Geant4 model. Right: A portion of an ORNL technical drawing illustrating the target hall, with pieces in our Geant4 model highlighted. The outer shaded cyan is the concrete monolith, with the inner indigo representing the steel containing the Hg target and moderators. In the bottom right corner, the shaded purple shows the location of Neutrino Alley relative to the target monolith [140].

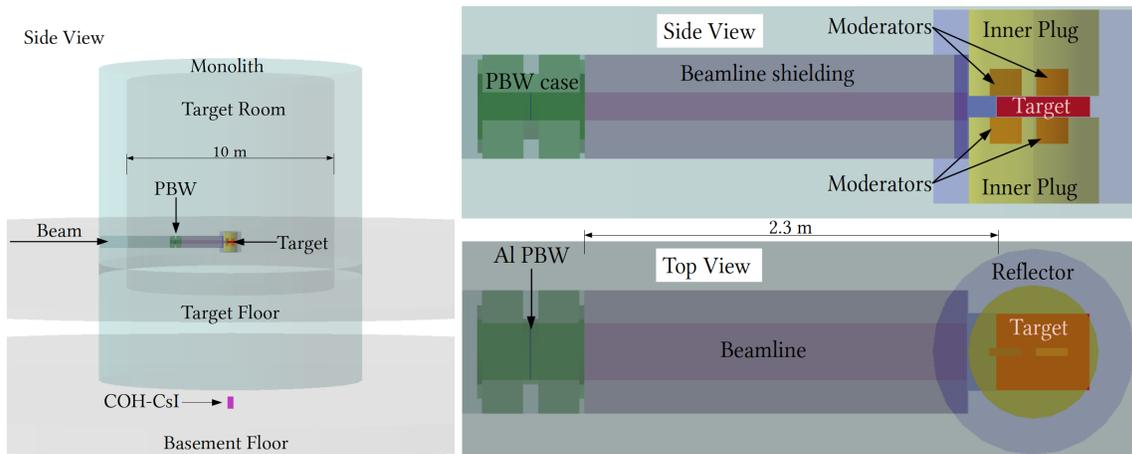


Figure 4.9: Our Geant4 model of the SNS, simplified from ORNL technical drawings. Left: The full Geant4 world, highlighting the monolith relative to the location of CsI in Neutrino Alley. Right: A view inside the outer monolith illustrating the target, neutron moderator suite, proton beam window, and beamline shielding.

the structures we build into our model. The details of our SNS model, along with their relative contributions to the overall π^+ production, are shown in Table 4.1, and the full visualization of our simple model is shown in Fig. 4.9.

Table 4.1: An overview of components in our Geant4 model that contribute to the overall pion-production. We also include the fraction of π^+ our simulations produce within each volume. We report the dimensions from the perspective of the beamline as either Width \times Height \times Depth or Diameter (\varnothing) \times Height. The depth of the Inconel proton beam window (PBW) is an approximation (indicated by an *) of 3 cm, which includes some amount of vacuum immediately before and after the window as a result of the curvature.

Component	Material	Dimensions	π^+ contributed	
			Aluminum PBW	Inconel PBW
Target	Hg	$39.9 \times 10.4 \times 50.0 \text{ cm}^3$	94.12%	(90.91%)
Target Casing	Steel	$40.9 \times 11.4 \times 51.0 \text{ cm}^3$	0.20%	(0.56%)
Inner Plug (2)	Be, D ₂ O	70.0 cm \varnothing , 45 cm	0.19%	(0.23%)
Moderator (4)	H ₂ (3), H ₂ O (1)	$4.0 \times 13.9 \times 17.1 \text{ cm}^3$	0.01%	(0.01%)
Reflector	Steel, D ₂ O	108 cm \varnothing , 101.6 cm	0.99%	(1.34%)
Beamline Shielding	Steel	$64.8 \times 54.6 \times 200.0 \text{ cm}^3$	0.93%	(1.72%)
Target Room	Steel	1002 cm \varnothing , 950.8 cm	0.00%	(0.14%)
Aluminum PBW	Al-6061, H ₂ O	$29.8 \times 14.6 \times 0.02 \text{ cm}^3$	2.77%	(—)
Aluminum Beamline	Air	$29.8 \times 14.6 \times 200.0 \text{ cm}^3$	0.79%	(—)
Inconel PBW	Inconel-718, H ₂ O	$26.7 \times 12.7 \times 3.0 \text{ cm}^{3*}$	—	(4.32%)
Inconel Beamline	Air	$26.7 \times 12.7 \times 200.0 \text{ cm}^3$	—	(0.77%)

Though most of the components we simulate are essentially unchanged during running despite routine maintenance and possible replacements, we must carefully consider the proton beam window (PBW) separating the vacuum of the accelerator from the target. Each proton must pass through the PBW, resulting in both proton energy loss and pion production as a result of interactions in the thin window. The PBW is routinely replaced due to radiation damage, and two different PBW designs have been in use during COHERENT’s live-time in Neutrino Alley. A two-layered film design using Inconel, a nickel-based alloy trademarked by the Special Metals Corporation [142], with water cooling between the films, was used from the initial SNS production runs until January 11, 2017. An aluminum plate design with 50 drilled pipes for water cooling was in place until the latest replacement reverted back to an Inconel PBW on April 7, 2020. Figure 4.10 illustrates both PBW designs as modeled in our Geant4 geometry.

The SNS accelerates protons into an accumulator ring, which ensures that a focused beam of monoenergetic protons is directed on to the target. This beam is magnetically spread to prevent overheating of the proton beam window and target casing [143]; we introduce a uniformly distributed source using a prior measurement of the beam profile

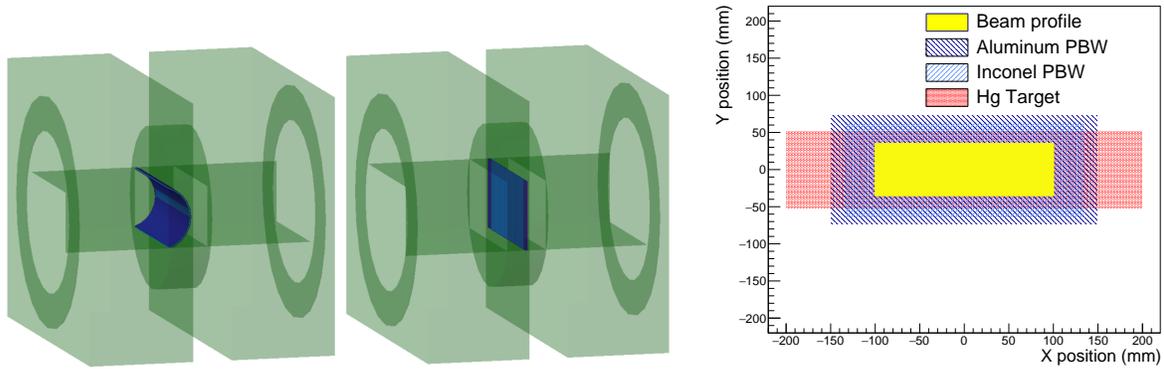


Figure 4.10: Left: Geant4 mockup of the dual-film Inconel PBW, with water cooling between the two films. The proton beam is incident from the left, and travels to the right. Center: Geant4 mockup of the aluminum plate PBW, with 50 vertical pipes for water cooling. Right: The position of incident protons shown relative to the profiles of the different PBW designs and Hg target.

at the target [144] to account for this. Our simulated profile is illustrated in the right panel of Fig. 4.10 to show its size relative to the beam window designs and target.

We specify the particles for the simulation to track, typically ν_x , π^\pm , μ^\pm , K^\pm , η , p , and n , to ensure that we do not truncate any possible neutrino production chain. Using the Monte Carlo framework of Geant4 and the QGSP_BERT physics model chosen in Section 4.2, we observe which particles and interactions are responsible for generating the SNS neutrino flux. These predictions depend on our chosen physics model; for example, the QGSP_BIC nuclear model predicts a small rate of η production for 1 GeV incident protons, while other lists do not.

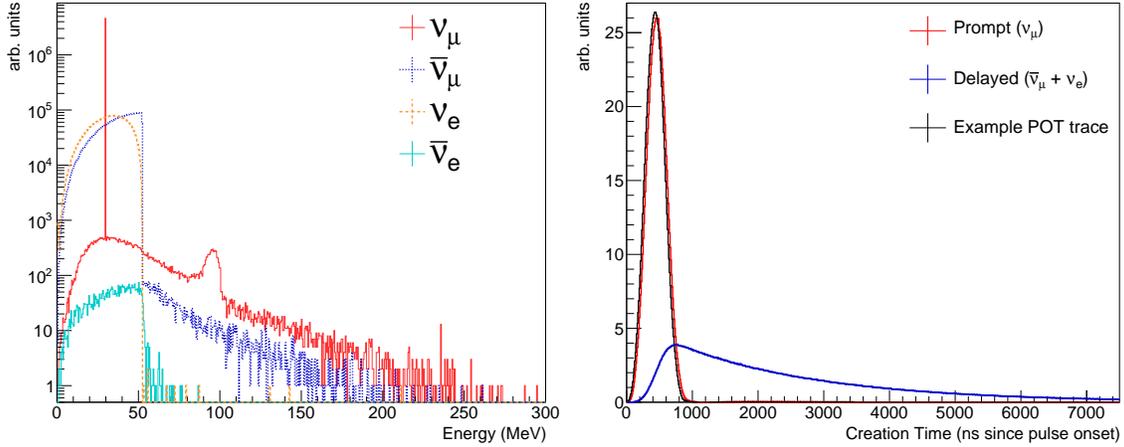


Figure 4.11: Distributions of neutrino energy (left) and creation time (right) produced at the SNS, using QGSP_BERT to model the interactions of 1 GeV protons incident on the aluminum PBW geometry. We convolve the single proton output of our simulations with the proton-on-target trace.

4.4 Neutrino Flux for the First Target Station

Figure 4.11 shows the energy and timing spectra for each neutrino flavor present in the simulation using the QGSP_BERT physics list to simulate incident protons with 1 GeV of kinetic energy on the SNS geometry with an aluminum PBW. We find that the SNS ν flux predictably demonstrates the characteristics of a pion decay-at-rest source such as the monoenergetic ν_μ at ~ 30 MeV from π^+ decay at rest and $\bar{\nu}_\mu$ and ν_e following the Michel spectra predicted from the three-body μ^+ decay at rest (DAR). Variations from these spectra include decays in flight (DIF), decays in orbit (DIO), and μ^- capture. We also observe some contribution from decay-at-rest kaons, notably in the ν_μ spectrum at ~ 240 MeV, but due to the small phase space available to produce these more massive particles, kaons have an almost negligible contribution to the SNS neutrino flux. Ultimately, this simulation predicts a decay-at-rest neutrino source with greater than 99% purity, with the exact creation process and parent particle breakdown shown for the aluminum PBW in Table 4.2. We also identify spectral features resulting from different creation processes in Figure 4.12.

Using 1 GeV protons incident on our SNS geometry from behind the beam window (PBW), our simulations predict 0.262 neutrinos per proton on target (POT). We find that our model of the SNS neutrino flux is primarily comprised of ν_μ , $\bar{\nu}_\mu$, and ν_e (each greater than $0.087 \nu_X/\text{POT}$, where $X = \mu, \bar{\mu}, e$) with a small contribution of $\bar{\nu}_e$ ($0.0001 \bar{\nu}_e/\text{POT}$,

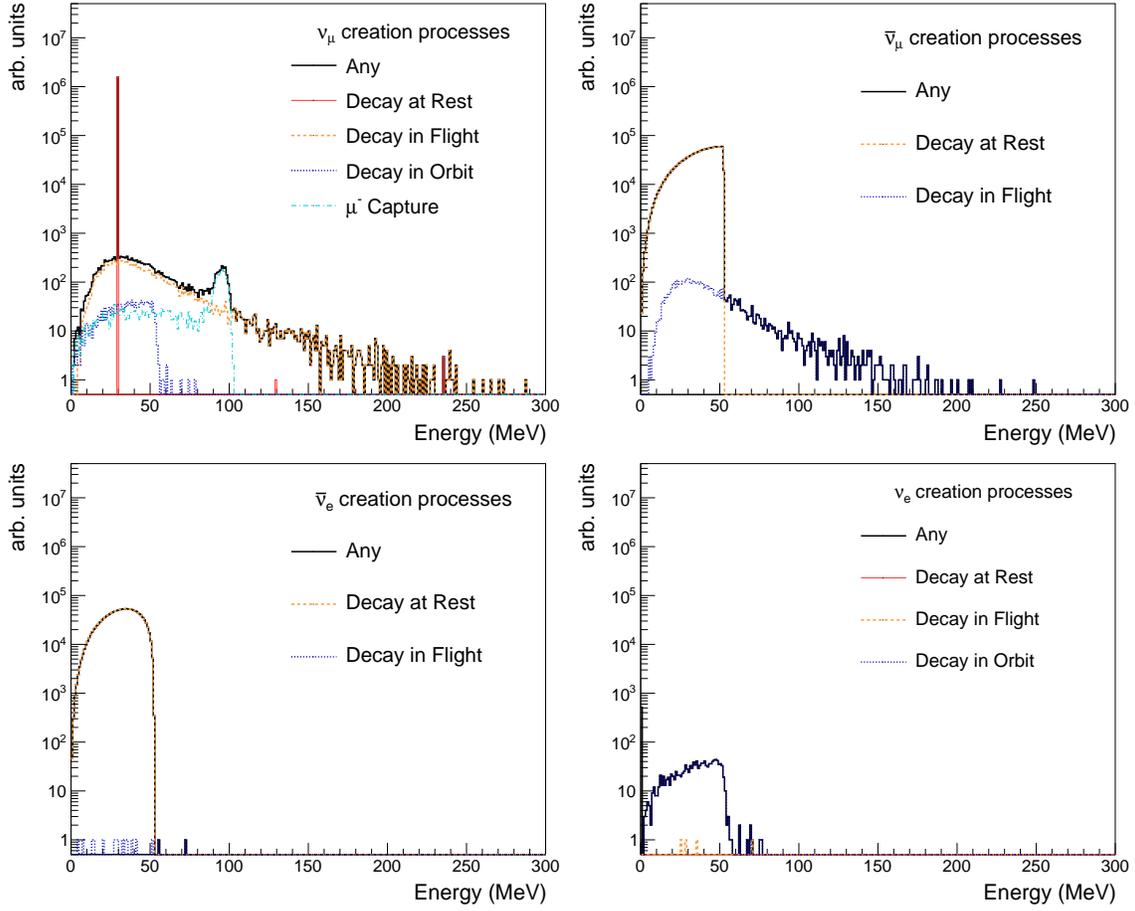


Figure 4.12: Distributions of neutrino energy broken down by creation process for each neutrino flavor produced at the SNS.

Table 4.2: A breakdown of the processes and parent particles which create neutrinos for 1 GeV protons at the SNS with an aluminum PBW. The creation processes are classified as decay at rest (DAR), decay in flight (DIF), μ^- capture, or decay in orbit (DIO). We include significant figures here to sum to 100% given the small contributions outside of the π^+ DAR chain.

	ν / POT	Creation Process				Parent Particle		
		DAR	DIF	μ^- Cap	μ^- DIO	π^+ or μ^+	π^- or μ^-	K^+
ν_μ	0.0875	98.940%	0.779%	0.196%	0.084%	99.7185%	0.2812%	0.0003%
$\bar{\nu}_\mu$	0.0875	99.718%	0.282%	—	—	99.7187%	0.2813%	—
ν_e	0.0872	99.999%	0.001%	—	—	99.9999%	—	0.0001%
$\bar{\nu}_e$	0.0001	—	0.331%	—	99.669%	—	100%	—

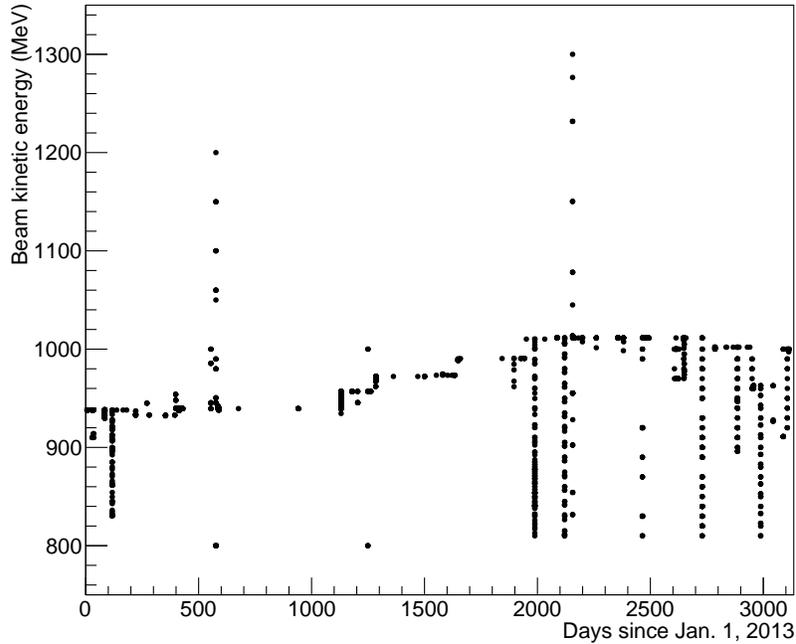


Figure 4.13: Incident proton energy at the SNS, shown from Jan. 1, 2013 – Aug. 1, 2021. Each point indicates a change to a new beam energy at the indicated timestamp.

not considering the activation of materials near the target). We also see a small flux of low-energy $\bar{\nu}_e$ from neutron β -decay that we neglect in this work, with a dedicated study of radioactive products produced as a result of SNS operations under construction.

COHERENT deployed detectors at the SNS prior to the accelerator systems reaching 1 GeV, so data taken at lower energies (~ 850 MeV) must also be understood. Fig. 4.13 shows the beam energy for COHERENT operations in Neutrino Alley, starting in January 2013 and extending through August 2021. The upcoming Proton Power Upgrade [89] will prepare the SNS for the planned Second Target Station (described in Section 4.5) by improving the accelerator. The upgrade will see the SNS operate at a more intense 2.0 MW, with 1.3 GeV incident protons by 2024. We use this simulation to study the dependence of the neutrinos produced on the incident proton energy and to develop an approach to account for changes to SNS operations over a run period. Figure 4.14 shows the energy dependence for both the total neutrino production and the fraction of neutrinos produced by the π^+ decay chain, and the parameters for each of the fits are listed in Table 4.3. This figure also demonstrates that while there are minimal differences in total neutrino production between the two PBW designs, the differences in the relative contribution of pion production resulting from interactions with the PBW (see Table 4.1) can impact

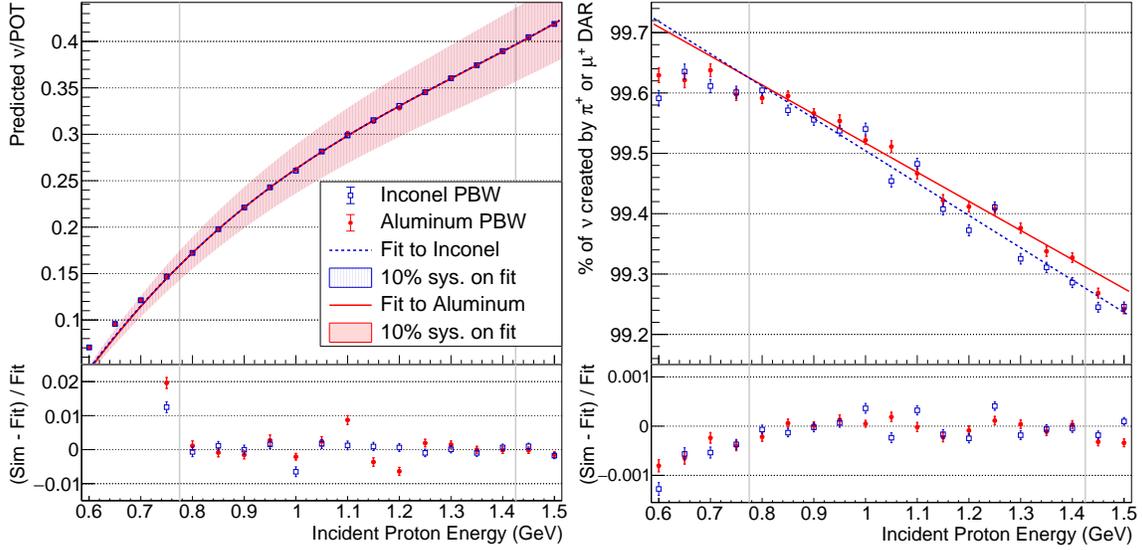


Figure 4.14: Left: The total neutrino flux from the SNS will depend on the incident proton energy, and each operational configuration demonstrates a cubic dependence on this parameter. Right: The fraction of neutrinos produced from decay-at-rest processes demonstrates a linear dependence on the incident proton energy above ~ 0.8 GeV. The fit range for both plots is $E \in [0.775, 1.425]$ GeV; this is the region between the vertical gray lines. The bottom panel in each plot shows the relative residuals, calculated as in the axis label from the simulation (“Sim”) and fit (“Fit”) predictions.

the stopping power of the SNS. The neutrino luminosity from the SNS given particular operating conditions can then be calculated as

$$\frac{\nu}{t} = \frac{\nu}{\text{POT}} \frac{\text{POT}}{t} = \frac{\nu}{\text{POT}} \frac{E_{\text{total}}}{E} \frac{1}{t} = F(E) \frac{P}{E}, \quad (4.2)$$

where E is the kinetic energy per proton, $F(E)$ is the fraction of ν produced per proton on target (POT) with incident kinetic energy E , E_{total} is the combined energy of all protons incident on the target in time t , and P is the SNS beam power (E_{total}/t). Figure 4.14 demonstrates that $F(E)$ can be estimated as a cubic polynomial in E with parameters defined in Table 4.3, for E between 0.775 and 1.425 GeV. Plugging this into Eqn. 4.2, we find a general expression for the SNS neutrino luminosity:

$$\frac{\nu}{t} = P \left(p_3 E^2 + p_2 E + p_1 + \frac{p_0}{E} \right). \quad (4.3)$$

Using this functional form and typical pre-upgrade operational parameters of 1.4 MW (7.0 GWhr/yr) and incident protons with 1 GeV of kinetic energy, we calculate 2.36×10^{15}

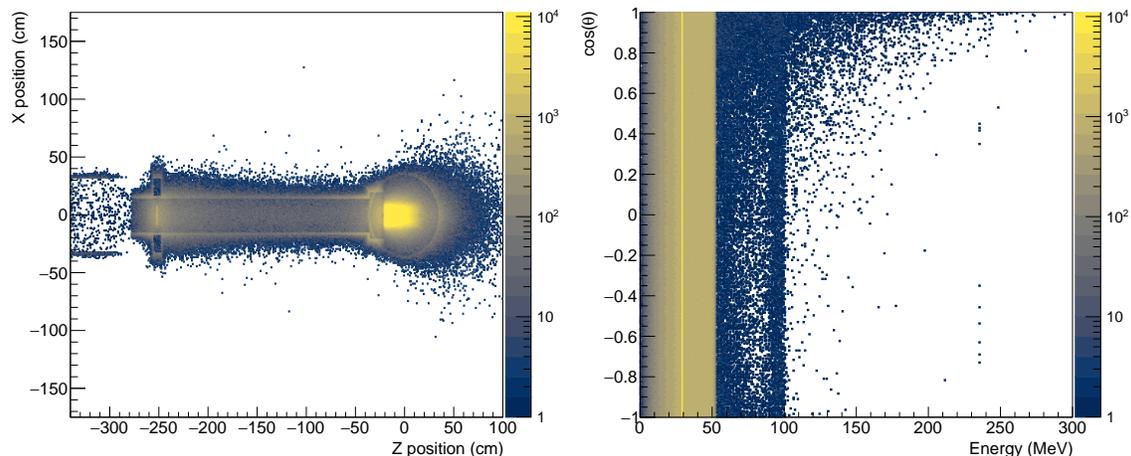


Figure 4.15: Left: A top-down view of the neutrino creation positions. Right: Distribution of the kinetic energies and production angles (relative to the beamline axis without convolving the creation position information from the left panel) of all neutrinos. The proton beam is incident from $-z$, and θ refers to the scattering angle relative to the z -axis.

Table 4.3: Fit parameters for the proton-energy dependence studies using both beam window designs. The three parameters for the cubic fits used in Eqn. 4.2 ($F(E) = p_3 E^3 + p_2 E^2 + p_1 E + p_0$) are illustrated in the left panel in Fig. 4.14, while the two parameters for the linear fits ($mE + b$) are illustrated in the right panel. The fit uncertainties do not consider the overall 10% systematic.

Design	p_3 [GeV $^{-3}$]	p_2 [GeV $^{-2}$]	p_1 [GeV $^{-1}$]	p_0	b	m [GeV $^{-1}$]
Aluminum PBW	0.28(2)	-1.12(6)	1.79(6)	-0.68(2)	99.99(1)	-0.48(1)
Inconel PBW	0.27(2)	-1.09(6)	1.75(6)	-0.67(2)	100.04(1)	-0.53(1)

neutrinos produced per second while the SNS is running. Estimating this production as an isotropic point source, we calculate a neutrino flux of $4.7 \times 10^7 \nu \text{ cm}^{-2} \text{ s}^{-1}$ at 20 meters from the target center (the approximate location of the first CEvNS measurements in CsI). Using the nominal SNS running time of 5000 hours per year, the SNS sees 1.58×10^{23} POT per year, with a ν luminosity of $4.25 \times 10^{22} \nu$ per year, or a flux of $8.46 \times 10^{14} \nu \text{ cm}^{-2} \text{ yr}^{-1}$ at 20 m from the target.

We also study the creation positions and momenta of the neutrinos, shown in Fig. 4.15. The volumes and materials which create the pions were listed in Table 4.1; the neutrinos are primarily produced after the short movements of pions and muons coming to rest. The spread of the beam and the movements of the particles result in a radial spread from

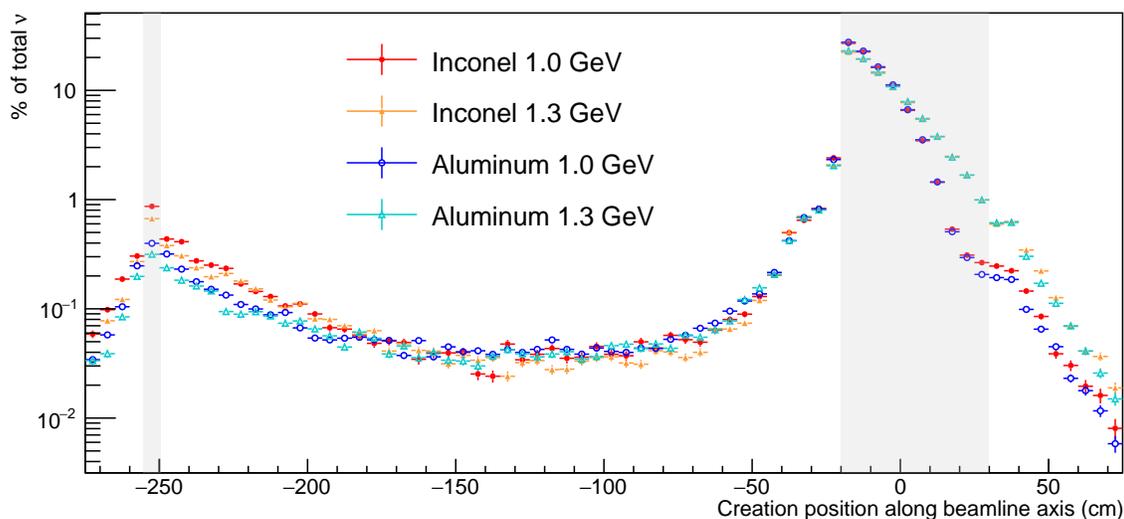


Figure 4.16: A comparison of neutrino production along the beamline for different SNS configurations and beam energies. The gray shading to the left indicates the position of the proton beam window, and the shading to the right indicates the position of the Hg target.

the beamline axis. Over 86% of the neutrinos are produced within 10 cm of the beamline axis, and almost all production (>99%) occurs within 0.5 m of the beamline axis. Along the beamline axis, we find that over 90% of the neutrino production occurs within the target and less than 5% of the neutrinos are produced at the PBW location 2.5 m upstream of the target. Because the π^+ and μ^+ decay at rest, we also have almost fully isotropic production of neutrinos up to about 50 MeV. We do note visible anisotropy in the right panel of Fig. 4.15 for $E_\nu > 60$ MeV that is consistent with neutrinos boosted in the forward direction from pions decaying in flight.

We find that both PBW designs cause some neutrino production outside of the target regardless of the incident proton energy as illustrated in Fig. 4.16. However, our detectors are deployed ~ 20 m from the target center, with the PBW placement only 2.5 m upstream of the target. To quantify the effect of this non-point-like neutrino production, we project the neutrino flux onto a 20 m sphere centered on the Hg target and determine an effective production angle based on the neutrino's projected location.

We can then determine how big of an impact the non-point-like SNS neutrino production has on the flux in Neutrino Alley. Using the histograms from Figure 4.17, we calculate

the total anisotropy as

$$a = \sqrt{\sum_{i=0}^{i=10} (S_i - P_i)^2}, \quad (4.4)$$

where S_i represents the counts in the i th bin for a spherical projection and P_i represents the counts in the i th bin for the point-source approximation. For an aluminum PBW and 1 GeV incident protons, this anisotropy is $\sim 5\%$. The dominant contribution is an excess near $\cos \theta \approx -1$, where θ refers to the polar angle upon arrival on the spherical projection, consistent with neutrino production within the PBW. There is also a secondary excess near $\cos \theta \approx 1$ consistent with neutrinos produced by decays in flight. For a small detector at the CsI location 19.3 m from the target center and at $\cos \theta \approx 0$, we predict less than a 1% deficit of the neutrino flux compared to the isotropic point-source approximation. The contributions to the neutrino flux error from geometric considerations are small, and add negligibly in quadrature to the 10% to the overall neutrino flux incident on our detectors in Neutrino Alley. The anisotropy depends on the relative contributions of the different materials in our SNS geometry outlined in Table 4.1, and emphasizes the need for new pion-production measurements such as those discussed in Section 4.7.

In general, we provide this global picture of the SNS neutrino production to individual detector systems in COHERENT. However, with precise knowledge of the detector geometry and location, we can also simulate the neutrino production specific to a subsystem. This does, however, require *significantly* more computational resources to generate results to an informative level of statistical precision.

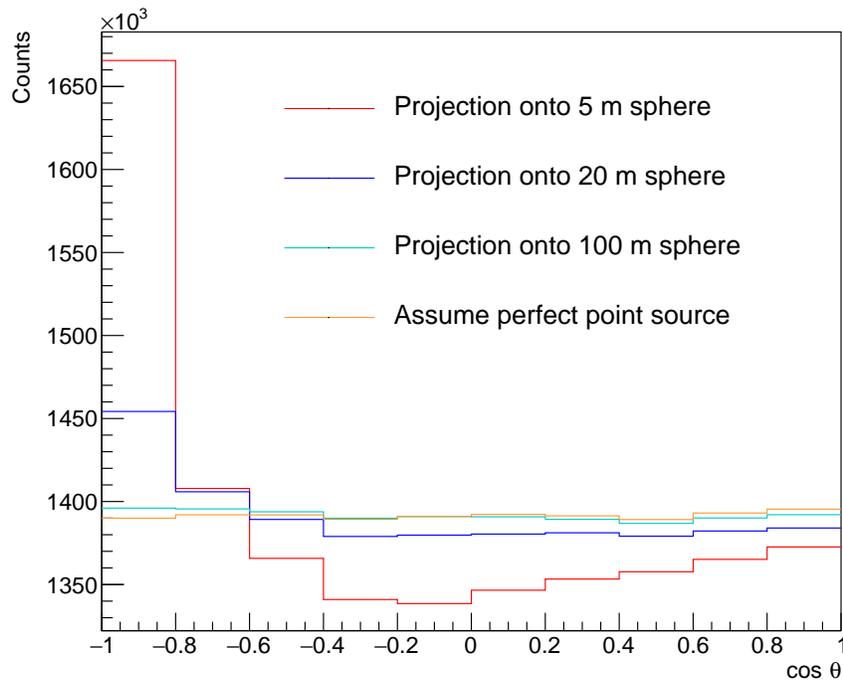


Figure 4.17: Examining projections of the SNS neutrino flux onto spheres of different sizes, and comparing each projection to the point-source approximation to determine the anisotropy.

4.5 Neutrinos at the Second Target Station

We also created a model geometry to estimate the neutrino production at ORNL’s planned Second Target Station (STS) [90]. With a projected completion in the early 2030s, COHERENT is engaged with the design phase of this facility to optimize location and shielding with the aim to deploy 10-ton-scale detectors for CEvNS and other physics. Using preliminary details about the planned target provided at the Workshop on Fundamental Physics at the Second Target Station in 2019 [145], we modeled 21 tungsten wedges surrounded by thin layers of tantalum and water, evenly spaced in an assembly with a 1.1 m diameter. We also modeled neutron moderators above and below the active target wedge, centered along the beamline axis. We simulated a 6 cm (width) \times 5 cm (height) beam profile to ensure that the profile is smaller than that of a single tungsten wedge and included the aluminum PBW and beamline shielding as implemented in our First Target Station (FTS) geometry. This target geometry is illustrated in Fig. 4.18 and is centered inside a 5 m vacuum box, then enclosed in a steel box (10 m outer edge, 5 m inner) to mimic pion production in typical shielding materials without assuming the geometry of the STS target

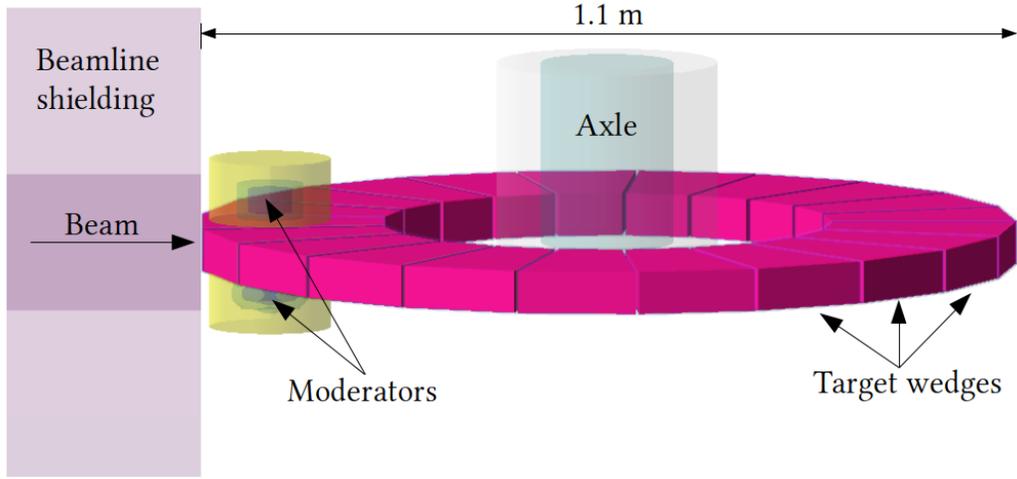


Figure 4.18: Geant4 implementation of the Second Target Station target and moderators.

surroundings.

With this simple geometry and 1.3 GeV incident protons, our simulations predict $0.13 \nu_X / \text{POT}$ for $\nu_\mu, \bar{\nu}_\mu$, and ν_e from the π^+ decay chain, resulting in an approximate total $0.39 \nu / \text{POT}$. This estimate is larger than the predictions for the FTS operating at 1.3 GeV due to the increased density of a solid tungsten target. We cannot accurately discuss the decay-at-rest fraction of neutrinos or relative impact of the PBW since the shielding surrounding the target remains unknown. However, we note that protons do escape the end of the 25-cm thick active target wedge with enough energy that nearly 10% of the simulated pion production in our simple geometry occurs downstream of the target.

The STS will receive one of every four pulses from the SNS linear accelerator and will operate as a 15 Hz, 0.8 MW facility. Possible locations for 10-ton scale COHERENT detectors at the Second Target Station have been identified within a few tens of meters from the planned target location [146]. With a tungsten target rather than mercury, the uncertainty on the overall pion production is expected to be nearer to 15% (see right panel of Figure 4.5), though efforts to normalize the neutrino flux at the FTS using D_2O and new hadron production experiments using a range of targets at lower beam energies will help to constrain our predictions (see Section 4.7).

4.6 Light Dark Matter Production at the SNS

This work was focused on understanding the neutrino fluxes but also explored the creation of other interesting particles. In particular, π^0 , η^0 , and π^- production are relevant to dark

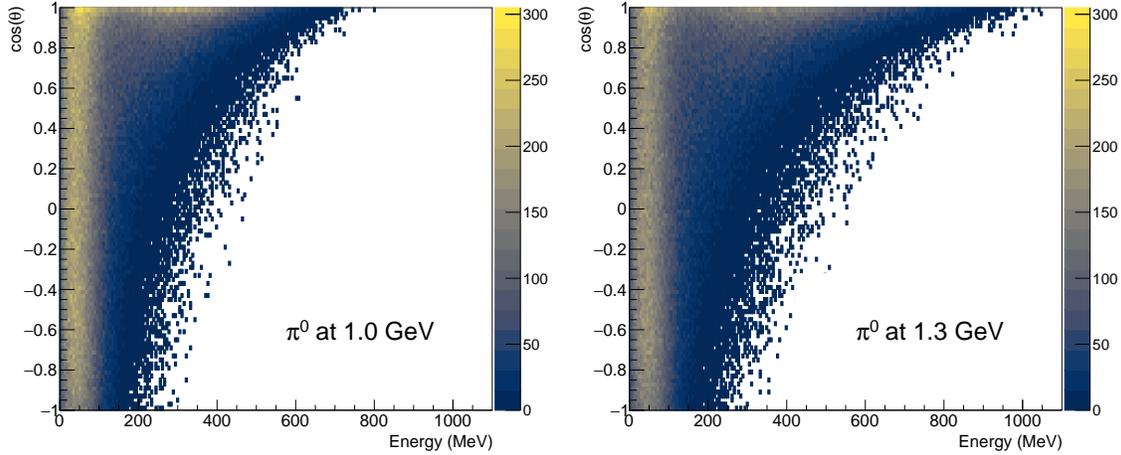


Figure 4.19: Distribution of production angles and creation energies of π^0 at the SNS-FTS assuming an aluminum PBW. Left: Using 1 GeV incident protons to mimic the current operating conditions of the SNS. Right: Using 1.3 GeV incident protons to mimic the operating conditions following the upgrade.

matter searches using the SNS as an accelerator [147, 148]. Here, we present some findings regarding the production of such particles using QGSP_BERT, noting that no effort was made in this work to specifically validate the production of any hadrons other than π^+ . As mentioned in Section 4.2, η production is excluded from our discussion here because it is not predicted by QGSP_BERT. Predictions with QGSP_BIC have previously been used with this simulation geometry to predict η flux for sensitivity studies [148].

Figure 4.19 shows the scattering angle as it relates to the creation energy for SNS-produced π^0 from 1 GeV incident protons on the left and 1.3 GeV incident protons on the right. We observe strong forward production for both, but note that we will have a small flux directed towards Neutrino Alley ($\cos\theta \approx 0$). This is relevant primarily for the π^0 which could decay in flight into dark-matter particles that cause an observable nuclear recoil in our CEvNS detectors. For π^- , dark matter could be produced in an absorption process or in a charge-exchange process; both are more efficient at non-relativistic energies, and each would emit particles isotropically and negate any impact of forward production.

Assuming an aluminum PBW, the SNS produces 0.11 π^0 /POT and 0.05 π^- /POT for 1 GeV incident protons. We also predict that the upgraded 1.3 GeV incident protons will produce 0.17 π^0 /POT and 0.09 π^- /POT. This study also demonstrates the potential gain of the STS for dark matter searches, particularly in aiming for forward-positioned detectors and a reduced distance to the target.

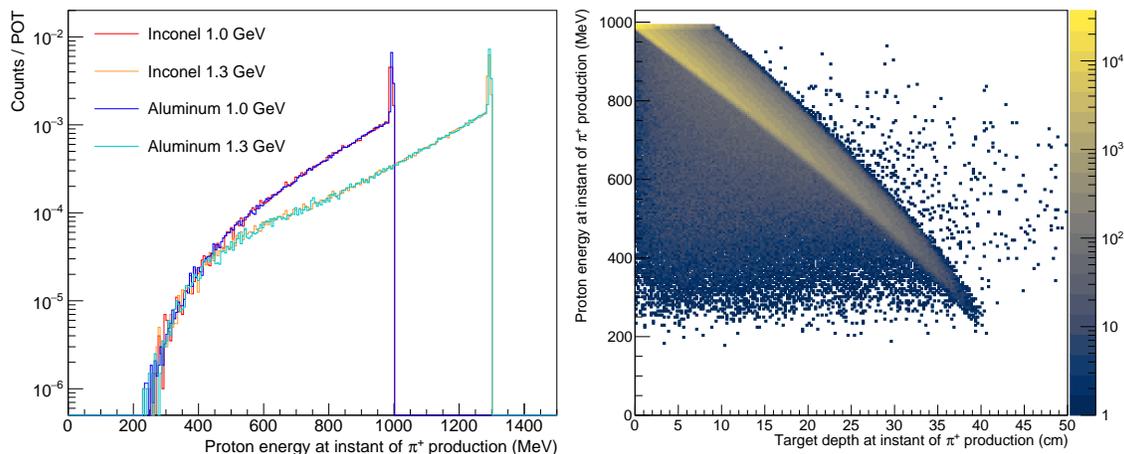


Figure 4.20: Top: A histogram of the proton energies which produce π^+ at the SNS. Bottom: A closer look at how protons lose energy in the Hg target before creating π^+ .

4.7 Ongoing Efforts

In the absence of pion-production data for protons incident on Hg at energies up to 1.3 GeV, the $\sim 10\%$ uncertainty assigned to our neutrino flux is a robust estimate that cannot be significantly improved through simulation. For the moment, the cross-section results on Ar are still dominated by statistical uncertainties [107], but the uncertainty on the neutrino flux is now the dominant uncertainty in the final results from CsI [108]. Two types of experimental measurements could further reduce this uncertainty.

Pion-production measurements with thin Hg targets would allow us to validate our simulation against interactions on the same material as the SNS target. The proposed EMPHATIC experiment at the Fermilab Test Beam Facility could measure differential pion-production cross sections on Hg at proton energies as low as 2 GeV with expected uncertainties less than 10% [149]. Meanwhile, the NA61/SHINE collaboration [150], which has measured pion production on both thin and replica targets for a variety of accelerator neutrino experiments, is investigating the possibility of reducing the energy of the CERN SPS H2 proton beamline to 1 GeV for low-energy pion-production studies [151]. These measurements will benefit neutrino experiments at the SNS and at other pion decay-at-rest neutrino sources with GeV-scale protons incident on a mercury target such as the JSNS² sterile-neutrino search at the Japan Spallation Neutron Source [152]. As a decay-at-rest source is insensitive to the production angle of the pion, new pion-production measurements should ideally cover as close to a 4π acceptance as possible.

Thin-target data at ≥ 1 GeV, however, cannot account for the effects of proton energy

loss in the SNS target and from scattering in the PBW as shown in Fig. 4.20. Reducing our uncertainty using only thin-target data would require convolution of the simulated proton energy-loss profile with the observed pion production from thin-target data at varied energy and on a variety of target materials. A separate approach to reducing neutrino-flux uncertainties would directly measure the total neutrino production at the SNS target. A D₂O detector, deployed at the SNS, would measure the charged-current interaction

$$\nu_e + d \rightarrow p + p + e^-. \quad (4.5)$$

The cross section of this reaction is well understood; theoretical calculations, taking several disparate approaches, have converged to the 2–3% level [153, 154]. A moderately sized detector, about 680 kg, could achieve similar statistical precision in about four SNS beam-years of operation. The observed ν_e flux from the SNS target could then be multiplied by three to obtain the total flux of all three neutrino flavors generated by π^+ decay. The COHERENT collaboration plans to build such a detector to directly normalize the simulated SNS neutrino flux [155]; we discuss the design details in Chapter 5. We note that if the neutrino flux can be independently measured to high precision, one can in principle use neutrino data to validate models of hadron and neutrino production.

Chapter 5

Design Studies for a Flux Normalization Detector at the SNS

With the 10% uncertainty on the SNS neutrino flux from a lack of π -production data as described in Chapter 4, the most effective way to reduce this systematic is to take new data. While new, large-acceptance thin-target pion production data will be incredibly informative, the thick and dense FTS and STS targets quickly cause protons to lose energy, such that pion production will take place for proton energies less than the incident beam energy. Because of this, a direct normalization of the SNS neutrino flux is more immediately useful to the collaboration. Thin-target data can then be used to better inform the model predictions relating to the individual contributions of the target, beamline, and moderators we discussed in the previous chapter.

The collaboration is currently commissioning a D₂O detector to normalize the neutrino production at the SNS; this concept is outlined in [156]. This chapter outlines the detection strategy, simulation efforts, and design decisions that have led to the current and future deployments. The work presented in this chapter pulls from the efforts and ideas of the D₂O working group within the COHERENT collaboration, especially those of Igor Bernardi, Yuri Efremenko, Karla Tellez-Giron-Flores, Matthew Heath, Jon Link, Jason Newby, Diana Parno, and Kate Scholberg.

5.1 D₂O Detector Overview

The cross sections for neutrino interactions in the tens-of-MeV range are not well known experimentally [28], and theoretical calculations for neutrino interactions with large nu-

clei can be quite complicated due to the nuclear physics that must be taken into account. However, the cross sections for small nuclei, such as those of hydrogen or deuterium, can be precisely calculated with effective field theories since there is no bound final state following an interaction. In particular, the charged-current cross section for

$$\nu_e + d \rightarrow p + p + e^-, \quad (5.1)$$

which we'll refer to as "CC-D" throughout this work, is of particular interest to theorists given its importance to the Sudbury Neutrino Observatory (SNO) which confirmed the existence of neutrino oscillations [19], and given that the inverted process $p + p \rightarrow d + e^+ + \nu_e$ is more commonly known as the first branch of the pp -chain, a nuclear fusion process which dominates the energy production in stars like the Sun [157].

The uncertainty on the cross section for CC-D is estimated at only 2-3% by examining the agreement between calculations utilizing different theoretical approaches [158, 159]. As laid out in [28, 160], these approaches include:

- Standard nuclear physics¹: The nucleus is treated as a particle with assigned quantum numbers. A model of the potential (typically parameterized from data) is used to generate nuclear wave functions. These wave functions are then acted on by operators which parameterize the transformation properties of nuclear states observed in experiment [163].
- Hybrid effective field theory (EFT*): Transition operators are taken from chiral perturbation theory and wave functions generated by a potential model are used to calculate the nuclear transition [164].
- Pionless effective field theory (π EFT): Nucleon-nucleon interactions and two-body currents are described as point-like in an expansion of q/m_π given the low characteristic momentum of the interaction q relative to the pion mass m_π [165–167].

The 2-3% uncertainty envelope between these approaches can likely be further reduced by implementing experimental constraints (such as the decay width of tritium) on the theories, and there is active work in this sector. In particular, the effective field theory approaches described above depend on only one two-body coupling (generally $L_{1,A}$ in literature), and experimentally constraining this axial exchange-current operator reduces uncertainty on both the EFT* and π EFT approaches [168].

¹This strategy was first used for μ -capture calculations [161, 162].

With low backgrounds and a large volume of deuterium, measuring the flux of ν_e neutrinos at the SNS could reach a statistical precision at the few-percent level and enable us to reduce the largest systematic for all CEvNS detectors from the COHERENT collaboration. As we describe in Chapter 4, the majority of the ν_e flux that is expected to be coincident with the beam signal is produced by the decay of μ^+ , which themselves are produced by the π^+ decay. The ν_e flux, as demonstrated in the neutrino-flux simulations of the SNS, accounts for 1/3 of the total SNS neutrino flux since there are only sub-percent contributions (0.22% for 1 GeV incident protons) on the total neutrino flux from the π^- chain. In measuring the ν_e flux, we ultimately normalize the total SNS flux ($\Phi_\nu \approx 3\Phi_{\nu_e}$).

5.1.1 Signal

To amass a chemically stable volume of deuterium able to quickly reach statistical precision, we use “heavy water”, or D_2O , which features an extra neutron on each hydrogen nucleus compared to H_2O . Deploying this detector off-axis reduces the decay-in-flight contributions to the SNS neutrino flux, so our ν_e will primarily be produced with energies between 0 and 50 MeV. This flux of ν_e will then produce e^- through the well-understood CC-D process inside our detector, with an average energy in the tens-of-MeV range. These energetic electrons traveling through the heavy water will emit Cherenkov radiation with the emitted frequency distribution per unit length given by the Frank-Tamm result [169]

$$\frac{\partial^2 E}{\partial x \partial \omega} = \frac{q^2}{4\pi} \mu(\omega) \omega \left(1 - \frac{1}{\beta^2 n^2(\omega)} \right), \quad (5.2)$$

where E is the energy of the emitted photon spectrum, x represents the distance traveled in the medium, $\beta = |\vec{v}|/c$ describes the speed of the electron, ω is the frequency of the emitted radiation, and $\mu(\omega)$ and $n(\omega)$ are the frequency-dependent permeability and index of refraction of the medium, respectively. This frequency distribution is not, however, fully observable, as our photomultiplier tubes (PMT) are limited by their frequency-dependent quantum efficiency (QE).

The physical constraints of Neutrino Alley limit the detector size and shape (more in Section 5.2), so we cannot have full photocathode coverage of the volume and do not plan to use reconstructed Cherenkov rings to identify events. Rather, we focus on maximizing the overall light collection and choose to use a reflective layer behind the PMTs.

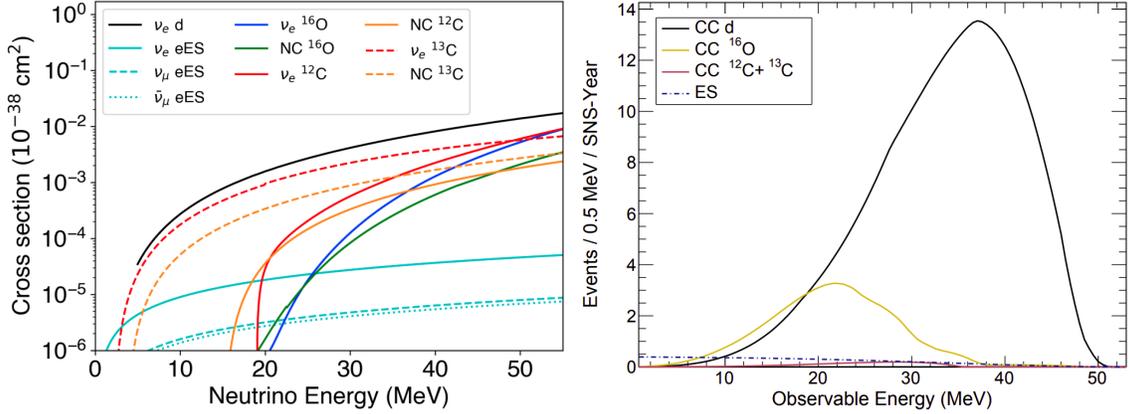


Figure 5.1: Left: A plot comparing the predicted cross sections for possible neutrino interactions in a D₂O detector at the SNS, including charged-current (CC), neutral-current (NC), and elastic-scattering (ES) interactions. Right: A plot showing the predicted event rates in a D₂O detector at the SNS as a function of observed electron energy for the highest-probability processes. Both figures reproduced from [156].

5.1.2 Background

Several backgrounds have been considered so far, including a mix of steady-state and beam-related backgrounds. Ambient backgrounds in Neutrino Alley (like the 511-keV γ s from the hot-off-gas pipe) can be mitigated with a moderate amount of lead shielding. The low overburden (8 m.w.e.) of our surface-level experiment means that we will also need to characterize the detector’s response to cosmic rays and other backgrounds that are not associated with the beam. Muons passing through the detector will also emit Cherenkov radiation, but cosmic-ray showers in our modest overburden and in the detector surroundings can also produce γ s that are observable in the detector. We do plan to have muon veto panels surrounding the detector to help identify these events in our analysis, and 2” lead shielding (inside the muon veto) to offer some protection from cosmic-ray-induced γ -rays, but Michel electrons have an energy distribution in the tens of MeV similar to our signal flux. The production of Michel electrons within our detector or in the detector surroundings will be important to characterize and distinguish from our signal. These steady-state backgrounds will be measured during the anti-coincidence periods.

Beam-related backgrounds do include neutrons, which can interact in the surrounding shielding or capture on O to produce γ s that are observable by our PMTs. With the range in arrival times for ν_e and the time it takes for a prompt neutron to thermalize and capture in the detector, a timing cut is not an easily viable option. Thus, we must directly simulate

the neutron background to develop a strategy to mitigate this background which primarily appears within our signal window. However, these neutrons are not the dominant beam-related background for this D₂O detector, given that both d and O will be possible target nuclei for incident neutrinos. The charged-current interaction

$$\nu_e + {}^{16}\text{O} \rightarrow e^- + {}^{16}\text{F}^*, \quad (5.3)$$

which we'll refer to as "CC-O" in this work, similarly produces an electron observable through Cherenkov radiation. Cross sections of interest for ν_e+d and ν_e+O are shown in Figure 5.1, and demonstrate that the energy distributions for the electrons produced from these charged-current interactions are similar and that we must have a detector capable of resolving these signals.

5.1.3 Energy Reconstruction and Resolution

In simulation, we know the "true" energy of the electron, E_{src} . However, in data, the observed energy of an event is reconstructed from the measured number of photoelectrons (PE) for the event. Our simulation demonstrates that more PE correspond to higher energy events in a roughly linear manner ($E_{\text{src}} \propto N_{\text{PE}}$) as shown in Figure 5.2. Our energy reconstruction, then, determines the energy scale (PE / MeV) by measuring the slope s of the line illustrated in Figure 5.2, and defining $N_{\text{PE}} = sE_{\text{rec}}$.

In Figure 5.3 left, we show the distribution of the number of optical photons recorded for simulated electrons with 30 MeV isotropically emitted uniformly throughout the detector volume. We observe a Gaussian-like peak with a long low-energy tail. In our analysis, we perform a two-staged fit to identify the central value of the peak to determine the expected number of photoelectrons per MeV. The first stage of the fit is a Gaussian fit, with no restrictions on the fit parameters. Typically, this finds a mean μ_1 that is not representative of the distribution, but consistently falls within the peak region. We then perform a second fit, restricting the fitting to the range $\mu_1 \pm 100$ PE. The best-fit parameters μ and σ from the second fit are used to characterize the energy reconstruction. The point (μ, E_{src}) is expected to fall on the red line illustrated in Figure 5.2, so by varying E_{src} , we produce enough points to accurately determine the slope s .

We also define the fractional energy resolution for a given E_{src} as σ/μ , again utilizing the best-fit parameters from the two-staged fit shown in Figure 5.3 left. This energy-dependent resolution is plotted in the right panel of Figure 5.3. We characterize this dependence with the modified exponential curve

$$f(p_0, p_1, p_2) = p_0 e^{p_1 E_{\text{src}}} (E_{\text{src}})^{p_2}, \quad (5.4)$$

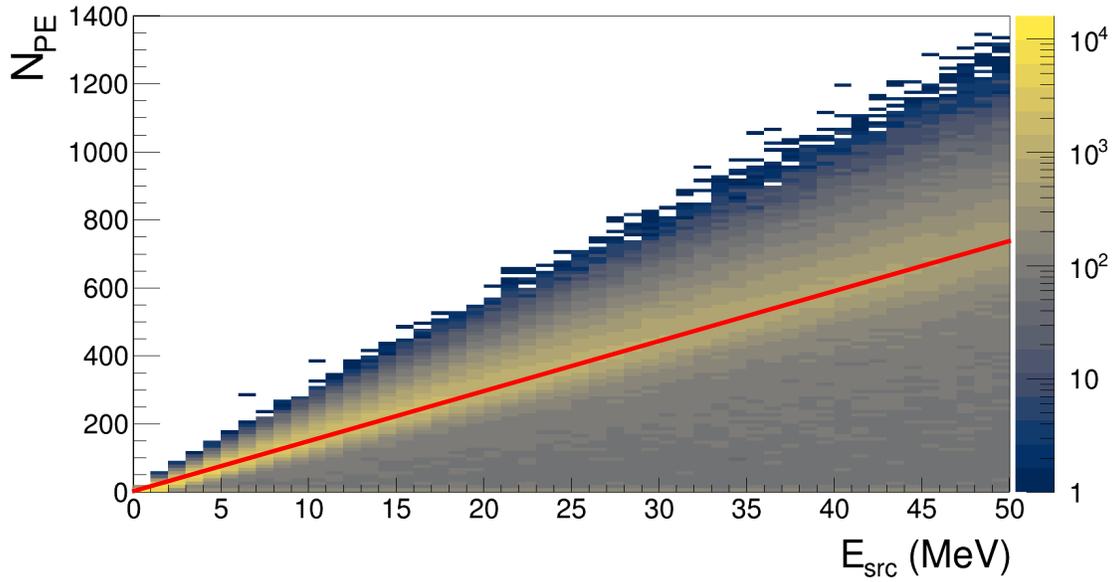


Figure 5.2: The simulated distribution of the number of PE (N_{PE}) observed for electron events with initial energy E_{src} at random angles and from uniformly sampled positions within the D_2O and H_2O volumes in the Tank geometry (see Section 5.2.3). The red line superimposed onto the distribution illustrates the linear dependence which defines our energy reconstruction E_{rec} . Deviations away from that line demonstrate the effect of our energy resolution.

where p_i are parameters determined by the fitter. These curves describe the general trends of the energy resolution and are used to compare different design options and to quickly estimate the energy resolution for use in our fiducial analyses (more in Section 5.5). The simulations demonstrate that we will have better than 20% energy resolution in this geometry above 10 MeV, or for the majority of of our signal region. As an example, this means that for a population of 30 MeV electron events, about 67% of the events will be reconstructed between 24 and 36 MeV.

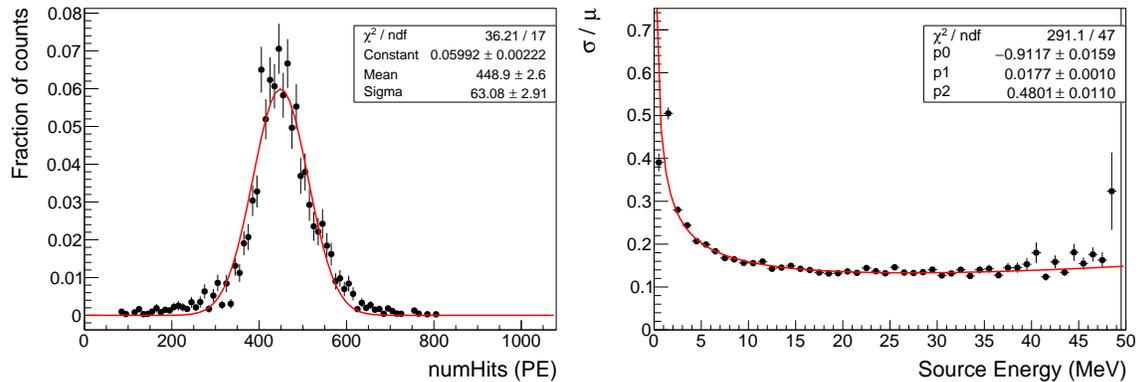


Figure 5.3: Left: The distribution of optical photons recorded for simulated 30 MeV electrons with uniformly sampled position and direction in the Tank geometry. We perform a two-staged Gaussian fit (described in the text) to the peak to obtain a measure of the conversion from MeV to photoelectrons; the red curve illustrates the second-stage fit. Right: The simulated fractional energy resolution calculated from the fit parameters μ and σ from Gaussian fits as in the left panel. We then model the behavior using the modified exponential from Equation 5.4. The large changes and high uncertainties at higher energies are associated with low statistics in those bins.

5.2 Detector Geometries

Adapting a Geant4 simulation originally designed by Matthew Blackston, we have modeled different geometric options, optical properties, and design parameters for the planned D₂O detector. In this section, we discuss the simulation comparisons which contributed to the design decisions outlined in [156]. In all cases considered in this section, the material definitions and optical properties are identical; we change only the geometry of the volumes we implement and the locations of the PMTs. We will discuss optical changes separately in Section 5.6.3. As such, we'll present the material descriptions only in Section 5.2.1 alongside the geometric description. The geometric descriptions we offer here have cartesian dimensions Width \times Height \times Depth, or cylindrical dimensions which will have a Diameter \times Height along a given axis.

5.2.1 Box Geometry

The neutrino program at the SNS was, in truth, an afterthought — the original design of the target facility did not intentionally include accommodations for neutrino physics. Neutrino Alley, the experimental home of COHERENT detectors, was designed as an ad-

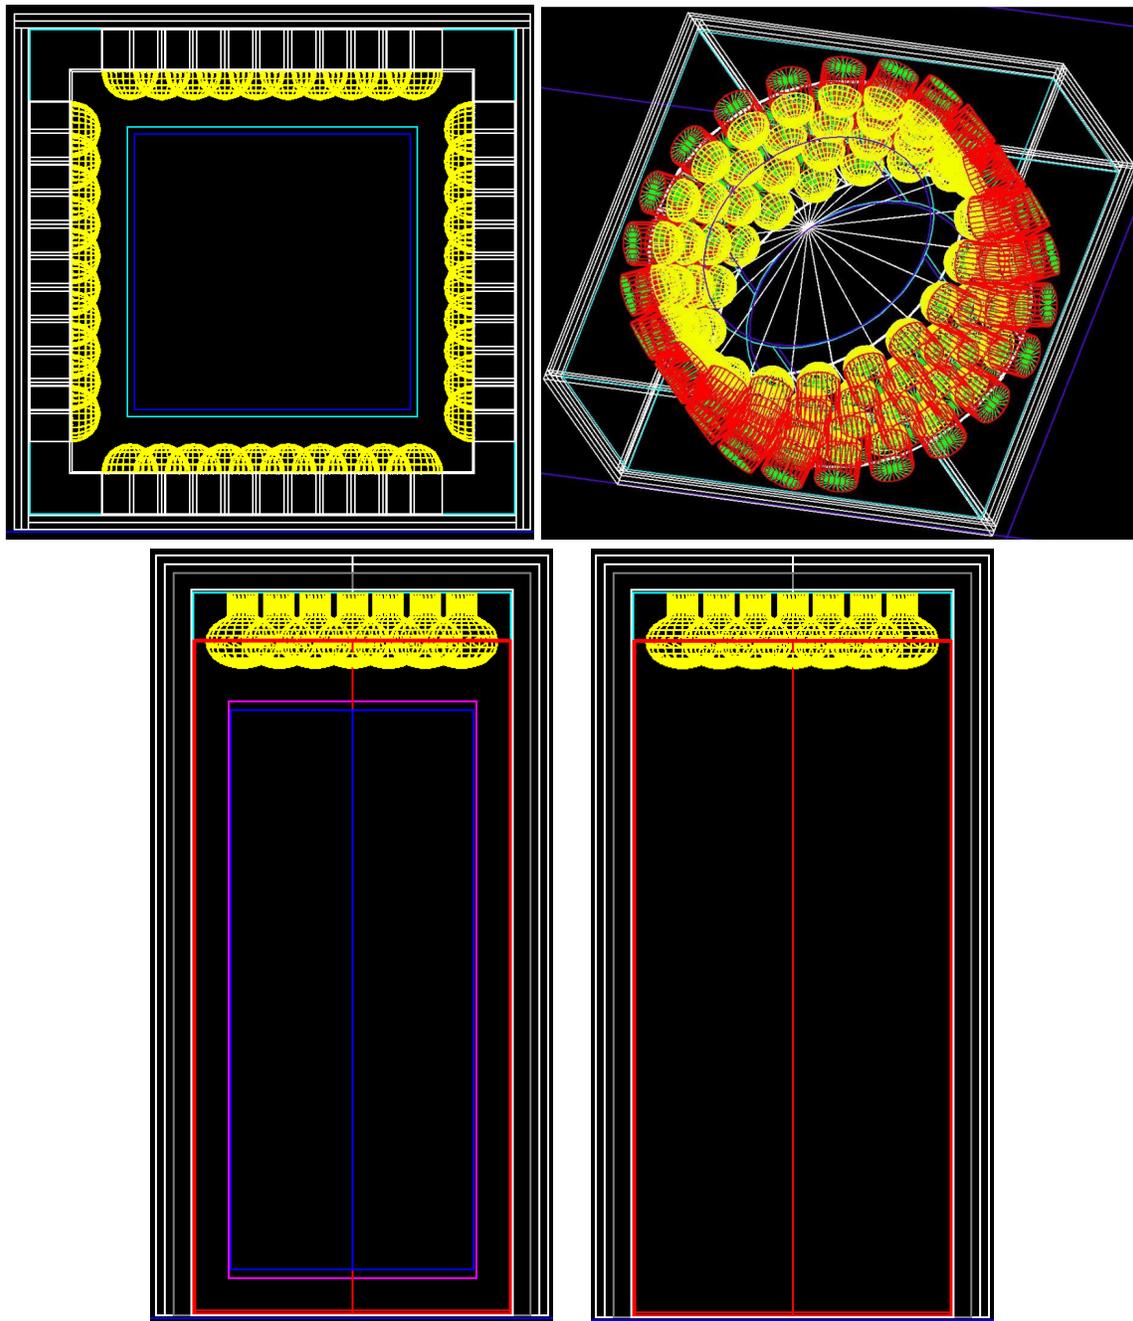


Figure 5.4: Screenshots of the Geant4 models for the “Large”-sized Box geometry from Section 5.2.1 (top left) and Axle geometry from Section 5.2.2 (top right), and the Tank and No-D₂O geometries described in Section 5.2.3 (bottom left) and Section 5.2.4 (bottom right), respectively.

ditional fire-safety exit, and as such, it is mandatory to abide by the safety restrictions of Neutrino Alley and leave an unimpeded 3 ft walking corridor. Thus, the total detector depth, or the dimension expanding from the wall into the hallway of Neutrino Alley, must be less than a meter.

The “box” geometry arose from the goal to produce the largest possible tank given the space constraints of Neutrino Alley. We have a height restriction of 2 m for deployment in the chosen location, and we expect to need a few inches of lead shielding on all sides of the detector to limit ambient backgrounds in Neutrino Alley, as well as muon veto panels surrounding the detector to mitigate cosmic-ray backgrounds.

Early simulation work suggested that the overall energy collection of this detector design improves with the inclusion of a tail-catcher region, or a volume designed to ensure the electrons will pass through enough water to radiate all their energy; all the D₂O designs we describe here use a 10 cm H₂O tail-catcher surrounding a transparent acrylic vessel containing the D₂O. The PMTs will be deployed within the H₂O volume, which itself will be contained by a stainless-steel outer tank. A reflective layer will be used to improve the overall light collection, and is placed just behind the photocathode and on all the steel surfaces so that light cannot scatter out of the sensitive region of the detector. To allow for a 10 cm tail-catcher, PMTs on both top and bottom of the detector, and a acrylic vessel with only flat sides for simplicity, we use a 1 inch thick acrylic vessel to contain the inner volume of D₂O with dimensions $140 \times 140 \times 60 \text{ cm}^3$, totaling 1.3 (metric) tons of heavy water.

We explored two iterations of this geometry: the “Large” detector, which would use 1.3 tons of D₂O and 112 PMTs and occupy the full allowed region in Neutrino Alley, and a “Half”-sized version to reduce costs by limiting the amount of D₂O to 0.67 tons and using only 80 PMTs. A Geant4 screenshot of the “Large”-sized detector is shown in the top left panel of Figure 5.4. For our preliminary investigations, we aimed to investigate the light collection and energy resolution of each geometry for electrons distributed uniformly in both the D₂O and H₂O volumes to model the response to CC-D and CC-O events; our ability to correctly reconstruct the energy of a given event will be vital for reducing backgrounds (discussed in Section 5.4).

5.2.2 Axle Geometry

In the box geometry, our simulations observed an abundance of CC-O events in the corners of the H₂O volume. As this is the dominant background for our normalization of the neutrino flux, we transitioned to a more complex geometry which minimized the amount

of H_2O in the volume observed by the PMTs without sacrificing performance (e.g., by reducing the 10 cm tail-catcher). In this model, we use a cylindrical acrylic tank, placed on its side like a wheel. We again organize the PMTs with minimal spacing between them and angle each to point inwards towards the central axis at all points on the cylinder, preventing corners or geometrically unavoidable blindspots. As in the box geometry, the reflector is placed directly behind the photocathodes of the PMTs and on the flat steel surfaces so that events behind the PMTs cannot be observed, and light cannot scatter out of the sensitive volume of the detector.

Like the box geometry, the axle geometry also had two iterations; a “Large” detector, which would contain 1.02 tons of heavy water and require 100 PMTs, and a “Half”-sized version to reduce the amount of D_2O to 0.67 tons and use only 84 PMTs. A Geant4 screenshot of the “Large”-sized detector is shown in the top right panel of Figure 5.4.

Figure 5.5 illustrates the expected signals from CC-D and CC-O (the calculation of which we’ll describe in detail in Section 5.4) for the “Half”-sized geometries. Note the reduction in CC-O counts for the axle geometry in the region above 10 MeV for the same amount of live time (more obvious in the excess of low-energy events since the overall normalization of the curves are identical). This is illustrative of the decrease of H_2O counts in the corners of the tail catcher; these events still exist, but are now emitted behind the PMTs so they do not reconstruct with any precision. The axle geometry, then, increases the precision of our CC-D measurement because there will be fewer events from the dominant beam-coincident background, CC-O, above a given threshold. We’ll examine the specifics of this threshold setting in Section 5.4.

5.2.3 Tank Geometry

The engineering required to build something like the wheel geometry would certainly be more costly than the “simple” box geometry, and funding for this detector was not obvious at the time of these design studies. To reduce costs even further than the half-sized geometries we posed earlier, we imagined a third geometry for the detector which would use 0.6 tons of D_2O (which a donation by Hank Sobel of UC Irvine could supply) and only 12 PMTs to reduce purchasing costs for the detectors, high-voltage supplies, and acquisition system. We chose not to explore a Cartesian geometry to avoid the “corner” problems we observed with the initial box design. Simulations from Matthew Heath improved the description of the PMTs and indicated that muon-veto panels on the bottom of the detector do not add significantly to the physics performance of any of the geometries, and we adopted both changes in this model. A Geant4 screenshot of the detector is shown in the

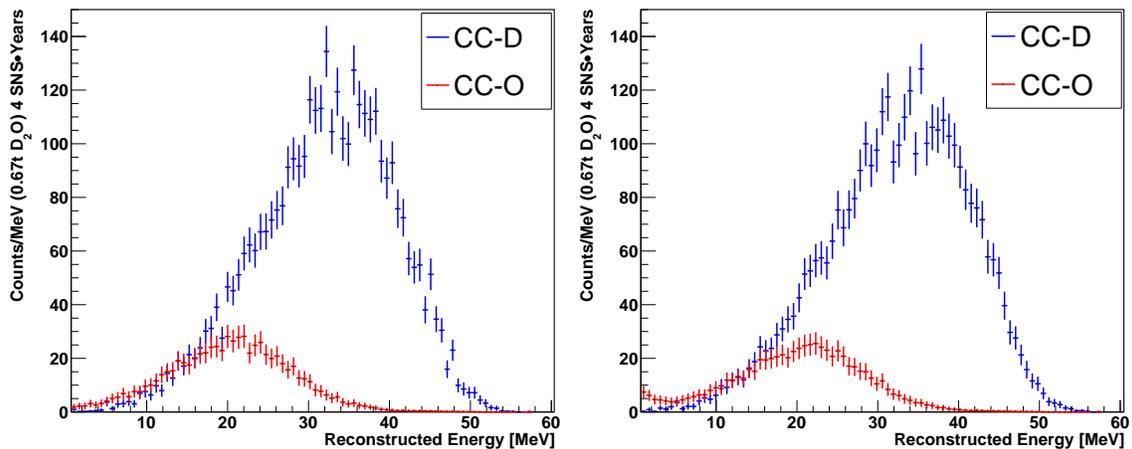


Figure 5.5: A comparison of the simulated response of the “Half”-sized box (left) and axle (right) geometries to CC-D signal and CC-O background. These curves are calculated using the strategy outlined in Section 5.4.

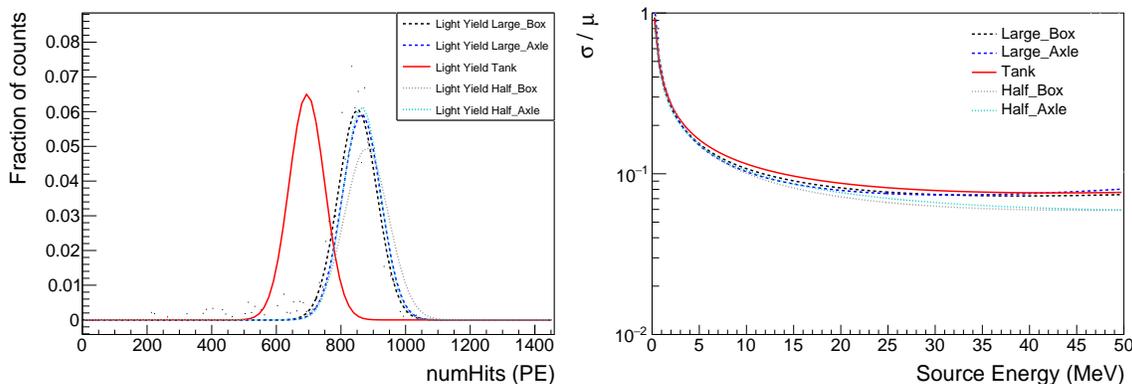


Figure 5.6: Left: A comparison of the best-fit light collection (see Figure 5.3 left) for uniformly distributed 30 MeV electrons in each geometry. Right: A comparison of the best-fit fractional energy resolution (see Figure 5.3 right). We find that the “Half”-sized geometries with a large number of PMTs have the best performance, but the low-cost Tank design does not sacrifice much in the way of performance since we are primarily interested in events with $E_{\text{src}} > 20$ MeV. We use these studies to determine that the Tank design is the best path forward for the collaboration.

bottom left panel of Figure 5.4.

These comparisons demonstrate that the lowest-cost “Tank” design does not significantly sacrifice performance for events of interest ($E_{\text{src}} > 20$ MeV) compared to the other options we considered, and this design has become our baseline model. An engineering

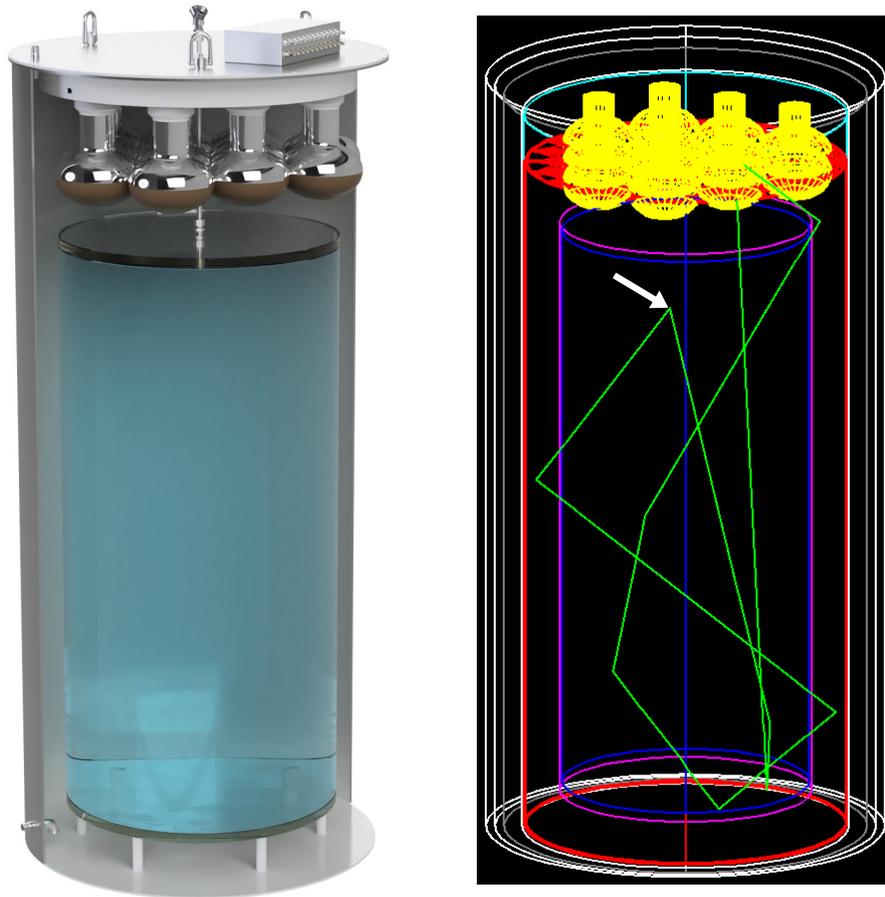


Figure 5.7: The chosen geometry described in Section 5.2.3, with a preliminary engineering diagram from Eric Day (CMU) on the left (reproduced from [156]), and a Geant4 screenshot illustrating a sample event on the right. The white arrow in the event snapshot indicates the position of the electron at one instant during the event, and the green lines represent the trajectory of example γ s emitted through Cherenkov radiation at that instant.

diagram is shown in Figure 5.7 left, with an example snapshot from an event shown in Figure 5.7 right. Due to the (relatively) low-cost of this design, the collaboration is ultimately pursuing two independent modules with similar geometry to pursue the flux normalization, so the additional studies we'll describe using this geometry will also be useful to pursue interesting design options for the second module.

5.2.4 No-D₂O Tank Geometry

Due to a supply-chain delay as a result of the COVID-19 pandemic and the immense benefit of commissioning the electronics and taking data prior to the SNS beam-energy changes in 2023 [89], we are currently moving forward with the deployment of a detector without an acrylic vessel (and thus, no deuterium). A Geant4 screenshot of the detector is shown in the bottom right panel of Figure 5.4. This geometry will be used for a commissioning run of the electronics, and we run simulations of this geometry to inform an initial measurement of the CC-O cross-section (our dominant background) which can inform supernova detection in water-based detectors using this cross section [24].

5.3 Simulation Physics

Similar to the neutrino-flux simulations described in Chapter 4, this simulation is written in Geant4. However, it uses a modular physics list which pulls selected interaction models from several different options rather than implementing a standard combination of interaction models (such as QGSP_BERT). In this way, we allow Geant4 to track all particles, but we prevent it from modeling every interaction of every particle to improve the computational efficiency. Electromagnetic interactions are handled with the standard Geant4 implementation, and we also add in the “extra” electromagnetic interactions, including synchrotron radiation and electro- and photo-nuclear processes. Particle-decay processes are always simulated, but ion-decay processes are generally ignored due to run-time errors; we do model ion-decay processes for some sources, such as for studying the effects of neutron capture de-excitation. The hadronic interactions are built using the binary cascade model, and can use the high-precision neutron models when appropriate. Finally, we model the production of optical photons for Cherenkov processes, and allow for Rayleigh scattering in addition to absorption and boundary processes. We prevent the simulation from generating more than 300 optical photons at any given step, and we prevent the production of particles which do not have enough energy to travel more than 10 μm in the given material. With these conditions, we use the UNIFIED model [170] to simulate the surface processes (reflection, refraction, absorption) of optical photons through our multiple volumes; we will extensively discuss the optical model and our description of the detector optics in Section 5.6.3.

With these processes initialized, we reduce the computational load by simulating electrons within the D₂O and H₂O volumes, and convolving the results with the theoretical

cross sections for CC-D and CC-O neutrino interactions to obtain a measure of our expected signal. Before convolution, our simulated signal source is uniform in energy from 0 to 60 MeV, and emitted isotropically from positions uniformly distributed throughout both the D₂O and H₂O volumes. The SNS flux-weighted doubly-differential cross sections of ν_e+d and ν_e+O interactions were generated by Jes Koros and Kate Scholberg respectively using SNOwGLoBES [171], and we convolve these probabilities of a neutrino generating an electron with a given energy and angle with the simulated response to such an event.

Since the simulation manages the creation and tracking of optical photons, we apply the quantum efficiency (QE) of our PMTs, or the probability that the photon will create a PE in the PMT and thus be detectable, at the step which generates the optical photon. This saves computation time because we do not waste cycles following the track through many reflections/materials only to ignore it because of the QE cut, and applying the cut at this early stage does not lose accuracy because we do not use any wavelength-shifting materials in the detector.

We implement the QE of our PMTs as a wavelength-dependent array taken from [172], and plotted in Figure 5.8 (left). If an optical photon is outside the wavelength range, it is immediately discarded. If the photon passes the QE cut for its wavelength, we track its path through the detector volume and record the energy and timestamp of the photon if its trajectory enters a PMT volume. An optical photon that hits a PMT is treated as equivalent to an observed photoelectron. The number of PE in each PMT, along with information about the initial characteristics (energy, position, direction, etc.) of the simulated electron, are recorded in output for later analysis. As an example, the Cherenkov spectrum for 30 MeV electrons traveling through D₂O is plotted in Figure 5.8 (right), showing both the emission spectrum and the expected response of a PMT.

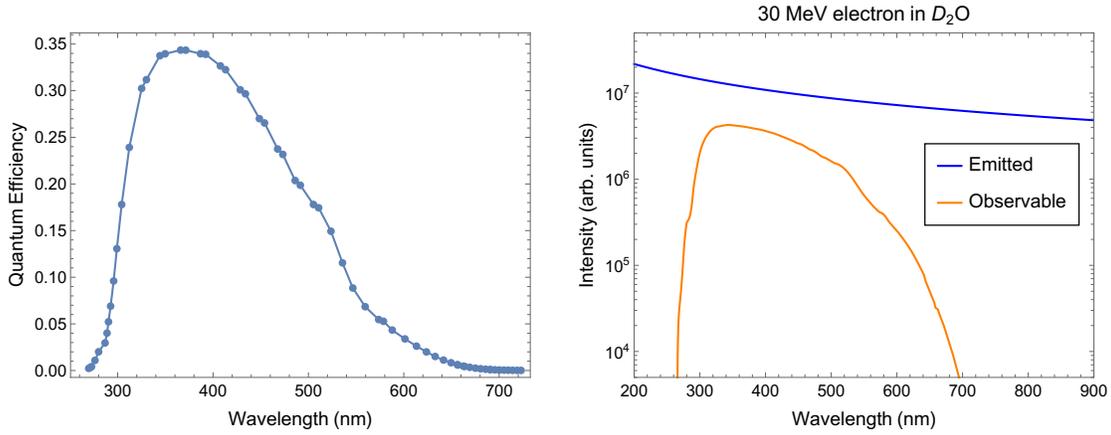


Figure 5.8: Left: A plot of the wavelength-dependent QE we apply in simulation. Right: A plot convolving the Cherenkov spectrum for a 30 MeV electron traveling through pure D₂O with the QE from our PMTs to compare the emitted radiation (blue) to the observable radiation (orange) in our detector.

5.4 Expected Response to Signal and Background

In this section, we simulate each of the following sources to inform our separation of signal and background events:

- CC-D (signal): incident ν_e will interact with deuterium to produce electrons in the tens-of-MeV inside D₂O volume (see Equation 5.1)
- CC-O (background): incident ν_e will interact with oxygen to produce electrons in the tens-of-MeV inside D₂O and H₂O volumes (see Equation 5.3)
- BRNs (background): incident neutrons will downscatter in the detector surroundings and capture on nuclei in the detector volumes (ex: H, O), producing a shower of γ s.
- Cosmics (background): incident high-energy cosmic rays can produce showers of γ s along with neutrons and muons
- Michel electrons (calibration): incident muons can come to rest in the detector, decay, and produce electrons inside the D₂O and H₂O volumes

With simulation, we model the response to each source separately, then normalize the distributions by the expected event rate to make predictions of the signal our detector will observe.

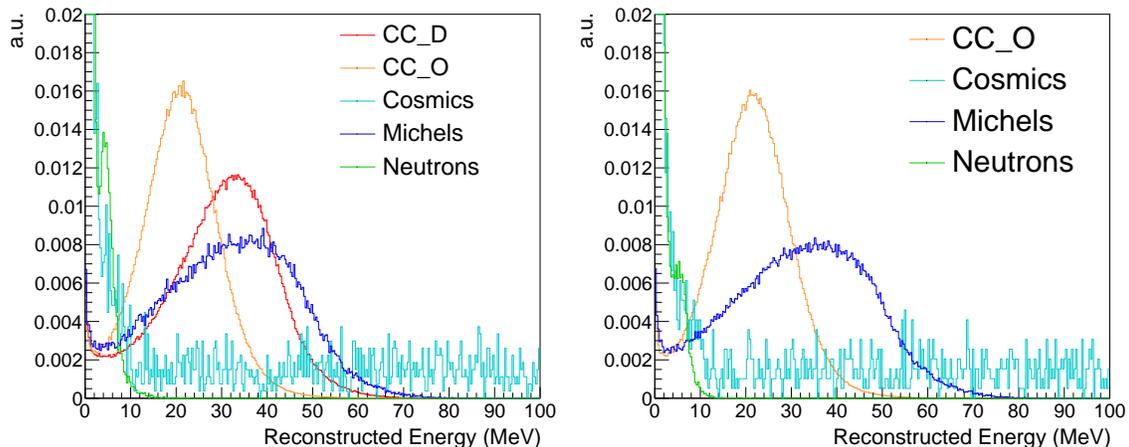


Figure 5.9: Energy distributions for events from simulated sources using the chosen Tank geometry of the D₂O detector (left) and using the modified commissioning geometry without the acrylic vessel (right). The integral for each curve is normalized to unity for $E_{\text{rec}} \in (0, 100)$ MeV to first examine the differences in spectral response (we examine the relative normalizations in Figure 5.10). Note the distinct similarities in the distributions for each geometry, which illustrate our discussion that the acrylic vessel and optical differences between D₂O and H₂O can be (largely) ignored. This also establishes that data collected with the commissioning run can be used to directly study backgrounds *in situ* for the full detector.

In the cases of CC-D and CC-O, if we have an incident ν_e flux $\phi_e[\text{cm}^{-2}\text{s}^{-1}]$, a total cross section for the specific interaction $\sigma_i[\text{cm}^2]$, a volume with N_t possible interaction targets, and a detector efficiency $\epsilon = 0.97$ to account for expected dead-time from the veto, the expected event rate (R_i in counts / s) can be defined as

$$R_i = \phi_e N_t \sigma_i \epsilon. \quad (5.5)$$

Then, the distributions we model with the simulations will each integrate to the expected event rate for that source, such that we can also explore the relative contributions to our total signal for given time periods. This strategy does not account for simultaneous different interactions, but we do not anticipate significant pileup.

For both the signal and background ν_e interactions, we simulate electrons uniform in energy, direction, and position, then convolve the results with the calculated cross-section to model the events that we would see in Neutrino Alley². We do not allow events to be

²In this way, we simulate both D₂O and H₂O events simultaneously; we apply a weighting in analysis that saves us from needing to re-simulate statistically significant results for both types of interactions.

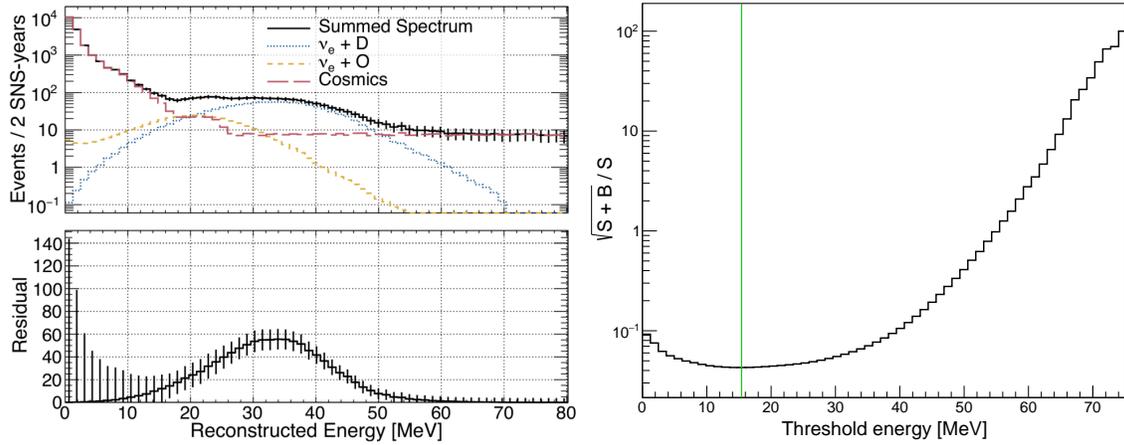


Figure 5.10: Left: The expected signal from the Tank geometry after 2 years of beam-on at the SNS (total 10000 hours of runtime), excluding the neutron contributions (which are expected to be quite small relative to the cosmics component). The bottom panel shows the background-subtracted CC-D spectrum with statistical errors. Figure reproduced from [156]. Right: The statistical uncertainty of the signal excess using a set threshold, given by Equation 5.7, for the Tank geometry after 2 SNS-years. This analysis demonstrates that a statistical uncertainty of 4.2% can be achieved using a threshold of 15.4 MeV, marked with the vertical green line.

simulated within the acrylic vessel, because the cross sections for carbon interactions are negligible as shown in Figure ?? right. Our signal, the electrons from $\nu_e + d$ events, peaks at around 35 MeV, and the majority of the signal spectrum is between 20 and 40 MeV. The background $\nu_e + O$ events peak lower in energy – nearer to 20 MeV, with the majority of the spectrum between 10 and 30 MeV. By applying a threshold in the analysis, we will increase the sensitivity of our measurements despite sacrificing some of the signal statistics.

For the cosmic-ray muon background, we simulate cosmic rays through the Neutrino Alley overburden using the Cosmic-ray shower generator (CRY) module for Monte Carlo transport codes [173]. With these simulations, we study the interactions of cosmic particles (n, e^-, μ, γ) generated above the Neutrino Alley overburden with energy and direction distributed by CRY using the geographic location of the SNS and an arbitrarily chosen date (July 1, 2019). We then observe the response of the detector to the cosmics, as well as to the showers of secondary particles the cosmics can create in the concrete surroundings of our detector. We then use the built-in elapsed time from the CRY module to determine an expected number of simulated events for a fixed time interval to inform

our background model.

Muons pose a unique problem for us, as muons will lose energy through Cherenkov radiation. Some of them come to rest in the detector, and ultimately decay at rest into electrons that will also leave an observable energy deposit according to the Michel spectrum. We will deploy muon veto panels to reduce the background of muons passing through the detector, and the background we consider results from “untagged” muons, or background events which are not detected by the veto panels. The physics model in Geant4 which controls the decay of a muon into an electron causes run-time errors for our simulation when used with the CRY module. As no solution has yet been identified, we separately model the detector’s response to Michel electrons by sampling from the well-known energy distribution

$$\frac{d\Gamma}{dx} = \frac{G_F^2 m_\mu^5}{192\pi^3} x^2 \left(6 - 4x + \frac{\alpha}{\pi} f(x) \right), \quad x = \frac{2E_e}{m_\mu}, \quad 0 < x \leq 1, \quad (5.6)$$

where E_e is the energy of the electron, m_μ is the mass of the muon, $\alpha \approx 1/137$ is the fine structure constant, and $f(x)$ represents radiative corrections, to understand the detector response to a very useful, steady-state calibration source. We observe in Figure 5.9 that a simple threshold will also serve to reduce the impact of cosmic-ray backgrounds on our main signal.

Finally, we simulate beam-related neutrons using the spectral model developed from measurements for the CsI detector [85]. This D₂O detector will be deployed at the same location in Neutrino Alley, albeit with different shielding requirements since this detector has very little sensitivity to small nuclear recoils. In [85], an incident neutron-energy distribution (in MeV) following the power law E^α was used, where $\alpha = 1.6$ was observed. We used this as our simulation input, and modeled the interactions of neutrons up to 400 MeV as they travel through the concrete walls of Neutrino Alley, through the lead shielding and veto panels of our D₂O detector, and finally through the sensitive volume of our detector. As shown in Figure 5.9, the neutron spectra will be almost completely below any threshold used to separate the signal/background electron events.

With these rates, we plot our expected signal in Figure 5.10 (left) and use this to determine an energy cut which optimizes the precision of our rate estimate, or minimizes the Poisson uncertainty on the residual spectrum. The statistical precision is calculated from the uncertainties on the signal S (CC-D) and background B (CC-O + Cosmics) counts above a given threshold E . As an example, $S(E) = \int_E^{75 \text{ MeV}} R_S(E') dE'$, where E is the threshold, and $R_S(E')$ is the measured signal with reconstructed energy E' . We estimate

Table 5.1: Expected event rates using Equation 5.5 for CC-O and CC-D and the CRY module normalization for untagged cosmic-ray events. We also investigate the expected event rates in the region of interest $E_{\text{rec}} \in [22, 68]$ MeV. We neglect contributions from Michel electrons (because their rate and spectra are well-known) and neutron backgrounds (because events are shown to be below threshold). Reproduced from [156].

Process	Total Events	Events Passing Cuts
CC-D	1070	910
CC-O	390	160
Cosmics	21150	315

$B(E)$ similarly. Then,

$$s(E) = \frac{\sqrt{\sigma_S^2 + \sigma_B^2}}{S} = \frac{\sqrt{S(E) + B(E)}}{S(E)}. \quad (5.7)$$

We plot $s(E)$ in Figure 5.10 (right), and identify an uncertainty of 4.2% at 15.4 MeV, meaning that using a threshold of 15.4 MeV will achieve 4.2% statistical precision in 2 SNS-years (10000 hours of 1.4 MW operations) using the Tank geometry. The value of the analytical threshold and precision listed here are different from the predictions in [156] (4.7% statistical precision in 2 SNS-years using a 22 MeV threshold); ultimately, the threshold will be determined by our analysis with the commissioning geometry and a preliminary characterization of the CC-O cross section.

The rate estimates are shown in Table 5.1 using the threshold used in [156], calculated for CC-O and CC-D using Equation 5.5 and for cosmic-ray events using the elapsed time provided by the CRY module. We exclude any estimation of BRN rates since all events will be below the energy range of interest³, and we also exclude the estimation of Michel electron rates since we will be able to fully characterize their spectra and rates in the anti-coincident periods.

5.5 Fiducialization Analyses

One of the most notable problems with the chosen tank design is the lack of photocathode coverage; we use only 12 PMTs. We depend on the energy resolution, or our ability to observe all of the energy deposited from a given electron and correctly infer the initial

³In addition, the flux and spectra are not yet well-known (see Chapter 6).

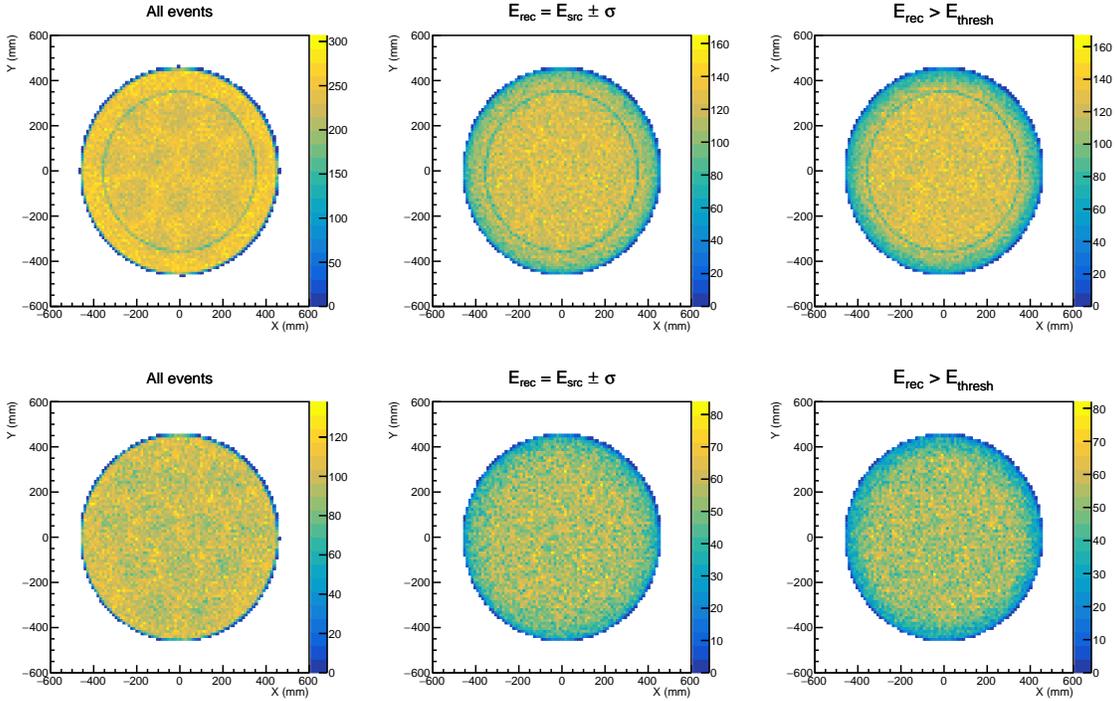


Figure 5.11: Distributions of the position of simulated electron events projected as if looking down from the top of the detector for the full D₂O detector (top) and the version without the acrylic vessel (bottom). Left: Every simulated event; note the slight discoloration which indicates the locations of the PMTs, and the discolored ring in the Tank geometries (top) indicating the location of the acrylic vessel. Center: Events which can be reconstructed within one Gaussian sigma of the perfectly reconstructed energy deposit. Right: Events which fall above a threshold set at 15.6 MeV, or 340 photoelectrons.

electron energy from this observation, and less photocathode coverage means less ability to reconstruct the full energy deposited. Because the PMTs are all at the top of the detector, we need to understand the impact the geometry will have on the observations. Ultimately, this is an exploration of understanding the tails of the light-yield distribution we illustrated in Figure 5.3 left, which have not been well-accounted for in our estimation of the energy resolution using a Gaussian fit, and identifying possible strategies to reduce poorly-reconstructed events in our analysis.

In Figure 5.11, we show the projection of the position distribution for simulated electron events (both CC-D and CC-O) onto the XY plane, corresponding to a top-down view of the detector volume. We examine, in the different panels, the effects of two different cuts compared to the results with no cut. Both cuts we use make an effort to isolate a

region of the detector for which the reconstructed energy of an event falls within an acceptable range around the truth-value E_{src} available in simulation.

The first strategy uses all truth values from the simulation, and demands that the energy reconstruct within one standard deviation of the true electron energy ($E_{\text{rec}} = E_{\text{src}} \pm \sigma(E_{\text{src}})$), where σ is calculated from the best-fit of Equation 5.4 as in Figure 5.3 (right). If the response of the detector to a simulated event passes that cut, the position of that event is included in the histogram shown in the center panels. The application of this cut causes an observable increased reduction of counts near the edges of the detector. This is fully expected, as electrons created near the edges can more easily escape the detector and leave lower energy deposits.

The second strategy attempts to mimic the reduction of counts near the edges with the use of a simple threshold (E_{thresh}) that could reasonably be applied to data. If $E_{\text{rec}} > E_{\text{thresh}}$, the position of that event is included in the histogram shown in the right panels of Figure 5.11. The impact of this cut depends on the value assigned to E_{thresh} , which is illustrated in Figure 5.12. Here, we examine the efficiency ϵ of the threshold cut as it relates to the D₂O and H₂O volumes (assuming the Tank geometry). The efficiency is calculated as

$$\epsilon(E_{\text{thresh}}) = 1 - \frac{\int_{E_{\text{thresh}}}^{55} p(E)dE}{\int_0^{55} p(E)dE}, \quad (5.8)$$

where E is the reconstructed energy of the event and $p(E)$ is the energy-dependent light-yield distribution for electron events uniformly distributed from 25 to 45 MeV and isotropically emitted from uniformly distributed locations in the selected volume. We use this efficiency to determine the threshold setting, and aim to maximize the efficiency of the threshold cut for the H₂O volume while minimizing the same efficiency for events in the D₂O volume – this is to say, we want to cut out H₂O events without sacrificing sensitivity to the D₂O volume.

Using Figure 5.12, we identify the value $E_{\text{thresh}} = 15.6$ MeV as the highest threshold setting we can set without cutting more than 5% of events from the D₂O volume, while providing the largest relative loss of H₂O events. The bottom panel of Figure 5.12 illustrates that this also nearly maximizes the separation of the efficiencies; the true maximum occurs at $E_{\text{thresh}} \approx 18$ MeV. As we discussed in Section 5.4, the analysis produces the best statistical precision (4.2%) at $E_{\text{thresh}} \approx 15.4$ MeV, which improves the ratio of CC-D events to CC-O events.

We now aim to quantify the fiducial volume, or the fraction of the geometry within which we can accurately reconstruct the electron energy. In the sections that follow, we

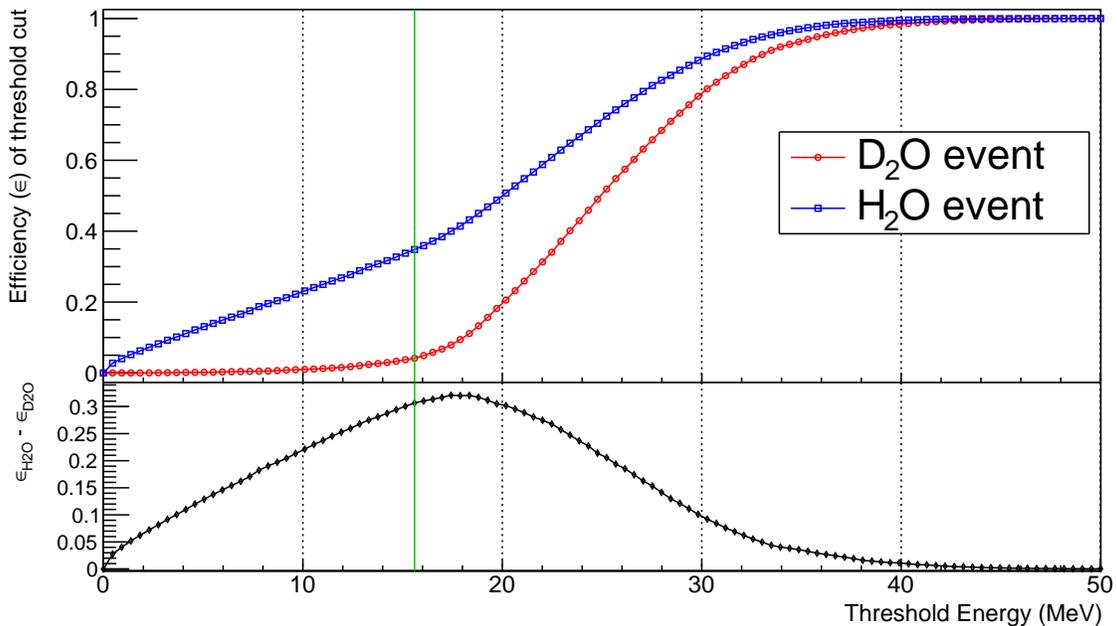


Figure 5.12: Top: A plot of the efficiency of the threshold cut, calculated from Equation 5.8, for electron events sampled from 25 to 45 MeV and uniformly populating the D₂O and H₂O volumes. Bottom: A measure of the difference in the cut efficiency for the two volumes. The green line superimposed onto the plots illustrates our chosen threshold value, $E_{\text{thresh}} = 15.6$ MeV.

quantify the fiducial volume for the Tank geometry and the soon-to-be-deployed No-D₂O commissioning geometry, again noting that they should be consistent but will display different features due to the inclusion of the acrylic vessel in the Tank geometry. We separately investigate two dimensions of these cylindrical geometries: the distance from the central axis (Section 5.5.1) and the distance from the PMTs (Section 5.5.2).

5.5.1 Radial Fiducialization

To quantify the portion of the volume which responds in an easily characterizeable way, we calculate the distance from the detector’s central axis and plot the information from Figure 5.11 in a more compact way in the left panel of Figure 5.13. The content of each bin is scaled by the total area that would be associated with a ring of the given radius and a thickness consistent with the bin width. In this way, we observe a flat distribution when we have a uniform population of events in this axial dimension, but deviations from that flatness when we start to lose events in those bins. This gives us a way to identify the

fiducial volume of a geometry, which represents the fraction of the full detector which cleanly reconstructs the initial energy of the electron; we search for the first bin with a greater than 15% deviation away from a baseline characterized by the average of the first 20 cm. The vertical lines in Figure 5.13 (left) mark this location, which is consistent with the size of the acrylic vessel.

5.5.2 Vertical Fiducialization

With the axial volume quantified, we also want to explore cuts in the vertical dimension. We observe poorer reconstruction for events generated closest to the PMTs⁴, and our threshold cut will eliminate events which might have occurred behind the PMTs. Here, though, we begin to observe true differences in using simulation truth and potential experimental cuts. Simulation suggests that the energy reconstruction is best at the bottom of the detector, but events closest to the PMTs have a higher likelihood to leave a larger energy deposit and pass the threshold cut.

These results show a similar fall-off to the axial plots at the very bottom of the detector, where the tail catcher begins, demonstrating that the oxygen contributions will again be reduced with the simple energy cut. However, we do not observe a definitively “good” reconstruction region for the vertical dimension. Even using all truth-values available in simulation, we find that there is a significant correlation between vertical position and effective event reconstruction. Notably, events near the bottom of the detector are more likely to be accurately reconstructed than events near the PMTs; we believe this is due to the angular distribution of the simulated events. The threshold cut, however, does not mimic this behavior. Events closest to the PMTs are the most likely to have high light collection and will generally pass the threshold. More work is needed to fully quantify a fiducial cut and understand how the vertical position of an event will affect the reconstruction. A planned LED flasher system may help us to untangle these effects in the commissioning geometry, as this will produce calibrated events at well-defined positions within the detector.

Based on Figure 5.13, we find that for the Tank and No-D₂O geometries, the fiducial volume associated with an analysis cut $E_{\text{thresh}} = 15.6$ MeV is roughly consistent with the volume contained within the acrylic vessel of the Tank geometry. We can then estimate the fiducial volume as a 70 cm $\varnothing \times 1.5$ m vertical cylinder coaxial to the full detector, starting just below the PMT locations and extending to the floor.

⁴This effect is still under investigation, but is believed to be an effect of the angular distribution of events.

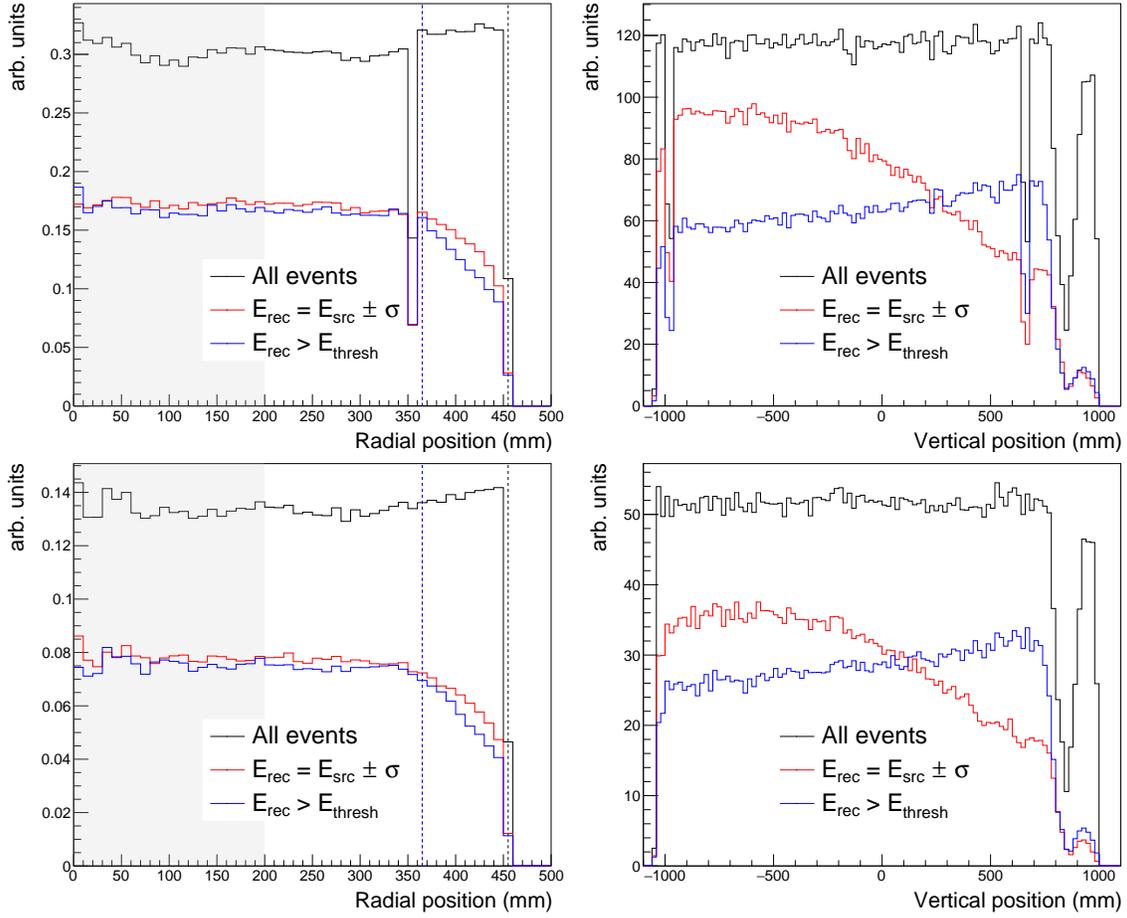


Figure 5.13: Event position distributions which pass cuts for the Tank geometry (top) and the commissioning No-D₂O geometry (bottom). In all plots, $E_{\text{thresh}} = 15.6$ MeV. Left: Distribution of the distance from the central axis for simulated electron events. The vertical lines indicate the positions at which the applied cut reduces the events by more than 15%. The dip in the top-left plot at 350 mm indicates the location of the acrylic vessel for the D₂O geometry, where events are not generated; we ignore these bins in the determination of the fiducial volume since they will have artificially lower event rates. Right: Distribution of the vertical position of simulated electron events. The valley at 900 mm indicates the location of the PMTs in both geometries; the two dips in the top-right plot at roughly -950 and $+650$ mm indicate the location of the acrylic vessel.

5.6 Optical Measurements

The optical properties we include in simulation can greatly impact the detector performance. For example, Figure 5.14 illustrates the change in overall light collection and en-

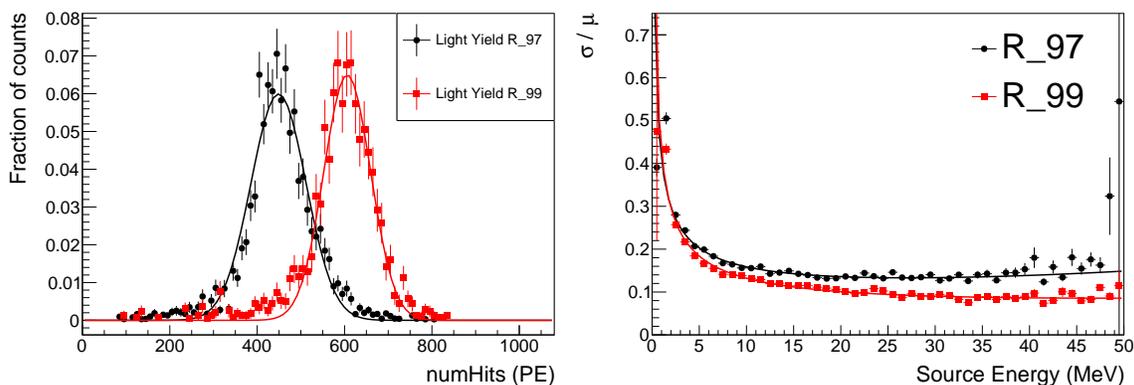


Figure 5.14: Left: The raw light collection for 30 MeV electrons with total reflectivity R set to the indicated amount in percentage. Right: A comparison of the fractional energy resolution for the Tank geometry indicating that higher reflectivities will have better energy reconstruction.

ergy resolution associated with different reflectances.

Given the significance of these changes, we will need to match these optical properties as closely as possible in simulation to make the best predictions. In the discussion that follows, I overview a series of optical measurements and an analysis methodology which should be repeated with the final choice of reflective coating, and with water-soaked coating to understand the change of properties during the detector deployment.

These measurements were conducted in the ORNL Physics Division’s Chemistry Support Laboratory and were only possible with the help, guidance, and forethought of Mike Febrarro and Brennan Hackett. Both of the measurements we detail utilized a laser to emit a focused beam of light at a specified wavelength, and a silicon photodiode [174, 175] to measure the wavelength-dependent intensity (or flux) of detected light. This work was completed with the laser wavelength at 516 nm — a visible, bright green laser. In future, we’ll want to understand the optical properties of our reflective surface at wavelengths appropriate to the conditions of D_2O (from Figure 5.8, our QE maximizes near 366 nm). There are two main optical properties we care about in our Geant4 simulation: the total reflectance of a material and the specularity, which is a measure of the angular distribution of light reflected from a material that can depend on properties such as the roughness of a surface.

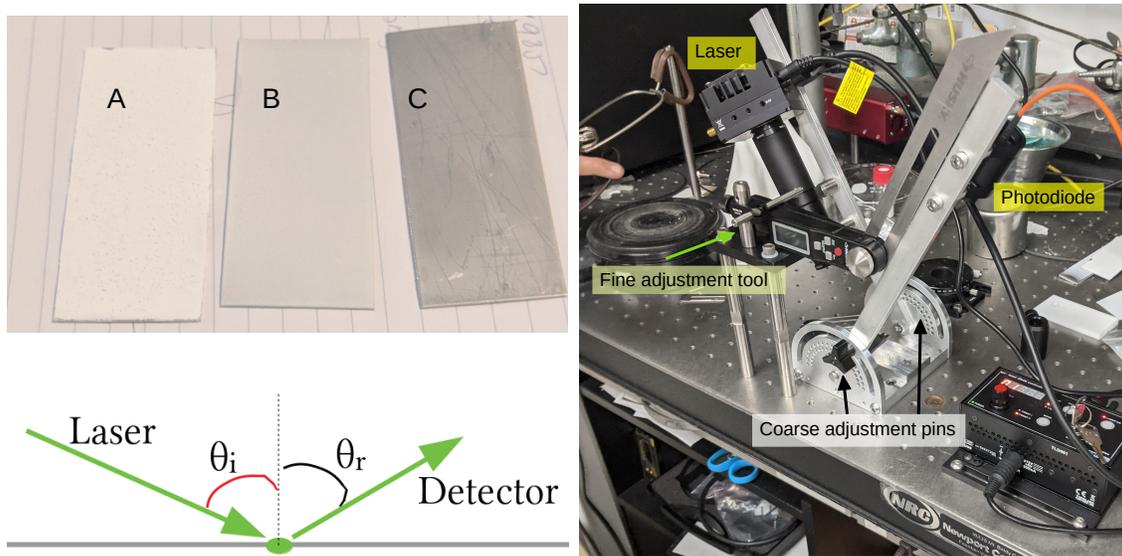


Figure 5.15: Top Left: A color-corrected photo of three test samples for the reflectivity studies outlined in Sections 5.6.2 and 5.6.3. Sample A is a barium sulphate paint, Sample B is Performix “PLASTI DIP”, and Sample C is an unpolished stainless steel sample. Bottom Left: An illustration of the tool shown in the photo to the right identifying the relevant laser angle (θ_i) and detector angle (θ_r). Right: The optical tool described in Section 5.6.3 for measuring the surface roughness of a sample. The laser is affixed to the left/back leg, and the photodiode (with the orange cabling) is affixed to the right/front leg. The front/back distinction of the legs refers only to the pin the user would move in order to move the given leg, as the laser and photodiode are aligned on-axis. Both legs can be coarsely adjusted in 5 degree increments by moving the pins shown on the bottom of the tool, and the laser can be more finely adjusted using the attachment shown at the center of the tool.

5.6.1 Material Samples

Our D₂O tank will feature multiple layers of Tyvek tape glued to the entire interior surface area of the stainless steel outer tank. As such, we want to test the optical properties of these reflective coatings adhered to stainless steel. ORNL staff machined thin 1×2 in² “coupons” for us, as shown in the top left panel of Figure 5.15, that could be painted or covered with tape to serve as our test surfaces. We study three samples in this work (labeled as they appear in Figure 5.15): A) two coats of 6080 barium sulphate (BaSO₄) white reflective paint (diluted and sprayed onto the surface), B) two coats of Performix brand “PLASTI DIP” multi-purpose rubber coating and C) a “control” coupon which we

did not modify from the machined steel.

The samples we study here are *not* planned for use in the D₂O detector due to the high cost (A) or insufficient performance (B, C); instead, we explore the available resources and develop a procedure that can be used in the future to both measure and validate the optical properties we use in simulation, largely following the discussion laid out in [176].

5.6.2 Total Reflectance

We make use of a four-port ThorLabs integrating sphere [177] to measure the total reflectance of our sample. The photodiode and laser are affixed to the top and rear of the integrating sphere, respectively. The side port uses the provided port plug, which is coated with the same PTFE as the sphere itself. The sample is then affixed to the open front port, such that the laser light will reflect off the sample and within the integrating sphere until detection by the photodiode.

As a calibration, we use a provided PTFE port plug to measure the intensity recorded by the integrating sphere. The port plug has a known reflectance of $R_0 = 99\%$ [177]. The average intensity the photodiode observes over a 2-second period then represents the measurement for a sample with total reflectivity of 99% for the incident wavelength. Deviations above or below this calibration measurement then indicate material reflectivity above or below 99%, quantified as

$$R = \left(\frac{I - I_0}{I_0} + 1 \right) R_0 = \left(\frac{I}{I_0} \right) R_0, \quad (5.9)$$

where I is the observed intensity for the sample, I_0 is the observed intensity from the calibration, and R_0 is the total reflectance of PTFE. This is a wavelength-dependent quantity, and we recommend performing the Tyvek measurement for several different incident laser wavelengths to fully characterize the reflectance we input to simulation. As a procedural test for $\lambda = 516$ nm, the total reflectance for 516 nm light incident on Sample A was near 97%.

5.6.3 Specularity

While the integrating sphere tells us the total reflectivity of a material, we must also measure the angular distribution of reflected light. For the D₂O detector, we plan to use a material with diffuse reflectivity; light will bounce around at all angles, eventually making its way to the PMTs at the top of the geometry. Diffuse reflectors are generally characterized

with a cosine-like behavior described by Lambert’s law

$$I = I_0 \cos \theta, \tag{5.10}$$

where θ describes the angle of reflected light relative to the surface normal, I_0 is the incident intensity of light, and I is the reflected intensity of light at the given outward angle. Specular reflectors, on the other hand, are defined by their mirror-like reflectance; the angle of reflectance is dependent on the incident angle. Our Geant4 simulation uses the UNIFIED model, which allows for both diffuse and specular reflectance based on the user-defined parameters for a given material and surface. Figure 5.16 demonstrates the possibilities for reflection, and the user must define the probability for each type of reflection. In order to accurately represent the material of interest, we must measure these properties for our specific implementation.

In cases where the microscopic roughness of a surface is not perfectly polished, such that the normal for any microfacet is different from the average surface normal, the specularity of a material is defined by two separate components: the specular spike and the specular lobe. Specular spike reflectance is the specular angle of scatter relative to the surface normal; it has no dependence on any microfacets. Specular lobe reflectance is the specular angle of reflection relative to the facet normal; it will depend on the specific facet, and it will have a broader distribution in angle as a result. We illustrate both types in Figure 5.17, and we use this modified version of Figure 5.15 (bottom left) to define the angle between facet normal and surface normal as α . The distribution of α s for a surface is the standard metric to quantify the roughness of a surface; the parameter σ_α (with units of radians) is defined to represent the standard deviation of a Gaussian distribution describing the angles of rough facets relative to the surface normal. This parameter can then be measured experimentally and applied in simulation.

The tool photographed in Figure 5.15 was used to directly measure the intensity of light reflected from the sample to a particular angle, and thus measure the specularity and surface roughness of a material. Using the ThorLabs photodiode, we measure the background-subtracted spectral intensity across the full wavelength range – an example result is shown in Figure 5.18 for intensities averaged over a 10-second window. The units on the recorded intensity are treated as arbitrary; the photodiode is expected to be linear, so our measurements only depend on the relative differences which do not depend on units. The uncertainties on the observed intensity in each wavelength bin were not recorded, but are known to be small; we estimate a flat uncertainty of 0.001 [arb. units] to the spectral intensities to account for the baseline fluctuations of the measured intensity

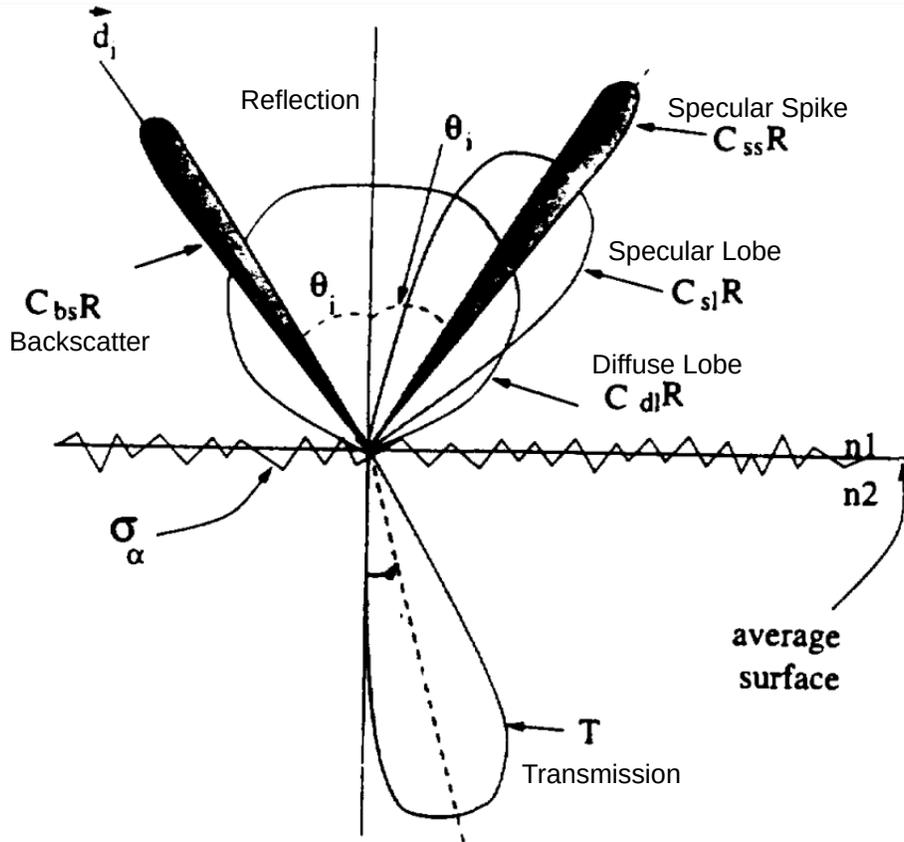


Figure 5.16: A diagram of the UNIFIED model's approach to modeling surface reflection, reproduced from [170]. In Geant4, the user specifies both the reflection and transmission probabilities of a surface, but also the coefficients C_X which describe the four possible reflections of incident light \vec{d}_i : backscatter, specular spike, specular lobe, and diffuse lobe. In our Geant4 model of the D_2O detector, we choose $C_{dl} = 1$ and all other $C_X = 0$.

from 800 to 1000 nm which are expected to be consistent with zero intensity⁵. In future, we recommend taking multiple time-averaged runs at each incident and outgoing angle, and using the spread of the observed measurements to characterize the uncertainties on the spectral intensities.

At each angular position combination of the incident laser and observing photodiode, we fit a Gaussian to the peak in the observed intensities near the incident laser wavelength.

⁵The 800-1000 nm region was chosen to be away from the fit region and attempted to avoid visible-light backgrounds from the next room.

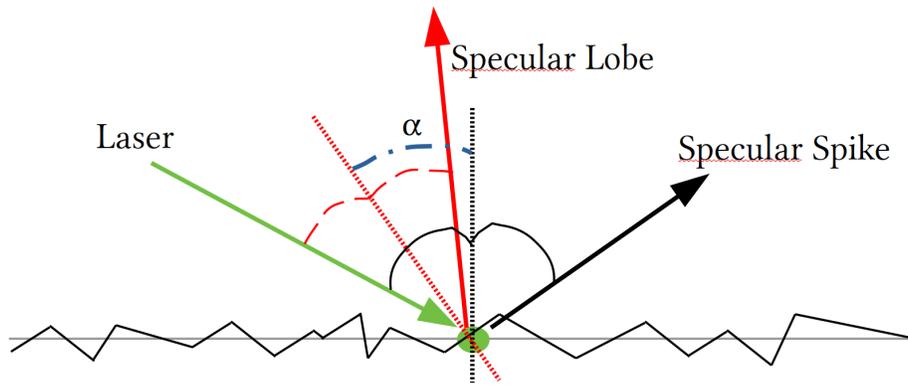


Figure 5.17: An illustration of the different types of specular reflectance for incident laser light shown as the green arrow. The black arrow representing specular spike reflectance ignores the roughness of a surface; light will be reflected at an equal angle relative to the average surface normal (shown as the dotted black line). The red arrow representing specular lobe reflectance will depend on the roughness of a surface. In this example, the laser is incident on a microfacet with a normal vector that deviates from the average surface normal by an angle α . The specular lobe reflectance will reflect incident light at an equal angle relative to the microfacet normal (shown as the dotted red line).

The amplitude of the Gaussian is a measure of the observed intensity for this angle, and the uncertainty on the fit parameter is carried into the next step. We restrict these fits to the range 500 - 530 nm to prevent the fitter from being impacted by random backgrounds; an example fit is also shown in Figure 5.18.

For a given incident laser angle $\theta_i = 30^\circ$, we examine the spectral intensity at increasing photodiode angle θ_r in the left panel of Figure 5.19. We observe that Sample A has primarily diffuse reflectance, while Sample B has apparent contributions from diffuse (cosine-like behavior for 0-10 and 45-60 degrees), specular spike (large increase in intensity at 30 degrees), and specular lobe reflectance (symmetric and non-cosine behavior for angles above and below 30 degrees). To further characterize the relative contributions, we plot the Gaussian amplitude from each combination of incident and detector angles that we measured in Figure 5.19 right. We fit this data with a Gaussian restricted to $\mu = 0$, and find more than a 20° spread, indicating that the specular lobe reflectivity dominates over the specular spike. However, this strategy does not fully describe the data, as we have likely underestimated the errors (see χ^2 / ndf in Figure 5.19). An unconstrained Gaussian ($\mu \neq 0$) can better describe the data, but it does not have a clear physical interpretation. It is also worth noting that during these tests, a misalignment of the laser and photodi-

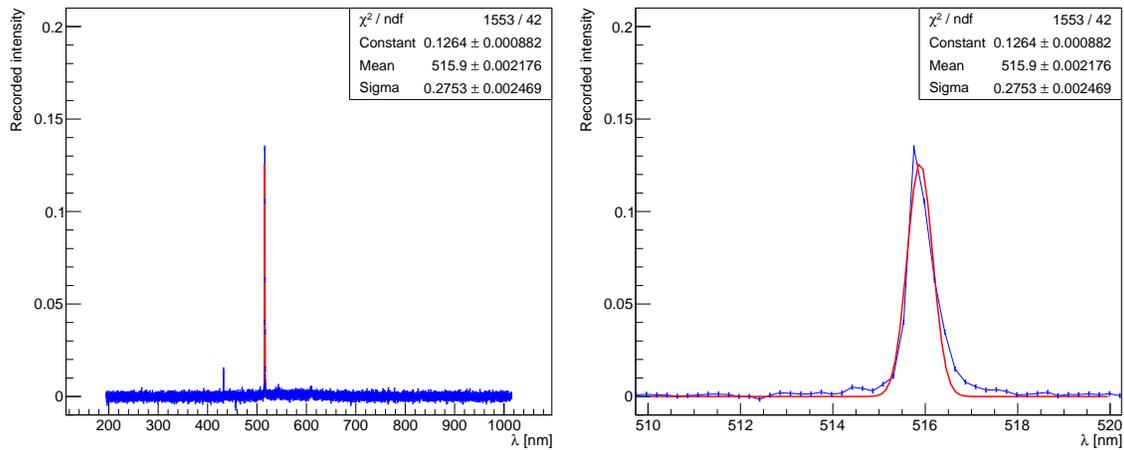


Figure 5.18: An example spectral intensity measurement using the tool for measuring the surface roughness, showing the full observation (left) and a zoomed-in region around the laser wavelength (right). In this example, the incident angle (laser leg) and the reflected angle (photodiode leg) were both 40 degrees. In both plots, the blue points represent the intensity recorded by the photodiode (in arbitrary units with artificial uncertainty), while the red line illustrates the best-fit Gaussian.

ode was uncovered; the shift could be a result of this issue. For any measurements of the specularly and surface roughness for our final choice of reflector, we recommend dedicated runs at finer-binned angles and at multiple locations on the surface to improve our understanding of this behavior.

During our testing, there was a misalignment in the laser-photodiode axis. The focused laser-light would reflect off the sample, but the photodiode was shifted slightly off-axis from the focus point. To compensate, we used a slightly larger / unfocused laser, such that there was a 2 mm disk on the sample where the incident light could reflect. This compensated for the misalignment, but we recommend future testing of how focusing the laser can impact the extraction of the surface roughness. We would also recommend repeating measurements at different locations on the sample to test for uniformity of the surface roughness.

Both the total reflectance and specularly measurements we describe were done to develop a procedure to collect and analyze the data for these optical measurements. We used test samples to work through the procedure, and identified several avenues for improvement in our eventual characterization of the final choice for our D₂O detector's reflective coating. Tests with water-soaked materials are now in development by COHERENT col-

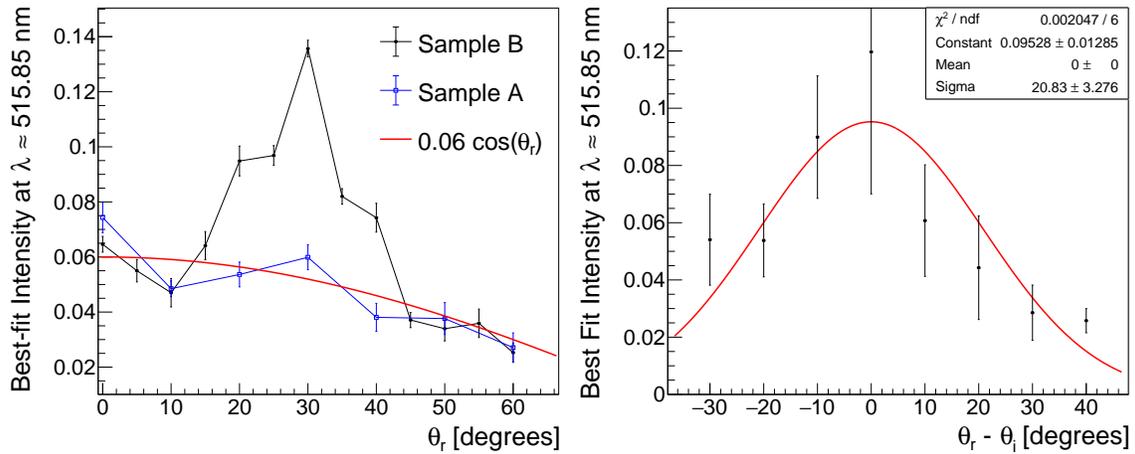


Figure 5.19: Left: The amplitudes of the best-fit Gaussians to the measured spectral intensity (in arbitrary units) for an incident laser at $\theta_i = 30^\circ$ over varied photodiode angle for samples A and B. The red line is a cosine curve with amplitude selected to offer a visual cue of Lambert’s law — it is not a fit to the data. We observe clear excess near $\theta_r = 30^\circ$ in the measured intensity for the glossy Sample B, indicating that it is a more specular reflector. Sample A roughly follows the cosine behavior of a diffuse reflector. Right: The amplitude of the best-fit Gaussian to the measured spectral intensity (in arbitrary units) from each combination of incident and detector angles for Sample A. We average the amplitudes of like-combinations (e.g., (θ_i, θ_r) combinations of $(30^\circ, 10^\circ)$ and $(45^\circ, 25^\circ)$ both give $\theta_r - \theta_i = -20^\circ$) to better estimate the uncertainty on our observations. We fit this data with a Gaussian with fixed $\mu = 0^\circ$, such that the width of the fit describes the surface roughness $\sigma = \sigma_\alpha$, corresponding to spectral lobe reflectance. In future, finer-binned data for $\theta_r = \theta_i \pm 10^\circ$ can enable examination of the relative contributions of the specular spike and specular lobe.

laborators at ORNL, as we need to understand the impact of long-term submersion on the optical properties of the selected reflective coating.

5.6.4 Impact of Optical Studies on Simulation

With simulation, we can understand the effect of changes in the total reflectivity or the specularity of our reflective surfaces. In Figure 5.14, we observed that an increase from our baseline $R = 97\%$ to a total reflectivity of $R = 99\%$ leads to increased light collection; this makes sense, as higher reflection (lower absorption) means that more optical photons will make it to the PMTs. We also note the improved fractional energy resolution with higher

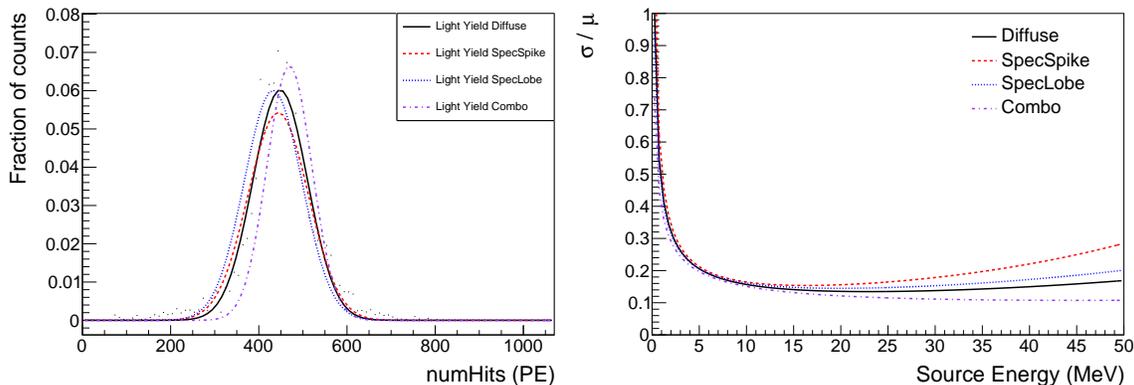


Figure 5.20: A comparison of the light yield for 30 MeV electrons (left) and fractional energy resolution (right) uniformly distributed throughout the Tank geometry, implemented with varied reflectance types. The Diffuse, SpecSpike, and SpecLobe labels indicate the relevant coefficient is 1 with others set to zero. The Combo label is a simulation in which $C_{ss} = C_{sl} = 0.33$ and $C_{dl} = 0.34$. Our baseline simulations implement the Diffuse curves.

overall light collection — a high-quality reflector can resolve signals of interest to better than 15%. However, more reflective surfaces can significantly increase financial costs for the detector, and a 97% reflectivity is observed to do well enough that these additional costs are not necessary.

Our planned reflective coating is multiple Tyvek layers glued to the inner surface of the stainless steel tank; this setup is known to be primarily diffuse reflectance. In our baseline simulations, we choose $C_{dl} = 1$ and all other $C_X = 0$. Nonzero specular components can change the performance of the detector, as shown in Figure 5.20. Since we’ve simulated a perfectly diffuse reflector, any changes to the surface roughness predictably have no effect on the detector performance. To summarize, accurate measurements of the optical properties of our chosen reflector using the strategies we discussed in this section will benefit the accuracy of our model of the detector response, but generally only the total reflectance is required to accurately model a detector using a diffuse reflector.

5.7 Conclusion

The normalization of the neutrino flux is the dominant uncertainty that stands in the way of a precision CEvNS program from the COHERENT experiment. We observed in Chapter 4 that simulation alone cannot reduce this uncertainty; it is imperative that new data be collected to better understand the neutrino flux from the SNS. Thin-target pion-

production data will greatly benefit our work, but a more direct normalization through the use of the D₂O detector we described in this chapter will account for the thick SNS target and be immediately useful to the COHERENT program. Our work in this section details several design studies which aim to improve the performance of the detector given the detector size and funding constraints we faced in the attempts to rapidly deploy the detector before the long shutdown of the SNS in 2023 for the planned beam upgrades.

Chapter 6

Neutrons in Neutrino Alley

While the previous chapters focused on improving our understanding of the SNS neutrino production, this is not the sole focus of my work with the COHERENT collaboration. As we mentioned in Chapter 2, the model of the expected detector response to both signal and background events is of utmost importance to interpret any experimental observation. Our efforts in understanding the neutrino flux aim to reduce the uncertainty of the signal model in our interpretation of CEvNS data. We now turn our attention to the background model, which is dominated by the neutron flux in Neutrino Alley.

The discussion in the chapters that follow pulls heavily from the discussions and efforts of the Backgrounds and MARS working groups within the COHERENT collaboration. In particular, the joint investigations of the MARS detector response and the monitoring of BRN backgrounds in Neutrino Alley reported in this work have been conducted with the guidance and efforts of Belkis Cabrera-Palmer, Jing Liu, Jason Newby, Yuri Efremenko, Max Hughes, Justin Raybern, and Conan Bock.

6.1 Neutrons as a CEvNS Background

As discussed in Chapter 2, the only observable signature of a CEvNS interaction is the small nuclear recoil. However, an incident neutron will cause the same result; *neutron-nucleus scattering* is indistinguishable from CEvNS in our detectors and has a higher cross section than its neutrino counterpart. Passive high- Z materials are used in CEvNS detectors to shield from γ backgrounds (and can block low-energy neutrons), but will also serve to convert high-energy neutrons into a shower of lower-energy neutrons. In particular, fast neutrons with tens of MeV in energy will be the dominant background

for CEvNS searches in Neutrino Alley, based on simulations of the nuclear recoils within shielded CEvNS detectors [178].

The Spallation Neutron Source is the most intense neutron-production facility in the world, with an estimated 20-30 beam-related neutrons (BRNs) produced each time an Hg nucleus is spalled by a 1 GeV proton [179]. Knowing that neutrons will be a background for CEvNS detectors, the COHERENT collaboration launched a campaign to study the BRN rates at a variety of locations in the SNS target building (see Section 6.2.1). This included locations above, behind, and beneath the target. Ultimately, this effort identified a location with a manageable BRN flux; Neutrino Alley is, at its closest point, 19 m from the Hg target inside the SNS shielding and moderating monolith, but is protected from BRNs by a concrete/gravel fill occupying almost all of that distance.

The spallation process which produces the SNS neutrinos also produces an intense flux of prompt neutrons. These BRNs then lose energy in the target, moderator suite, and shielding monolith geometries illustrated in Figure 6.1. The majority of neutrons will be stopped by this amount of material, yet some can still make it into the alley. It is impractical to simulate the neutron flux in the same way we study the neutrino flux; between the required number of events to observe statistically significant results at detector locations in Neutrino Alley and the uncertainties on neutron interaction models, there is limited, if any, information to gain from simulating the complete details of the neutron propagation through shielding materials.

Instead, we use experimental data in Neutrino Alley to build a model of the BRN flux at various detector locations in Neutrino Alley. The paths neutron take to reach the alley are of particular interest to COHERENT, as trenches, hallways, changes in beamline shutter positions, and other such gaps in the shielding affect the observed BRN rates in the alley (see Section 6.2.3.1). There are also ongoing efforts to develop a simplified neutron propagation model to predict the results of these experimental efforts starting from the primary neutron production within the Hg target and modeling the attenuation of the flux rather than individual neutron propagation [104].

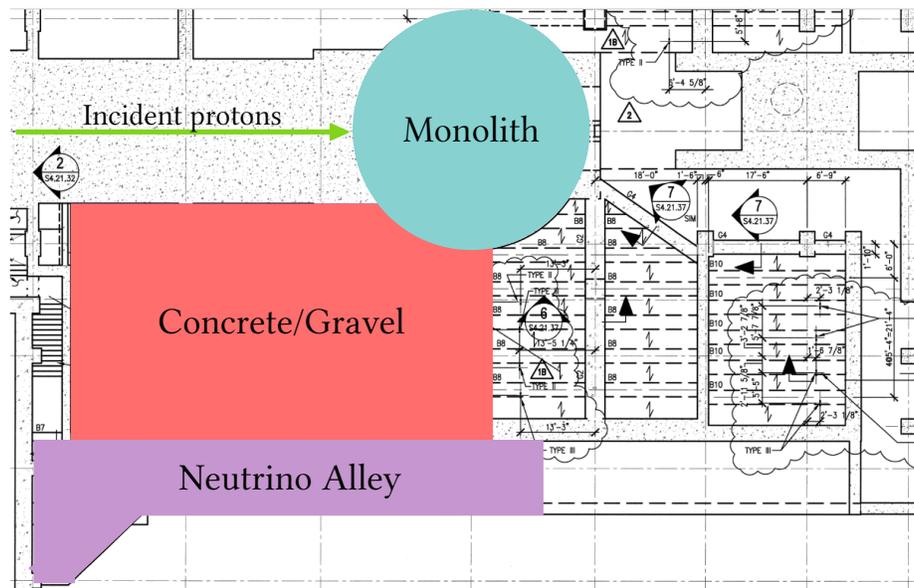
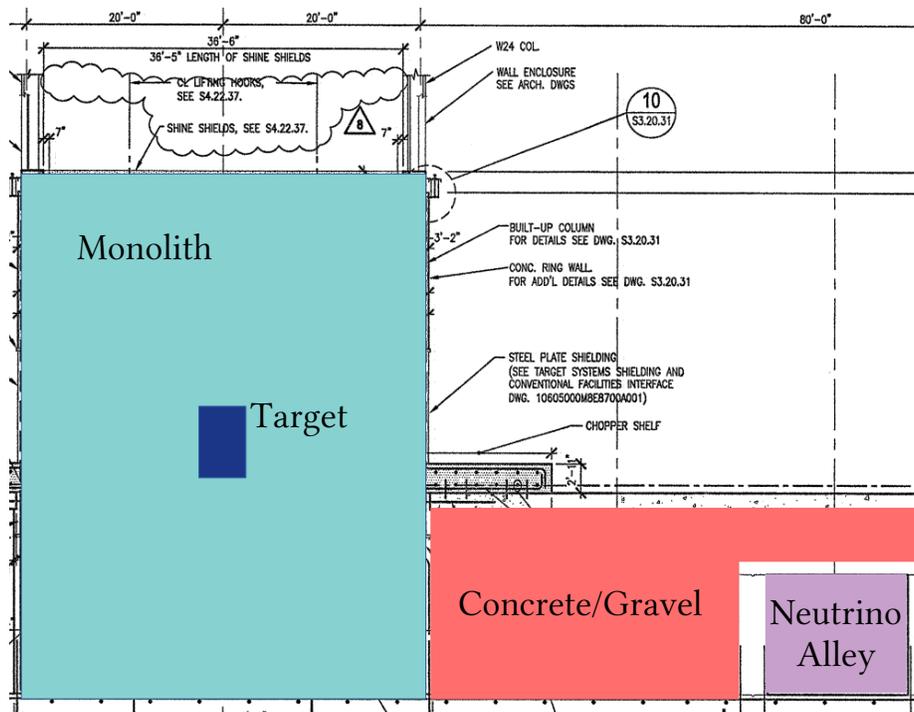


Figure 6.1: Top: A duplicated diagram from Figure 4.8, here with additional shading in red, showing the location and size of the concrete/gravel fill relative to the target location and the position of Neutrino Alley. Bottom: A top-down view of the SNS target hall, illustrating the size and location of the concrete/gravel fill relative to the target monolith and Neutrino Alley.

6.2 Strategies to Measure BRN Backgrounds *In Situ*

Because of the importance of understanding the neutron rate in Neutrino Alley, there have been multiple approaches to measuring this quantity; some measurements are detector-specific approaches which account for the exact shielding configuration, while others have been made by dedicated neutron detectors. In this section, we overview the collaboration's efforts to measure and monitor the BRN rates near COHERENT detector locations. Figure 6.2 shows the relative locations of BRN flux measurements in Neutrino Alley.

6.2.1 Historical Monitoring Systems

A dedicated BRN-measurement campaign occurred in 2013 to identify the neutron-quiet location that we now know as Neutrino Alley. Several positions were identified as possible deployment locations, both on the main target floor and in the basement.

6.2.1.1 Neutron Scatter Camera

Following the formation of the COHERENT collaboration in 2013 and the decision to deploy at the SNS, a background measurement campaign identified that Neutrino Alley had a much lower neutron flux than other possible deployment locations. The concrete/gravel fill between the target and Neutrino Alley clearly attenuates the neutron flux to a manageable level for the very sensitive CEvNS detectors. However, there is still a flux of BRNs in the alley, and understanding the rate and spectrum of neutrons throughout Neutrino Alley is of high importance to the collaboration.

The Neutron Scatter Camera (NSC) provided by Sandia National Laboratories [183] is pictured in Figure 6.3. The general strategy of a NSC affixes PMTs to two independent scintillating volumes to search for coincident neutron scatters in both volumes. The energy of the neutron can be estimated as the sum of the energy deposited by the initial scatter and the energy calculated from the time-of-flight between scatters, and the scattering kinematics are sensitive to the incident neutron direction. In the case of the COHERENT background campaign, we know the approximate direction of the incident BRNs, and were primarily focused on the spectral details and observable flux.

Neutrino Alley, labeled as “basement C2.5” and “basement C4” in the NSC results shown in Figure 6.4 (left) and “Basement 8 m.w.e.” in the right panel of the same figure, features orders-of-magnitude reduction of the BRN flux compared to the other locations investigated at the beamline level in Fall 2013. We immediately notice that the

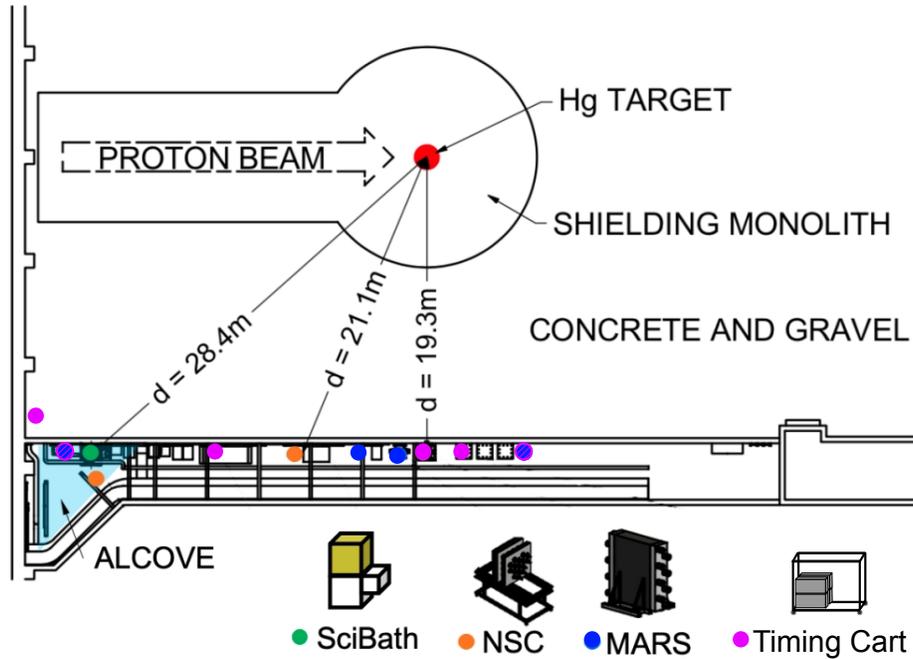


Figure 6.2: An illustration of Neutrino Alley, highlighting the location of various BRN measurements by different subsystems at various times. Figure updated from [180].

BRN spectrum is dominated by the few- to tens-of-MeV range, which helps inform the CEvNS detectors about the type and amount of shielding to use in their designs. The BRNs generally arrive within $2 \mu\text{s}$ of the beam spill, which corresponds to the “prompt” analysis window associated with the ν_μ arrival in Neutrino Alley, suggesting that the “delayed” analysis window will be an excellent tool for any subsystems which are especially neutron-sensitive. The NSC observed a prompt BRN flux of $5.3 \text{ n m}^{-2} \text{ MWh}^{-1}$ in the center of Neutrino Alley (near the CsI deployment) and a prompt BRN flux of $272.1 \text{ n m}^{-2} \text{ MWh}^{-1}$ in the alcove (near the LAr deployment). The measured fluxes from the NSC have very large uncertainties ($>50\%$) due to poor statistics and limited understanding of the double-scatter detection efficiency [180, 184, 185]; precision measurements of the BRN flux and spectra are needed to build confidence in the neutron background model for COHERENT detectors.

6.2.1.2 SciBath

While the NSC identified Neutrino Alley as the best possible deployment location in the SNS target hall, SciBath was deployed to take precision measurements in the higher-

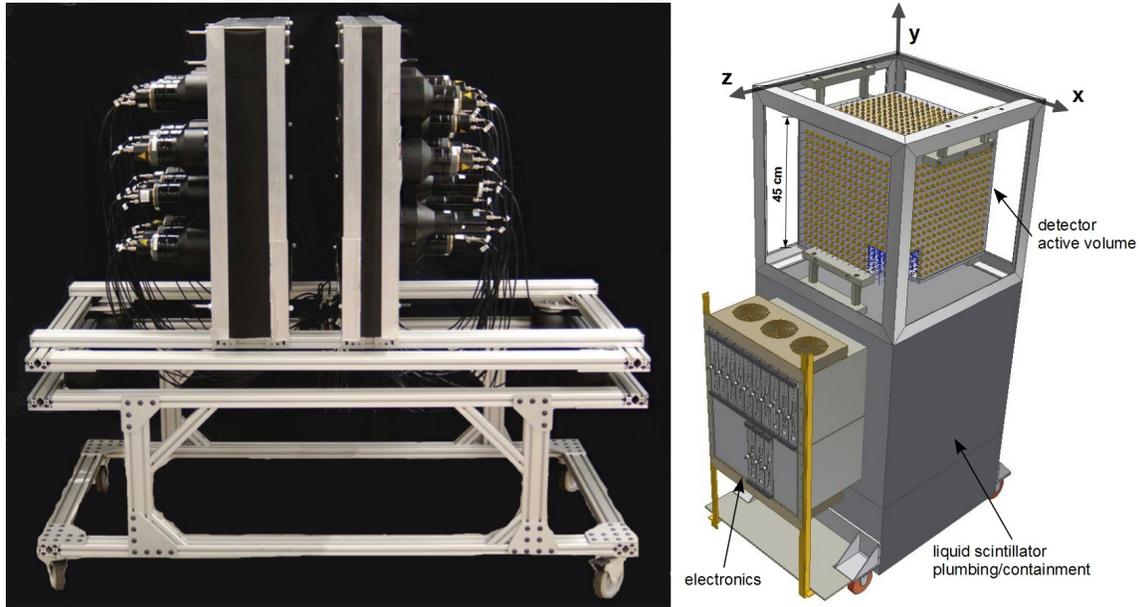


Figure 6.3: Left: A photo of the Neutron Scatter Camera, reproduced from [181]. Incident neutrons scatter in both the left and right scintillating volumes; the energy collected in the initial scatter and the time between observed scatters enable measurements of the incident neutron energy. In this configuration, electronics are deployed in the bottom of the frame, and the cart is transported to the relevant deployment location. Right: A labeled diagram of the SciBath detector, reproduced from [182]. Incident neutrons thermalized via proton recoils before capturing on hydrogen (similar to the MARS detection strategy).

background alcove due to the additional space the alcove provided for the LAr cryogenic systems. A three-dimensional array of 768 wavelength-shifting fibers in an 80-liter liquid scintillator (LS) bath, SciBath had been previously used by COHERENT collaborators working on Fermilab neutrino experiments [186], and the deployment characterized both the neutron-flux and the cosmic-muon-flux backgrounds simultaneously. Scintillation light from charged particles (e.g., cosmic muons or protons recoiling from elastic neutron scatters) was captured by the fibers and guided to PMTs. Incident neutrons experienced several scatters in the scintillator, ultimately thermalizing before capturing on hydrogen $n + {}^1\text{H} \rightarrow {}^2\text{H} + \gamma$ and producing a characteristic 2.2 MeV gamma-ray. Events with a coincident signal of the neutron thermalization and capture γ could then be tagged as stopped neutron events and used to infer the neutron flux. SciBath observed a prompt BRN flux of $(2.1 \pm 0.4) \times 10^{-5}$ neutrons per square meter each beamspill, or, using consistent units to the NSC results, $13.3 \pm 2.5 \text{ n m}^{-2} \text{ MWh}^{-1}$, and a delayed BRN flux consistent

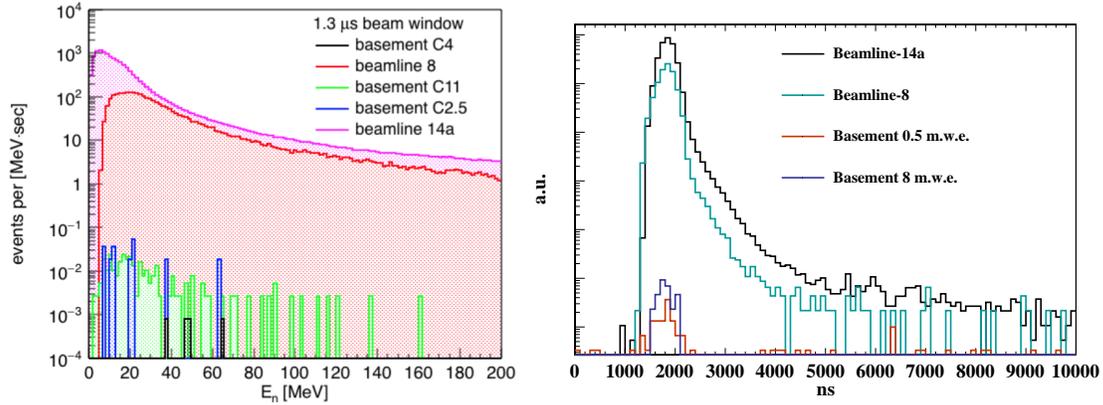


Figure 6.4: Left: The observed energy spectra of BRNs (in a $1.3 \mu\text{s}$ analysis window coincident with prompt neutron arrival following the beam-spill) by the NSC at the indicated locations in the SNS target hall. The “basement C4” spectrum corresponds to the NSC deployment in the center of Neutrino Alley, while “basement C2.5” corresponds to the deployment in the alcove. Reproduced from [185]. Right: The observed arrival time of BRNs at the indicated locations in the SNS target hall. This demonstrates that for Neutrino Alley (Basement 8 m.w.e.), BRNs are expected in the “prompt” analysis window ($< 2\mu\text{s}$ following a beam spill) coincident with ν_μ arrival times, but are significantly reduced in the “delayed” analysis window associated with ν_e and $\bar{\nu}_\mu$ arrival times (μs). Reproduced from [184].

with zero.

6.2.2 Detector-specific Analyses

COHERENT’s CEvNS detector groups have used their own strategies to characterize the observed BRN rates in their specific shielding configurations. We do not produce BRN flux measurements from these detector-specific studies, largely because the reduction of the neutron fluxes through given shielding configuration are not the priority of these studies; only the expected rates after shielding contribute to the final analyses.

The initial CsI[Na] detector deployed LS cells inside the chosen shielding configuration to directly measure the observed BRN rate during beam-on operations [85]. The LAr detector performed a three-week “no-water” run and removed the water shielding to directly measure neutron rates, in addition to utilizing the SciBath measurements at their alcove deployment location [107].

The NUBEs measurement of the neutrino-induced-neutron cross section σ_{NIN} also re-

quired knowledge of the BRNs that make it through the shielding configuration. Simulated models of the BRN flux based on neutron production models convolved with attenuation lengths of different concrete compositions were explored for the NUBES, along with a synthesis of the available Timing Cart data (see Section 6.2.3.1) near the end of the concrete/gravel fill [103, 104].

6.2.3 Currently Deployed Monitoring Systems

The physics goals of COHERENT are impacted by the BRN flux at all detector locations, yet our most precise measurement (from SciBath) has a 19% uncertainty and is confined to a single location. The subsystems discussed in this section are currently taking data in Neutrino Alley in pursuit of an accurate characterization of BRN fluxes throughout Neutrino Alley. The derivation of flux estimates from these systems are not finalized, due to ongoing detector characterization efforts and measurements. We present here a preliminary discussion of the goals of these independent subsystems.

6.2.3.1 Timing Cart

The “Timing Cart” is a collection of four, unshielded LS cells with a single PMT per cell, shown in Figure 6.5 left. These cells are placed on a small cart for mobility, and can be deployed anywhere in Neutrino Alley to obtain a quick estimate of the BRN rate at that location. Due to the small mass of this subsystem, the main goal of the Timing Cart is to track the BRN arrival time at different locations in Neutrino Alley to inform our model of BRN propagation through the SNS target hall.

As an example, the results of the Timing Cart helped to determine that the activities of the beamline experiments at the SNS impact the neutron rate in Neutrino Alley. In Figure 6.5, we show a top-down view of the SNS target hall and highlight the locations of the neutron beamlines which transport neutrons to the relevant experimental halls. Each beamline is generally accompanied by a “trench”, which is a hollow space dug into the surroundings to provide enough space to perform maintenance tasks. These trenches can provide an alternate path for neutrons to propagate to Neutrino Alley, and the Timing Cart has preliminarily shown increased BRN rates at the locations of these trenches. Similarly, the Timing Cart has shown that beamline “shutters” impact the observed BRN rates in Neutrino Alley. If a beamline shutter is closed, neutrons are heavily moderated to prevent their propagation down the beamline. In the alcove, the Timing Cart observed a decrease in BRN counts correlated with closed shutters for the beamlines above the alcove.

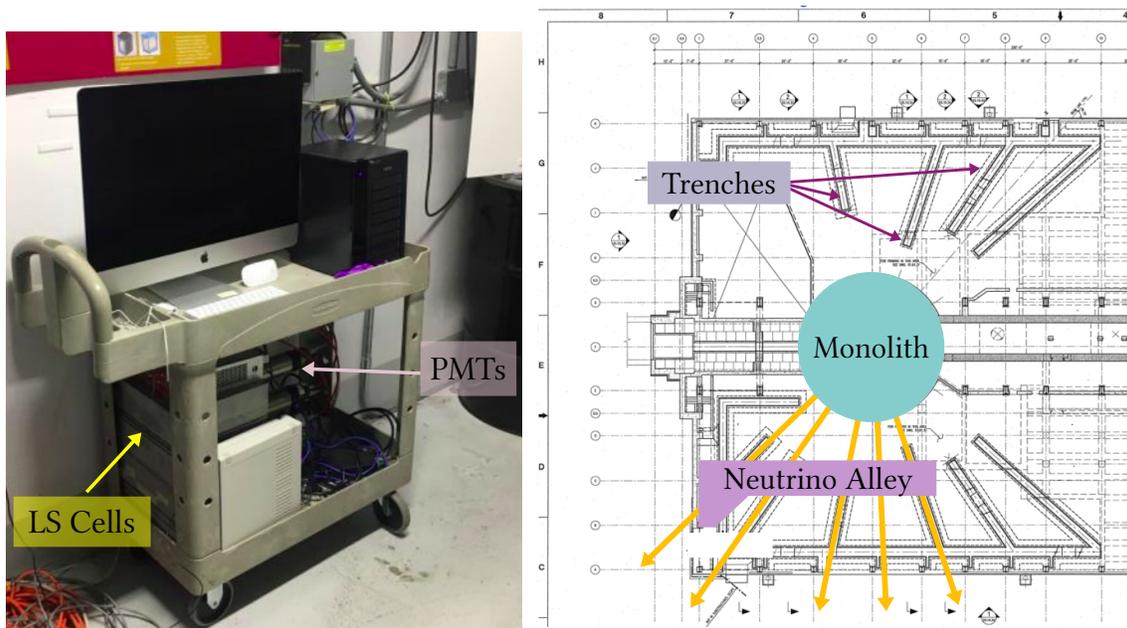


Figure 6.5: Left: A photo of the Timing Cart deployed near the center of Neutrino Alley. Right: A top-down view of the SNS target hall, showing the position of Neutrino Alley (purple) relative to neutron beamlines of interest (yellow arrows). Most of the beamlines shown feature “trenches”, which give maintenance access to the beamlines but introduce gaps in the shielding and possible alternate paths for neutron propagation.

The mobility of the Timing Cart is also of incredible benefit to the collaboration. It is small enough to be deployed inside the stairwell nearest to the alcove of Neutrino Alley. The observations of the Timing Cart in the stairwell helped to determine that neutrons can avoid the concrete/gravel fill by propagating down the stairwell and into Neutrino Alley. The quick estimates of BRN arrival times and rates help to identify locations or details that require further investigation by the collaboration. Figure 6.6 shows the results from each of the Timing Cart runs throughout Neutrino Alley. These efforts are ongoing; the Timing Cart is currently deployed in the center of Neutrino Alley and has preliminarily observed a change to the BRN flux following the decommissioning of the NUBES in June 2022, which further demonstrates the need for constant monitoring of the BRN rates in Neutrino Alley.

Location	Hall coordinate (m)	Integrated neutron flux $\frac{\text{Neutron}}{\text{m}^2 * \text{MWhr}}$	Power-law Power
CsI location	19.4	$1.39 \pm 0.43 \pm 0.14$	Fixed to 1.24
Alcove take 1 shutter 1 closed	2.1	$2.49 \pm 0.68 \pm 0.26$	Fixed to 1.24
Alcove take 1 shutter 1 open	2.1	18.72 ± 2.23	1.08 ± 0.30
In front of rad monitor with no lead	25.8	54.73 ± 2.29	1.20 ± 0.08
In front of rad monitor with no lead on floor	25.8	46.94 ± 2.46	1.14 ± 0.09
In front of rad monitor with lead	25.8	46.28 ± 1.22	1.27 ± 0.04
Alcove take 2 production running	2.1	13.25 ± 0.56	1.61 ± 0.08
Next to trench/timing rack	12.1	$0.68 \pm 0.19 \pm 0.04$	Fixed to 1.24
Alcove take 4	2.1	25.84 ± 1.24	1.41 ± 0.09
Alcove on top of NUBEs shield	2.1	19.60 ± 0.98	1.25 ± 0.10
In stairwell production running	0.1	261.84 ± 3.89	1.17 ± 0.03
Next to trench/timing rack take 2	12.5	$0.58 \pm 0.014 \pm 0.03$	Fixed to 1.24
CsI Location take 2	19.4	0.50 ± 0.16	0.93 ± 1.25

Table 5: Summary of the fluxes

Figure 6.6: A summary table, provided by Max Hughes, of the Timing Cart measurements throughout Neutrino Alley. Reproduced from an internal COHERENT technical note. The column “Power-law Power” relates to the spectrum of incident BRNs (α in Equation 8.3).

6.2.3.2 MARS

The Multiplicity and Recoil Spectrometer (MARS) is a gadolinium-doped plastic scintillator originally designed to study the small fast-neutron fluxes (tens to hundreds of MeV) in underground experiments [187], and is shown in Figure 6.7. This mobile detector has been re-configured to meet the space restrictions of Neutrino Alley and monitor the BRN rates in our surface-level experiment. Part of this refurbishing involved the removal of a neutron detection mechanism, the so-called “multiplicity” mode. The original design featured two stacked modules sandwiching a lead and steel table, as shown in the left panel of Figure 6.7, such that neutrons could interact in the first module, produce secondary neutrons in the lead (hence “multiplicity”), and those secondary neutrons would be observed in the second module.

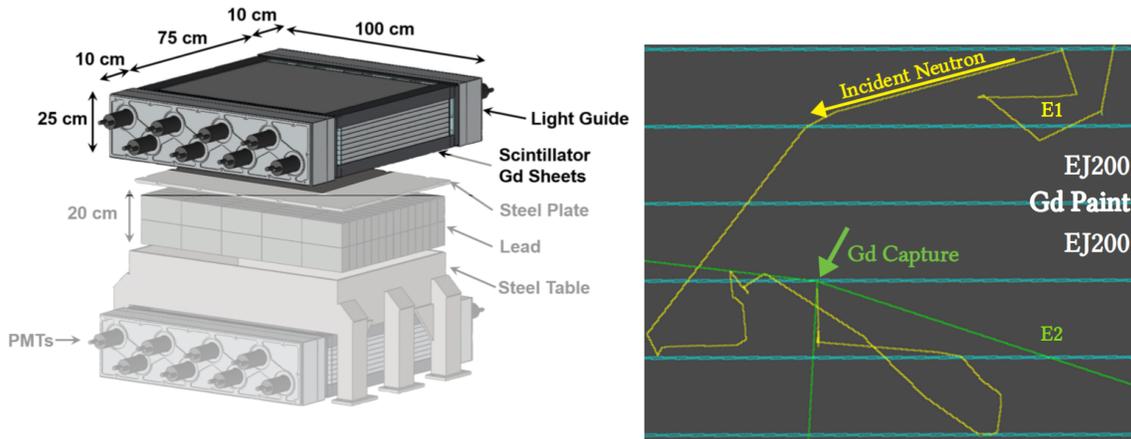


Figure 6.7: Left: An exploded diagram of the full MARS detector, reproduced from [187]. The unshaded region indicates the module deployed to Neutrino Alley. Right: A simulated MARS capture-gated event. The incident neutron (yellow) down-scatters in the scintillator, then the thermal neutron captures on Gd. The de-excitation of the Gd produces an ~ 8 MeV shower of gammas (green).

The MARS detector as deployed in Neutrino Alley consists of twelve 2-cm-thick EJ-200 plastic scintillator layers interleaved with Gd-coated Mylar sheets¹. The dimensions we discuss here are given in Width \times Height \times Depth coordinates. The overall dimensions of the detection volume (scintillator + Gd sheets) are $75 \times 100 \times 25$ cm³. This module stands vertically, held by a Unistrut frame, with each of the left and right faces covered by a $10 \times 100 \times 25$ cm³ acrylic light guide. On each side, eight 5-inch-diameter ADIT B133D01 photomultiplier tubes (PMTs) are coupled to these light guides by silicon grease. Black tape covers the scintillator and acrylic volumes for light-tightness, and the scintillating volume is further wrapped in an aluminum sheet for fire safety. We include a photo of the MARS deployment in Neutrino Alley in Chapter 7 (see the left panel of Figure 7.1).

Without the lead table, MARS analyses cannot operate in “multiplicity” mode, and we conduct so-called “capture-gated” analyses. When a fast neutron elastically interacts with hydrogen nuclei in plastic scintillator, the recoiling protons generate a prompt scintillation pulse. If the neutron does not escape, it will quickly thermalize while bouncing around the detector. In the presence of a material with a large thermal-neutron capture cross section like Gd, there is a high probability that de-excitation γ -rays will be produced. In Gd specifically, the maximum total energy of the emitted γ -rays is ~ 8 MeV from the

¹We assume the Gd contains natural abundances; it is not isotopically pure.

following interaction on naturally occurring Gd isotopes:



The capture-gated mode searches for two coincident signals: the prompt scintillation pulse from proton recoils associated with the neutron thermalization, and a delayed pulse from the γ shower signifying a neutron capture and subsequent Gd de-excitation. An example capture-gated event is illustrated in the right panel of Figure 6.7.

This process in MARS has a time constant τ_{nCapt} that depends on the Gd concentration and neutron thermalization time; it was measured to be $18.7 \mu\text{s}$ for the full-MARS configuration [188]. Thus, using the known timing and energy of the Gd gamma-ray scintillation pulses, the population of fully thermalized and absorbed neutrons can be isolated. Neutrons with energies in the few to tens of MeV will primarily transfer all their kinetic energy to the scintillator in multiple elastic scatterings, such that the integrated prompt pulse is a direct measure of the incoming neutron energy. Since the fraction of neutrons thermalizing purely via elastic interactions decreases with increasing neutron energy, being of the order of a few percent for ~ 50 MeV neutrons, the capture-gated mode is expected to provide accurate rate and spectral information on BRN energies no larger than a few tens of MeV [187].

Data readout, diagrammatically shown in Fig. 6.8, is done with two 14-bit, 250 MHz Struck SIS3316 sixteen-channel waveform digitizers [189] mounted in a VME crate. The first digitizer is used to record the observations of the sixteen PMTs, and the second asynchronously records two SNS-provided timing signals, arbitrarily named “E39” and “E61” triggers (where “E” stands for “event”), which are synchronized with the SNS beam-spill. The E39 triggers are provided uninterrupted at 60 Hz during beam-on and beam-off periods, while the E61 triggers are supplied during beam-on operation whenever there are protons sent to target. These beam spills also occur with a 60 Hz frequency but with one event skipped every ten seconds (averaged frequency 59.9 Hz)

The signals from a pair of neighboring PMTs on one side of MARS and the corresponding pair at the same height on the opposite side are joined into the same digitizer channel group, for a total of four non-overlapping groups. These groups are illustrated in the sample event of Figure 6.9, such that each plot in the left panel corresponds to a trigger group in the manner labeled in the right panel. A trapezoidal Finite Impulse Response (FIR) filter² with peak and gap times equaling 4 samples (16 ns) is applied on each

²A trapezoidal filter is a moving average which compares the difference integrals of two “peak” time regions separated by the “gap” time. The FIR threshold requires that the difference between the “peak”

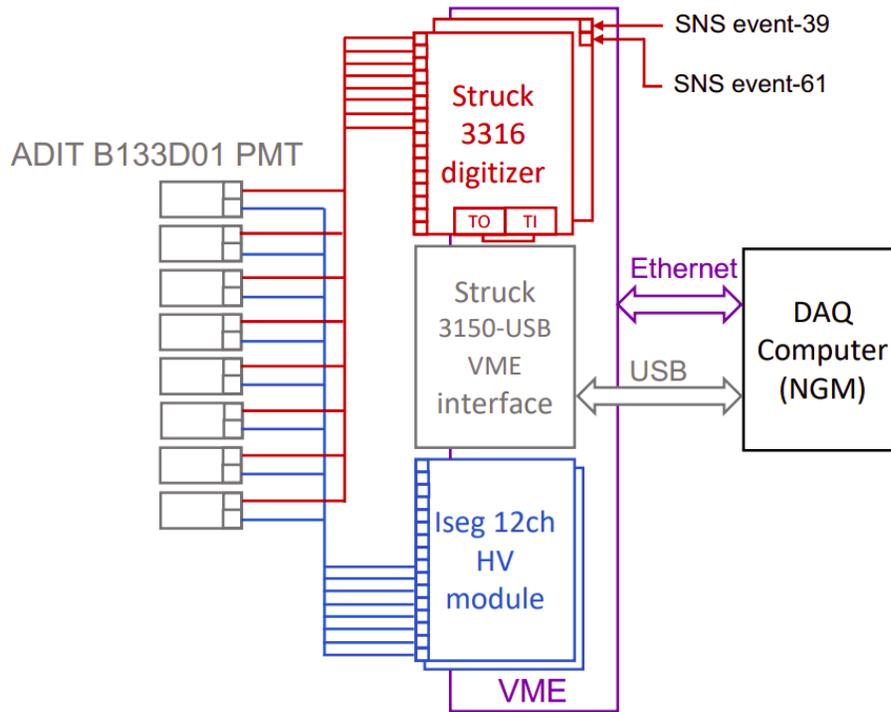


Figure 6.8: A diagram of MARS’ electronic read-out and high-voltage setup, representing only 8 of the 16 PMTs for simplicity. TO and TI represent the Trigger Out and Trigger In connectors, respectively. Figure reproduced from [180].

four-channel sum signal to generate a trigger signal [189]. This FIR trapezoidal threshold value — 300 ADC — reduces dead time while still recording low-energy signals. The internally generated trigger is routed to the external trigger input to start the sampling of all sixteen channels, and 150 samples (600 ns) are recorded for each trigger: 50 baseline pre-samples, and 100 samples containing the majority of the scintillation pulse. A sample event is shown in Figure 6.9.

For each recorded event, the raw channel data contain the channel ID, timestamp, peak height value and sample index³, and sums from each of six accumulators. Each accumulator stores the integral of 25 samples, with accumulator 0 integrating samples 0-24, accumulator 1 integrating samples 25-49, etc. At the first stage of processing, each PMT channel integral is calculated as the sum of accumulators 2 to 5 (signal region — 100 samples) minus the twice the sum of accumulators 0 and 1 (baseline region — 50 samples, integrals be greater than a threshold value.

³The peak height refers to the maximum of the waveform; we track the maximum voltage associated with the event and the sample-time at which that voltage occurred.

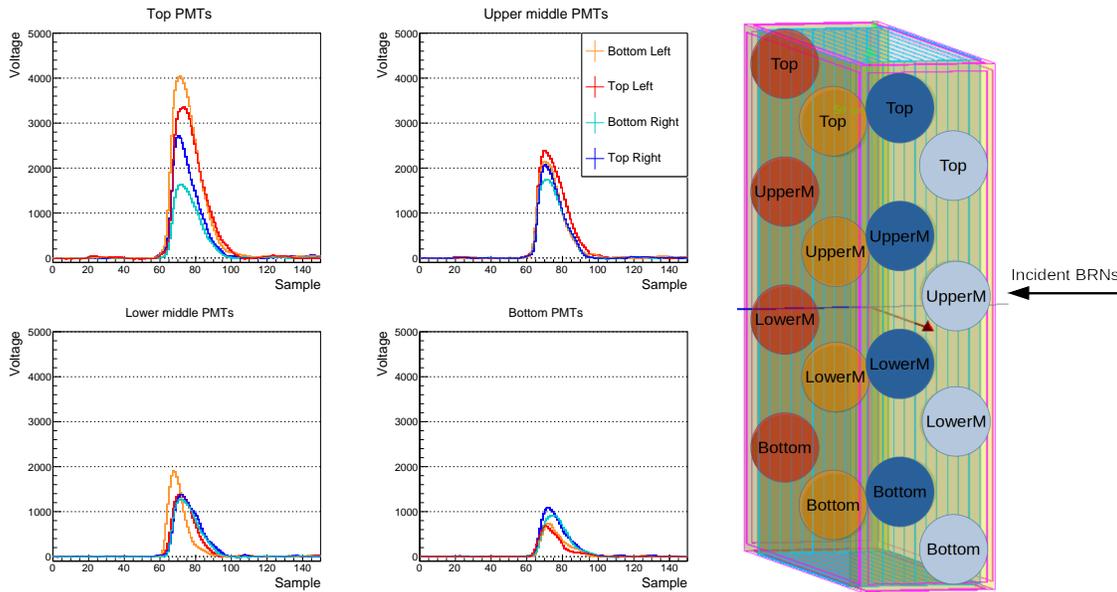


Figure 6.9: Left: A sample event, showing a recorded waveform from each of MARS’ 16 PMT channels. Right: A Geant4 screenshot (see Chapter 7 for details) highlighting the PMT locations and trigger groups for the MARS detector. The color schemes are connected across both figures; for example, the red curve of the upper left panel of the waveform trace corresponds to the “Top Right” PMT in the “Top” trigger group; this corresponds to the shaded red “Top” circle of the Geant4 screenshot.

doubled for equivalence to the signal region). The sum of the sixteen baseline-subtracted integrals is taken to represent the total energy of the scintillator pulse in the digitizer output units (“ADC” units). The times elapsed since the previous E39, the previous E61, and the previous scintillator pulse are also recorded.

MARS’ deployment in Neutrino Alley is focused both on the monitoring of BRN rates and measurements of the incident neutron spectrum. The Timing Cart has identified key locations of interest where the more-massive MARS can collect the statistics to perform a spectral unfolding; this work contributes to the development of a simulation response matrix and the preliminary analyses of data taken throughout Neutrino Alley. We detail the efforts to design and validate a simulation to model the MARS response in Chapter 7 and make a preliminary demonstration of our ability to interpret spectral information from MARS data. We then report on measurements of the BRN rate at three distinct locations in Neutrino Alley in Chapter 8 and describe these measurements in the context of other COHERENT neutron-monitoring efforts.

Chapter 7

Developing a Simulation of MARS to Unfold the BRN Spectrum

The work in this chapter focuses on the introduction to the MARS detector and the development of a simulation response matrix, with the ultimate goal of extracting a neutron flux and energy spectrum from BRNs observed in Neutrino Alley. The Geant4 simulation model I describe is my own work, but it was greatly benefitted by validations against an independent model built by Jing Liu and Conan Bock using GEARS [190].

7.1 Geant4 Simulation of MARS

7.1.1 Geometry

We model the MARS geometry (described in Section 6.2.3.2) as accurately as possible with the intent to study energy depositions and particle interactions *without* performing optical simulations, pulling extensively from Caleb Roecker’s Geant4 simulation of the original MARS detector [188]. The scintillator material is defined using the listed fractions of natural hydrogen and carbon with atomic ratios for EJ-200 [191]. The Gd-paint is simulated as a mixture of Gd and acetone, with the mixture left as a tuneable parameter for matching the observed detector response (see Section 7.2). All other volumes (steel, air, etc.) are defined using default material descriptions provided in Geant4, which implement databases of physical responses provided by the National Institute of Standards and Technology (NIST) [192].

Our simulated detector is also shown in Figure 7.1 (right), and features a solid $75 \times 100 \times 25 \text{ cm}^3$ scintillator volume with 13 layers of Gd-paint mixture as $75 \times 100 \times 0.07$

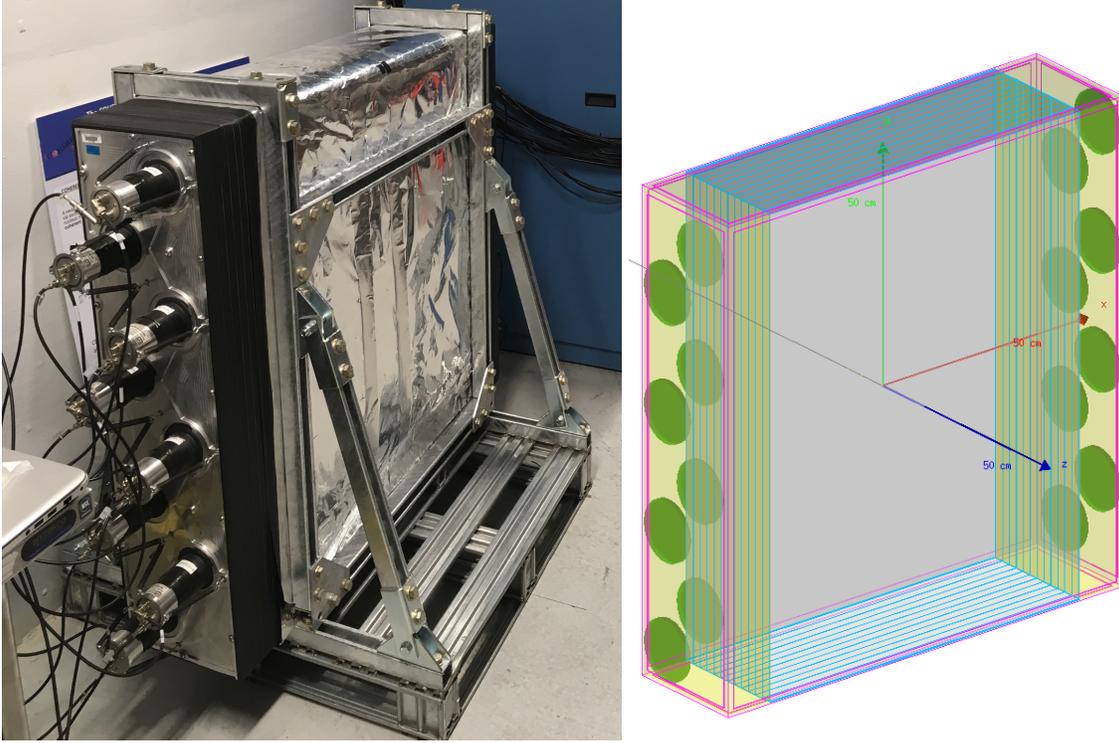


Figure 7.1: Left: A photograph of the MARS detector in Neutrino Alley. Right: A screenshot of the Geant4 model of the MARS detector.

cm^3 boxes inserted into the scintillator with 2 cm spacing. On the top and bottom of the detector, the sides with no photocathode coverage, we add additional Gd-paint layers with dimensions $75 \times 0.05 \times 25 \text{ cm}^3$. On the edges of the detector which house the PMTs, acrylic volumes of dimension $10 \times 100 \times 25 \text{ cm}^3$ serve as our model of the light guides. We simulated each PMT as a 5 in $\text{Ø} \times 5 \text{ mm}$ quartz cylinder, but we do not simulate any optical greases or other more specific couplings of the PMTs to the detector.

7.1.2 Physics Model

We use our simulation model to predict the behavior of the detector for low-medium energy particles. No incident particle is expected to exceed 1 GeV; we use calibration sources in the few MeV range, though BRNs from the SNS could have energies up to a few hundred MeV. We use the standard Geant4 physics list QGSP_BERT_HP to model all interactions in our detector, though at present we have not validated the neutron thermalization model – we primarily focus on the light collection and response to gammas in this work. The interactions of γ s demonstrate consistency with other low-energy physics lists, as shown in

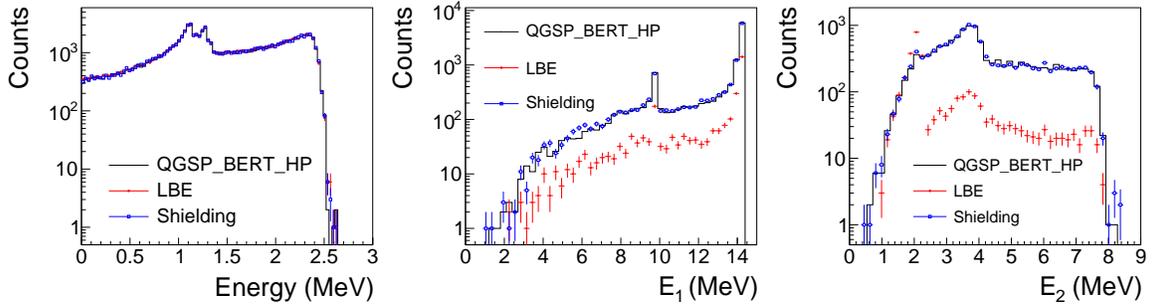


Figure 7.2: Left: A comparison of the calculated energy deposits from 1.1 and 1.3 MeV γ s simulated inside MARS using three different physics lists: QGSP_BERT_HP, Shielding, and LBE. All models are demonstrated to be consistent in predicting the MARS response to γ s. Center/Right: A comparison of the calculated energy deposits from 14.1 MeV neutrons incident on MARS from outside the detector. The center plot shows the energy deposited by the neutron thermalization (E_1), while the right plot shows the energy deposited following the Gd de-excitation (E_2). LBE predicts an efficiency that is a factor of 4 lower than the other models, in addition to predicting a different spectral response to the Gd de-excitation from either QGSP_BERT_HP or Shielding.

Figure 7.2, though we do note discrepant response in the Gd de-excitation models. Given the agreement between QGSP_BERT_HP and the Shielding physics list, which uses data-driven approaches to neutron interactions below 20 MeV, we trust that these models are more appropriate than LBE (the Low Background Experiment physics list).

The work with the original MARS detector used Geant4’s standard Shielding [193] physics list in combination with the MENATE_R [194, 195] description of neutron scattering on carbon [187, 188]. Given our focus in this work on the response to low-energy neutrons and calibration sources in the few-MeV range, we do not currently initialize the MENATE_R description as it is primarily important for neutrons above 20 MeV, but we may re-evaluate this choice in future. We use Geant4’s photon evaporation model to handle the de-excitation physics following a neutron capture, though future efforts will incorporate the FIFRELIN model of γ emission following a gadolinium de-excitation to further improve our model of the MARS energy response [196, 197].

We choose *not* to perform an optical simulation of this detector, but instead create a detector response model by measuring the optical effects and matching simulation to that data. In this way, we do not need to model any optical greases which couple the PMTs to the light guides, any potential optical coupling between the paint layers and scintillating layers (particularly in the outer edges of MARS), etc. We simply validate our simulations

by modeling a source and comparing the raw energy deposition from simulation to the measured response of the detector. These calibrations inform the development of a full detector characterization which would enable the unfolding of the incident neutron spectrum as follows.

7.1.3 Event Processing

To interpret our simulation data and compare our calculations to the observed energy observables (which will be extensively discussed in Chapter 8), we process simulated events to create analogous structures in simulation and data. The raw observation of MARS is the integrated pulse recorded in all 16 PMTs (a sample event is shown in Figure 6.9, with the integration window illustrated), which is proportional to the total energy deposited in the detector. Each event represents a group of energy deposits, and its height and width will depend on the strength of the interaction, the geometry and optical properties of MARS, and the operating voltages and gains of the PMTs themselves. We thus use the simulation as a tool to unpack each of those effects, and build a model of the detector which helps us to better understand the MARS response on an event-by-event basis.

Our simulation processing is outlined in this section, and undergoes five main steps as visualized in Figure 7.3:

1. Collecting Deposited Energy into Pulses — We group the energy deposited by single particles in time to develop a model of the waveform pulse that we would observe in Section 7.1.3.1.
2. Quenching — The scintillation light yield by different particles will depend on the properties of both the incident particle and the surrounding scintillating material. We build a model of the differences between particle interactions in EJ200 in Section 7.1.3.2.
3. Smearing — The light collection of MARS is dependent on both the placement of PMTs in the geometry and the amplification of signal within each individual PMT. We describe the development of the smearing model in Section 7.2, but we'll discuss its general implementation in our simulation model in Section 7.1.3.3.
4. Threshold — We measure the signal acceptance fraction and implement a threshold cut in Section 7.1.3.4.
5. Pairing — Finally, we pair all pulses which pass threshold in Section 7.1.3.5.

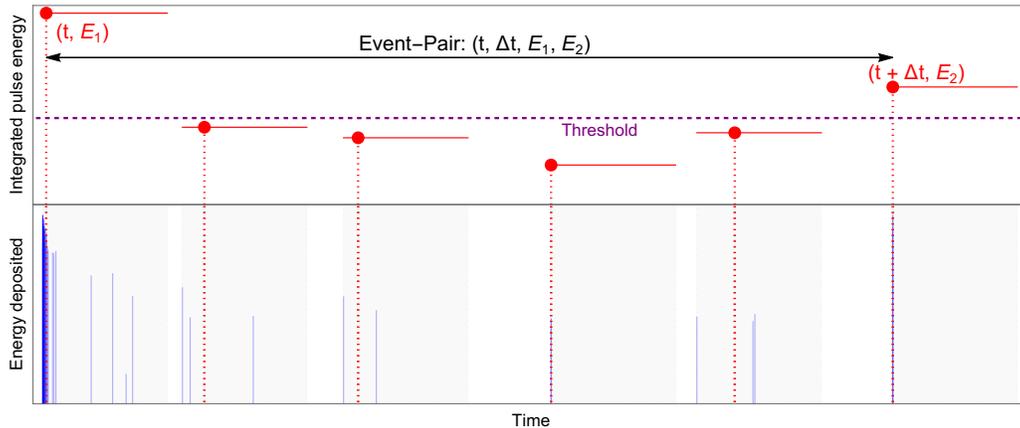


Figure 7.3: An illustration of the processing for a single simulated neutron event. Each blue line in the bottom panel corresponds to an interaction in the MARS detector; its position along the x-axis represents the time of the interaction while its height represents the amount of energy deposited. We group energy depositions in fixed-width time bins, shown as both the gray shaded regions on the bottom panel and the red horizontal lines on the top panel, and title each group a “pulse”. The red points of the top panel demonstrate the energy-weighted vertex of each pulse, with energy and time calculated using Equations 7.2 and 7.3 respectively. All pulses which have a total energy exceeding the threshold setting are paired as described in Section 7.1.3.5. The only possible event-pair in this example is then defined by four variables: t (the global time in simulation, relative to the emission of the incident particle), Δt , E_1 , and E_2 .

7.1.3.1 Collecting Deposited Energy into Pulses

Energy depositions in the MARS scintillating volume are gathered in bins of constant width to mimic the physical scintillating pulse (see Chapter 6); these groupings of energy depositions are what we’ll refer to as our simulation “pulses”. We structure the pulse objects to contain an array of the time and energy deposit for each constituent interaction, and allow the object to update relevant data members throughout the next several steps. The analysis favors time bins on the order of 400 ns¹, but there was no significant difference in simulation using 400 ns vs. 1 μ s time bins, so we choose the larger of the two values for computational efficiency — larger time bins limit the number of pulses that can be created per event.

¹We record 150 samples per event, totaling 600 ns. However, the first two accumulators are baseline, so the signal is recorded for 100 samples, or 400 ns.

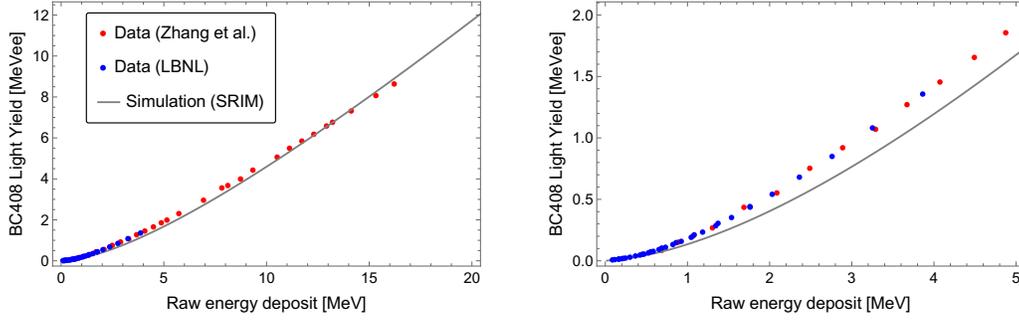


Figure 7.4: Left: A comparison of the proton quenching model parameterization we implement in simulation to the central values of the measurements by Zhang *et al.* [198] and LBNL [199]. Right: The same comparison, zoomed into the 0 - 5 MeV region for a closer look at the low-energy discrepancy. We do not show the experimental uncertainties on the measurements in either panel, but note that they are described to be less than 10% of the central values.

7.1.3.2 Quenching

The low-energy light yield of proton recoils in plastic scintillators is of high importance to a variety of experiments, so measurements on specific scintillator compositions are prevalent in the literature. For our purposes, we want to create a quenching model in simulation to convert the raw energy deposited (MeV) to the observable light yield (MeVee, or “electron-equivalent” energy) of the MARS EJ-200 scintillator². Recent proton quenching data has been taken for proton recoils less than 20 MeV [198, 199]. We illustrate these datasets in Figure 7.4.

Rather than extrapolate these datasets to extend to all energies of interest³, we use the Stopping and Range of Ions in Matter (SRIM) toolkit [200] to calculate the energy loss of particles in the scintillator ($\frac{dE}{dx}$), and use the Chou parameterization to calculate the total light yield $L(E)$ as[201]

$$L(E) = S \int_0^E \left[1 + kB \frac{dE}{dx} + C \left(\frac{dE}{dx} \right)^2 \right]^{-1} dE, \quad (7.1)$$

²Quenching accounts for the tendency of nuclear recoils to induce a smaller response than electrons of the same energy; nuclear recoils are less ionizing than electronic recoils.

³For example, a high-energy neutron could cause a proton to recoil with greater than 20 MeV, but since we are interested in neutrons below 100 MeV, the majority of neutrons will create cascades of low-energy proton recoils.

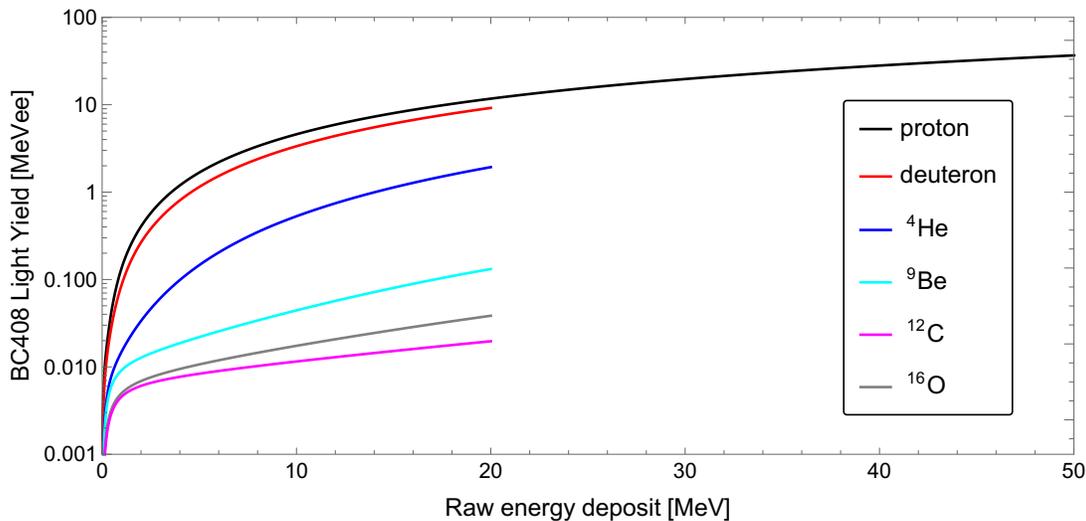


Figure 7.5: A comparison of the light yield models implemented for different particles in MARS. The proton model ranges from 0 - 100 MeV, while all other ions are limited to 0 - 20 MeV.

where S is the absolute scintillation efficiency, and kB and C are adjustable parameters that are empirically fitted parameter to the model. Using the SRIM results in combination with the $kB = 1.15 \times 10^{-2} \text{ g cm}^{-2} \text{ MeV}^{-1}$ and $C = 0.84 \times 10^{-5} \text{ g}^2 \text{ cm}^{-4} \text{ MeV}^{-2}$ parameters found by Zhang *et al.* [198] to generate a lookup table to apply in simulation⁴. Thus, in simulation, we use the total energy (in MeV) of a recoiling proton to look up the appropriate light yield (MeVee). We show the parameterization we apply in Figure 7.4 relative to the central values of the LBNL and Zhang measurements, and note that this offers a rough parameterization of the available data. We do, however, overquench in the low-energy range; we artificially reduce the light collection in our simulation model by roughly 25%. We will return to this point in Section 7.2.3, but generally accept this uncertainty given the power of this approach: we can generate a light-yield lookup table for any particle of interest.

Inelastic interactions in MARS can produce recoiling ions, and we use this SRIM-parameterization approach to generate quenching models for deuterium, tritium, helium, beryllium, carbon, and oxygen isotopes in addition to the proton quenching. We show the quenching models as applied in simulation in Figure 7.5, and note that we only model the ion light yield for energy deposits up to 20 MeV.

⁴The SRIM calculations were performed by Ana Paula Vizcaya Hernandez, and Belkis Cabrera-Palmer developed the lookup table.

The “vertex” of each simulated pulse is then calculated by summing over the individual quenched energy deposits E_i (at time t_i and position \vec{r}_i) to obtain a measure of the energy-weighted average parameters of the pulse. We define the pulse energy (E_p), time (t_p), and position (\vec{r}_p) as:

$$E_p = \sum_i E_i, \quad (7.2)$$

$$t_p = \sum_i t_i \frac{E_i}{E_p}, \quad (7.3)$$

$$\vec{r}_p = \sum_i \vec{r}_i \frac{E_i}{E_p}. \quad (7.4)$$

All calculations beyond this point in the simulation processing ignore the constituent energy depositions of the pulse and are performed on the pulse vertex.

7.1.3.3 Smearing

Just as for our D₂O simulations, we use the MARS simulations to generate events and describe probability distributions for given events to produce observable signals in the detector. The energy resolution, or the width of the Gaussian peak that describes the observation of a given energy deposit⁵, is characterized by performing calibrations with sources of known energy and matching our simulated model of the energy deposited to experimental observation.

The specifics of the smearing map we apply are described in great detail in Section 7.2, and are found to have a position dependence⁶. In the simulation processing, we use the energy-weighted position of the pulse, \vec{r}_p , to determine the expected smearing parameters, and apply the smearing to the quenched energy deposit of the pulse. We *must* use the quenched energy deposit here, as the smearing map is developed using γ calibrations; we have only measured the light response (quenched, MeVee) rather than the raw response (recoil energy, MeV).

We do not construct a smearing model that is continuous in position; we determine the smearing parameters only at specific locations in MARS and use a linear interpolation⁷ between measured positions to extract the smearing parameters $s(\vec{r})$ (scaling⁸ — units

⁵This is the same concept of resolution we used in Chapter 5; we merely apply it to different physical interactions.

⁶Note that the MARS PMTs are on the sides — this position dependence simply means that events close to the PMTs will have a larger observed light yield than events far from the PMTs.

⁷This is a three-dimensional linear interpolation, and is described in detail in Section 7.2.

⁸We may also use the term “slope” in this work interchangeably with “scale”.

MeVee / ADC) and $\beta(\vec{r})$ (resolution $\sigma = \beta\sqrt{E} - \beta$ in units $\sqrt{\text{MeVee}}$) at the particular location of interest. We then sample from a distribution S_p of possible observed energies in MARS based on the quenched pulse energy E_p (units MeVee)

$$S_p = \frac{s(\vec{r}_p)}{\sqrt{2\pi\beta^2(\vec{r}_p)E_p}} \exp\left(\frac{-(E - E_p)^2}{2\beta^2(\vec{r}_p)E_p}\right), \quad (7.5)$$

such that the pulse energy in ADC units will be the random sample from S_p .

7.1.3.4 Threshold

As we described in Section 6.2.3.2, the MARS detector triggers only if the summed energy observed by a group of four PMTs exceeds a set threshold; this is done using the digitizer’s trigger logic and applying a “Finite Impulse Response” (FIR) trapezoidal filter with peaking and gap times of 16 ns. However, this does not lead to a strict cutoff in our energy spectrum – rather, there is some increased probability to observe events of a given energy with a defined threshold setting. For beam operations, we set an FIR = 300 ADC threshold, which is the lowest possible setting that produces negligible dead time in MARS during nominal SNS operations. We measure this directly by lowering the detection threshold below the nominal operating value and comparing the observed spectra. Taking the ratio of the observed counts per energy bin for the operating threshold compared to a setting which allows lower-energy events through, we characterize the signal acceptance fraction ϵ as

$$\epsilon = \frac{1 + \text{erf}(p_0 + p_1 E)}{2}, \quad (7.6)$$

where p_0 and p_1 are free parameters, to estimate the probability that an event with a given energy E in ADC units will pass the threshold cut, assuming that at some energy that probability will be effectively unity. Sample measurements of the detection threshold are shown in Figure 7.6, and we repeat the measurement in time to confirm the stability of the threshold function we implement in simulation. The bottom panel of Figure 7.6 shows the comparison of all measurements in time. Because there is no clear trend of the observed threshold in time, we create an “envelope” of the best fits, as shown by the black points, and use the best fit to this envelope as our threshold function in simulation. The parameters from this best fit are $p_0 = -3.47 \pm 0.03$ and $p_1 = (3.01 \pm 0.04) \times 10^{-4} \text{ ADC}^{-1}$. For each simulated pulse, we then evaluate the threshold curve at the value of the smeared energy deposit, and use that signal acceptance fraction to determine whether or not the pulse would be observed.

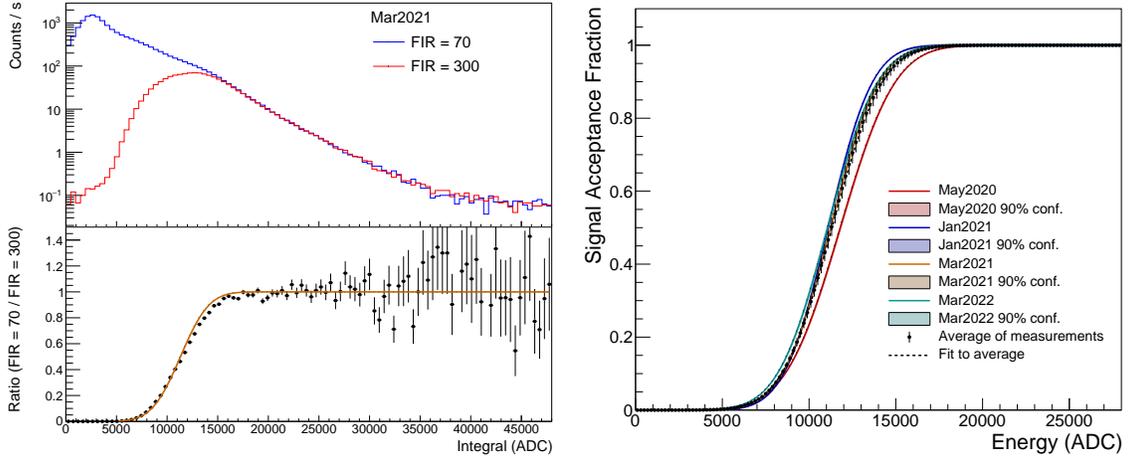


Figure 7.6: Left: An example measurement of the signal acceptance fraction for FIR = 300 operations of MARS in March 2021, using a separate run with FIR = 70 operations as a indicator of the signal that does not pass the higher threshold cut. Right: The best-fit sigmoidal threshold curves from each threshold measurement, along with the 90% confidence intervals on the fit. We take the average of the fits, and perform a fit to that average to determine the signal acceptance fraction we apply in simulation.

7.1.3.5 Pairing

At this point in our simulation processing of an event, we have a list of pulses which would pass the detection threshold. The pairing procedure is done just as for the analysis; we pair all pulses which occur within a $200 \mu\text{s}$ window of the first pulse's timestamp, maintaining time-ordering of all pairs. So, for a simulated event with three pulses (p_1, p_2, p_3) within $200 \mu\text{s}$, we would create three event-pairs $(\{p_1, p_2\}, \{p_1, p_3\}, \{p_2, p_3\})$.

These pulse-pairs then populate the output of our simulation processing; we track the raw energy, quenched energy, and smeared energy of both pulses in the pair, the energy-weighted time of the first pulse, and the elapsed time between pulses. This replicates the event-pairing in data, so with this procedure we generate distributions of E_1 , E_2 , t (where $t = 0$ occurs $800 \mu\text{s}$ prior to the corresponding E39 signal of the event), and Δt for particular sources incident on the MARS volume.

7.2 Characterizing the MARS Response

To understand the response of the MARS detector to incident neutrons, we must first understand its response to gamma-rays. The E_2 signal of all neutron events is the de-

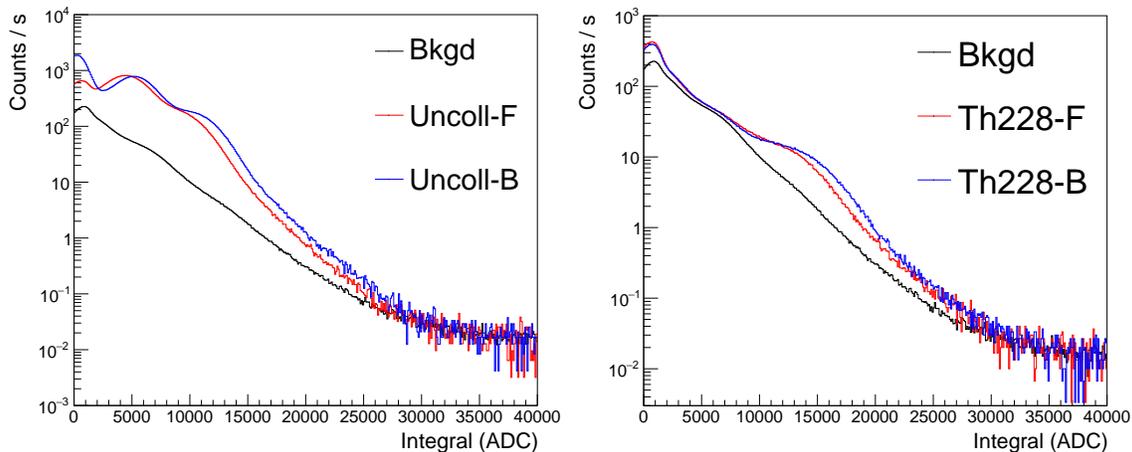


Figure 7.7: MARS data measured with uncollimated ^{60}Co (left, ~ 1 MeV γ s) and ^{228}Th (right, ~ 2.6 MeV γ) positioned on the alley-side (front) and wall-side (back) of the detector, compared to background (no-source) measurements during beam-off. The consistent shift of the observed spectrum to higher energies when the source is on the wall-side of the detector indicates that different smearing parameters should be considered.

excitation of gadolinium following a capture — a process which releases a shower of γ s summing to approximately 8 MeV. In this section, we overview efforts to characterize the position-dependence of the MARS response, develop the smearing model we apply in simulation, and validate the simulation processing we outlined in Section 7.1.3. We note that the majority of this effort is characterizing the detector while the SNS is not producing neutrons (e.g., maintenance days, scheduled shutdowns), but Section 7.2.4 briefly discusses observed changes to the MARS response during beam-on operations.

7.2.1 Gamma Reponse and Calibration

Cobalt-60 is a common radioactive nuclide that emits two γ -rays with energies of 1.17 MeV and 1.33 MeV after its decay to an excited state of Nickel-60. These incident γ s generally undergo either the photoelectric effect or Compton scattering. Both processes result in fast-moving electrons in a scintillating material — the electrons ionize nearby atoms, and the resulting light can be observed with the PMTs affixed to the edges of MARS. A ^{60}Co event does not produce a pulse-pair; the gammas quickly lose energy in the scintillator, so single pulses contain the full energy deposit.

Sample spectra for an uncollimated source are shown in the left panel of Figure 7.7. Note that the main peak of this signal is just below 10000 ADC, while the measured sig-

nal acceptance fraction from Section 7.1.3.4 is only $\sim 20\%$ in this range. In order to see the signal from this calibration, we must lower the detection threshold. We perform the calibrations discussed in this section at FIR = 30, which accepts near-100% of the signal above 3000 ADC^{9,10}.

MARS is a large detector, and has only 16 PMTs on two edges of the scintillating volume with which to measure the interactions. Thus, an understanding of the position-dependence of the MARS response is crucial to developing a full response matrix. We use a calibration grid of 20 positions (5 rows, 4 columns) as illustrated in Figure 7.8 to measure changes in the light collection of MARS as a function of position. After identifying differences in uncollimated measurements with different calibration sources¹¹ taken on the front (closest to the alley) and back (closest to the wall) of the detector, illustrated in Figure 7.7, we began repeating this 20-position grid on the front and back of the detector (40 total calibration scans) in early 2022.

For each position in the grid scan, we simulate the response of MARS to the cobalt-60 source, including the attenuation of the incident γ -rays as they interact in the collimator. This allows us to account for the geometric effects on the raw spectrum separately from the effects on the light collection; for example, events near the corners have a higher probability to escape the detector before depositing all their energy and will shift the raw response spectrum lower, but these events are also closer to the PMTs and will shift the observed response spectrum higher.

We calculate the simulated pulse energy for each event in MeVee using the processing we described in Section 7.1.3, shown in the right panel of Figure 7.9 for an example scan position near the center of MARS. Each event from this simulated distribution has a raw energy E which is smeared to create a model of the expected observation in data. This smearing is performed by constructing a Gaussian distribution with mean $\mu = E$ and a variance $\sigma^2 = \alpha^2 E^2 + \beta^2 E + \gamma^2$ [202–205] to model the detector resolution. The smearing parameters α , β , and γ each account for different parts of the signal production: α relates

⁹We do not apply a different threshold curve in simulation, because the smeared pulse energy is not used in this section – we derive the position-dependent map of smearing parameters in this section.

¹⁰As we previously noted, our threshold setting was chosen because it prevented dead-time in the DAQ. We can safely lower the threshold during beam-off operations, but we do introduce dead-time if the SNS is producing neutrons (see Section 7.2.4).

¹¹Thorium-228 is an unstable isotope with a long decay chain; it undergoes several α and β decays before reaching the stable isotope Lead-208. At low energies, the emission spectrum is more complicated than ^{60}Co , but the beta decay of Thallium-208, in particular, emits a 2.6 MeV gamma that is a clean observable above the MARS threshold.

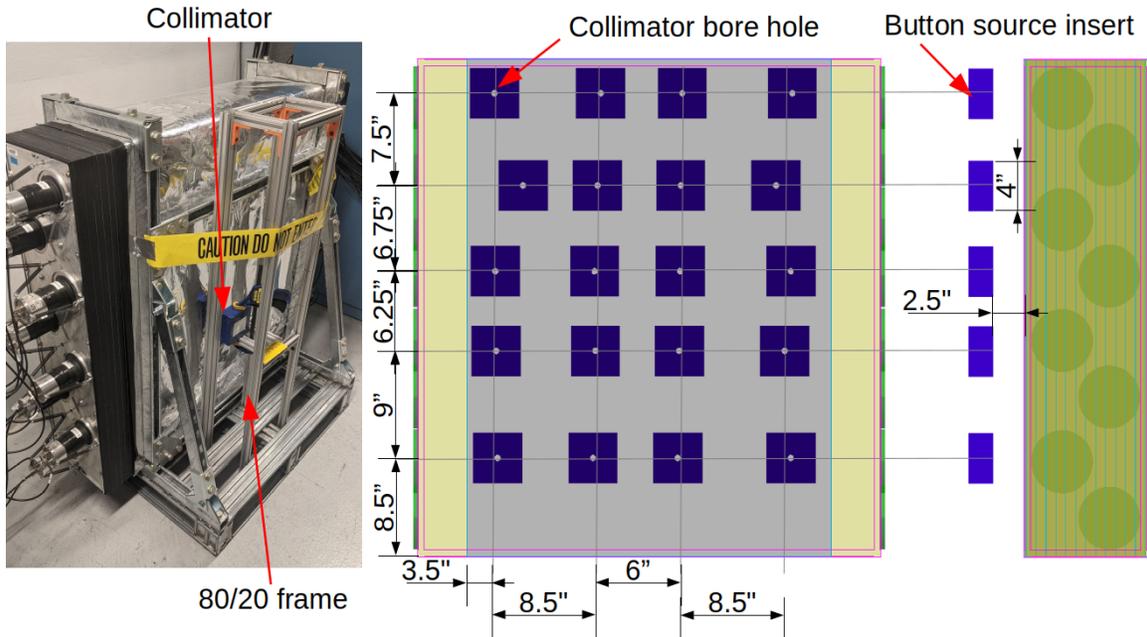


Figure 7.8: Left: A photo of the configuration for taking data with a cobalt-60 source. We place the source within a $4'' \times 4'' \times 2''$ collimator, and control the position of the collimator using a fabricated test stand using spare 80/20 aluminum framing available around the lab. The button source extends approximately $1/8''$ into the $2''$ thickness of the collimator. Center: A front-view diagram showing the position of the collimator relative to the MARS scintillating volume for each of the 20 grid positions on the alley-side of the detector. Right: A side-view diagram showing the locations of the collimator relative to the PMT positions (dark green).

to the propagation of light from the interaction point to the photocathode, β accounts for statistical variation in the light production and electronic amplification (the Fano Factor), and γ accounts for electronic noise [203]. We then apply a linear scale factor to convert the smeared energy deposit from MeVee to ADC units ($E'[\text{ADC}] = sE[\text{MeV}]$). In some cases, it is appropriate to approximate $\alpha = \gamma = 0$ [206]; prior MARS work followed this strategy [188, 207], and we follow suit in this work but note that we will need to reevaluate this in future as exploratory studies using nonzero α and γ parameters demonstrate a better fit to the data. The full smearing model mimics the distribution expressed in Equation 7.5, and we sample from that distribution to fill the simulation histogram shown in the right panel of Figure 7.9.

We use a procedure developed by Sam Hedges [103] to determine the values for the resolution β and scale s which best fit the data at each position in the scan. This strategy

samples through the user-defined parameter space using *emcee* [208] and uses RooFit [209] to minimize the negative log likelihood (NLL) of the model.

The likelihood function is a measure of how effective a predictive model is at predicting each data point. As an example, this could be done with a Gaussian estimation: for a dataset of N measurements Y_i with known errors σ_i , and fit parameters $\vec{\theta}$ such that the function $f(X_i, \vec{\theta})$ is a model of the dataset, the likelihood function is given by

$$L(Y_i|\vec{\theta}) = \prod_{i=1}^N \frac{1}{\sqrt{2\pi\sigma_i^2}} \exp\left(\frac{-(Y_i - f(X_i, \vec{\theta}))^2}{2\sigma_i^2}\right). \quad (7.7)$$

Note that this function will return its maximum value when the model, using parameters $\vec{\theta}$, most accurately describes the dataset. The NLL, referring to $-\ln L(Y_i|\vec{\theta})$, would then be minimized for the given $\vec{\theta}$.

We find the best-fit smearing parameters $\vec{\theta}(\vec{r}) = \left\{ \frac{N_s}{N_b + N_s}, s, \beta \right\}$ for each position in our ^{60}Co scan of MARS using a predictive model of the data

$$f(\vec{\theta}) = N_s S(\vec{\theta}) + N_b B, \quad (7.8)$$

which is a sum of the probability distribution functions (PDF) from smeared simulated signal $S(\vec{\theta})$ (calculated from Equation 7.5) and measured backgrounds B . We require N_b to be within 2% of the expectation from measured background rates¹², but allow N_s to vary by up to 15% of the expectation from background-subtracted source data. We restrict β to take values between 0.15 and $0.5 \sqrt{\text{MeVee}}$ (corresponding to an energy resolution between 15 and 50% for 1 MeV γ s), and the energy scale s to take values between 3500 and 7000 $\frac{\text{ADC}}{\text{MeV}}$.

We perform a binned fit for computational efficiency¹³, so the N measurements Y_i then represent the measured rate in each energy bin X_i , and we use the model $f(X_i, \vec{\theta})$ to predict the expected rate, F_i , in that energy bin. Rather than a Gaussian estimation, we use a Poissonian measure of the fit quality (to account for potentially low statistics), $\frac{F_i^{Y_i}}{Y_i!} e^{-F_i}$. The NLL we minimize is given by

$$-\ln L(Y_i|\vec{\theta}) = -\sum_{i=1}^N [Y_i \ln(F_i) - F_i], \quad (7.9)$$

where we neglect the term $\ln(Y_i!)$ because it has no dependence on the fit parameters $\vec{\theta}$.

¹²We take 30-minute background scans, but 5-minute source scans, so we allow the normalization of the source to float given its higher statistical uncertainty. Generally, the background rates do not change during beam-off calibration scans.

¹³Unbinned fits took a full week to process. Binned fits took an hour.

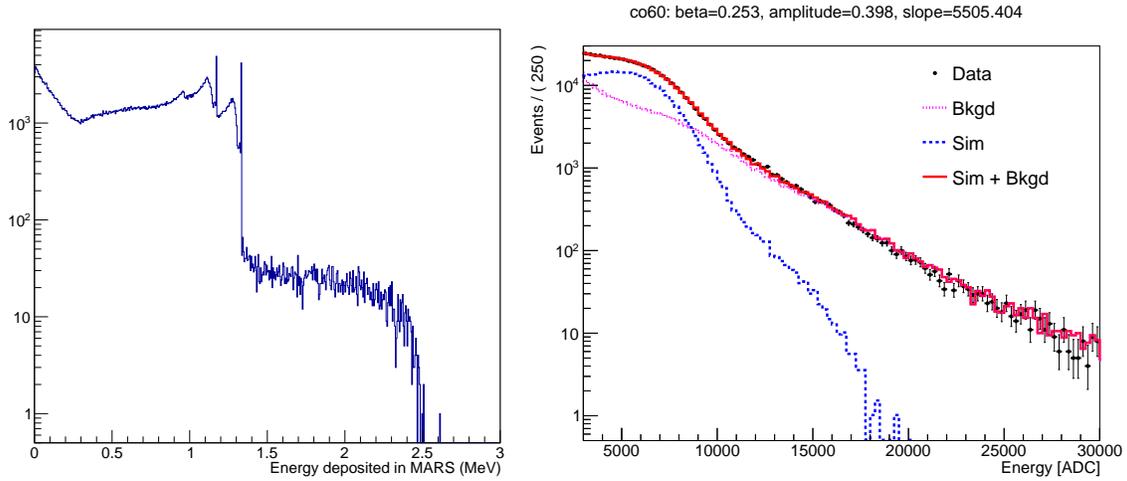


Figure 7.9: The fitting procedure for an example calibration position (Row 4 Column 2). Left: The distribution of raw simulated pulse energies observed in MARS from a collimated cobalt-60 source. Right: The best-fit result for a model of the smeared simulation plus measured background to the measured data in the presence of the cobalt-60 source. The blue curve of the right panel represents the smeared energy spectrum MARS will realistically observe from cobalt-60. The values listed atop the figure refer to the best-fit smearing parameters “beta” = $\beta(\vec{r})$, “amplitude” = $\frac{N_s(\vec{r})}{N_s(\vec{r})+N_b(\vec{r})}$, and “slope” = $s(\vec{r})$. The parameters $\beta(\vec{r})$ and $s(\vec{r})$ feed back into Equation 7.5.

With best-fit smearing parameters from each of the 40 locations, we produce a position-dependent light response model. We show the results for the roughly-annual calibrations performed on the front (alley side) of MARS in Figure 7.10. May 2022 was the first calibration to take data on the front and back of MARS, which identified a $\sim 7\%$ increase in the light collection on the wall side of MARS compared to the alley side. The exact cause is not yet known, but we are studying the possibility of air-gaps between layers on the front of MARS and are improving our simulation model to include possible reflected events from the wall.

To apply these maps in simulation, we manually extrapolate our experimental observations to the edges of the scintillating volume so that we can perform a linear interpolation between the observations for any possible pulse position. We studied several strategies for the extrapolation to the edges, and found minimal differences between approaches to the vertical extrapolation. We choose a method that fits the data from each of the four columns in the scan with a constant, and assign the best fit value to the top and bottom edges of MARS ($y = \pm 50$ cm) in that x position as illustrated in the center panel

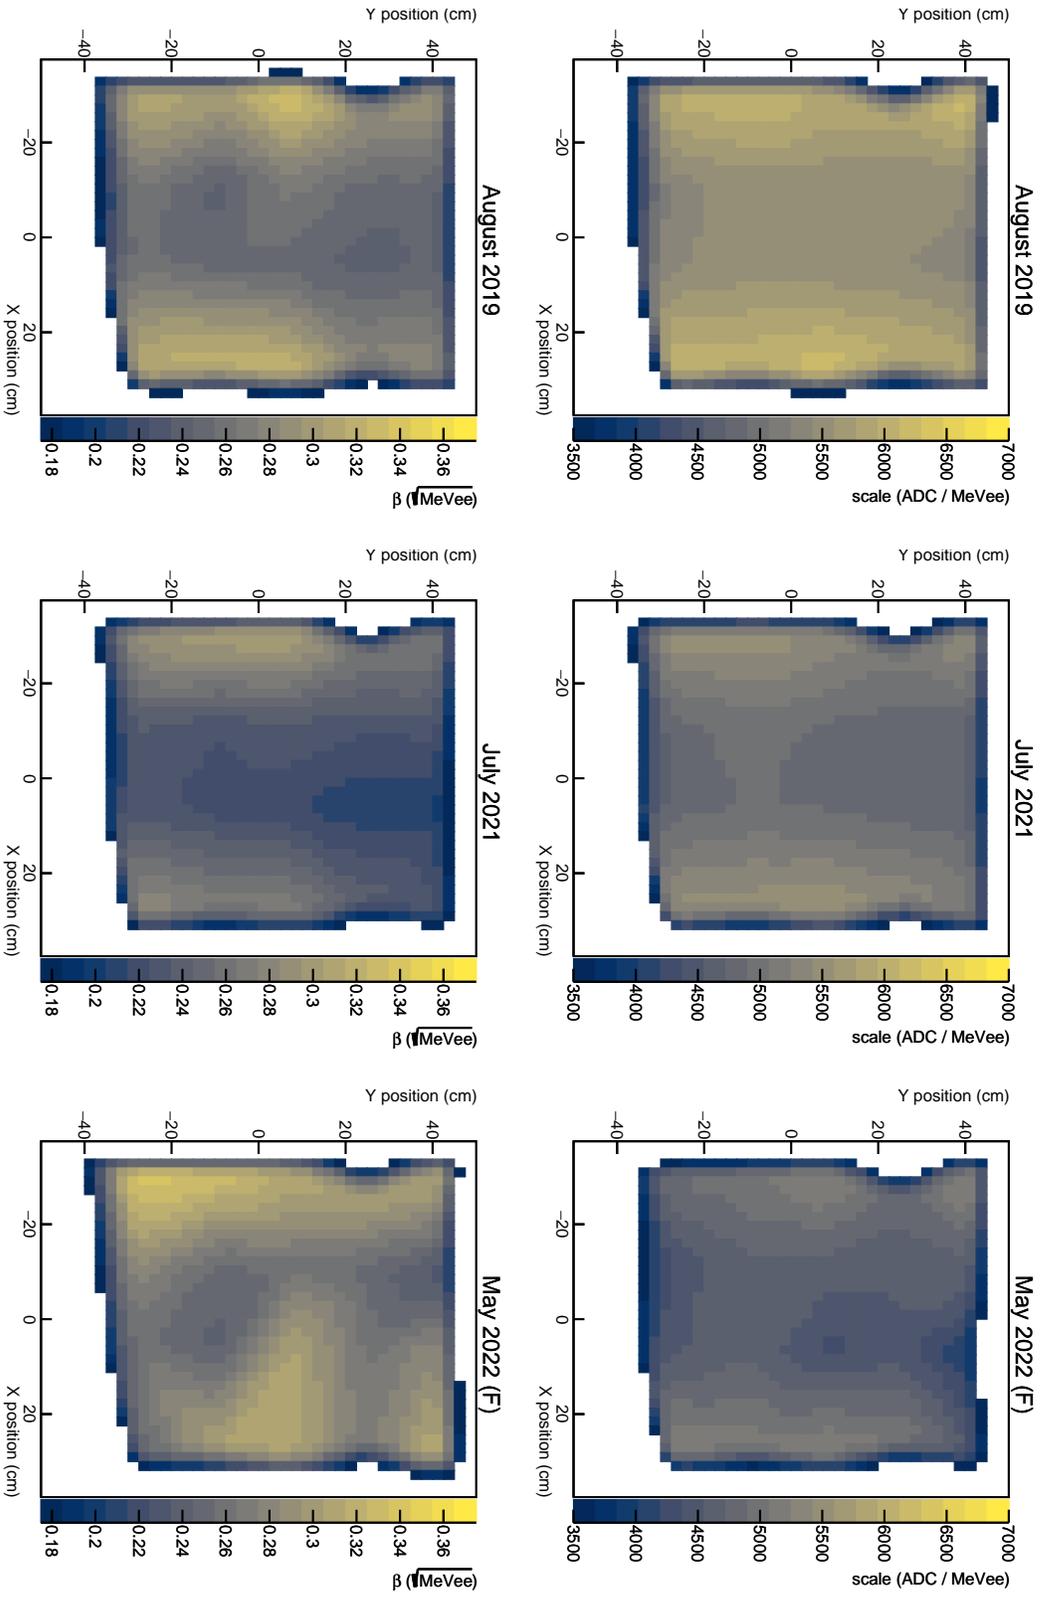


Figure 7.10: Two-dimensional maps of the smearing parameters collected for calibrations on the alley-side of MARS at the indicated date. The top plots show the values for the energy scale s , while the bottom plots show the values for the energy resolution β . The z-axis color scales are consistent, such that variations of color between the plots indicate the differences in the detector performance measured by the calibration. The border of the plots indicate the edges of the MARS scintillating volume to illustrate the fraction of the volume that is affected by the extrapolation procedure described in the text.

of Figure 7.8.

The extrapolation to the horizontal edges is more complicated, as the locations of the PMTs heavily impact the observation. We choose to extrapolate each row using a quadratic fit, expecting that the light yield will be higher near the edges of MARS. Then, we evaluate the best-fit quadratic at $x = \pm 35$ cm (the left/right edges of the scintillator) for each of seven rows: the five rows of the calibration scan (at y position illustrated in the center panel of Figure 7.8) and the two artificial rows at $y = \pm 50$ cm that we developed in the vertical extrapolation. This strategy ensures that there are values assigned to the corners of the scintillating volume.

To summarize, we extrapolate the smearing parameters as constant in the vertical dimension and quadratic in the horizontal dimension. Including the points we add to the map from the extrapolation, we then have 42 points (20 measured, 8 from the vertical extrapolation of four columns, 14 from the horizontal extrapolation of seven rows) describing the two-dimensional smearing map on either the front or back of MARS. These maps span the full scintillating volume of MARS, so for any pulse position $\vec{r}_p = \{x_p, y_p, z_p\}$ in simulation, we can determine the smearing parameters ($s_p = s(\vec{r}_p), \beta_p = \beta(\vec{r}_p)$) to use. We use the built-in Delaunay interpolation from ROOT's TGraph2DErrors object to determine the smearing parameters at $\{x_p, y_p, \pm 12.5 \text{ cm}\}$ on the front (s_f, β_f at $z_f = 12.5$ cm) and back (s_b, β_b at $z_b = -12.5$ cm) of the detector. We then perform a linear interpolation between these front and back parameters, such that

$$s_p = s_f + (z_p - z_f) \frac{s_f - s_b}{z_b - z_f}, \quad (7.10)$$

$$\beta_p = \beta_f + (z_p - z_f) \frac{\beta_f - \beta_b}{z_b - z_f}. \quad (7.11)$$

The fully extrapolated results from May 2022 are shown in Figure 7.11. For older scans, where we had not yet identified the discrepancy between measurements on the front and back of MARS, we assume a crude scaling of $s_b = 1.07s_f$ and $\beta_b = 0.93\beta_f$ based on the results from May 2022. The smearing parameters s_p and β_p can then be used in Equation 7.5 to generate a distribution of possible observations for the given simulation event.

7.2.2 Michel Electron Validation

To validate the calibration and the implementation of this light map in simulation, we target other sources of interest. In particular, Michel electrons are produced when cosmic-ray muons come to rest and decay within the MARS volume ($\mu^- \rightarrow e^- + \nu_\mu + \bar{\nu}_e$), leaving

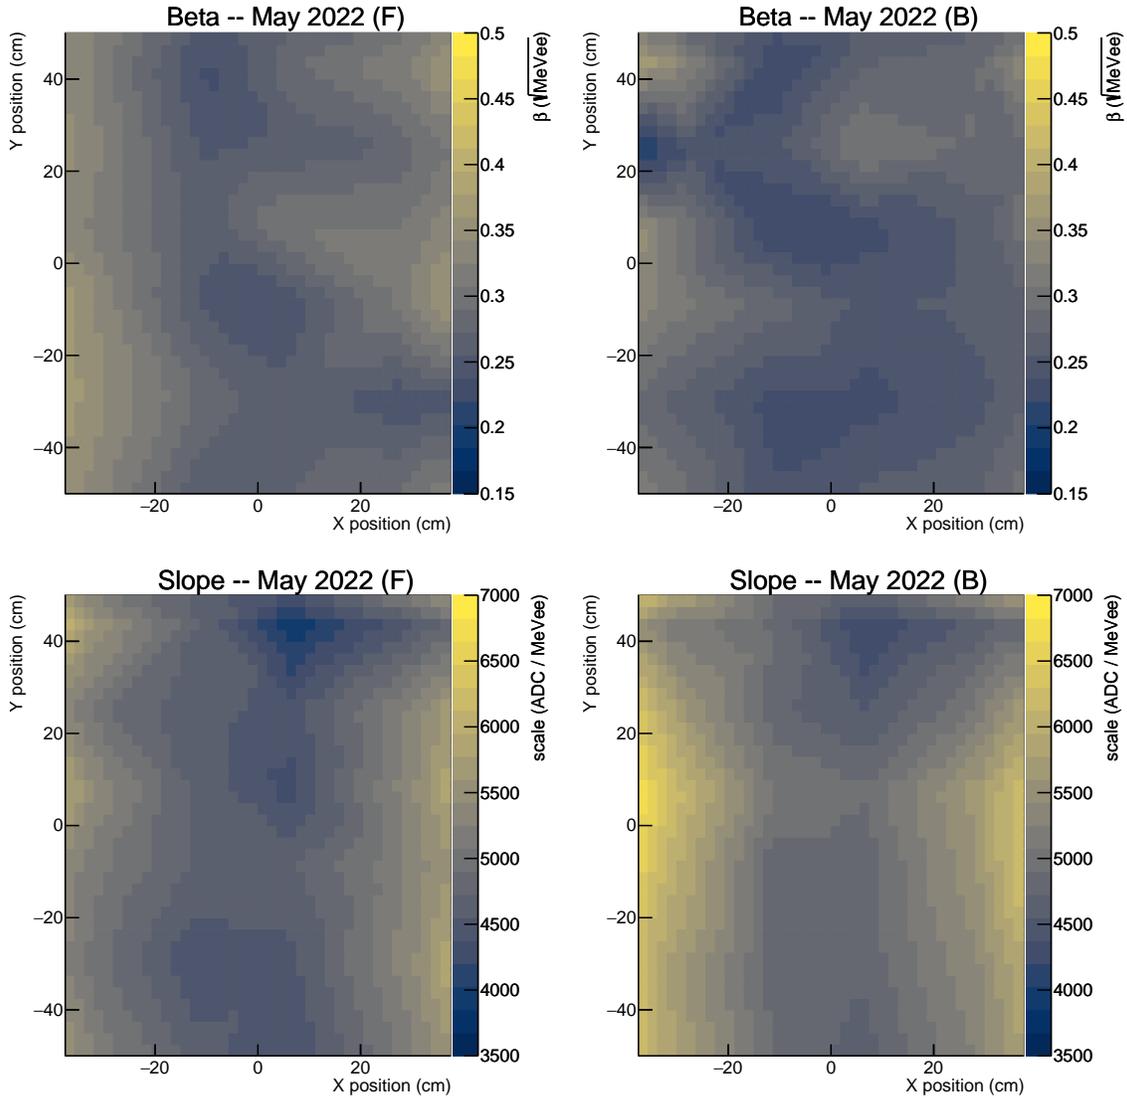


Figure 7.11: A comparison of the fully extrapolated front (left) and back (right) maps of the resolution β (top) and scale s (bottom) for the ^{60}Co calibrations conducted in May 2022.

an event-pair signature: the energy deposited by the muon as it comes to rest, and the energy deposited by the electron after the muon decays. With a spectrum in the tens-of-MeV range (see Equation 5.6 for expression), Michel electrons provide a test of the linearity of our light map, and since electrons are not quenched, the MARS response to this calibration source can directly validate our conversion from MeVee to ADC units.

To identify Michel electrons in the data, we place a cut $E_1 > 300$ kADC to identify

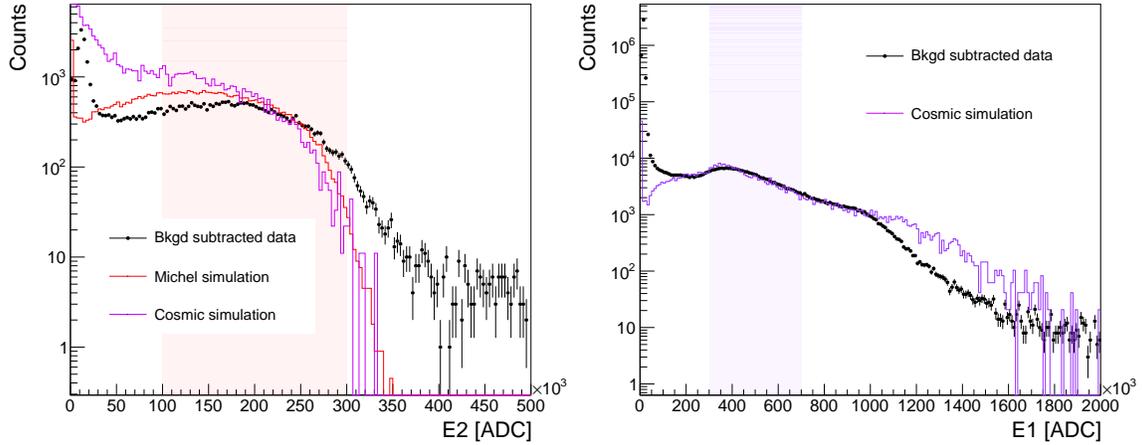


Figure 7.12: Left: The Michel electron feature of the MARS E_2 spectrum, compared for our two simulation cases and background-subtracted data. Right: A comparison of the observed E_1 spectrum in MARS during beam-off operations, which is believed to be largely associated with cosmic rays, to a simulation of the MARS response to cosmic rays using the CRY module.

the large energy deposit of muons and $\Delta t < 10\mu s$ to identify muon-decay events (this is nearly 5 times larger than the muon lifetime $\tau_\mu \approx 2.2\mu s$). With these cuts, the E_2 spectrum is expected to describe the energy deposited from Michel electron events. We apply the same cuts to both simulation and data, and compare the observed energy deposit in MARS to the fully smeared energy deposit from simulation.

In simulation, we model Michel electrons through two independent methods. Our first strategy uses the CRY package [173] to model cosmic rays as they pass through the overburden of Neutrino Alley, identically to the D₂O simulations in Chapter 5, and simulate the full event of a muon coming to rest in MARS and producing a Michel electron. The second strategy entails simulating electrons distributed throughout the MARS volume, in random directions, with energies sampled from the Michel spectrum (see Equation 5.6). In both cases, we fail to accurately predict the data – there is about a 15% discrepancy between the predicted endpoint of the E_2 spectrum in simulation and the observed endpoint from data, illustrated in Figure 7.12 (left).

The benefit of the CRY simulation of cosmic rays is that we can also examine other spectral features, as shown in the right panel of Figure 7.12. For example, a clear observable in the MARS E_1 spectrum is the so-called “30°” bump above 300 kADC; this is where muons pass through MARS at a roughly 30° angle, thus passing through the majority of

the scintillator and leaving a measureable energy deposit as they come to rest. Our simulation clearly predicts this peak, but fails to accurately model the higher-energy behavior of MARS (believed to be saturation of the observed signal). Further studies into what the feature around 1100 kADC represents and whether there are ambient backgrounds to subtract from the raw MARS spectrum we use as a comparison point are underway.

The discrepancies with the Michel electron data, however, suggest that our smearing model could be incomplete¹⁴ or could point to nonlinearity in the MARS response. We are actively investigating both possibilities. A complete model might require α or γ parameters, which could smear the data towards higher energies; preliminary studies have shown that the inclusion of these additional smearing parameters may better describe the data. To examine the nonlinearity, we pursue calibration with additional sources. We have conducted ²²⁸Th calibrations, which produce 2.6 MeV γ s, which preliminarily show a consistent discrepancy between simulation and data. We have also performed new calibrations with an AmBe source, which produces a 4 MeV γ (along with neutrons), that have not yet been analyzed. Finally, we explore the Gd de-excitation spectrum from our neutron calibrations as described below.

7.2.3 Neutron Response and Calibration

During the deployment and commissioning of MARS in late 2017, we calibrated the detector using a californium-252 source, which produces multiple neutrons of a few MeV each. Measurements with ²⁵²Cf provided preliminary information to the analysis about the E_2 spectrum following a neutron capture, prior to the development of both the Geant4 simulation described in this chapter and our current analysis methodology described in Chapter 8. In support of the data we had previously collected, we built a simulation model for this calibration, and identified that the majority of the neutron spectrum is not observable with MARS using its typical operating threshold; the measured E_1 of the event-pair results from γ s produced in the fission, which will help to calibrate the neutron detection efficiency *if* we know that our simulated response to neutrons is accurate. To validate our model of both the neutron response and the efficiency, this source was not the practical choice.

As such, we turned to a deuterium-tritium (DT) neutron generator, which initiates fusion reactions between deuterium and tritium that produce a monoenergetic 14.1 MeV

¹⁴It was examinations of this discrepancy that ultimately uncovered that the front and back of MARS respond differently to calibration sources.

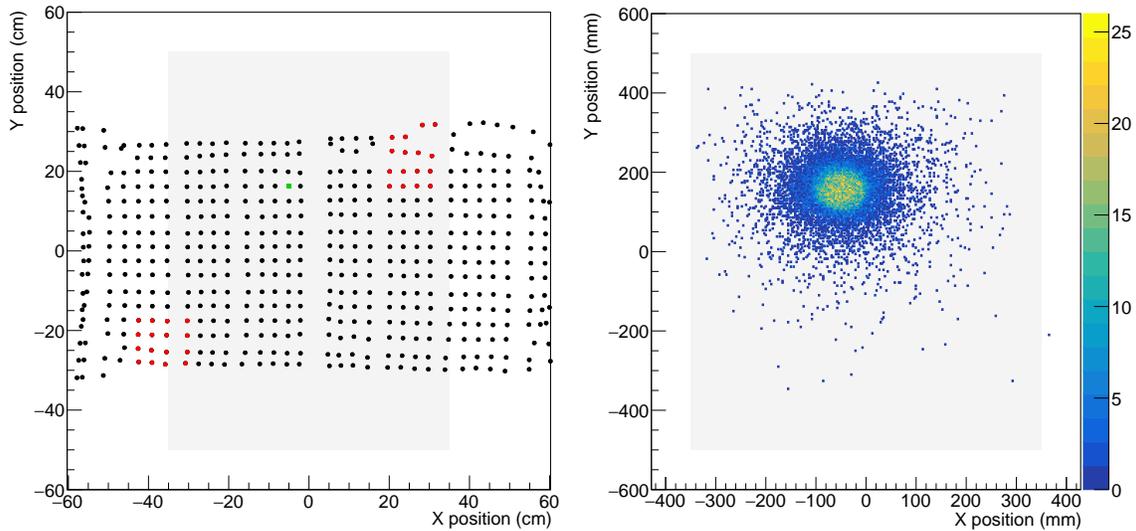


Figure 7.13: Left: The simulated neutron direction associated with the center of each backing detector pixel onto the MARS scintillating volume (gray rectangle). The red points indicate segment 5 (top right) and segment 25 (bottom left), which are the example comparisons used throughout this section. Right: The simulated distribution of E_1 pulse positions associated with pixel 110. The projection of this pixel onto MARS is indicated by the green point in the left panel.

neutron ($D + T \rightarrow n + {}^4\text{He}$). This calibration is further benefitted by the emitted 3.5 MeV $\alpha = {}^4\text{He}$, which must be in the opposite direction of the neutron to conserve momentum, and many DT generators provide a “backing” detector to tag the α emission time and direction. These characteristics provide a fantastic calibration source of monoenergetic neutrons with measured time and incident direction. Data was taken with an API DT generator provided by ORNL placed across the alley from MARS (46 inches from the front of the detector) in October 2019¹⁵.

While our validation against Michel electrons is a useful check, this DT calibration is the true test of the validity of our simulation model. Comparisons of our simulations to experimental DT data enable validations of our model of the Gd concentration in MARS, the Gd de-excitation physics and emission spectrum, the application of energy resolution parameters from the cobalt-60 calibrations, the neutron thermalization and our quenching model, and the predicted neutron-capture efficiency. In this section, we overview some preliminary comparisons between simulation and data for the DT calibration.

¹⁵The analysis of that data was performed by Belkis Cabrera-Palmer [180] and Justin Raybern [207].

We simulate the DT source as an isotropic, monoenergetic 14.1 MeV neutron emitted from 46 in away from the front face of MARS. The DT generator provided by ORNL features a backing- α detector with 512 pixels (32×16), each with a 4.5° FWHM angular acceptance¹⁶. We can then associate each experimental event with the triggered pixel to create a position-dependent validation of our simulation model. In Figure 7.13 (left), we show the projection of the central neutron direction associated with a trigger in each backing detector pixel. In simulation, we calculate the angle between the neutron direction and the projected direction from every backing detector pixel to determine if the pixel could have triggered. For example, the right panel of Figure 7.13 illustrates the distribution of E_1 pulses which are associated with a single pixel of the backing detector. We group pixels by their position, such that we compare “segments”, or non-overlapping groups of 16 pixels (4×4) which are all very near each other, to increase the statistics without sacrificing the position-dependent analysis.

In Figure 7.14, we show comparisons of simulation to data for the event-pair Δt (the time between E_1 and E_2) given that $E_2 \in [23, 57]$ kADC, which is expected to depend on the Gd-concentration and the thickness of the paint layers¹⁷. We use the DT calibration data to tune our simulation to match the observation. In this work, we use a Gd concentration of 15% and 0.07-cm thick paint layers, which demonstrates consistency with the measured capture time-constant of data across all pixels. However, the overall normalization, which is a measure of the neutron detection efficiency, is not correct. Ultimately, we will tune the simulated efficiency to this calibration measurement by altering the Gd concentrations and paint-layer thicknesses, as well as by re-examining the underlying physics model to ensure its validity. To further tune our simulation, however, we must first validate the spectral response.

We compare the spectral observations of E_1 and E_2 in Figure 7.15 by normalizing the simulation curves to have the same integral as the measured data. There are noteworthy shape differences in both the E_1 and E_2 spectra. Focusing on the E_2 spectrum shown in the right panels, we note a sharp excess near 20000 ADC, followed by a deficit in the tails of the spectrum. This result further indicates that we do not effectively smear our simulation, but it is worth noting that the endpoints differ by roughly 15% — a consistent

¹⁶This backing detector is a 100 μm thick YAP:Ce crystal mounted on a fiber-optic faceplate. A Hamamatsu H13700 multi-anode flat panel PMT coupled to a resistor network is used to collect the scintillation light from the crystal, and the four scaled outputs enable position reconstruction using Anger-logic [180, 210].

¹⁷Higher amounts of Gd provide less opportunity for a neutron to escape the paint layer and bounce around the scintillator before finding the next paint layer.

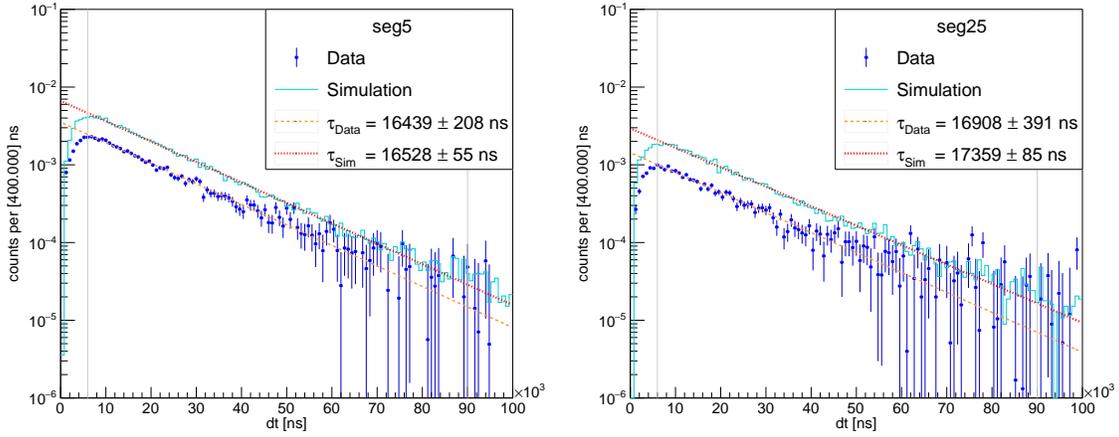


Figure 7.14: A comparison of the simulation prediction of the time between events in a pair to the experimental observation for segments 5 (top) and 25 (bottom). Segment 5 is near the top-left of the detector, while segment 25 is near the bottom-right of MARS. These segments include the pixels highlighted in red from Figure 7.13.

result with the deficit of the predicted endpoint for the higher-energy Michel electron spectrum. Note that the disagreement of the spectra is also the limiting factor for further tuning of the Gd concentration, as we cannot accurately describe the number of counts meeting a specific energy cut unless we trust our prediction of the energy spectrum.

The comparisons to E_1 are significantly worse; visually, we can see that the overall spectrum is shifted towards lower ADC values. We expect that this is due to the lower light yield of the quenching model we implement in simulation compared to the observed data. The peak of the E_1 spectrum is roughly 25% lower in simulation than in data, which is consistent with the deficit for low-energy proton recoils demonstrated in Figure 7.4. We will resolve this in future work, but we must first ensure the smearing model accurately predicts the E_2 spectrum, which is unaffected by quenching¹⁸.

7.2.4 SNS-ON Characterization Efforts

The presence of the hot-off-gas (HOG) background in Neutrino Alley during beam-on operations has an observable impact on MARS. There are measurable differences to known energy spectra, such as the Michel electron spectrum shown in Figure 7.16. In particular, during beam-on operations, the presence of the HOG background provides a constant flux

¹⁸The energy deposits which contribute to E_2 are electronic recoils caused by γ s from the Gd de-excitation.

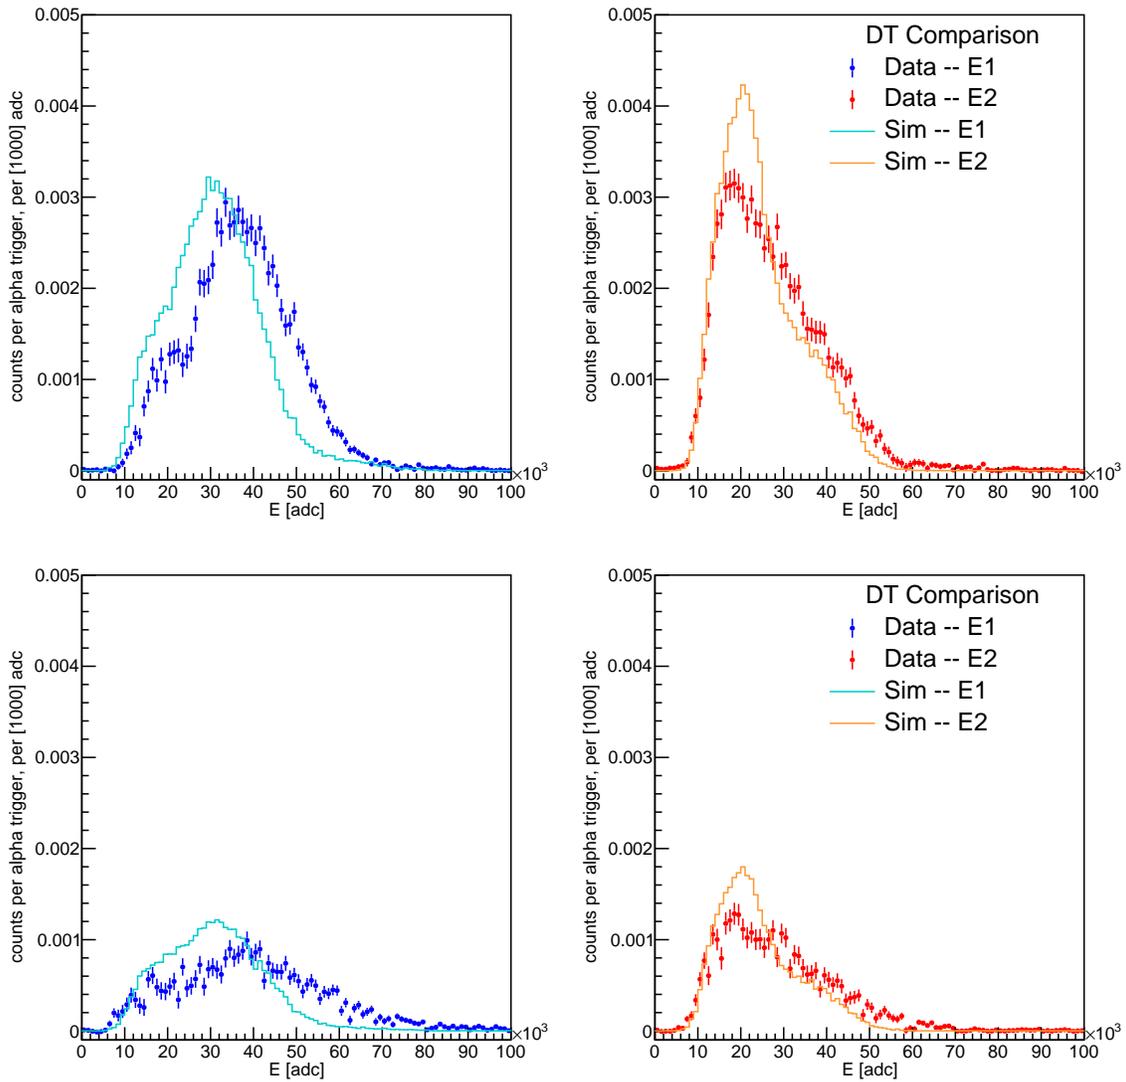


Figure 7.15: A comparison of the simulation prediction of the E1 and E2 spectra to the experimental observation for segments 5 (top) and 25 (bottom). Segment 5 is near the top-left of the detector, while segment 25 is near the bottom-right of MARS. These segments include the pixels highlighted in red from Figure 7.13.

of 511 keV γ s. These are below threshold and will not independently trigger MARS, but they do deposit energy in MARS and increase the measured integral for typical events. The collaboration is deploying 2" thick lead shielding around the HOG pipe to reduce the impact of this background, which will improve our future measurements.

To characterize the response of MARS during beam-on operations for past data, we

have performed ^{60}Co calibrations during nominal beam-on operations. These analyses are still ongoing due to dead-time issues lowering the FIR threshold to observe the spectral features of our ^{60}Co calibrations, along with quickly-changing HOG rates that have a measurable impact on the background spectrum in the region of interest.

The normalization of the observed spectra is an important input to our fitting procedure, as discussed in Section 7.2. The total run-time is not an accurate normalization parameter when the DAQ suffers from dead time during low threshold operations in high HOG background rates. We improved our calculation of the run duration using an external 60 Hz timing signal provided by the SNS (E39) to monitor the detector throughout the calibration; rather than normalizing by the DAQ run duration, we can instead normalize by the duration when MARS was able to collect data.

We also deployed an independent 2" $\varnothing \times 2$ " NaI cylindrical crystal with a single PMT atop MARS to monitor the HOG background rates in the summer of 2021. The right panel of Figure 7.16 shows the observed differences between 24-hour run periods; in general, we are able to monitor changes in the HOG rates minute-by-minute and will be able to use this to characterize observed changes in the MARS energy spectra. Given the considerable challenges we've faced in the characterization of MARS during beam-on operations, we have also begun a dedicated waveform study to search for small 511 keV γ events in the MARS response.

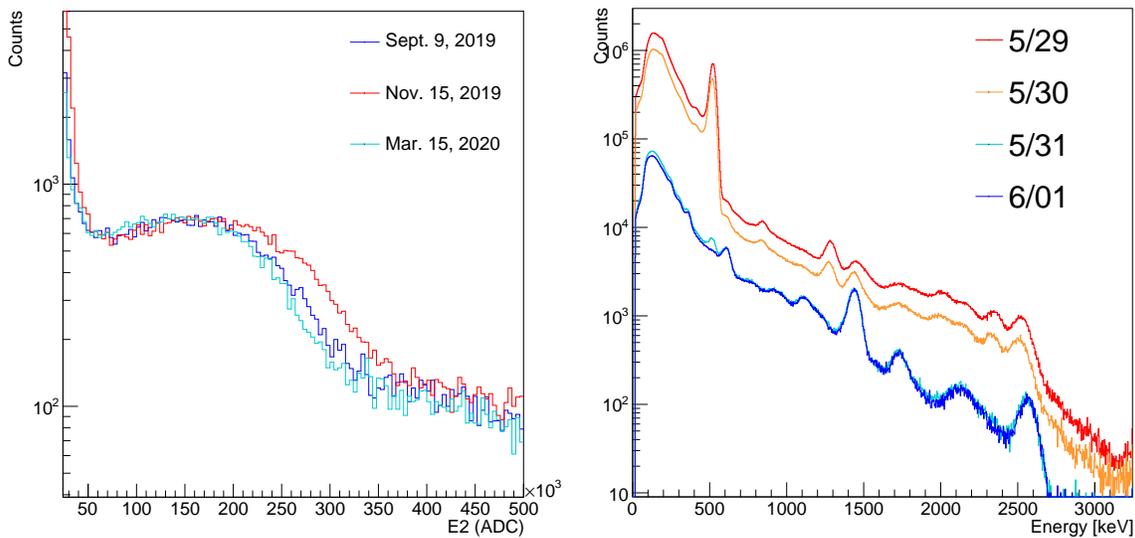


Figure 7.16: Left: A plot of the observed E_2 spectrum during 24-hour run periods on the given calendar dates. The SNS was running at a nominal 1.4 MW on November 15, 2019, but was shut down for maintenance on September 9, 2019 and March 15, 2020. We note an observable shift in the observed Michel electron spectrum during beam-on operations, in addition to the gradual decrease of the detector gain over long periods of time (see Chapter 8). Right: The daily recorded spectrum in late May and early June 2021 from an independent $2''\varnothing \times 2''$ NaI cylindrical crystal affixed to the top of MARS. The SNS shut-down started on 5/31/2021, and we notice an immediate decrease of the 511-keV peak associated with the HOG background.

7.3 Efforts towards Unfolding Neutron Spectra

The overall normalization of the simulated DT calibration is a measure of the simulated neutron capture efficiency. As we observed in Figure 7.14, simulation consistently over-predicts the capture efficiency, which is related to the expected number of events per bin. Ultimately, we aim to demonstrate consistency between simulation and data for 14 MeV neutrons by predicting the E_1 and E_2 spectra to within a few percent. Then, we can use simulation to derive a full energy-dependent efficiency, tuned to the experimental measurement of the efficiency at 14 MeV. We measure the energy deposited by BRNs in Neutrino Alley in ADC units, and this procedure will convert this spectrum back to the “raw” incident neutron energy in MeV units. With our efforts so far, we can make some preliminary comments about the BRN spectrum.

In the left panel of Figure 7.17, we show the observed BRN spectrum from Run 2

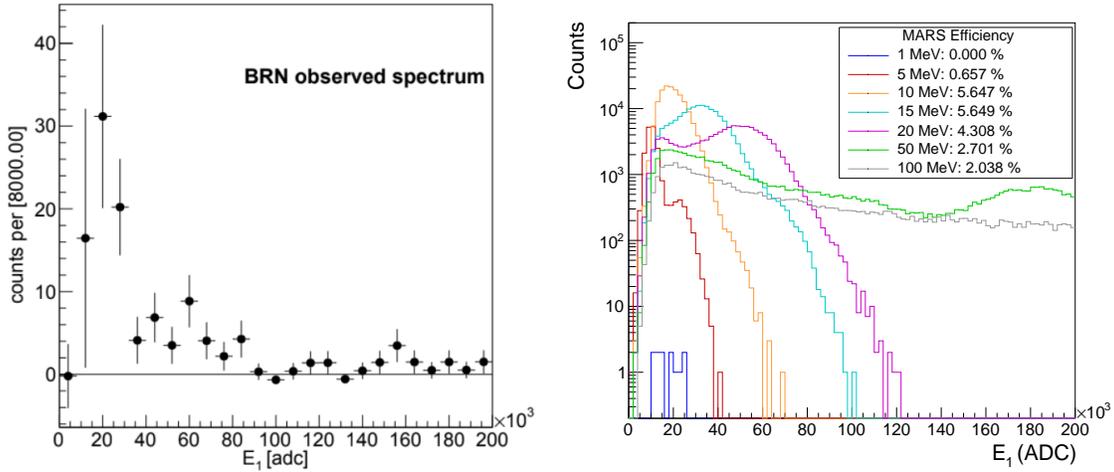


Figure 7.17: Left: The observed spectrum of BRNs in MARS from Neutrino Alley, reproduced from [180]. Further details about the measurement of the E_1 spectrum of incident BRNs is contained in Chapter 8. Right: The simulated spectrum of neutron events with given energy; we apply all analysis cuts and thresholds to predict the MARS efficiency for each neutron energy. We can infer from this that the likelihood of observing high-energy neutrons in Neutrino Alley flux is quite low given the lower capture efficiency at higher neutron energies; the BRNs of Neutrino Alley are most likely to have few to tens of MeV of kinetic energy given the observed excess in the sub-100 kADC range.

(more details in Chapter 8) in ADC units. The right panel of Figure 7.17 shows a series of simulation curves for monoenergetic neutrons. Even with the overprediction of the capture efficiency in simulation, we do not expect MARS to observe a measurable rate of BRNs below 5 MeV. Turning our attention to the peak in the data at 20000 ADC, we observe that ~ 10 MeV incident neutrons will generate a similar signal. There are very few BRN counts observed above 100000 ADC, suggesting that neutrons above 20 MeV do not comprise a significant fraction of the BRN flux in Neutrino Alley.

In future, we will use a Markov-Chain Monte-Carlo (MCMC) strategy to take all of the characterization information (capture efficiency, spectral response, quenching model, etc.) we've outlined in this chapter and unfold the incident spectrum. We have made significant progress towards this goal, and uncovered several avenues of exploration to further validate our simulation model. In particular, our immediate efforts must focus on the development of our smearing model; the efficiency measurements and spectral unfolding are dependent on this piece, which we cannot yet trust. Further calibrations using higher-energy γ sources, during both beam-off and beam-on operations, are being

explored, and we also plan to reexamine our choices about smearing parameters. Yet, the detection efficiency and spectrum of BRNs in Neutrino Alley is only half the battle; we now turn our attention to the analysis of BRN measurements throughout the alley with MARS.

Chapter 8

Monitoring BRN Rates with MARS

Although our main goal with this detector is to measure the neutron flux and spectrum in Neutrino Alley, we can provide useful partial information to the collaboration while we work towards the complete characterization of MARS. Observational differences in MARS rates can indicate fluctuations of the BRN background, and we present the rate-analyses conducted to date in this chapter. We separate MARS operations into 5 run periods based on the conditions of the MARS operations and the deployment location in Neutrino Alley. We will briefly discuss Runs 1 and 2, which were analyzed primarily by COHERENT collaborators Belkis Cabrera-Palmer [180] and Justin Raybern [207] in the context of BRN measurements throughout Neutrino Alley, but the focus of this work will be the ongoing analyses of Runs 3, 4, and 5.

8.1 Data Quality

To ensure the quality of each day of data, we monitor the detector using Grafana [211] to display and track the changes in characteristics of our choosing. We extract a daily average of the total detector trigger rate, the trigger rates of individual PMTs, the PMT baselines, and the rates of the E39 and E61 timing signals from the data for the health monitoring of the detector¹. Due to the frequency of inconclusive health checks in 2021 and early 2022, we transitioned to a closer monitoring of the detector and opted to perform all health checks in real time. We now extract the relevant rates at the end of every data file, each of which is generally ~ 1 GB in size and corresponds to about 1 hour of beam-on

¹Timing outages disrupt our analysis flow, so we monitor E39. Similarly, beam-off maintenance periods can sometimes include sporadic beam-testing, so monitoring of E61 can help us track down unexpected rate excesses.

data or 1.5 hours of beam-off data². This has helped us to localize issues to a single binary file rather than a full 24-hour run period.

We do not include data in the analysis which have a trigger rate “spike”, meaning that the daily trigger rate is significantly changed from the surrounding date. Typically, we exclude data only if the trigger rate is more than double the rates on surrounding days, but we do allow for exceptions in cases close to this cutoff if the cause of the spike can be identified. The average trigger rate depends on both beam conditions and ambient backgrounds, so we make decisions about which data to include by examining the rate in context with the detector performance in surrounding days and with the SNS operations. As an example, during Run 3, we saw increasing background rates from the HOG pipe, which increased the trigger rates in MARS by at least a factor of 4 over the course of the run³. The trigger-rate spikes we observe during overall lower rate beam-on operations are similar to healthy data from months with very high beam-on rates⁴, so our exclusion criteria are chosen on a case-by-case basis, detailed in Appendix A. Ultimately, it is the pre-processing (see Section 8.2.1) which solves any borderline cases; if the pre-processing takes longer than 24 hours for a single day of data (healthy data with high rates can take up to a few hours in pre-processing), we manually stop the analysis and exclude the data.

We also use the pre-processed file size as an indicator of the health of the data. As an example, data when the E39 timing signal from the SNS is interrupted do not correctly create event-pairs in the pre-processing, and such data are not used for our typical analyses. A typical, healthy day of beam-off (-on) data from Run 3 will have a pre-processed file size near 300 (750) MB, which is a significant reduction of the many-GB of binary data taken during each 24-hour run. We generally exclude data with unreasonably small (<10 MB) or large (>1.5 GB) file sizes relative to the surrounding data, but again allow for manual exceptions in borderline cases.

Figure 8.1 shows all the daily rates of data we include in our analysis of Run 3 data in blue, and data rejected by these quality checks in red. Ultimately, we end up removing about 16% of the data collected during Run 3, 6% of the data from Run 4, and 12% of the

²This expectation is for stable, nominal operations. During a flush of the HOG pipe, for example, the background rates are considerably higher and fill up the ~ 1 GB files much faster.

³We do not know with any certainty what caused the increasing backgrounds, though suspiciously a beamline component was found to have more radiation damage than expected during the long shutdown following our Run 3.

⁴As an example, a trigger rate of 2000 Hz would generally be classified as a spike in July 2019, but is a typical healthy day of data-taking during July 2021 when the nearby hot-off-gas (HOG) background rates were much higher.

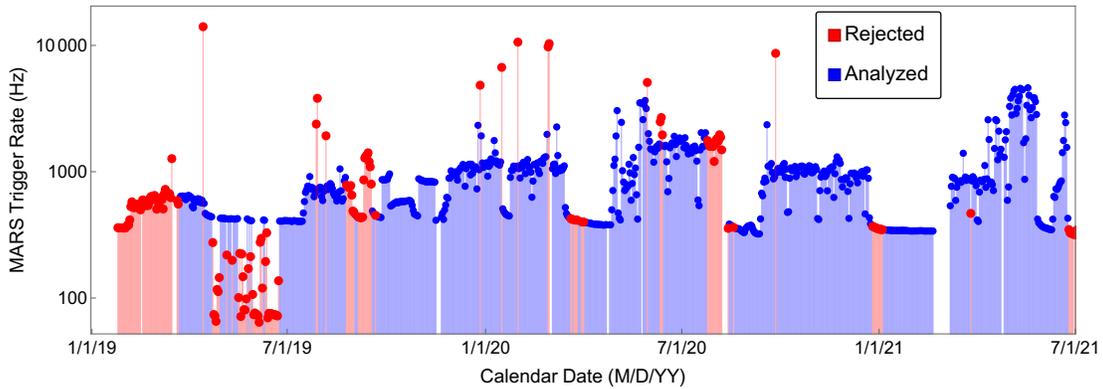


Figure 8.1:]

A plot of the daily MARS trigger rate for Run 3. The days highlighted in red are rejected by the data quality checks. The observed increase to the detector trigger rate (note the logarithmic y -axis) during beam-on operations in 2020 and 2021 is correlated to increasing HOG backgrounds; the detector health is trusted throughout the run because the trigger rate consistently returns to stable beam-off operations (400 Hz).

data during the portion of the ongoing Run 5 we analyze in this work; we report the specific inclusion/exclusion criteria for each day of data in each run in Appendix A.

8.2 Analysis Strategy

8.2.1 Preprocessing

To prepare the data for analysis, we associate any two scintillator pulses separated by no more than $200 \mu\text{s}$ into a pair, maintaining the time-ordering of the pulses. Any single pulse can be a part of an unconstrained number of pairs. Each pair is described by the variables $t, E_1, \Delta t, E_2$, where t is the time relative to the preceding E39 trigger (which occurs at $t = 800 \mu\text{s}$)⁵, E_1 represents the first pulse's energy in ADC units, Δt is the inter-pulse time and E_2 is the energy of the second pulse, also in ADC units. These parameters are illustrated in Figure 8.2. When a capture-gated neutron detection occurs, (t, E_1) represent the prompt neutron-pulse time and energy, while $(\Delta t, E_2)$ represent the Gd neutron-capture time and the energy of the Gd de-excitation γ -ray pulse. The comparisons between DT data and simulation in Section 7.2.3 show that the characteristic neutron capture time in MARS is $\tau < 17 \mu\text{s}$, so the $200 \mu\text{s}$ maximum bound on Δt will miss a negligible number of neutron-

⁵This requires that $t < 1/60 \text{ s}$ for all event-pairs, since the frequency of the E39 trigger is 60 Hz.

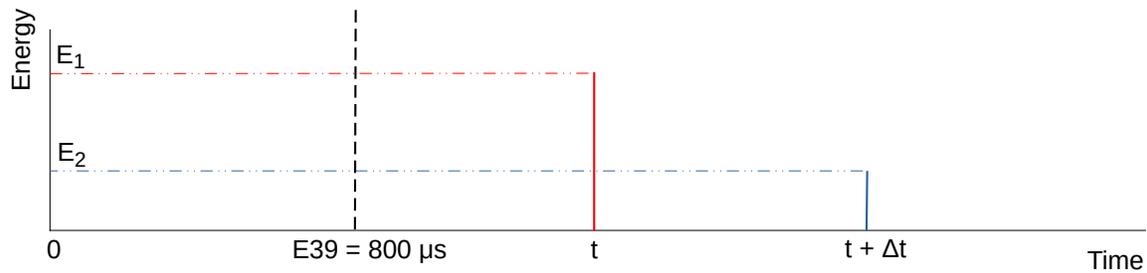


Figure 8.2: An illustration of a single event-pair in the MARS analysis. The variables E_1 and E_2 are the energies of the pulses in the pair, t is the time of the first pulse (E39 occurs at $800 \mu s$), and Δt is the elapsed time between the first and second pulses.

capture events.

We plot the raw spectra from a sample week of beam-off and beam-on data in Figure 8.3. We immediately observe a feature at roughly 400000 ADC units; this is associated with cosmic-ray muons passing through the detector at a less-than 30° angle (thus depositing a substantial amount of energy). By placing the muon cut $E_1 > 300000$ ADC, we can isolate the population of muon events to observe a feature in the E_2 spectrum corresponding to Michel electrons. Placing this cut with the opposite time-ordering (muon is associated with the E_2 pulse) gives a measure of the expected background for this estimation. This is the strategy for all the analyses that follow; we identify cuts on the four event-pair variables to isolate interesting populations of events.

8.2.2 Identifying BRNs Using Cuts on E_2 and Δt

To isolate BRN events in MARS, we place restrictions on the allowed values of Δt and E_2 , which have well-understood behavior associated with the thermalization time (Δt) and gadolinium de-excitation spectrum (E_2). We overview a baseline strategy which applies the same cuts to all data in Section 8.2.2.1, and a more rigorous approach to derive appropriate cuts given the changing detector performance in Section 8.2.2.2. In both cases, we then use timing windows to examine signal compared to random backgrounds. These timing windows are the prompt signal window $t \in [800, 802] \mu s$, the same-width background window $t \in [1000, 1002] \mu s$, and the much longer background window $t \in [1, 15] ms$ ⁶. Our rate calculation is then an estimation of the excess in the signal region over

⁶The values of these cuts should seem a bit weird; we mentioned in Chapter 6 that the BRNs should arrive within $2 \mu s$ of the beamspill! However, we shift E39 by $800 \mu s$ in the acquisition to allow for some

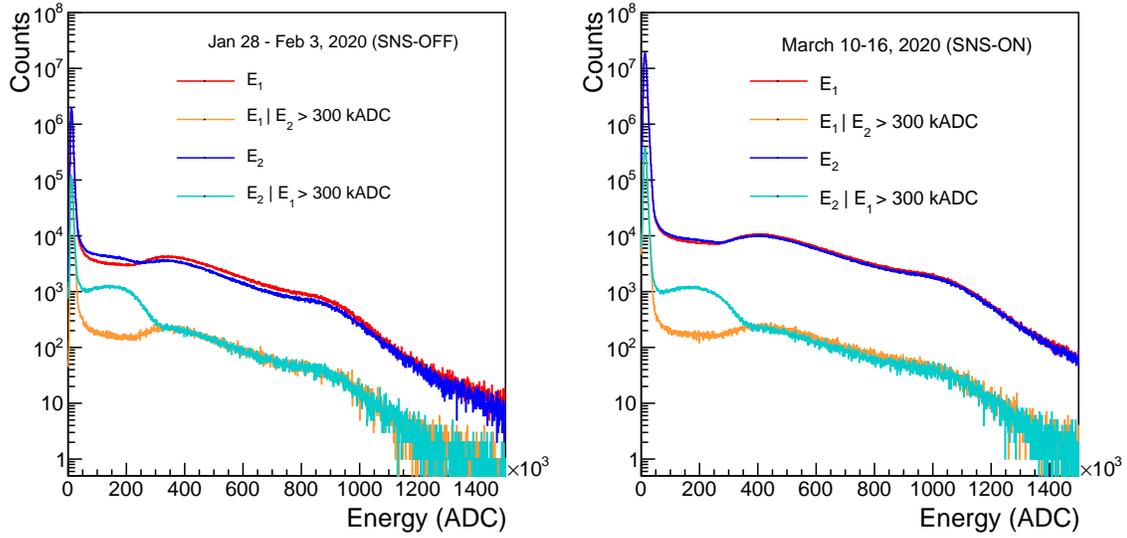


Figure 8.3: A sample week of data collected in MARS during beam-off (left) and beam-on (right) periods, illustrating the total measured spectra of E_1 and E_2 , along with the easily-identifiable Michel electron peak when a large E_1 deposit is observed (consistent with a muon passing through the detector).

the time-normalized long background. The short and long background windows are expected to be fully consistent, as both are anti-coincident with the beam spill and represent a characterization of the steady-state background.

8.2.2.1 Constant Cuts

The cuts described in this section were developed for use in the Run 2 and DT calibration analyses, along with simulation, and are applied to Runs 3, 4, and 5. We use the cut $E_2 \in [26, 60]$ kADC to describe the gadolinium de-excitation⁷, and apply a cut $\Delta t \in [6, 48.1]$ μ s based on results from Run 2 and our knowledge of the relevant capture times in the detector; Michel electrons dominate the 0 - 6 μ s region given that the muon lifetime is 2.2 μ s, and the measured capture constant from the DT calibration was $\tau < 17 \mu$ s, so going beyond 50 μ s starts to add more background events than signal events into the

characterization of the steady-state background before and after each beam-spill.

⁷We note from our work in Section 7.2.3 that this does not describe the full peak (which would be roughly $E_2 \in [10, 70]$ kADC); there was a dedicated effort in Run 2 to determine the optimal cuts on E_2 and Δt which maximizes the significance of the signal excess over background. Essentially, focusing on the high-energy tail of the E_2 spectrum minimizes the impact of the high rate of background events near threshold.

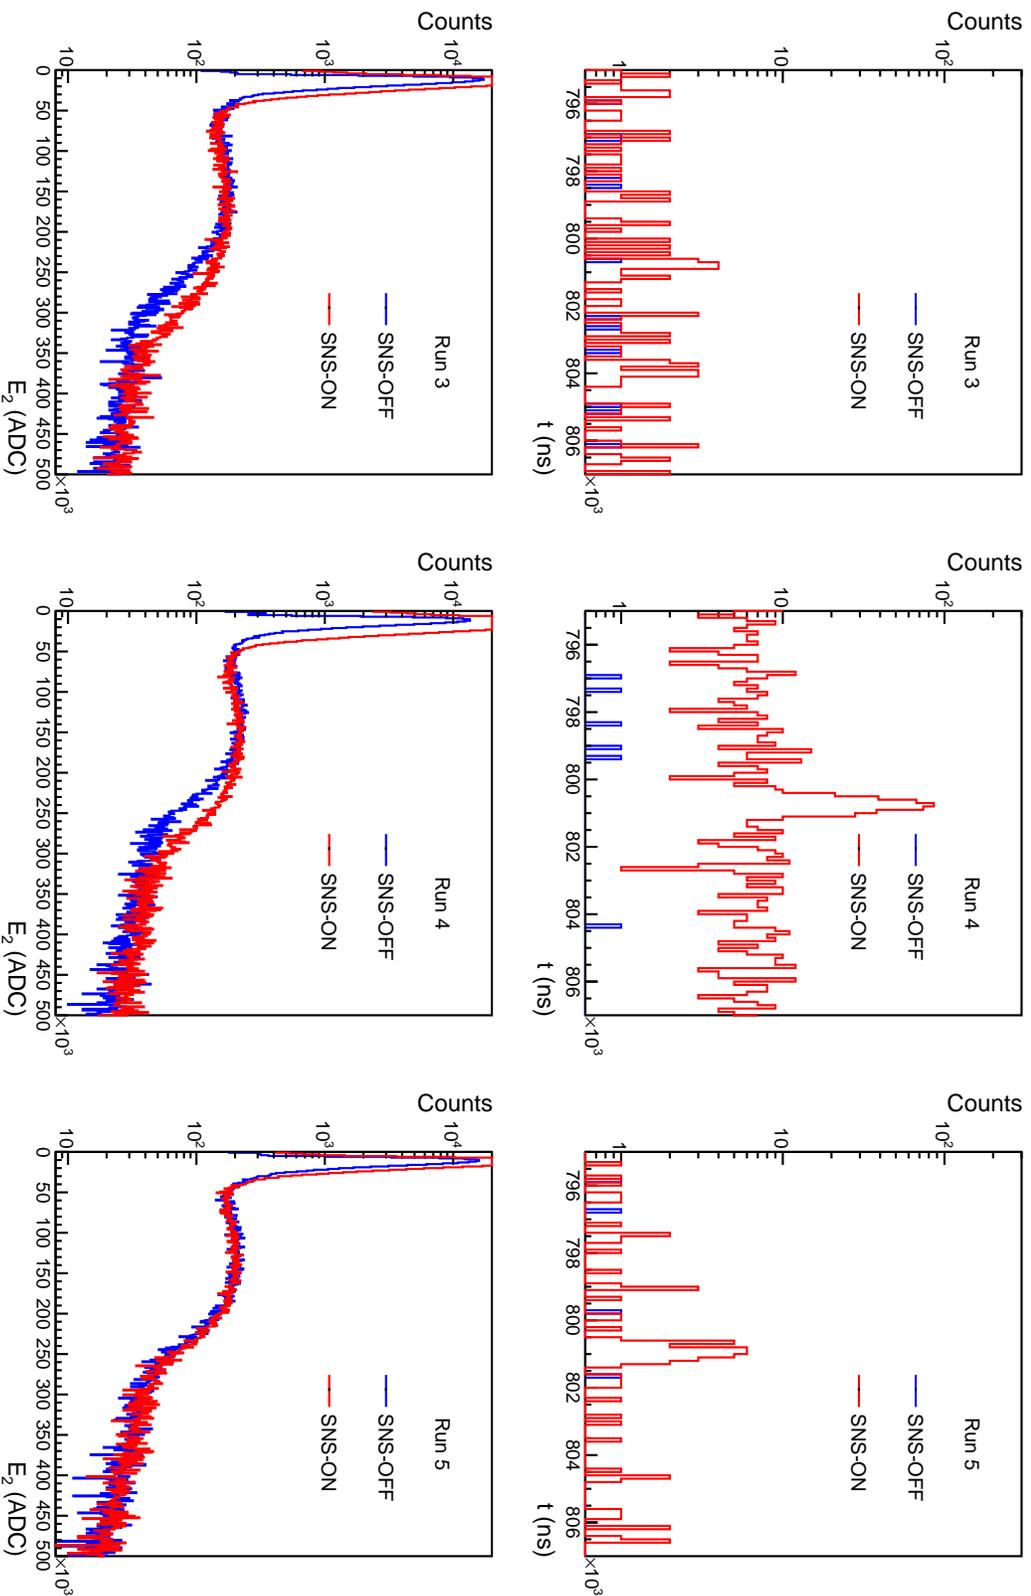


Figure 8.4: A comparison of t for BRN-like events (top) and the E_2 spectrum for Michel electron-like events (bottom) in observed pulse-pairs in MARS for example 24-hour periods of beam-on and beam-off data during the different production runs.

analysis.

The top panels of Figure 8.4 illustrate the distribution of arrival times for events in a 24-hour period which pass the neutron-capture cuts on Δt and E_2 used in Run 2. We have zoomed in to show the region surrounding our expected signal window, $t \in [800, 802]\mu s$, and observe clear excess of counts in this window during beam-on operations in the high BRN-rate areas of Neutrino Alley investigated by Runs 4 and 5 (more in Section 8.3). Run 3 took place in an especially neutron-quiet location of the alley, so there is no pronounced excess for a single day of data. These plots establish that our knowledge of the detector operations during Run 2 can indeed be applied to the more recent runs.

However, we also know from Run 2 that MARS has an observable spectral difference for beam-off and beam-on operations. The bottom panels of Figure 8.4 illustrate the Michel electron feature in the E_2 spectrum for the same 24-hour run periods in Runs 3, 4, and 5. The shift to higher energies in Runs 3 and 4, associated with pile-up of HOG 511 γ s, affects our interpretation of the excess in our signal region. The lack of significant shift in the Run 5 data is likely because MARS is farther from the HOG pipe than any of the other runs. We can also immediately note that these rough cuts demonstrate relatively consistent MARS operations across all run locations.

Cosmic-ray muons passing through MARS can generate large pulses and represent the main contribution to the high-energy ends of the E_1 and E_2 spectra. Muons decaying within the detector produce highly energetic Michel electrons with a decay time constant $\tau_{\mu^-} = 2.2 \mu s$ [212] and a kinetic-energy cutoff at ~ 50 MeV [213]. These events are included in our (E_1, E_2) analysis, as the event-pair will have pulses corresponding to the muon's ionization E_1 as the muon penetrates the detection medium and stops and the Michel electron deposition E_2 . A sample week of data from beam-off and beam-on periods is shown in Figure 8.3, and a double-exponential fit to Δt illustrated in Figure 8.5 demonstrates our ability to identify these Michel electrons. This fit is given by

$$f = n_0 \exp\left(\frac{-\Delta t}{\tau_{\mu}}\right) + n_1 \exp\left(\frac{-\Delta t}{\tau_{\text{Gd}}}\right), \quad (8.1)$$

where $\tau_{\mu, \text{Gd}}$ refer to the muon lifetime and Gd capture-time constant, respectively, and n_i are the normalizations for each component. The best-fit values of $\tau_{\mu} = 1.99 \pm 0.01 \mu s$ and $\tau_{\text{Gd}} = 15.49 \pm 0.41 \mu s$ inform our analysis cuts on Δt .

The cosmic muons (which produce the Michel electrons) will also knock off neutrons in the detector or in the concrete surroundings of Neutrino Alley. These muon-induced neutrons (μIN) will be a steady-state background in our paired analysis, because there is an appreciable rate of muon ionization in the detector followed by Gd de-excitation.

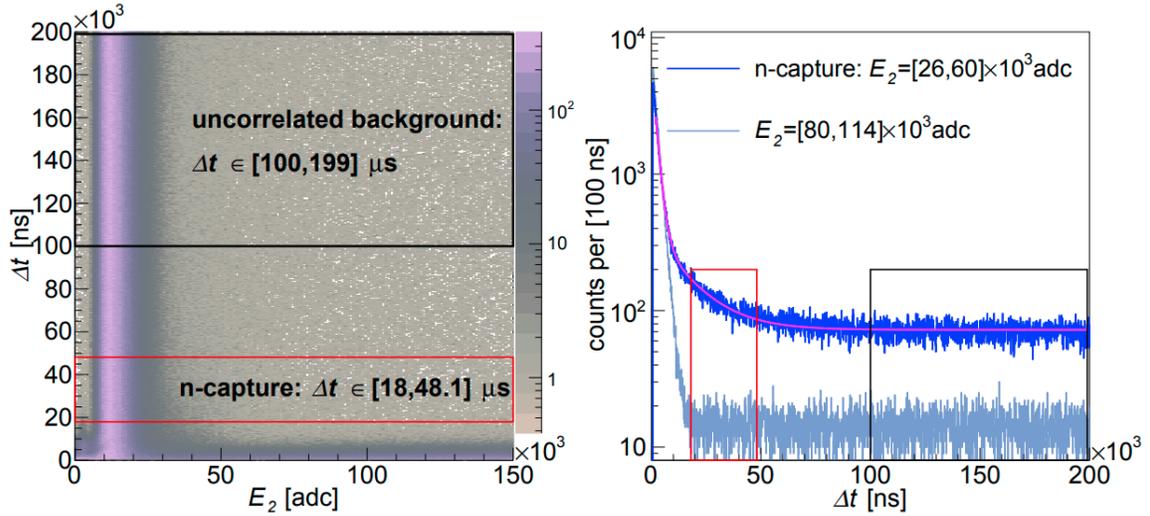


Figure 8.5: Left: The distribution of Δt and E_2 given that $E_1 > 300$ kADC for a five-week beam-off period in Spring 2018. Right: The Δt projections for two E_2 ranges. $E_2 \in [26, 60]$ kADC corresponds to a Gd de-excitation, while $E_2 \in [80, 114]$ kADC is a measure of the background. We fit the neutron-capture region, illustrated by the pink curve, with a double-exponential function (Equation 8.1) to identify the Michel electron ($\tau_\mu = 1.99 \pm 0.01 \mu\text{s}$) and neutron capture ($\tau_{\text{Gd}} = 15.49 \pm 0.41 \mu\text{s}$) populations. Figures reproduced from [180].

Triple events, where we might observe the muon, neutron, and capture, are not observed. During Run 2, we used a stable five-week period of beam-off data in Spring 2018 to identify the gadolinium de-excitation spectrum using μIN . Figure 8.5 shows the distributions of Δt and E_2 for all event-pairs with $E_1 > 300$ kADC in that period. Note that this is the same cut we apply to identify Michel electrons, and we do observe that irreducible population of events in the right panel of Figure 8.5. Using the E_2 range appropriate to neutron captures, however, can isolate the smaller rate of the muon-induced neutron capture events.

We observe in the right panel of Figure 8.5 that $\Delta t \in [0, 48.1] \mu\text{s}$ will contain the majority of neutron capture events, but will be dominated by the substantial rate of Michel electron events. However, by restricting our search to $\Delta t \in [18, 48.1] \mu\text{s}$, we can isolate the pure muon-induced neutron population. We examine this population in Figure 8.6, and observe that, just as with the shifted feature associated with Michel electrons we showed in Figure 8.4 (bottom), there is both a shift in the spectrum associated with the HOG background during SNS beam-on operations and a gradual decrease of the overall detector gain.

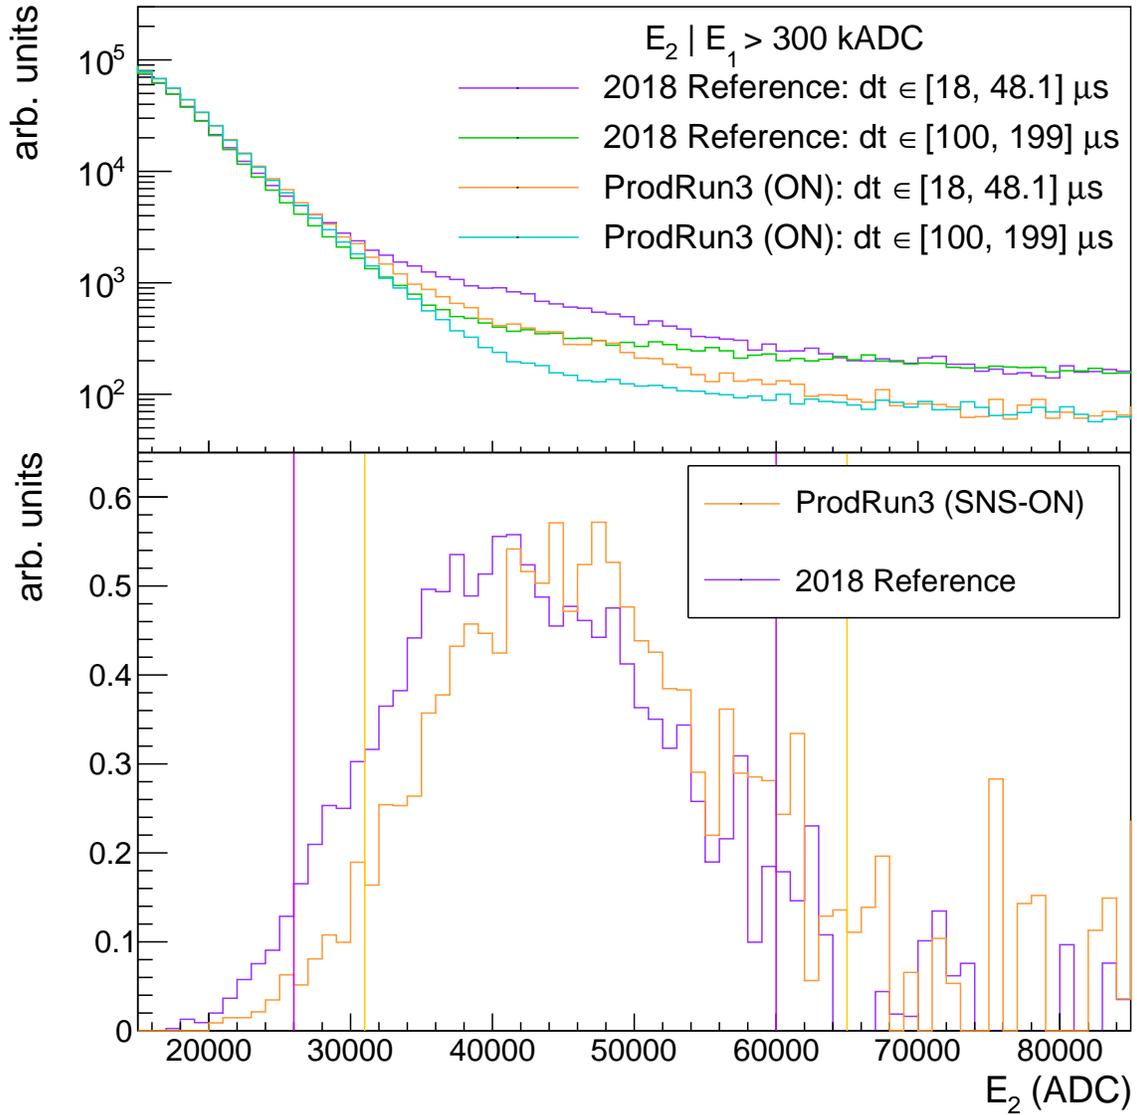


Figure 8.6: Top: Distributions of the E_2 spectrum for neutron capture and background Δt cuts for two weeks of data during beam-off and beam-on periods in Run 3, shown relative to the five-week beam-off reference period from April 2018, scaled by the duration of the Δt range. Bottom: The residual spectrum, defined as signal ($\Delta t \in [18, 48.1] \mu\text{s}$) minus background ($\Delta t \in [100, 199] \mu\text{s}$), which signals μIN capture events. The vertical lines indicate areas of approximately equal integral.

In Figure 8.7, using the constant cuts on Δt and E_2 we observe the correlated fluctuations of both muon-induced neutrons and Michel electrons in time. An independent measurement of the muon flux throughout the alley indicates that there are no major spa-

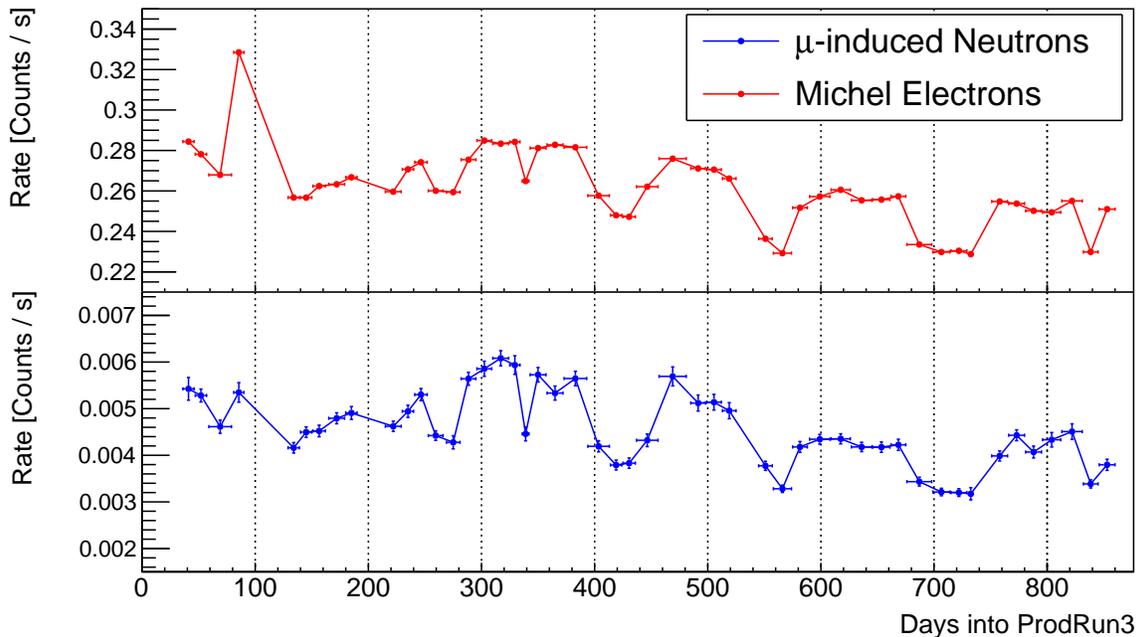


Figure 8.7: Michel electron (top) and muon-induced neutron rates (bottom) identified using the constant cuts from Run 2 ($\Delta t \in [18, 48.1]\mu\text{s}$ and $E_2 \in [26, 60]$ kADC).

tial changes in the Neutrino Alley overburden, which means that the instability of these rates is caused by our approach.

8.2.2.2 Variable Cuts Using Muon-induced Neutrons

Since we know that the flux of cosmic-ray muons, and thus the rate of muon-induced neutrons (μIN), is expected to be constant in time⁸, we can use this population of events to derive variable cuts on E_2 and account for the spectral shift of the HOG backgrounds. We keep Δt fixed, as the time constants are only expected to be impacted by major detector changes (PMT failures, changes to the detection efficiency, etc.), which are not caused by background fluctuations. We also keep the width of our E_2 region fixed; we identify the minimum value of an E_2 range for which the rate of μIN is equal to the observed rate during the five-week beam-off reference period from April 2018. Mathematically, this

⁸We ignore the seasonal fluctuations of cosmic-ray muon rates, given that the slow-variation is not correlated with the beam and thus cannot be the cause of our observations in Figure 8.7, which shows the rates as they depend on the calendar date.

means we want to equate

$$\int_{26 \text{ kADC}}^{60 \text{ kADC}} E_{2,\text{ref}} dE_{2,\text{ref}} = \int_{E_{\text{min}}}^{E_{\text{min}}+34 \text{ kADC}} E_{2,i} dE_{2,i}, \quad (8.2)$$

where $E_{2,i}$ refers to the measured μIN spectrum for the i th aggregation period using a constant cut $\Delta t \in [18, 48.1]$, leaving E_{min} as the only variable.

This procedure for deriving E_2 cuts depends heavily on the statistical fluctuations of single bins in the μIN spectrum. As such, we aggregate data with similar beam status (i.e., our aggregation periods do not mix beam-off and beam-on data) and monitor the BRN rates in Neutrino Alley in roughly two-week intervals. In future, we will develop alternate strategies, such as fitting the Gd deexcitation spectrum to the residual μIN spectrum to determine the appropriate energy scaling, so that we may monitor the daily BRN rates throughout Neutrino Alley and study the potential impact of shutter status for each of the beamlines which pass over the alley. The details of the aggregation periods for the Runs 3, 4, and 5 are given in Appendix A.

As a visual cue for this procedure, the bottom panel of Figure 8.6 illustrates the identified E_2 ranges for two-week periods of data in Run 3; we identify the μIN population in the data, and match the integral of the residual counts to the integral of the 2018 reference period for $E_2 \in [26, 60]$ kADC. Figure 8.8 shows measured μIN rates for the chosen aggregation periods of Run 3, 4, and 5. Note that while the Constant E_2 cuts fluctuate with beam conditions and in time as we previously demonstrated, the Variable E_2 cuts result in a constant rate of μIN by design. We derive these variable E_2 cuts for each data-aggregation period, and apply them in our search for BRNs.

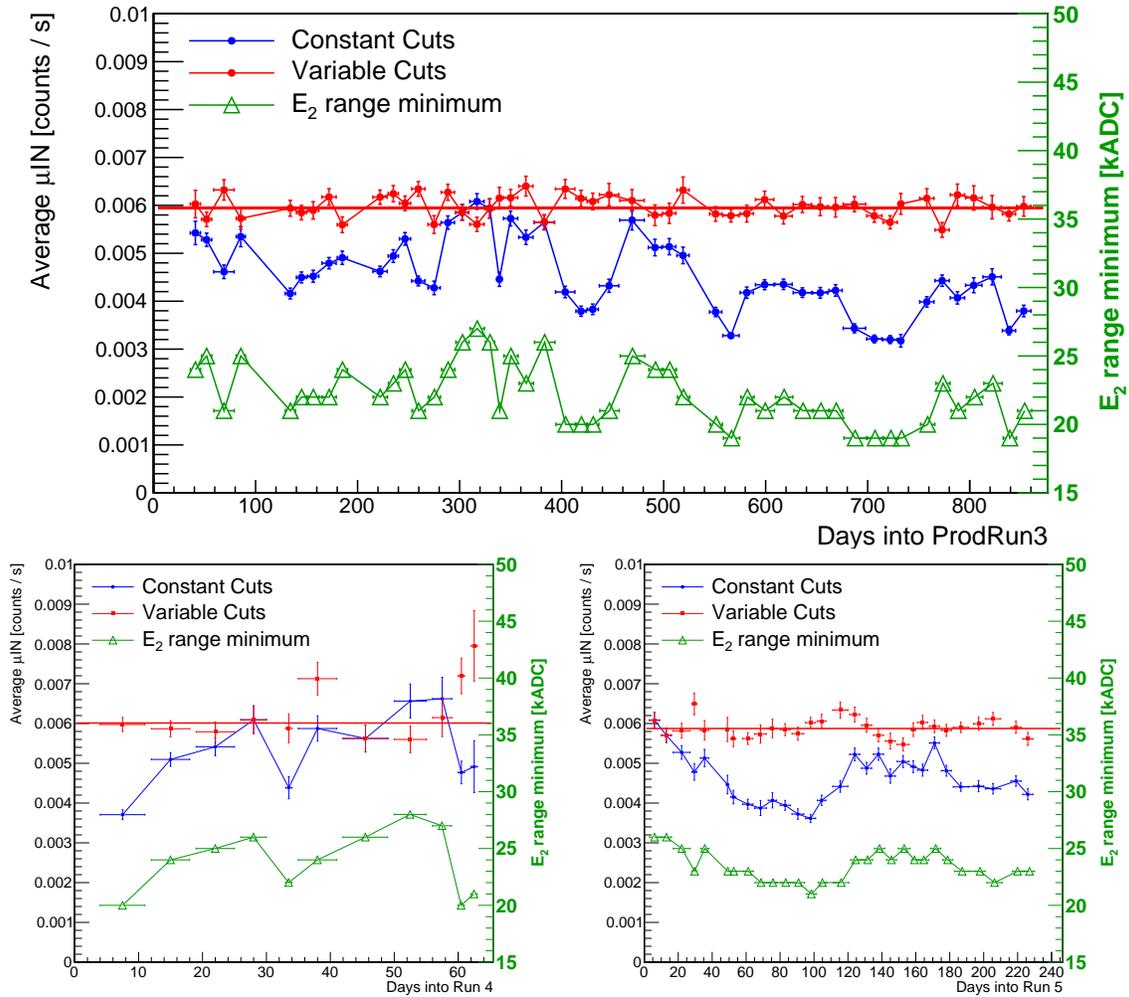


Figure 8.8: A comparison of the observed μIN rates during Run 3 (top), 4 (bottom left) and 5 (bottom right) using the constant E_2 cut found by the April 2018 reference versus the variable E_2 cuts derived by getting as-close-as-possible to the observed μIN rate from April 2018. The minimum E_2 value for the range of the derived cuts is shown in green. The rate of μIN using Variable E_2 cuts is fit with a constant, which is consistent across production runs by design.

8.3 Production Run Results

In this section, we report on the full observations of the MARS detector, and compare the measurements of MARS with results from SciBath, the Neutron Scatter Camera (NSC), and the Timing Cart. Recall that the locations of each of the MARS runs discussed below were illustrated relative to other COHERENT neutron-monitoring subsystems in Figure

6.2.

The SNS provides databases of beam operations, including the time-dependent beam power and energy. We integrate these databases into our analysis using the timestamps of each 24-hour run period to calculate the total beam delivered during each day of data and, subsequently, in each aggregation period. Using the variable E_2 cuts, and the fixed cut $\Delta t \in [6, 48.1]\mu s$, we examine the t and E_1 distributions of the BRN population in each run. We subtract the total number of counts in the signal region $t \in [800, 802]\mu s$ from the time-scaled number of counts in the long background region $t \in [1, 15]$ ms to measure the rate of the BRN-like excess in each aggregation period. We then normalize this by the total beam delivered to report our results as observed BRN / GWh. We'll then convert these observations to a neutron flux that can be compared to other systems in Section 8.3.5.

We perform two separate measures of the observed BRN / GWh. The first aggregates the entire run, and takes the total number of BRN-like events observed in all aggregation periods divided by the total beam delivered to estimate an “Average” rate, shown in red in the figures that follow. The second strategy fits a constant to the results from each aggregation period separately, and the “Best Fit” is shown in orange in the figures that follow. We expect the “Average” to be the more accurate measure of the flux, given the higher statistical significance of any BRN-like excess over the background, while the “Best Fit” serves as an indication of any atypical operations (e.g., if a significant excess above the fit is observed, we might look for a change in shutter positioning).

8.3.1 Runs 1 & 2: 2017 – 2018

Run 1 was a commissioning run from deployment near the CsI detector in Fall 2017 to Spring 2018, during which we identified several issues with the DAQ and found the FIR threshold setting which would limit dead time associated with the high HOG-background trigger rate. During this run, one of the PMT channels was unusually noisy and caused a feature in the E_1 spectrum where none should have existed. Ultimately, a reset of the high-voltage system returned this channel to standard operations⁹.

Run 2 was our first true search for BRNs, and MARS took data near the CsI detector from May - December 2018 (154 beam-on days, 3.872 GWh). The full analysis details are

⁹The specific PMT was channel 0, which is oriented at the top left of the detector when viewed from Neutrino Alley. We have frequently observed atypical behavior from this channel compared to other PMT operations, especially during low-threshold operations. In June 2022, we identified a potentially faulty read-in cable coming from this PMT; it remains unclear if this cable issue dates as far back as 2017.

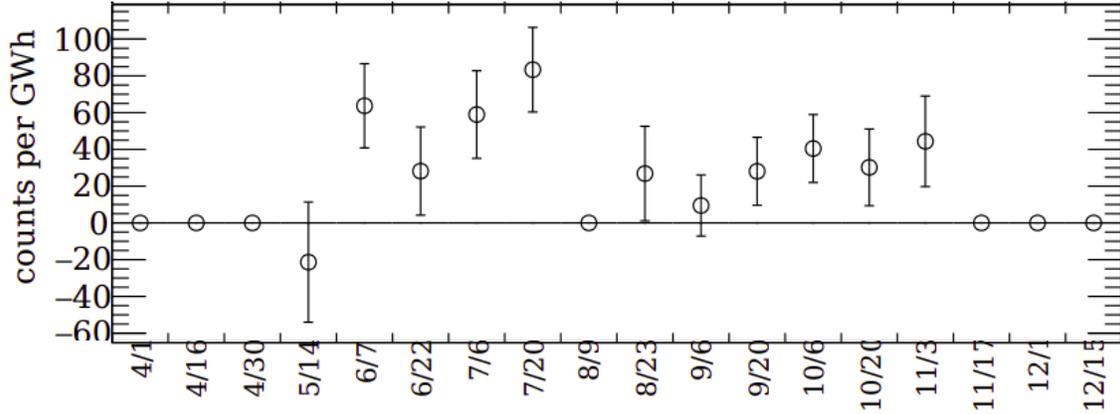


Figure 8.9: Observed BRN counts per GWh in the Run 2 data, aggregated over two week (335 hour) periods. The calendar date of the start of the aggregation period, given in MM/DD format (since all Run 2 data was collected in 2018). These results are normalized by beam delivered, such that the plot estimates the BRNs observed by MARS per GWh. The points falling at exactly 0 BRN per GWh represent data taken during beam-off and are assumed to be 0 (so that we do not divide by 0 beam-delivered). Figure reproduced from [180].

published in [180], and average rates of 38 ± 6.7 BRN per GWh, (0.96 ± 0.17 BRN per beam-on day) were observed. We reproduce the results per aggregation period in Figure 8.9.

8.3.2 Run 3: Jan 2019 – June 2021

For Run 3, MARS was also in the center of Neutrino Alley, but was moved about a meter closer to the alcove compared to the Run 2 location. We ran into a wide variety of data-quality problems over the 2.5 years at this location, causing $\sim 16\%$ of the data to be rejected prior to analysis. These problems ranged from SNS timing signal outages (meaning, no E39 or E61 triggers), to disk-space write errors, to spikes in trigger rate with presently unknown causes. In spite of these losses, MARS operations were generally very stable during this time.

We chose to aggregate the data in two-week periods, as in Run 2, due to the low rate of BRNs at this location. Our analysis includes 514 beam-on days and 9.39 GWh of beam delivered from January 2019 to June 2021. We measure an average 18.2 ± 1.5 BRN / GWh, with rates shown per aggregation period in Figure 8.10. We plot the arrival time t and

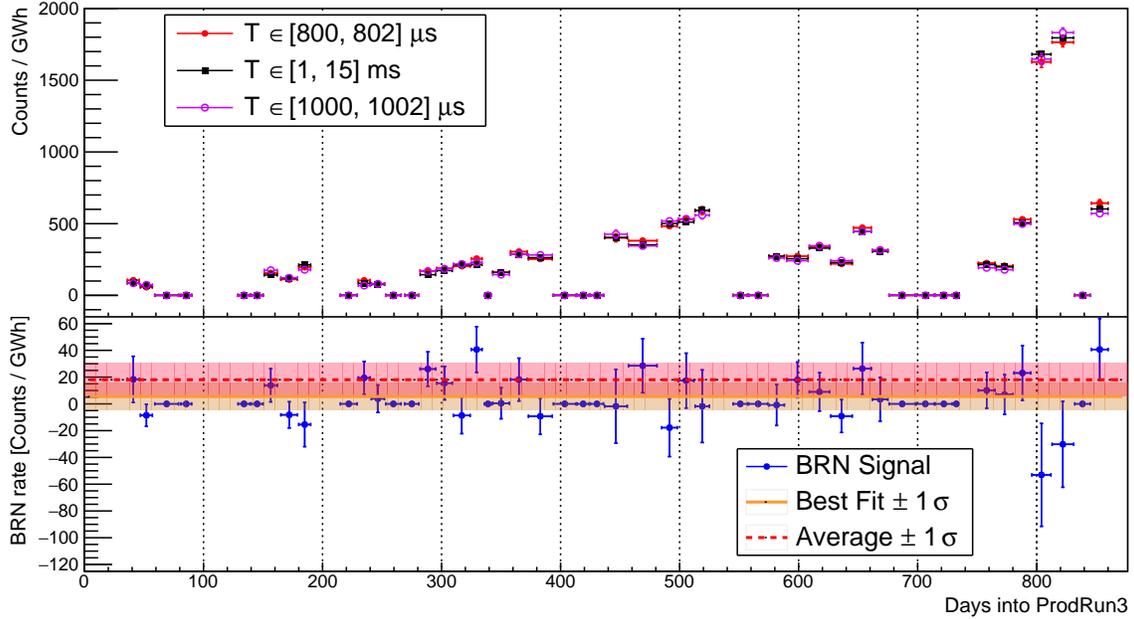


Figure 8.10: Top: Comparison of Run 3 data within expected region of BRNs (800 - 802 μs), an anti-coincident time window of equal length (1000 - 1002 μs), and a long-anti-coincident window (1 - 15 ms) to increase precision of the measured background. Bottom: Subtracting the background from the BRN expected window to estimate the signal. All results are normalized by beam delivered, such that the bottom plot estimates the BRNs observed by MARS per GWh. The points falling at exactly 0 BRN per GWh represent data taken during beam-off and are assumed to be 0 (so that we do not divide by 0 beam-delivered). We fit all data during beam-on days with a constant, shown as the horizontal orange line, to determine the “Best-Fit” BRN per GWh. The red line indicates the flat “Average” over all beam-on days (total BRN counts / total beam delivered).

energy spectrum E_1 of the observed BRN counts in Figure 8.11. We fit the arrival time with a Gaussian and find that the distribution peaks at 800982 ± 0.45 ns with a FWHM of 412 ± 97 ns; this is consistent with the POT trace (FWHM 350 ns). We fit the E_1 spectrum with a power law

$$R_{\text{BRN}} = N E_1^{-\alpha} \quad (8.3)$$

on the range $E_1 \in [30, 100]$ kADC to avoid any threshold effects, and find $\alpha = 2.07 \pm 0.45$ with normalization $N = (4 \pm 18) \times 10^{10}$ counts per GWh. This fit to the observed spectrum gives a rough estimate of the expected power law describing the flux of incident BRNs, but should *not* be interpreted as accurate for the raw neutron energy until a full spectral

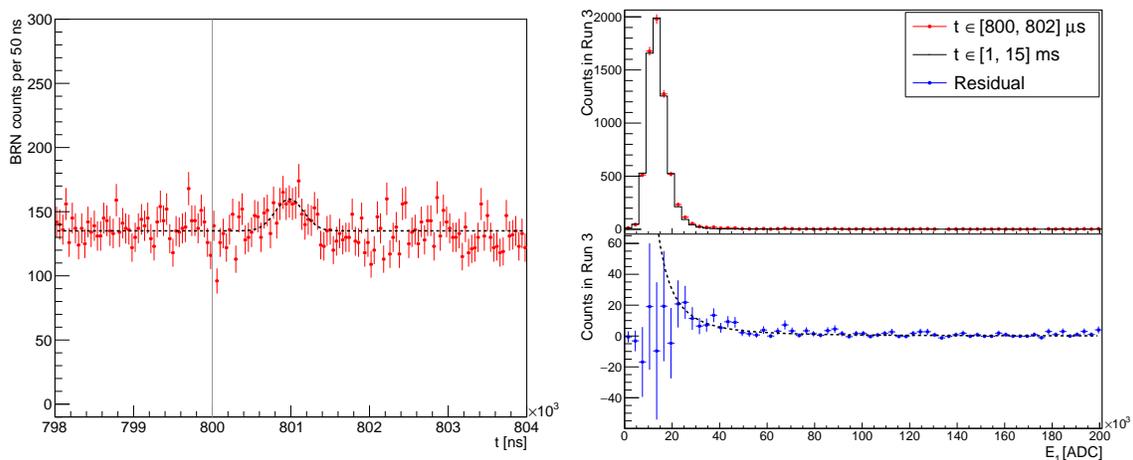


Figure 8.11: Distributions of t (left) and E_1 (right) for the BRN-like excess in Run 3. The dashed black lines represent the fit results.

unfolding is performed.

8.3.3 Run 4: July 2021 – September 2021

Run 4 was our first move to a location at the edges of the alley, where there is not as much of a concrete/gravel fill to attenuate the BRN flux. We generally refer to Run 4 as the “rad monitor” location, so named because MARS was placed in front of the safety system monitoring radiation levels in the area behind the wall. The proximity to this monitor and the geometric space that MARS occupies limited our ability to stay in this location (ORNL staff needed to be able to access that monitor), so Run 4 was only 2 months. A shorter run was still useful in this location because we observed much higher rates than in the center of the alley.

We used one-week aggregation periods for Run 4, which enable the use of our μIN approach to deriving the cuts¹⁰ while also allowing MARS to search for any changes in the measured rates. During the 63 beam-on days and 1.11 GWh of beam delivered in July and August 2021, we measured an average 10608 ± 98 BRN / GWh, shown in Figure 8.12.

We plot the arrival time t and energy spectrum E_1 of the observed BRN counts in Figure 8.13, and again fit the distributions. The arrival time is described by a Gaussian with $\mu = 800730 \pm 3$ ns and FWHM of 360.2 ± 4.2 ns, and the E_1 power-law fit (using Equation 8.3) returns $\alpha = 1.28 \pm 0.11$ and $N = (2.6 \pm 3.1) \times 10^7$ counts per GWh.

¹⁰In future, we will search for BRNs with daily monitoring by using the simulation-predicted Gd de-excitation spectrum to perform fits to the E_2 spectrum rather than using our current integral-approach.

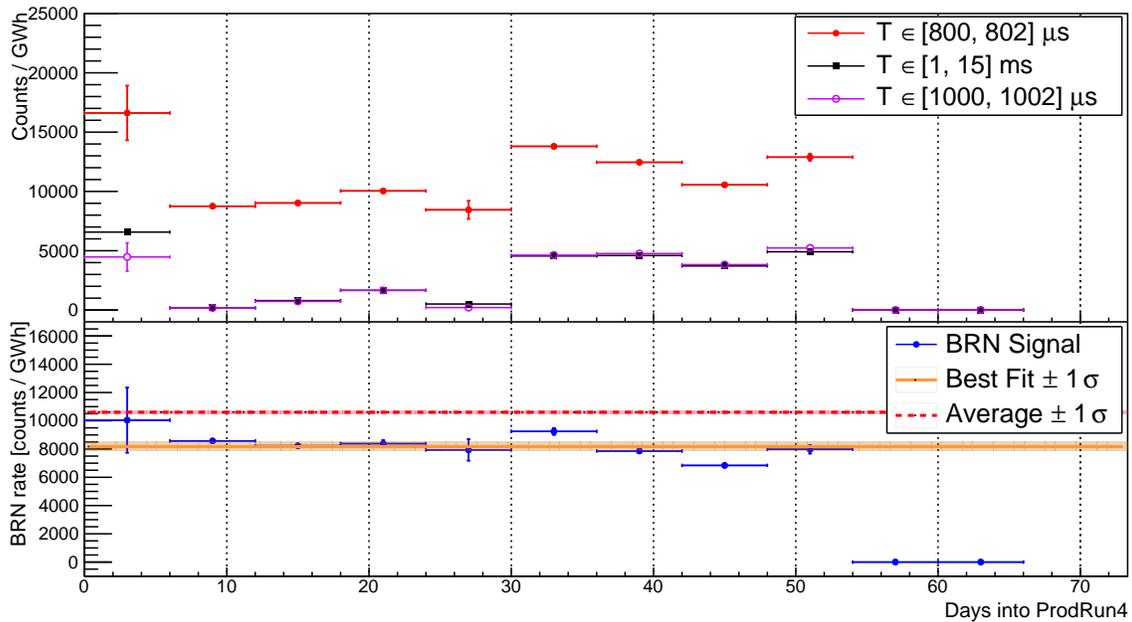


Figure 8.12: Top: Comparison of Run 4 data within expected region of BRNs ($800 - 802 \mu s$), an anti-coincident time window of equal length ($1000 - 1002 \mu s$), and a long-anti-coincident window ($1 - 15 ms$) to increase precision of the measured background. Bottom: Subtracting the background from the BRN expected window to estimate the signal. All results are normalized by beam delivered, such that the bottom plot estimates the BRNs observed by MARS per GWh. The points falling at exactly 0 BRN per GWh represent data taken during beam-off and are assumed to be 0 (so that we do not divide by 0 beam-delivered). We fit all data during beam-on days with a constant, shown as the horizontal orange line, to determine the “Best-Fit” BRN per GWh. The red line indicates the flat “Average” over all beam-on days (total BRN counts / total beam delivered).

Again, this fit to the observed spectrum gives a rough estimate of the expected spectrum of incident BRNs, but will require a full spectral unfolding to confirm.

There is a feature in the observed E_1 spectrum at roughly 40 kADC that we do not expect and are actively investigating. Our current aim is to determine if this feature is contained to a particular aggregation period or if it prevails throughout the run. We will then examine the behavior of our other event-pair variables to isolate the source of the feature.

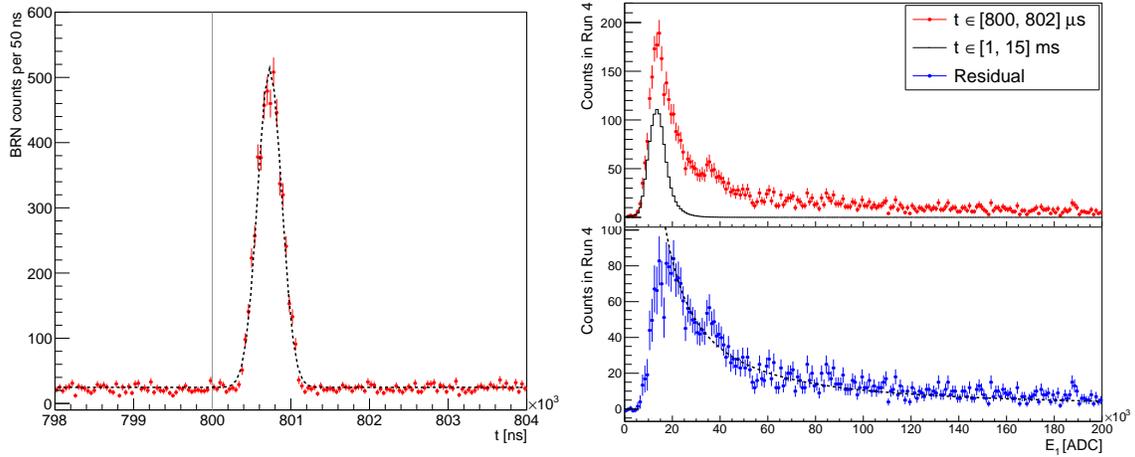


Figure 8.13: Distributions of t (left) and E_1 (right) for the BRN-like excess in Run 4. The dashed black lines represent the fit results.

8.3.4 Run 5: September 2021 – Present

Run 5 is at the other end of Neutrino Alley compared to Run 4; MARS is currently deployed in the alcove. The Timing Cart has demonstrated that this interesting location next to LAr likely observes additional neutrons from the stairwell and beamline trench. Deploying MARS at this location, for a long period of time, allows us to look for different neutron propagation paths and other differences relating to shutter positioning; we can then provide direct information to the ongoing LAr analyses. For the 101 beam-on days and 2.16 GWh of beam delivered between September 2021 and April 2022, we observe an average 1664 ± 28 BRN / GWh, illustrated in Figure 8.14.

We plot the arrival time t and energy spectrum E_1 of the observed BRN counts in Figure 8.15, and fit the observed distributions. The arrival time is generally described by a Gaussian with $\mu = 800893 \pm 4$ ns and FWHM of 410.2 ± 7.0 ns. However, there is a small excess preceding the larger excess that is consistent with the beam spill. This unexpected excess can be characterized by a secondary Gaussian, such that $\mu = 800254 \pm 6.8$ ns and FWHM 79.1 ± 10.9 ns. We do not know the origin of this excess, so we follow a similar procedure to the E_1 excess from Run 4: we must first determine if this excess is associated with any particular aggregation period, then we can begin to isolate the characteristics of the event-pairs described by this excess.

The E_1 power-law fit (using Equation 8.3) returns $\alpha = 1.68 \pm 0.12$ and $N = (1.2 \pm 1.6) \times 10^9$ counts. As a reminder, the fit to the observed spectrum is only a rough estimate of the incident BRN spectrum.

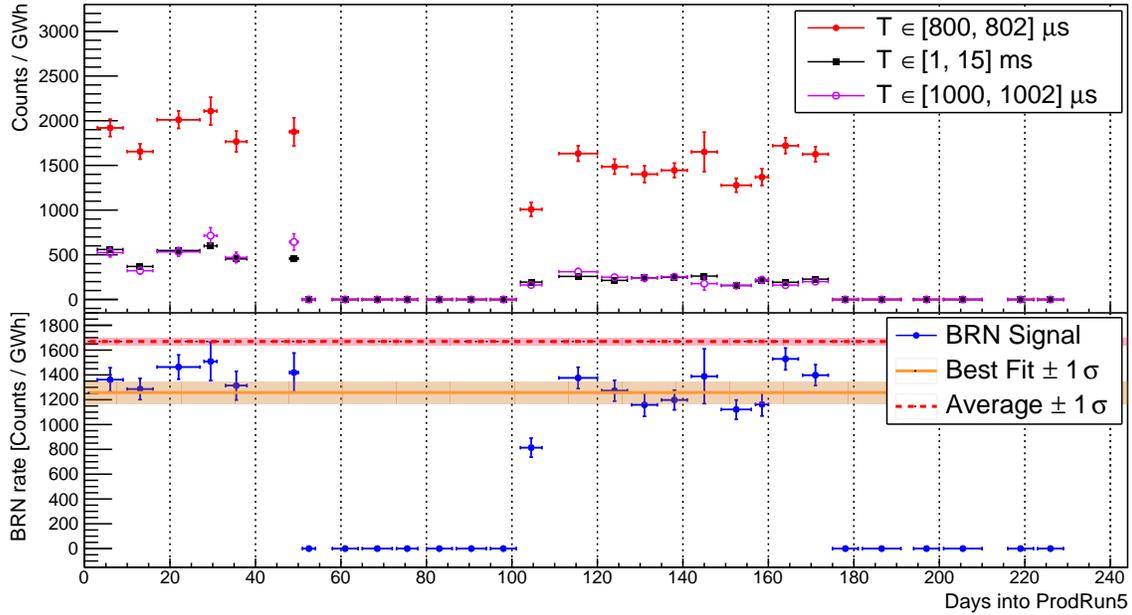


Figure 8.14: Top: Comparison of Run 5 data within expected region of BRNs ($800 - 802 \mu\text{s}$), an anti-coincident time window of equal length ($1000 - 1002 \mu\text{s}$), and a long-anti-coincident window ($1 - 15 \text{ ms}$) to increase precision of the measured background. Bottom: Subtracting the background from the BRN expected window to estimate the signal. All results are normalized by beam delivered, such that the bottom plot estimates the BRNs observed by MARS per GWh. The points falling at exactly 0 BRN per GWh represent data taken during beam-off and are assumed to be 0 (so that we do not divide by 0 beam-delivered). We fit all data during beam-on days with a constant, shown as the horizontal orange line, to determine the “Best-Fit” BRN per GWh. The red line indicates the flat “Average” over all beam-on days (total BRN counts / total beam delivered).

8.3.5 Comparing BRN Measurements throughout Neutrino Alley

From each of the MARS production run analyses, we have measured a rate of BRNs per GW of SNS operations. From the DT calibration, we know the detection efficiency ϵ is $(4.2 \pm 1.8)\%$ for 14.1 MeV incident neutrons (details of that analysis can be found in [180, 207]). We also know that the efficiency is energy-dependent from our characterization work in Chapter 7, but we do not yet have a full model of this energy-dependent efficiency. In this work, we use the measurement and approximate a flat efficiency across all incident neutron energies to roughly estimate the BRN flux. As shown in Figure 7.17, this is likely an overestimate, because 14 MeV neutrons have near-maximal efficiency in

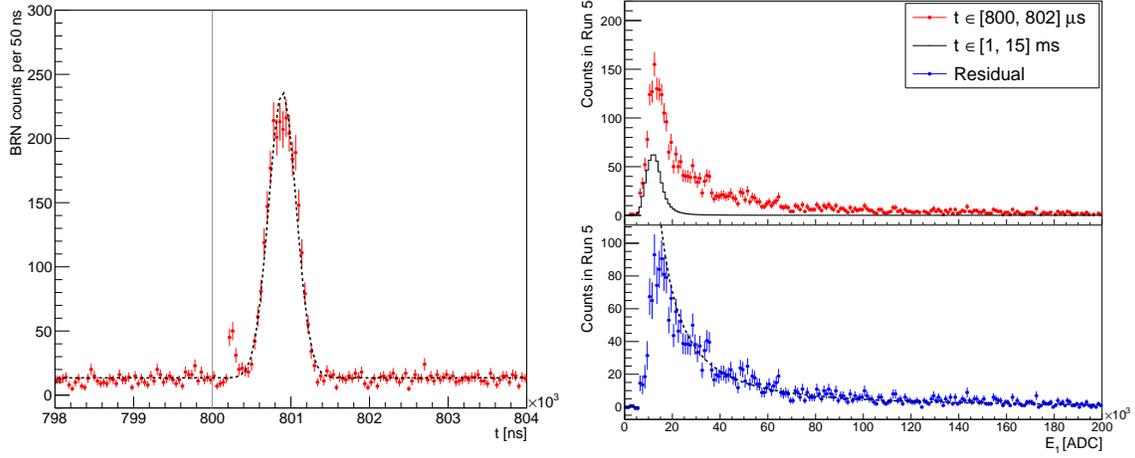


Figure 8.15: Distributions of t (left) and E_1 (right) for the BRN-like excess in Run 5. The dashed black lines represent the fit results.

MARS. Assuming a power-law spectrum with $\alpha \approx 1.5$, as in [85], we can estimate the uncertainty of the choice to use an energy-independent efficiency as

$$\sigma = \sqrt{\frac{\sum (\epsilon_{\text{sim}}(E_i) - \epsilon_{\text{DT}})^2 E_i^{-\alpha}}{\sum E_i^{-\alpha}}}, \quad (8.4)$$

where E_i is the simulated neutron energy, ϵ_{sim} is the simulated detection efficiency, and ϵ_{DT} is the measured efficiency. Using the seven points illustrated in Figure 7.17, we obtain a roughly 4% uncertainty, which is negligible compared to the 43% uncertainty from the calibration itself.

To convert our counts per GWh, R , to a flux estimate using the detection efficiency and cross-sectional area A as

$$\Phi = \frac{R}{\epsilon A} \longrightarrow \frac{R(E)}{\epsilon(E)A} = \frac{NE^{-\alpha}}{\epsilon(E)A}, \quad (8.5)$$

where the inclusion of Equation 8.3 after the arrow indicates our eventual goal to provide energy-dependent flux information to the COHERENT detector systems once a spectral unfolding is achieved. The total uncertainty will be the sum in quadrature of the statistical uncertainty reported in the previous sections and the uncertainty on the efficiency.

We report our estimation of the observed flux of prompt BRNs from all measurements in Neutrino Alley in Table 8.1. These observations, despite their significant uncertainties, demonstrate that there is at least an order of magnitude variation of the BRN flux in Neutrino Alley. We plot these based on the detector distance from the alcove wall in

Table 8.1: A comparison of all BRN measurements of the baseline neutron flux in Neutrino Alley, with position estimated relative to the alcove wall. This table excludes detector-specific rate estimates from finalized COHERENT analyses so that we ignore any shielding configuration dependencies. Lines of the table with a gray shading indicate atypical configurations of the Timing Cart, which explore interesting changes to the BRN flux caused by effects other than the alley position. We also report the uncertainty (σ) on each measurement as a percentage of the observed flux (Φ).

Detector	Deployment Notes	Position (m)	BRN $\text{m}^{-2} \text{MWh}^{-1}$	σ/Φ
Timing Cart	Stairwell	0.1	261.84 ± 3.89	1.5%
MARS	Alcove, Run 5	1.5	53 ± 23	43%
Timing Cart	Alcove, shutter 1 closed	2.1	2.49 ± 0.45	18%
Timing Cart	Alcove, shutter 1 open	2.1	18.72 ± 2.23	12%
Timing Cart	Alcove, take 2	2.1	13.25 ± 0.56	4.2%
Timing Cart	Alcove, atop NUBEs shield	2.1	19.60 ± 0.98	5%
Timing Cart	Alcove, take 4	2.1	25.84 ± 1.24	4.8%
SciBath	Alcove	3.7	13 ± 2.5	19%
NSC	Alcove, Basement C2.5	4	300	>50%
Timing Cart	Trench/Timing Rack	12.1	0.68 ± 0.23	34%
Timing Cart	Trench/Timing Rack, take 2	12.5	0.58 ± 0.14	25%
NSC	Center, Basement C4	13	5	>50%
MARS	Center, Run 3	15	0.58 ± 0.25	43%
MARS	Center, Run 2	18	1.20 ± 0.56	47%
Timing Cart	CsI Location	19.4	1.39 ± 0.45	33%
Timing Cart	CsI Location, take 2	19.4	0.50 ± 0.16	32%
Timing Cart	Rad Monitor, no lead	25.8	54.73 ± 2.29	4.1%
Timing Cart	Rad Monitor, lead	25.8	46.28 ± 1.22	2.6%
Timing Cart	Rad Monitor, floor	25.8	46.94 ± 2.46	5.2%
MARS	Rad Monitor, Run 4	26	335 ± 140	43%

Figure 8.16. As we previously mentioned, this suggests that there are alternate neutron paths that avoid the concrete/gravel fill since the edges of the alley have much higher rates than the center.

We have demonstrated the ability to measure the BRN rates throughout Neutrino Alley with MARS, and have obtained consistent results with measurements from other systems. Further characterization work is needed to reduce current uncertainties on the prompt BRN flux and to unlock the full potential of MARS as we focus on unfolding the incident neutron spectrum.

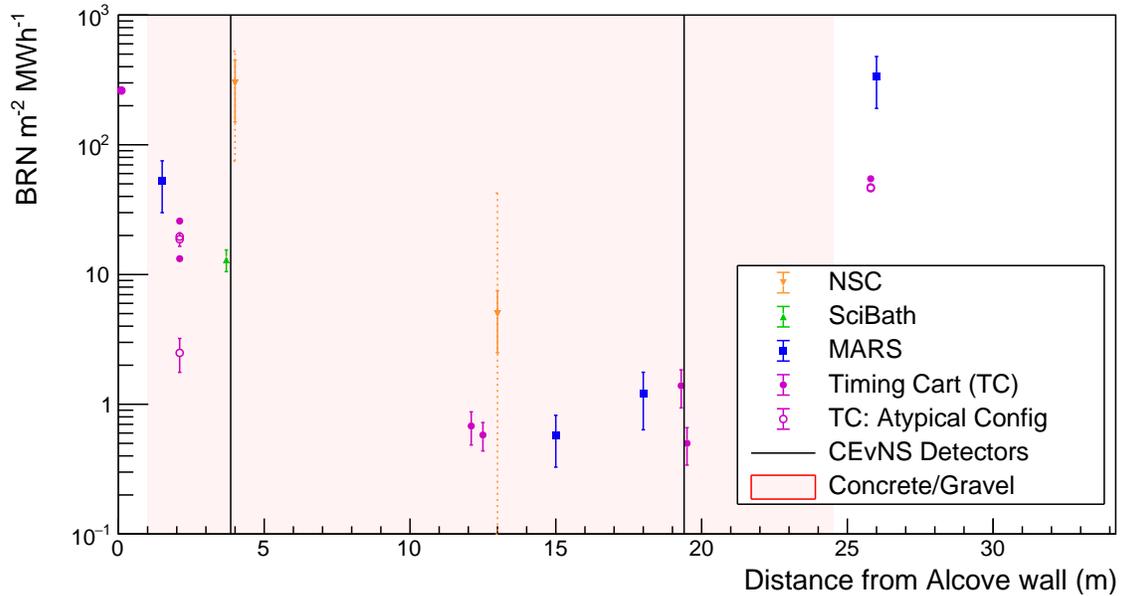


Figure 8.16: A comparison of all BRN measurements of the baseline neutron flux in Neutrino Alley, with a position estimated relative to the alcove wall. We show both 50% (solid) and 75% (dotted) uncertainties on the NSC measurements; the exact uncertainty is very large, but not precisely known. The “Atypical Config” label on some of the Timing Cart data correspond to the shaded-gray rows from Table 8.1. The shaded red region indicates the approximate location of the Concrete/Gravel shielding between Neutrino Alley and the SNS target, and the vertical lines indicate the approximate locations of the CsI and LAr subsystems.

Chapter 9

Conclusion

From Pauli’s proposition of the neutrino to present day, we have established that three active flavors of neutrinos exist; that while neutrinos are created and destroyed in flavor states, they can oscillate between flavors since the flavor states are superpositions of mass states; and that while they are among the most abundant particles in the universe, neutrinos have incredibly low interaction rates. With the COHERENT experiment, we contribute to measurements of low-energy neutrino scattering cross sections, with a specific emphasis on measurements of CEvNS.

Cross-section measurements generally require precise knowledge of the signal flux; this is even more important when searching for rare interactions, and all neutrino interactions can be considered rare. In this work, we have developed a full characterization of the signal flux of neutrinos produced at the SNS in Chapter 4, albeit with a large uncertainty associated with the production of pions from 1 GeV protons incident on a mercury target. This 10% systematic is the dominant uncertainty for all of COHERENT’s neutrino observables. We have also contributed to design efforts for a D₂O detector in Chapter 5 that will bring this systematic to the few-percent level. A commissioning version of this detector was deployed without D₂O in July 2022.

While understanding the expected signal is a necessary part of the cross section measurements, it is not the sole piece of the puzzle. Neutron events will cause similar nuclear-recoil signatures to CEvNS, and we must understand this incident flux of particles as well. Using MARS, we have made measurements of the prompt BRN flux throughout Neutrino Alley (Chapter 8) and are pushing towards a full detector characterization (Chapter 7) that can be used to unfold the incident spectrum for each measurement. Our flux measurements have a high uncertainty associated with a mismatch between calibration data and simulation, and it is of high importance to resolve these discrepancies in future work.

The efforts described in this work are far from the final story; each chapter has several avenues for further exploration, study, and improvement. However, the knowledge we have gained characterizing the fluxes at the Spallation Neutron Source will directly impact our efforts to interpret detector observations in Neutrino Alley and at the upcoming Second Target Station, supporting the next decade of COHERENT physics.

Appendix A

Daily Data Collected During MARS Production Runs

We begin a new MARS data run every 24-hours, allowing us to separate runs by date. This appendix presents the individual days of data recorded as live-time for each production run of MARS discussed in this work. Each table represents a calendar month of data, with each row in a table representing a single day of data-taking. A row highlighted in red in the tables below indicates that the data from that day was excluded from our search for BRNs in Neutrino Alley.

We use the alternating shading of gray / white to indicate the aggregation blocks in the data we include in the analyses; a change in color indicates the start of a new aggregation period. Aggregation periods can and do ignore calendar transitions, and we maintain consistent shading across the monthly tables. The tables included in these sections each contain the following columns:

- Date: The calendar date, based on the UTC Timestamp in the next column, on which the 24-hour run period began.
- UTC Time: The timestamp when the 24-hour run period began, formatted as a six-digit number reflecting the hour, minute, and second in UTC. It is important to note that for Production Run 3, the computer operating the DAQ had a 2 hour and 47 minute delay from local time, which modified the conversion to UTC time. We signify this with the asterisk “UTC*”. For data collected after July 2021, this is no longer an issue because we began using a different computer (with no observable delay) to operate the DAQ.

- Binary: The top-level folder for the binary data files associated with this 24-hour run on the ORNL computing cluster “hcdatal”. The specific file locations on the noted RAID server will generally be “/dataX/coherent/data/mars/prodRunNUMBER/rawData/”, where dataX is replaced by the string in this column.
- Size: The size of the ROOT file generated in the pre-processing step in megabytes.
- Rate: The average trigger rate of the MARS detector during this 24-hour period.
- E39 Rate: The average trigger rate of the E39 timing signal. Nominal operations are 60 Hz.
- E61 Rate: The average trigger rate of the E61 timing signal. During beam-off, the expected E61 rate is 0 Hz, though accelerator testing is indicated by non-zero values. During beam-on, the E61 rate fluctuates depending on how many protons were sent to target during the run period. Nominal 1.4 MW operations result in E61 rates near 60 Hz.

A.1 Production Run 3

Table A.1: Details of Run 3 data from January 2019. MARS was moved on January 23, 2019, but a DAQ crash was observed in late February 2019. We generally ignore all data from January 2019 in our analyses because the pre-processed file sizes are small relative to the observed trigger rate.

Date	UTC* Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2019-01-24	003219	data5	602	597.773	—	—
2019-01-25	003220	data5	32	591.681	—	—
2019-01-26	003221	data5	32	558.536	59.939	39.867
2019-01-27	003222	data5	30	538.948	59.948	52.074
2019-01-28	003223	data5	32	549.398	59.949	59.459
2019-01-29	003224	data5	32	539.071	59.951	36.511
2019-01-30	003225	data5	33	570.017	59.952	49.198
2019-01-31	003226	data5	32	621.674	59.933	59.069

Table A.2: Details of the Run 3 data from February 2019. A DAQ-crash was identified on February 20, 2019, and data-taking halted while the issue was resolved. Normal operations resumed on February 27, 2019, and we exclude all prior data from this month given the low preprocessed file size relative to the observed trigger rate.

Date	UTC* Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2019-02-01	003227	data5	32	637.741	59.939	59.614
2019-02-02	003228	data5	46	624.398	59.954	39.761
2019-02-03	003229	data5	33	606.550	59.951	45.496
2019-02-04	003230	data5	32	648.701	59.925	58.633
2019-02-05	003231	data5	45	599.984	59.945	25.091
2019-02-06	003232	data5	33	504.842	59.950	2.305
2019-02-07	003233	data5	28	607.630	59.943	59.087
2019-02-08	003234	data5	32	553.619	59.954	7.872
2019-02-09	003235	data5	30	588.228	59.939	54.834
2019-02-10	003236	data5	32	624.733	59.940	58.963
2019-02-11	003237	data5	32	602.437	59.943	31.318
2019-02-12	003238	data5	33	506.830	59.953	4.060
2019-02-13	003239	data5	29	669.322	59.940	55.209
2019-02-14	003240	data5	54	724.812	59.938	54.884
2019-02-15	003241	data5	49	650.746	59.935	55.576
2019-02-16	003242	data5	59	667.987	59.936	58.819
2019-02-17	003243	data5	62	676.202	59.894	59.492
2019-02-18	003244	data5	49	667.306	59.945	59.620
2019-02-19	003245	data5	48	622.473	59.943	36.489
2019-02-27	001704	data5	612	603.684	59.946	52.070
2019-02-28	001705	data5	642	637.804	59.939	59.828

Table A.3: Details of the Run 3 data from March 2019. We exclude data collected on March 21, 2019 due to the unreasonably high trigger rate and preprocessed file size. We also ignore data from March 30 and 31, 2019 due to the irregularity of the E39 rate.

Date	UTC* Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2019-03-01	001706	data5	644	640.414	59.947	59.603
2019-03-02	001707	data5	646	641.768	59.940	59.830
2019-03-03	001708	data5	638	630.904	59.939	54.517
2019-03-04	001709	data5	643	639.089	59.950	59.459
2019-03-05	001710	data5	618	600.752	59.951	32.972
2019-03-06	001711	data5	624	619.698	59.933	57.835
2019-03-07	001712	data5	650	647.570	59.911	59.567
2019-03-08	001713	data5	650	647.430	59.923	59.751
2019-03-09	001714	data5	630	623.230	59.937	45.123
2019-03-10	001715	data5	646	629.126	58.411	47.361
2019-03-11	001716	data5	541	533.008	59.946	22.110
2019-03-12	001717	data5	528	517.257	59.954	33.981
2019-03-13	001718	data5	591	584.027	59.930	57.927
2019-03-14	001719	data5	606	601.881	59.950	59.457
2019-03-15	001720	data5	582	576.146	59.945	42.684
2019-03-16	001721	data5	614	610.637	59.947	59.396
2019-03-17	001722	data5	611	609.104	59.942	59.625
2019-03-18	001723	data5	607	603.009	59.957	56.549
2019-03-19	001724	data5	581	567.258	59.951	33.677
2019-03-20	001725	data5	601	594.773	59.941	56.379
2019-03-21	001726	data5	3047	14062.479	54.036	9.882
2019-03-22	001727	data5	529	560.721	59.780	0.217
2019-03-23	001728	data5	486	467.774	59.955	0.002
2019-03-24	001729	data5	469	458.247	59.961	0.000
2019-03-25	001730	data5	467	453.141	59.964	0.063
2019-03-26	001731	data5	468	449.321	59.956	3.715
2019-03-27	001732	data5	457	444.812	59.956	0.000
2019-03-28	001733	data5	463	441.651	59.958	0.000
2019-03-29	001734	data5	452	439.790	59.950	0.000
2019-03-30	001735	data5	436	245.945	33.645	0.000
2019-03-31	001736	data5	346	65.970	9.059	0.000

Table A.4: Details of the Run 3 data from April 2019. We exclude nearly half of the data this month due to work occurring on the E39 timing signal during the long shutdown of the SNS.

Date	UTC* Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2019-04-01	001737	data5	344	65.295	9.035	0.000
2019-04-02	001738	data5	341	65.366	9.064	0.000
2019-04-03	001739	data5	381	107.733	14.955	0.000
2019-04-04	001740	data5	309	112.357	15.628	0.000
2019-04-05	001741	data5	333	145.077	20.218	0.000
2019-04-06	001742	data5	443	427.985	59.722	0.000
2019-04-07	001743	data5	442	428.155	59.960	0.000
2019-04-08	001744	data5	441	427.887	59.966	0.000
2019-04-09	001745	data5	441	427.133	59.960	0.000
2019-04-10	001746	data5	440	425.982	59.965	0.000
2019-04-11	001747	data5	439	424.721	59.960	0.000
2019-04-12	001748	data5	286	218.963	30.985	0.000
2019-04-13	001749	data5	437	422.562	59.959	0.000
2019-04-14	001750	data5	437	422.731	59.959	0.000
2019-04-15	001751	data5	437	422.353	59.959	0.000
2019-04-16	001752	data5	439	424.778	59.960	0.000
2019-04-17	001753	data5	429	198.917	28.095	0.000
2019-04-18	001754	data5	437	422.805	59.967	0.333
2019-04-19	001755	data5	437	422.083	59.966	0.999
2019-04-20	001756	data5	445	421.484	59.959	0.143
2019-04-21	001757	data5	438	422.789	59.960	0.000
2019-04-22	001758	data5	440	425.470	59.964	0.000
2019-04-23	001759	data5	349	100.799	14.153	0.000
2019-04-24	001800	data5	429	204.024	28.422	0.000
2019-04-25	001801	data5	319	71.366	9.907	0.003
2019-04-26	001802	data5	335	223.075	31.113	0.496
2019-04-27	001803	data5	395	147.470	20.718	0.000
2019-04-28	001804	data5	345	80.993	11.451	0.000
2019-04-29	001805	data5	342	80.342	11.426	0.000
2019-04-30	001806	data5	332	98.410	13.997	0.000

Table A.5: Details of the Run 3 data from May 2019. We exclude a large fraction of the data this month due to work occurring on the E39 timing signal during the long shutdown of the SNS.

Date	UTC* Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2019-05-01	001807	data5	440	422.843	59.963	0.292
2019-05-02	001808	data5	320	170.931	24.868	0.034
2019-05-03	001809	data5	430	411.271	59.953	0.167
2019-05-04	001810	data5	416	212.654	31.122	0.000
2019-05-05	001811	data5	425	409.472	59.950	0.000
2019-05-06	001812	data5	362	94.522	13.879	0.000
2019-05-07	001813	data5	193	64.313	9.459	0.002
2019-05-08	001814	data5	331	65.447	9.578	0.000
2019-05-09	001815	data5	322	64.639	9.449	0.000
2019-05-10	001816	data5	321	64.290	9.432	0.000
2019-05-11	001817	data5	319	63.655	9.399	0.000
2019-05-12	001818	data5	330	64.332	9.517	0.000
2019-05-13	001819	data5	427	276.293	40.596	0.000
2019-05-14	001820	data5	431	296.698	43.298	0.000
2019-05-15	001821	data5	374	119.789	17.348	0.000
2019-05-16	001822	data5	430	415.064	59.956	0.000
2019-05-17	001823	data5	428	412.065	59.952	0.000
2019-05-18	001824	data5	404	194.371	28.365	0.000
2019-05-19	001825	data5	291	328.512	48.142	0.000
2019-05-20	001826	data5	347	59.499	8.706	0.000
2019-05-21	001827	data5	329	65.267	9.542	0.000
2019-05-22	001828	data5	326	64.708	9.455	0.000
2019-05-23	001829	data5	324	63.845	9.318	0.000
2019-05-24	001830	data5	323	64.652	9.450	0.000
2019-05-25	001831	data5	322	64.346	9.432	0.000
2019-05-26	001832	data5	321	64.185	9.421	0.000
2019-05-27	001833	data5	322	64.163	9.424	0.000
2019-05-28	001834	data5	321	63.966	9.412	0.000
2019-05-29	001835	data5	195	63.674	9.397	0.006
2019-05-30	001836	data5	296	137.042	20.247	0.000
2019-05-31	001837	data5	424	406.155	59.941	0.275

Table A.6: Details of the Run 3 data from June 2019.

Date	UTC* Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2019-06-01	001838	data5	424	407.340	59.945	0.000
2019-06-02	001839	data5	422	406.135	59.950	0.000
2019-06-03	001840	data5	422	406.769	59.943	0.000
2019-06-04	001841	data5	425	409.792	59.949	0.000
2019-06-05	001842	data5	425	409.238	59.948	0.000
2019-06-06	001843	data5	425	409.193	59.950	0.166
2019-06-07	001844	data5	425	408.803	59.951	0.987
2019-06-08	001845	data5	427	408.109	59.949	0.666
2019-06-09	001846	data5	424	407.827	59.954	0.000
2019-06-10	001847	data5	425	408.277	59.956	0.131
2019-06-11	001848	data5	428	403.761	59.946	0.300
2019-06-12	001849	data5	431	405.403	59.944	0.637
2019-06-13	001850	data5	429	409.706	59.951	0.858
2019-06-14	001851	data5	434	408.923	59.943	0.122
2019-06-15	001852	data5	432	407.651	59.954	0.595
2019-06-16	001853	data5	426	405.274	59.950	0.934
2019-06-17	001854	data5	429	404.381	59.942	0.402
2019-06-18	001855	data5	430	404.030	59.943	0.085
2019-06-19	001856	data5	428	404.854	59.943	0.445
2019-06-20	001857	data5	430	404.728	59.942	0.000
2019-06-21	001858	data5	428	404.891	59.944	0.445
2019-06-22	001859	data5	483	470.919	59.949	37.068
2019-06-23	001900	data5	586	580.263	59.948	49.377
2019-06-24	001901	data5	685	685.239	59.941	55.326
2019-06-25	001902	data5	720	720.550	59.933	48.525
2019-06-26	001903	data5	767	768.141	59.896	54.542
2019-06-27	001904	data5	762	763.819	59.927	57.192
2019-06-28	001905	data5	833	915.204	59.916	55.622
2019-06-29	001906	data5	736	731.036	59.930	56.599
2019-06-30	001907	data5	755	751.285	59.933	56.496

Table A.7: Details of the Run 3 data from July 2019. We exclude the data on July 5, 6, and 13, 2019 due to the unexpectedly high trigger rates.

Date	UTC* Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2019-07-01	001908	data5	740	739.537	59.917	57.075
2019-07-02	001909	data5	630	633.376	59.943	38.905
2019-07-03	001910	data5	668	674.430	59.940	52.546
2019-07-04	001911	data5	1856	2382.479	59.790	56.586
2019-07-05	001912	data5	3200	3811.671	59.627	57.542
2019-07-06	001913	data5	710	716.827	59.909	50.769
2019-07-07	001914	data5	750	749.898	59.925	57.155
2019-07-08	001915	data5	752	752.298	59.926	59.585
2019-07-09	001916	data5	630	636.607	59.106	35.372
2019-07-10	001917	data5	589	590.900	59.947	35.779
2019-07-11	001918	data5	712	697.472	59.924	49.758
2019-07-12	001919	data5	732	731.396	59.933	57.827
2019-07-13	001920	data5	1716	1924.880	59.824	55.812
2019-07-14	001921	data5	770	771.101	59.921	56.527
2019-07-15	001922	data5	723	722.524	59.932	50.777
2019-07-16	001923	data5	738	740.543	59.898	57.015
2019-07-17	001924	data5	804	805.802	59.918	59.468
2019-07-18	001925	data5	784	785.050	59.923	59.776
2019-07-19	001926	data5	732	733.754	59.940	58.947
2019-07-20	001927	data5	675	670.599	59.940	57.004
2019-07-21	001928	data5	648	644.040	59.947	58.530
2019-07-22	001929	data5	644	639.097	59.947	59.425
2019-07-23	001930	data5	584	585.904	59.929	36.018
2019-07-24	001931	data5	620	621.003	59.882	52.351
2019-07-25	001932	data5	693	690.751	59.936	59.200
2019-07-26	001933	data5	920	1052.511	59.901	58.015
2019-07-27	001934	data5	699	707.152	59.936	58.050
2019-07-28	001935	data5	594	585.102	59.919	32.199
2019-07-29	001936	data5	605	611.225	59.939	45.336
2019-07-30	001937	data5	835	846.091	59.893	58.296
2019-07-31	001938	data5	890	901.537	59.922	54.848

Table A.8: Details of the Run 3 data from August 2019. Calibration data using a ^{60}Co source was collected on August 15, and normal operations resumed on August 16. A lot of the data taken in the early part of the month is excluded from our analysis due to the low file sizes in the preprocessed data.

Date	UTC* Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2019-08-01	001939	data5	1.1	784.042	59.922	50.955
2019-08-02	001940	data5	0.0022	768.727	59.937	37.391
2019-08-03	001941	data5	22	749.144	59.926	49.985
2019-08-04	001942	data5	0.0022	761.618	59.930	53.866
2019-08-05	001943	data5	0.0022	776.411	59.936	59.613
2019-08-06	001944	data5	0.0022	652.466	59.927	33.624
2019-08-07	001945	data5	0.0022	488.374	59.957	0.997
2019-08-08	001946	data5	0.0022	475.999	59.463	0.500
2019-08-09	001947	data5	472	457.291	59.960	0.000
2019-08-10	001948	data5	463	448.368	59.964	0.000
2019-08-11	001949	data5	301	441.594	59.958	0.454
2019-08-12	001950	data5	0.0022	437.926	59.951	0.606
2019-08-13	001951	data5	0.935	434.725	59.938	0.642
2019-08-14	001952	data5	0.0022	433.792	59.930	0.609
2019-08-16	233524	data5	9.6	436.689	59.957	0.000
2019-08-17	233525	data5	84	861.131	59.694	48.296
2019-08-18	233526	data5	0.0022	1284.906	59.555	57.048
2019-08-19	233527	data5	0.0022	1328.243	59.882	59.136
2019-08-20	233528	data5	0.0022	1347.993	59.876	56.165
2019-08-21	233529	data5	0.0022	1410.836	59.858	54.557
2019-08-22	233530	data5	16	1199.687	59.883	57.083
2019-08-23	233531	data5	0.0022	1093.407	59.898	50.416
2019-08-24	233532	data5	0.0022	798.039	59.922	18.180
2019-08-25	233533	data5	501	488.343	59.088	0.000
2019-08-26	233534	data5	483	468.390	59.953	0.000
2019-08-27	233535	data5	358	458.206	53.478	0.000
2019-08-28	233536	data5	5.3	451.695	59.791	0.000
2019-08-29	233537	data5	331	445.838	59.956	0.167
2019-08-30	233538	data5	413	441.925	59.952	0.100
2019-08-31	233539	data5	455	438.573	59.951	0.000

Table A.9: Details of the Run 3 data from September 2019.

Date	UTC* Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2019-09-01	233540	data5	452	436.170	59.955	0.000
2019-09-02	233541	data5	451	433.901	59.900	0.000
2019-09-03	233542	data5	451	433.212	59.906	0.000
2019-09-04	233543	data5	450	431.529	59.883	0.199
2019-09-05	233544	data5	453	430.641	59.832	0.180
2019-09-06	233545	data5	449	430.172	59.938	0.747
2019-09-07	233546	data5	450	433.692	59.946	0.000
2019-09-08	233547	data5	452	435.349	59.974	0.000
2019-09-09	233548	data5	483	471.563	59.954	44.220
2019-09-10	233549	data5	535	509.619	59.944	49.712
2019-09-11	233550	data5	545	534.548	59.950	58.527
2019-09-12	233551	data5	551	541.664	59.952	59.283
2019-09-13	233552	data5	560	548.994	59.952	59.439
2019-09-14	233553	data5	570	560.212	59.940	59.710
2019-09-15	233554	data5	576	565.874	59.943	59.587
2019-09-16	233555	data5	583	572.620	59.953	59.326
2019-09-17	233556	data5	588	578.444	59.952	59.435
2019-09-18	233557	data5	576	566.507	59.952	56.763
2019-09-19	233558	data5	586	576.113	59.953	54.640
2019-09-20	233559	data5	588	581.132	59.821	59.713
2019-09-21	233600	data5	589	581.658	59.767	47.207
2019-09-22	233601	data5	586	579.864	59.945	55.726
2019-09-23	233602	data5	591	581.958	59.879	53.173
2019-09-24	233603	data5	590	582.299	59.885	59.776
2019-09-25	233604	data5	590	582.783	59.895	53.468
2019-09-26	233605	data5	594	586.373	59.931	55.190
2019-09-27	233606	data5	594	588.244	59.934	59.758
2019-09-28	233607	data5	597	590.902	59.933	59.819
2019-09-29	233608	data5	596	590.729	59.937	59.638
2019-09-30	233609	data5	591	577.508	59.596	26.940

Table A.10: Details of the Run 3 data from October 2019. Calibration data using a DT generator was collected on October 25, 2019; the surrounding week of data was used for testing, and we did not record typical beam data.

Date	UTC* Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2019-10-01	233610	data5	554	541.167	59.944	30.700
2019-10-02	233611	data5	503	492.714	59.880	0.000
2019-10-03	233612	data5	473	461.638	59.927	0.000
2019-10-04	233613	data5	470	458.493	59.953	0.000
2019-10-05	233614	data5	465	452.856	59.960	0.000
2019-10-06	233615	data5	459	446.628	59.951	0.000
2019-10-07	233616	data5	454	440.911	59.653	0.000
2019-10-08	233617	data5	449	439.657	59.941	0.000
2019-10-09	233618	data5	444	432.846	59.887	0.000
2019-10-10	233619	data5	435	429.201	59.968	0.000
2019-10-11	233620	data5	438	424.790	59.962	0.000
2019-10-12	233621	data5	436	421.898	59.962	0.000
2019-10-13	233622	data5	433	420.126	59.954	0.000
2019-10-14	233623	data5	431	416.902	59.962	0.000
2019-10-15	233624	data5	432	415.623	59.967	0.000
2019-10-16	233625	data5	3.3	415.738	59.953	0.000
2019-10-17	233626	data5	434	416.637	19.982	0.000
2019-10-18	233627	data5	433	418.045	59.952	0.000
2019-10-19	233628	data5	432	417.702	59.959	0.000
2019-10-20	233629	data5	432	416.462	59.950	0.000
2019-10-21	233630	data5	432	415.533	59.952	0.113
2019-10-22	233631	data5	432	414.813	59.951	0.220
2019-10-23	233632	data5	385	412.385	59.957	0.639
2019-10-28	155829	data5	459	438.405	59.955	0.028
2019-10-29	155830	data5	471	467.652	59.951	0.247
2019-10-30	155831	data5	442	418.931	59.954	0.270
2019-10-31	155832	data5	504	497.967	59.955	39.086

Table A.11: Details of the Run 3 data from November 2019. Here, we started to notice increased trigger rates and file sizes during beam-on operations; we do include this data regardless because the increase is roughly constant.

Date	UTC* Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2019-11-01	155833	data5	540	545.165	59.940	33.142
2019-11-02	155834	data5	614	627.295	59.936	24.480
2019-11-03	155835	data5	583	615.152	59.937	33.638
2019-11-04	155836	data5	896	888.680	59.871	56.961
2019-11-05	155837	data5	891	890.111	59.920	52.039
2019-11-06	155838	data5	1028	1015.927	59.893	58.052
2019-11-07	155839	data5	944	935.780	59.911	55.542
2019-11-08	155840	data5	974	964.105	59.908	59.619
2019-11-09	155841	data5	949	937.539	59.908	53.921
2019-11-10	155842	data5	1014	998.005	59.936	59.621
2019-11-11	155843	data5	926	926.036	59.911	51.408
2019-11-12	155844	data5	834	856.893	59.918	42.334
2019-11-13	155845	data5	1026	1008.783	59.819	56.163
2019-11-14	155846	data5	1066	1043.883	59.892	57.601
2019-11-15	155847	data5	1101	1076.988	59.878	57.300
2019-11-16	155848	data5	1109	1077.658	59.899	56.952
2019-11-17	155849	data5	1174	1140.816	59.881	58.609
2019-11-18	155850	data5	1175	1122.997	59.879	48.387
2019-11-19	155851	data5	741	766.736	59.936	22.269
2019-11-20	155852	data5	1074	1057.216	59.898	59.413
2019-11-21	155853	data5	1133	1107.779	59.876	59.477
2019-11-22	155854	data5	1159	1134.546	59.897	59.347
2019-11-23	155855	data5	1181	1146.831	59.883	58.604
2019-11-24	155856	data5	937	942.507	59.769	42.066
2019-11-25	155857	data5	1167	1145.545	59.879	56.927
2019-11-26	155858	data5	949	971.848	59.906	45.387
2019-11-27	155859	data5	1064	1047.776	59.889	55.383
2019-11-28	155900	data5	1056	1043.499	58.858	58.589
2019-11-29	155901	data5	1115	1093.614	59.894	59.457
2019-11-30	155902	data5	723	730.366	59.920	44.389

Table A.12: Details of the Run 3 data from December 2019. We start to exclude data with exceptionally high rates compared to the normal beam-on operations, such as the high trigger rates on December 3 and 23, 2019.

Date	UTC* Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2019-12-01	155903	data5	1253	2329.982	59.806	12.097
2019-12-02	155904	data5	797	1036.922	59.619	30.362
2019-12-03	155905	data5	5602	4839.357	59.477	57.120
2019-12-04	155906	data5	1522	1922.273	59.805	59.111
2019-12-05	155907	data5	1196	1172.305	59.720	59.407
2019-12-06	155908	data5	1108	1081.562	59.899	55.368
2019-12-07	155909	data5	1125	1104.429	59.779	59.432
2019-12-08	155910	data5	1131	1109.738	59.866	59.068
2019-12-09	155911	data5	1214	1178.268	59.883	56.648
2019-12-10	155912	data5	880	921.210	59.908	42.127
2019-12-11	155913	data5	1126	1101.063	59.846	58.825
2019-12-12	155914	data5	1156	1128.049	59.898	56.878
2019-12-13	155915	data5	1306	1263.397	59.867	58.668
2019-12-14	155916	data5	1279	1241.141	59.891	59.589
2019-12-15	155917	data5	1234	1202.289	59.875	58.924
2019-12-16	155918	data5	1364	1765.140	59.813	44.224
2019-12-17	155919	data5	1151	1410.151	59.869	33.368
2019-12-18	155920	data5	1247	1218.500	59.886	54.334
2019-12-19	155921	data5	1176	1152.697	59.888	57.945
2019-12-20	155922	data5	1176	1154.198	59.888	59.623
2019-12-21	155923	data5	1155	1126.396	59.888	55.724
2019-12-22	155924	data5	1224	1194.128	59.885	59.402
2019-12-23	155925	data5	2856	6700.839	59.348	0.000
2019-12-24	155926	data5	519	508.428	59.956	0.000
2019-12-25	155927	data5	495	483.404	59.960	0.000
2019-12-26	155928	data5	485	468.817	59.952	0.446
2019-12-27	155929	data5	479	459.044	59.826	0.253
2019-12-28	155930	data5	475	453.344	59.956	0.246
2019-12-29	155931	data5	474	448.815	59.960	0.088
2019-12-30	155932	data5	468	446.477	59.954	0.529
2019-12-31	155933	data5	848	901.242	59.413	52.317

Table A.13: Details of the Run 3 data from January 2020. We exclude the data from January 7, 2020 due to the incredibly high trigger rate.

Date	UTC* Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2020-01-01	155934	data5	1047	1026.944	59.897	51.469
2020-01-02	155935	data5	1025	1023.393	59.902	49.946
2020-01-03	155936	data5	1148	1107.841	59.888	55.631
2020-01-04	155937	data5	1091	1071.929	59.846	54.738
2020-01-05	155938	data5	1094	1076.294	59.883	58.742
2020-01-06	155939	data5	1128	1106.157	59.886	55.889
2020-01-07	155940	data5	6068	10612.861	59.041	9.818
2020-01-08	155941	data5	1071	1056.622	59.898	58.991
2020-01-09	155942	data5	890	891.180	59.917	39.004
2020-01-10	155943	data5	1160	1131.780	59.902	59.106
2020-01-11	155944	data5	1080	1075.034	59.897	47.049
2020-01-12	155945	data5	1218	1187.059	59.886	58.657
2020-01-13	155946	data5	1220	1186.276	59.881	58.921
2020-01-14	155947	data5	934	965.309	59.911	40.225
2020-01-15	155948	data5	1193	1173.116	59.888	58.147
2020-01-16	155949	data5	1178	1248.230	59.878	59.016
2020-01-17	155950	data5	1111	1082.722	59.889	57.004
2020-01-18	155951	data5	1115	1090.259	59.901	59.310
2020-01-19	155952	data5	993	990.496	59.904	47.480
2020-01-20	155953	data5	615	628.003	59.938	0.715
2020-01-21	155954	data5	778	842.330	59.921	39.264
2020-01-22	155955	data5	1168	1140.870	59.878	56.684
2020-01-23	155956	data5	1161	1138.878	59.890	59.284
2020-01-24	155957	data5	1144	1122.631	59.894	58.976
2020-01-25	155958	data5	1195	1173.365	59.879	59.133
2020-01-26	155959	data5	1012	1022.976	59.901	39.620
2020-01-27	160000	data5	910	966.235	59.898	47.263
2020-01-28	160001	data5	1237	1210.171	59.880	59.698
2020-01-29	160002	data5	1242	1213.605	59.895	53.934
2020-01-30	160003	data5	1201	1164.170	59.886	57.674
2020-01-31	160004	data5	1322	1280.443	59.872	59.035

Table A.14: Details of the Run 3 data from February 2020. We exclude the data from February 4, 2020 due to the high trigger rate, which ultimately crashed the DAQ for the following days. Data-taking resumed on February 7, but the health monitoring of the detector did not resume until February 8. Due to E39 timing signal issues, we also exclude data taken on February 25, 26, 28, and 29, 2020.

Date	UTC* Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2020-02-01	160005	data5	1353	1305.487	59.873	59.630
2020-02-02	160006	data5	1359	1312.430	59.874	58.538
2020-02-03	160007	data5	1511	1976.945	59.811	45.899
2020-02-04	160008	data5	5901	9728.238	59.045	13.988
2020-02-07	160011	data5	1240	—	—	—
2020-02-08	160012	data5	1056	1054.406	59.905	49.185
2020-02-09	160013	data5	1163	1142.030	59.882	59.581
2020-02-10	160014	data5	1175	1150.468	59.888	56.892
2020-02-11	160015	data5	956	967.968	59.911	47.325
2020-02-12	160016	data5	2274	2260.879	59.787	59.393
2020-02-13	160017	data5	1250	1342.989	59.865	58.719
2020-02-14	160018	data5	1075	1050.813	59.894	56.648
2020-02-15	160019	data5	1029	1017.915	59.864	53.915
2020-02-16	160020	data5	1096	1073.345	59.891	55.898
2020-02-17	160021	data5	1058	1032.460	59.855	56.777
2020-02-18	160022	data5	1079	1065.911	59.904	59.561
2020-02-19	160023	data5	1034	1112.224	59.891	38.918
2020-02-20	160024	data5	535	520.778	59.949	0.370
2020-02-21	160025	data5	490	466.272	59.946	0.046
2020-02-22	160026	data5	466	448.053	59.955	0.000
2020-02-23	160027	data5	458	438.477	59.954	0.000
2020-02-24	160028	data5	450	429.988	59.908	0.000
2020-02-25	160029	data5	5.3	425.228	1.795	0.000
2020-02-26	160030	data5	0.0087	420.532	0.000	0.000
2020-02-27	160031	data5	447	416.495	59.949	0.000
2020-02-28	160032	data5	52	415.880	19.984	0.000
2020-02-29	160033	data5	0.0087	415.266	0.000	0.000

Table A.15: Details of the Run 3 data from March 2020. Data from March 1, 6, 7, and 8, 2020 are excluded as a result of E39 issues. Furthermore, data was not taken on March 2, 2020

Date	UTC* Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2020-03-01	160034	data5	0.0087	413.996	0.000	0.000
2020-03-03	160036	data5	431	409.424	59.955	0.000
2020-03-04	160037	data5	428	406.419	59.945	0.000
2020-03-05	160038	data5	426	404.514	59.937	0.000
2020-03-06	160039	data5	57	401.982	23.984	0.000
2020-03-07	160040	data5	0.0087	399.646	0.000	0.000
2020-03-08	160041	data5	0.000252	399.163	0.000	0.000
2020-03-09	160042	data5	420	397.447	59.730	0.000
2020-03-10	160043	data5	419	396.130	59.956	0.000
2020-03-11	160044	data5	418	395.244	59.955	0.000
2020-03-12	160045	data5	417	393.441	59.955	0.000
2020-03-13	160046	data5	418	391.624	59.955	0.000
2020-03-14	160047	data5	411	390.433	59.954	0.000
2020-03-15	160048	data5	410	389.667	59.955	0.000
2020-03-16	160049	data5	412	386.379	59.953	0.000
2020-03-17	160050	data5	406	384.521	59.951	0.000
2020-03-18	160051	data5	406	383.741	59.952	0.000
2020-03-19	160052	data5	407	383.547	59.952	0.000
2020-03-20	160053	data5	407	384.141	59.958	0.000
2020-03-21	160054	data5	405	382.911	59.962	0.000
2020-03-22	160055	data5	404	381.063	59.952	0.000
2020-03-23	160056	data5	404	380.788	59.949	0.000
2020-03-24	160057	data5	406	380.848	59.947	0.000
2020-03-25	160058	data5	405	380.119	59.951	0.000
2020-03-26	160059	data5	403	378.023	59.947	0.000
2020-03-27	160060	data5	406	377.631	59.950	0.094
2020-03-28	160061	data5	408	379.898	59.950	0.000
2020-03-29	160062	data5	403	380.979	59.956	0.894

Table A.16: Details of the Run 3 data from April 2020. Though the trigger rates are much higher than typical, they are consistently high during this beam-on period, which extends into May 2020, and we allow all of this data into our analyses.

Date	UTC* Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2020-04-01	194735	data6	405	379.813	59.946	0.313
2020-04-02	194736	data6	404	379.747	59.958	0.037
2020-04-03	200426	data6	491	499.479	59.952	18.646
2020-04-04	200427	data6	649	709.714	59.937	27.808
2020-04-05	200428	data6	1134	1120.392	59.881	50.315
2020-04-06	200429	data6	1194	1256.846	59.884	46.580
2020-04-07	200430	data6	2041	1917.200	59.795	54.528
2020-04-08	200431	data6	3386	3045.680	59.723	57.444
2020-04-09	200432	data6	492	471.667	59.952	0.177
2020-04-10	200433	data6	451	428.329	59.962	0.197
2020-04-11	200434	data6	440	414.858	59.958	0.212
2020-04-12	200435	data6	2121	2459.797	59.773	43.459
2020-04-13	200436	data6	967	1018.547	59.899	59.444
2020-04-14	200437	data6	824	827.910	59.923	55.102
2020-04-15	200438	data6	705	708.815	59.932	43.325
2020-04-16	200439	data6	733	732.028	59.929	53.610
2020-04-17	200440	data6	794	789.237	59.923	57.204
2020-04-18	200441	data6	768	763.338	59.934	59.783
2020-04-19	200442	data6	1163	1145.735	59.884	58.130
2020-04-20	200443	data6	950	978.132	59.907	48.122
2020-04-21	200444	data6	574	594.687	59.937	40.453
2020-04-22	200445	data6	792	807.121	59.917	55.479
2020-04-23	200446	data6	1329	1294.889	59.871	57.518
2020-04-24	200447	data6	948	948.294	59.909	55.577
2020-04-25	200448	data6	1074	1069.775	59.894	58.601
2020-04-26	200449	data6	655	657.406	59.927	42.562
2020-04-27	200450	data6	443	421.552	59.959	0.494
2020-04-28	200451	data6	938	1560.776	59.860	38.321
2020-04-29	200452	data6	4464	3506.144	59.652	59.190
2020-04-30	200453	data6	4454	3503.510	59.622	55.197

Table A.17: Details of the Run 3 data from May 2020. The trigger rate on May 6, 2020 was deemed atypically high, and this day is excluded from our analyses. We also exclude data from May 18, 19, and 20 due to higher-than-typical rates.

Date	UTC* Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2020-05-02	034731	data6	3043	2584.055	59.753	56.479
2020-05-03	034732	data6	4213	3393.888	59.668	58.383
2020-05-04	034733	data6	4906	3663.382	59.627	59.357
2020-05-05	034734	data6	3466	3161.519	59.707	49.969
2020-05-06	034735	data6	8410	5088.317	59.531	56.659
2020-05-07	034736	data6	2044	1999.413	59.812	59.486
2020-05-08	034737	data6	1820	1742.109	59.807	58.736
2020-05-09	034738	data6	1596	1517.654	59.846	59.415
2020-05-10	034739	data6	1570	1494.079	59.862	59.039
2020-05-11	034740	data6	1571	1494.060	59.876	58.969
2020-05-12	034741	data6	1122	1189.311	59.887	40.292
2020-05-13	034742	data6	1511	1435.898	59.865	63.149
2020-05-14	034743	data6	1646	1557.350	59.851	62.958
2020-05-15	034744	data6	1684	1590.810	59.856	59.517
2020-05-16	034745	data6	1628	1533.852	59.824	54.493
2020-05-17	034746	data6	1712	1656.036	59.847	54.541
2020-05-18	034747	data6	2853	2484.320	59.760	59.592
2020-05-19	034748	data6	3135	2683.469	59.649	51.417
2020-05-20	034749	data6	2142	1956.220	59.800	59.166
2020-05-21	034750	data6	1716	1600.083	59.848	57.034
2020-05-22	034751	data6	1697	1581.500	59.854	55.923
2020-05-23	034752	data6	1734	1618.391	59.862	57.314
2020-05-24	034753	data6	1104	1198.011	59.882	31.748
2020-05-25	034754	data6	683	694.817	59.933	40.421
2020-05-26	034755	data6	801	884.627	59.917	39.100
2020-05-27	034756	data6	1639	1544.512	59.848	57.179
2020-05-28	034757	data6	1740	1627.004	59.840	58.132
2020-05-29	034758	data6	1806	1692.598	59.798	59.618
2020-05-30	034759	data6	1770	1656.916	59.815	59.700
2020-05-31	034800	data6	2001	1906.219	59.828	59.356

Table A.18: Details of the Run 3 data from June 2020.

Date	UTC* Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2020-06-01	034801	data6	1740	1637.761	59.834	59.425
2020-06-02	034802	data6	1249	1318.641	59.879	45.673
2020-06-03	034803	data6	1833	1725.711	59.840	59.228
2020-06-04	034804	data6	1875	1739.238	59.837	53.003
2020-06-05	034805	data6	2015	1852.653	59.818	57.338
2020-06-06	034806	data6	1863	1741.511	59.823	59.255
2020-06-07	034807	data6	1916	1790.797	59.823	57.775
2020-06-08	034808	data6	1902	1783.246	59.829	59.655
2020-06-09	034809	data6	1356	1548.420	59.856	46.154
2020-06-10	034810	data6	1936	1857.591	59.817	55.422
2020-06-11	034811	data6	1676	1614.025	59.836	54.593
2020-06-12	034812	data6	1580	1502.384	59.868	53.969
2020-06-13	034813	data6	1361	1382.738	59.860	47.501
2020-06-14	034814	data6	1575	1488.767	59.856	55.172
2020-06-15	034815	data6	1580	1486.798	59.864	54.968
2020-06-16	034816	data6	1400	1407.051	59.858	49.578
2020-06-17	034817	data6	1652	1561.779	59.851	59.206
2020-06-18	034818	data6	1630	1541.763	59.840	59.675
2020-06-19	034819	data6	1692	1595.468	59.775	59.429
2020-06-20	034820	data6	1547	1467.561	59.864	53.401
2020-06-21	034821	data6	1102	1212.781	59.884	35.806
2020-06-22	034822	data6	589	597.022	59.941	35.242
2020-06-23	034823	data6	497	538.487	59.944	10.727
2020-06-24	034824	data6	1545	1485.293	59.864	54.575
2020-06-25	034825	data6	1912	2035.331	59.813	59.412
2020-06-26	034826	data6	1663	1578.823	59.843	59.594
2020-06-27	034827	data6	1667	1578.030	59.851	59.738
2020-06-28	034828	data6	1835	1852.590	59.828	59.001
2020-06-29	034829	data6	2043	2034.906	59.787	52.865
2020-06-30	034830	data6	1517	1764.610	59.852	49.154

Table A.19: Details of the Run 3 data from July 2020. A disk-space issue on the DAQ machine prevented normal data-taking in the first half of July. We stopped taking data on July 14, 2020 to diagnose and address this issue before resuming normal operations on July 20, 2020; the E39 rate on July 20 is atypically low (and was not a result of the DAQ restart), so this day is also excluded from our analyses. Data from July 25, 2020 are ignored as a result of the preprocessed file size.

Date	UTC* Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2020-07-01	034831	data6	0.0022	1763.794	59.831	59.286
2020-07-02	034832	data6	0.0022	1675.789	59.814	58.269
2020-07-03	034833	data6	0.0022	1586.934	59.836	59.432
2020-07-04	034834	data6	0.0022	1662.656	59.851	59.682
2020-07-05	034835	data6	0.0022	1626.501	59.855	59.496
2020-07-06	034836	data6	0.0022	1583.394	59.842	59.537
2020-07-07	034837	data6	0.0022	1204.832	59.856	39.213
2020-07-08	034838	data6	0.0022	1641.687	59.838	58.190
2020-07-09	034839	data6	0.0022	1800.704	59.841	59.457
2020-07-10	034840	data6	0.0022	1832.545	59.828	59.471
2020-07-11	034841	data6	0.0022	1746.713	59.819	59.504
2020-07-12	034842	data6	0.0022	1961.287	59.813	57.985
2020-07-13	034843	data6	0.0022	1875.957	59.752	56.386
2020-07-14	034844	data6	0.0022	1490.340	59.852	58.958
2020-07-20	214909	data6	416	355.919	48.164	0.000
2020-07-21	214910	data6	407	385.030	59.962	0.000
2020-07-22	214911	data6	402	379.990	59.957	0.000
2020-07-23	214912	data6	397	373.239	59.948	0.000
2020-07-24	214913	data6	391	367.520	59.963	0.000
2020-07-25	214914	data6	27	362.790	59.959	0.029
2020-07-26	214915	data6	384	359.821	59.963	0.000
2020-07-27	214916	data6	382	357.325	59.963	0.000
2020-07-28	214917	data6	380	354.696	59.957	0.000
2020-07-29	214918	data6	380	353.297	59.956	0.000
2020-07-30	214919	data6	381	354.050	59.958	0.000
2020-07-31	214920	data6	377	348.771	59.960	0.000

Table A.20: Details of the Run 3 data from August 2020.

Date	UTC* Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2020-08-01	214921	data6	367	342.155	59.952	0.000
2020-08-02	214922	data6	363	336.791	59.967	0.000
2020-08-03	214923	data6	358	331.227	59.962	0.000
2020-08-04	214924	data6	360	329.427	59.962	0.000
2020-08-05	214925	data6	378	350.052	59.960	0.000
2020-08-06	214926	data6	390	362.036	59.959	0.283
2020-08-07	214927	data6	401	368.534	59.902	0.129
2020-08-08	214928	data6	399	372.331	59.950	0.000
2020-08-09	214929	data6	400	373.620	59.974	0.000
2020-08-10	214930	data6	403	377.935	59.958	0.111
2020-08-11	214931	data6	399	370.031	59.937	0.077
2020-08-12	214932	data6	375	349.584	59.957	0.055
2020-08-13	214933	data6	363	335.416	59.966	0.832
2020-08-14	214934	data6	359	327.622	59.962	0.309
2020-08-15	214935	data6	357	324.652	59.965	0.300
2020-08-16	214936	data6	350	322.815	59.959	0.983
2020-08-17	214937	data6	356	321.258	59.964	0.165
2020-08-18	214938	data6	356	322.000	59.967	0.297
2020-08-19	214939	data6	436	422.968	59.956	45.389
2020-08-20	214940	data6	669	676.114	59.937	54.985
2020-08-21	214941	data6	618	644.857	59.939	41.464
2020-08-22	214942	data6	825	826.657	59.924	58.828
2020-08-23	214943	data6	888	889.527	59.920	59.693
2020-08-24	214944	data6	818	840.513	59.918	43.349
2020-08-25	214945	data6	1741	2353.917	59.801	57.517
2020-08-26	214946	data6	1112	1104.037	59.899	59.644
2020-08-27	214947	data6	1110	1099.596	59.890	59.607
2020-08-28	214948	data6	1145	1124.792	59.901	57.889
2020-08-29	214949	data6	1181	1160.624	59.770	59.170
2020-08-30	214950	data6	943	986.675	59.896	47.798
2020-08-31	214951	data6	834	876.635	59.919	34.545

Table A.21: Details of the Run 3 data from September 2020. Data from September 2, 2020 is excluded due to the exceptionally high trigger rate observed.

Date	UTC* Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2020-09-01	214952	data6	811	906.837	59.879	47.803
2020-09-02	214953	data6	3148	8649.874	57.944	25.684
2020-09-03	214954	data6	953	979.394	59.907	56.112
2020-09-04	214955	data6	1063	1057.522	59.894	59.634
2020-09-05	214956	data6	1063	1056.888	59.899	59.649
2020-09-06	214957	data6	1076	1069.437	59.893	59.156
2020-09-07	214958	data6	953	997.352	59.900	45.251
2020-09-08	214959	data6	996	988.303	59.897	59.185
2020-09-09	215000	data6	1139	1133.836	59.889	58.933
2020-09-10	215001	data6	1183	1162.686	59.841	59.577
2020-09-11	215002	data6	829	902.219	59.921	37.969
2020-09-12	215003	data6	991	1048.505	59.903	46.700
2020-09-13	215004	data6	487	475.670	59.948	29.505
2020-09-14	215005	data6	476	482.334	59.957	26.856
2020-09-15	215006	data6	847	887.634	59.908	56.037
2020-09-16	215007	data6	1076	1072.349	59.905	59.597
2020-09-17	215008	data6	1016	1019.823	59.902	53.658
2020-09-18	215009	data6	1060	1052.194	59.888	59.775
2020-09-19	215010	data6	1074	1064.850	59.893	59.436
2020-09-20	215011	data6	1110	1100.862	59.905	59.681
2020-09-21	215012	data6	864	913.647	59.919	38.160
2020-09-22	215013	data6	964	977.638	59.897	55.098
2020-09-23	215014	data6	1081	1070.699	59.620	59.169
2020-09-24	215015	data6	1098	1091.575	59.896	59.012
2020-09-25	215016	data6	1093	1086.002	59.898	58.604
2020-09-26	215017	data6	1117	1108.182	59.861	59.540
2020-09-27	215018	data6	1117	1105.856	59.876	58.999
2020-09-28	215019	data6	909	970.231	59.907	43.813
2020-09-29	215020	data6	924	955.380	59.912	54.882
2020-09-30	215021	data6	1059	1055.781	59.894	56.722

Table A.22: Details of the Run 3 data from October 2020.

Date	UTC* Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2020-10-01	215022	data6	1022	1021.030	59.888	57.638
2020-10-02	215023	data6	1067	1063.099	59.904	59.682
2020-10-03	215024	data6	1123	1114.364	59.892	59.319
2020-10-04	215025	data6	1117	1106.322	59.879	58.789
2020-10-05	215026	data6	964	986.202	59.909	44.355
2020-10-06	215027	data6	966	980.461	59.909	54.924
2020-10-07	215028	data6	996	997.048	59.898	57.264
2020-10-08	215029	data6	1046	1042.760	59.904	59.744
2020-10-09	215030	data6	1104	1096.257	59.900	58.940
2020-10-10	215031	data6	988	1041.037	59.898	46.877
2020-10-11	215032	data6	444	426.010	59.958	6.441
2020-10-12	215033	data6	435	418.104	59.957	24.176
2020-10-13	215034	data6	664	707.948	59.925	51.463
2020-10-14	215035	data6	900	906.514	59.907	59.415
2020-10-15	215036	data6	974	966.362	59.906	56.487
2020-10-16	215037	data6	925	940.227	59.877	53.385
2020-10-17	215038	data6	1010	1011.898	59.903	59.553
2020-10-18	215039	data6	904	914.865	59.920	51.934
2020-10-19	215040	data6	895	944.345	59.913	48.122
2020-10-20	215041	data6	679	706.323	59.901	58.911
2020-10-21	215042	data6	1022	1012.160	59.887	59.137
2020-10-22	215043	data6	1063	1060.610	59.899	59.554
2020-10-23	215044	data6	1064	1061.216	59.895	59.779
2020-10-24	215045	data6	1094	1092.211	59.917	59.749
2020-10-25	215046	data6	1091	1088.421	59.889	59.735
2020-10-26	215047	data6	589	625.431	59.943	17.443
2020-10-27	215048	data6	910	963.557	59.909	48.457
2020-10-28	215049	data6	1174	1198.455	59.884	57.741
2020-10-29	215050	data6	1341	1310.426	59.875	59.762
2020-10-30	215051	data6	1236	1222.191	59.889	59.650
2020-10-31	215052	data6	1258	1237.537	59.873	59.295

Table A.23: Details of the Run 3 data from November 2020.

Date	UTC* Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2020-11-01	215053	data6	1169	1162.678	59.891	56.056
2020-11-02	215054	data6	826	866.572	59.926	40.047
2020-11-03	215055	data6	731	758.118	59.930	48.263
2020-11-04	215056	data6	852	868.657	59.923	51.643
2020-11-05	215057	data6	862	864.385	59.909	53.479
2020-11-06	215058	data6	893	897.640	59.897	56.884
2020-11-07	215059	data6	866	883.906	59.911	45.857
2020-11-08	215100	data6	668	697.671	59.929	35.486
2020-11-09	215101	data6	447	426.906	59.764	6.062
2020-11-10	215102	data6	760	801.812	59.918	52.635
2020-11-11	215103	data6	971	972.495	59.909	59.151
2020-11-12	215104	data6	975	972.848	59.906	59.639
2020-11-13	215105	data6	959	958.530	59.906	59.703
2020-11-14	215106	data6	979	976.188	59.909	56.730
2020-11-15	215107	data6	957	954.700	59.899	59.269
2020-11-16	215108	data6	872	889.963	59.916	48.694
2020-11-17	215109	data6	791	816.400	59.928	54.050
2020-11-18	215110	data6	938	939.344	59.911	59.584
2020-11-19	215111	data6	964	957.293	59.920	57.637
2020-11-20	215112	data6	985	982.704	59.873	58.458
2020-11-21	215113	data6	1001	1003.680	59.901	59.324
2020-11-22	215114	data6	913	920.909	59.914	47.414
2020-11-23	215115	data6	969	975.344	59.907	58.487
2020-11-24	215116	data6	871	463.270	59.851	0.449
2020-11-25	215117	data6	934	1033.941	59.910	50.174
2020-11-26	215118	data6	1018	1033.239	59.859	58.351
2020-11-27	215119	data6	855	886.853	59.924	43.360
2020-11-28	215120	data6	438	431.666	59.961	0.000
2020-11-29	215121	data6	406	387.482	59.962	0.000
2020-11-30	215122	data6	303	374.987	44.262	0.000

Table A.24: Details of the Run 3 data from December 2020. The data in the early part of this month are excluded due to an E39 outage at the SNS.

Date	UTC* Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2020-12-01	215123	data6	0.0087	368.106	0.000	0.000
2020-12-02	215124	data6	0.0087	367.226	0.000	0.000
2020-12-03	215125	data6	0.0087	361.184	0.000	0.000
2020-12-04	215126	data6	0.0087	358.733	0.000	0.000
2020-12-05	215127	data6	0.0087	355.732	0.000	0.000
2020-12-06	215128	data6	0.0087	354.456	0.000	0.000
2020-12-07	215129	data6	0.0087	351.335	0.000	0.000
2020-12-08	215130	data6	0.000252	349.796	17.595	0.000
2020-12-09	215131	data6	0.0087	349.122	46.442	0.000
2020-12-10	215132	data6	12	348.021	51.387	0.000
2020-12-11	215133	data6	370	346.257	59.951	0.000
2020-12-12	215134	data6	370	346.324	59.966	0.000
2020-12-13	215135	data6	370	346.661	59.957	0.000
2020-12-14	215136	data6	369	344.896	59.952	0.000
2020-12-15	215137	data6	369	344.936	59.949	0.000
2020-12-16	215138	data6	369	345.233	59.955	0.000
2020-12-17	215139	data6	368	343.694	59.948	0.000
2020-12-18	215140	data6	367	342.845	59.952	0.000
2020-12-19	215141	data6	367	342.842	59.943	0.000
2020-12-20	215142	data6	367	342.898	59.946	0.000
2020-12-21	215143	data6	366	342.035	59.944	0.000
2020-12-22	215144	data6	366	341.523	59.949	0.000
2020-12-23	215145	data6	366	342.499	59.944	0.000
2020-12-24	215146	data6	367	342.434	59.940	0.000
2020-12-25	215147	data6	366	341.816	59.930	0.000
2020-12-26	215148	data6	366	341.888	59.936	0.000
2020-12-27	215149	data6	367	342.423	59.940	0.000
2020-12-28	215150	data6	367	343.191	59.949	0.000
2020-12-29	215151	data6	367	343.022	59.947	0.000
2020-12-30	215152	data6	367	342.706	59.965	0.000
2020-12-31	215153	data6	366	342.017	59.944	0.000

Table A.25: Details of the Run 3 data from January 2021. We stopped normal data taking on January 27, 2021 to prepare for new ^{60}Co measurements and various beam-off characterization studies over the coming weeks.

Date	UTC* Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2021-01-01	215154	data6	366	342.158	59.940	0.000
2021-01-02	215155	data6	366	341.431	59.930	0.000
2021-01-03	215156	data6	365	340.797	59.942	0.000
2021-01-04	215157	data6	364	339.800	59.920	0.000
2021-01-05	215158	data6	364	339.624	59.921	0.000
2021-01-06	215159	data6	364	339.384	59.923	0.000
2021-01-07	215200	data6	364	339.649	59.967	0.000
2021-01-08	215201	data6	364	339.668	59.927	0.000
2021-01-09	215202	data6	364	339.494	59.922	0.000
2021-01-10	215203	data6	364	339.925	59.928	0.000
2021-01-11	215204	data6	364	340.112	59.932	0.000
2021-01-12	215205	data6	364	339.635	59.932	0.000
2021-01-13	215206	data6	365	340.566	59.940	0.000
2021-01-14	215207	data6	365	340.826	59.929	0.000
2021-01-15	215208	data6	365	340.850	59.935	0.000
2021-01-16	215209	data6	365	340.541	59.931	0.000
2021-01-17	215210	data6	364	339.929	59.933	0.000
2021-01-18	215211	data6	364	339.461	59.919	0.000
2021-01-19	215212	data6	364	339.135	59.912	0.000
2021-01-20	215213	data6	364	339.277	59.926	0.000
2021-01-21	215214	data6	364	339.841	59.938	0.000
2021-01-22	215215	data6	364	339.564	59.925	0.000
2021-01-23	215216	data6	363	339.171	59.914	0.000
2021-01-24	215217	data6	364	339.556	59.928	0.000
2021-01-25	215218	data6	364	339.785	59.921	0.000
2021-01-26	215219	data6	364	338.870	59.830	0.000

Table A.26: Details of the Run 3 data from February 2021. We resumed normal operations following calibration and characterization studies in late January and early February 2021.

Date	UTC* Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2021-02-11	015251	data2	539	537.079	59.954	41.351
2021-02-12	015252	data2	761	766.093	59.931	50.735
2021-02-13	015253	data2	902	903.181	59.915	59.263
2021-02-14	015254	data2	896	893.447	59.914	56.980
2021-02-15	015255	data2	890	886.000	59.889	54.091
2021-02-16	015256	data2	570	592.603	59.938	27.609
2021-02-17	015257	data2	756	753.989	59.937	53.170
2021-02-18	015258	data2	836	836.633	59.915	57.678
2021-02-19	015259	data2	833	833.816	59.911	56.645
2021-02-20	015300	data2	845	846.163	59.927	59.415
2021-02-21	015301	data2	822	824.183	59.923	58.715
2021-02-22	015302	data2	847	849.724	59.919	57.619
2021-02-23	180706	data2	575	674.566	58.887	37.328
2021-02-24	180707	data2	856	852.779	59.612	57.525
2021-02-25	180708	data2	907	907.361	59.913	59.796
2021-02-26	180709	data2	845	854.188	59.924	52.375
2021-02-27	180710	data2	885	885.289	59.932	56.585
2021-02-28	180711	data2	308	870.941	59.916	52.014

Table A.27: Details of the Run 3 data from March 2021. We exclude the data from March 2, 2021 due to preprocessing issues. We stopped taking data on March 30, 2021 to prepare for beam-on characterization studies and calibrations using a ^{60}Co source.

Date	UTC* Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2021-03-01	182832	data7	879	876.860	59.598	51.049
2021-03-02	182833	data7	42	467.798	59.900	0.651
2021-03-03	020514	data7	844	845.161	59.922	58.543
2021-03-04	020515	data7	913	913.517	59.929	58.590
2021-03-05	020516	data7	892	880.482	59.910	52.546
2021-03-06	020517	data7	840	850.421	59.915	55.179
2021-03-07	020518	data7	735	775.314	59.927	40.158
2021-03-08	020519	data7	435	414.798	59.957	11.568
2021-03-09	020520	data7	426	403.729	59.946	14.066
2021-03-10	020521	data7	651	689.493	59.906	51.133
2021-03-11	020522	data7	855	857.298	59.927	59.432
2021-03-12	020523	data7	832	834.130	59.922	57.941
2021-03-13	020524	data7	836	841.254	59.916	59.084
2021-03-14	020525	data7	863	868.477	59.911	59.415
2021-03-15	020526	data7	1035	1119.737	59.891	59.425
2021-03-16	020527	data7	723	751.999	59.935	54.905
2021-03-17	020528	data7	830	877.935	59.933	55.337
2021-03-18	020529	data7	2756	2577.620	59.742	49.680
2021-03-19	020530	data7	1519	1797.209	59.747	53.216
2021-03-20	020531	data7	790	802.407	59.931	52.706
2021-03-21	020532	data7	824	828.037	59.747	59.527
2021-03-22	020533	data7	787	791.060	59.932	53.232
2021-03-23	020534	data7	660	709.247	59.933	33.706
2021-03-24	020535	data7	1461	1526.456	59.847	56.592
2021-03-25	020536	data7	2973	2584.846	59.680	56.998
2021-03-26	020537	data7	2676	2531.054	59.772	52.690
2021-03-27	020538	data7	1170	1243.373	59.878	53.378
2021-03-28	020539	data7	1907	2093.866	59.811	43.569
2021-03-29	020540	data7	1323	1378.807	59.868	56.882
2021-03-30	020541	data7	1403	1506.463	59.549	45.705

Table A.28: Details of the Run 3 data from April 2021. April 1, 2021 was a calibration day, and is not included in our beam analyses. We do not exclude data with extraordinarily high rates from this month due to the consistent excess; these rates were deemed physical and are included in our analyses.

Date	UTC* Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2021-04-02	003849	data7	1239	1299.330	59.878	58.646
2021-04-03	003850	data7	1824	1786.862	59.837	59.575
2021-04-04	003851	data7	2041	2074.509	59.801	42.749
2021-04-05	003852	data7	541	593.431	59.940	19.014
2021-04-06	003853	data7	664	757.301	59.922	36.332
2021-04-07	003854	data7	4137	3288.440	59.705	58.128
2021-04-08	003855	data7	5789	3823.046	59.640	57.897
2021-04-09	003856	data7	2304	2808.269	59.726	41.664
2021-04-10	003857	data7	5317	4055.163	59.627	58.423
2021-04-11	003858	data7	6580	4366.920	59.610	55.029
2021-04-12	003859	data7	6235	4493.660	59.555	58.326
2021-04-13	003900	data7	2853	2883.985	59.754	42.982
2021-04-14	003901	data7	3634	3179.152	59.700	57.050
2021-04-15	003902	data7	5158	3632.890	59.651	53.181
2021-04-16	003903	data7	5439	3956.598	59.485	58.579
2021-04-17	003904	data7	6525	4570.070	59.545	58.503
2021-04-18	003905	data7	6289	4469.593	59.575	57.655
2021-04-19	003906	data7	7258	4454.453	59.631	56.697
2021-04-20	003907	data7	1426	1696.431	59.851	45.376
2021-04-21	003908	data7	876	870.290	59.916	59.393
2021-04-22	003909	data7	1859	1821.667	59.230	59.119
2021-04-23	003910	data7	4210	3353.118	59.662	59.551
2021-04-24	003911	data7	6349	4618.898	59.544	58.224
2021-04-25	003912	data7	5667	4033.180	59.596	52.391
2021-04-26	003913	data7	3725	3087.209	59.714	54.321
2021-04-27	003914	data7	2997	2674.349	59.748	39.002
2021-04-28	003915	data7	4163	3217.126	59.225	58.940
2021-04-29	003916	data7	5071	3835.003	59.657	59.466
2021-04-30	003917	data7	5073	3846.813	59.626	59.164

Table A.29: Details of the Run 3 data from May 2021. We choose May 31, 2021 as the final day of Run 3 arbitrarily; beam-off data taking continued through June 2021, but we used that time primarily to explore options for moving MARS, perform additional calibrations, and confirm that the detector behavior was stable and healthy despite the exceptionally high rates during the end of Run 3.

Date	UTC* Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2021-05-01	003918	data7	4706	3584.304	59.681	59.357
2021-05-02	003919	data7	3207	2836.779	59.767	33.158
2021-05-03	003920	data7	438	421.742	59.957	0.190
2021-05-04	003921	data7	422	394.864	59.959	0.096
2021-05-05	003922	data7	405	386.232	59.964	0.000
2021-05-06	003923	data7	397	377.075	59.957	0.000
2021-05-07	003924	data7	391	370.310	59.942	0.000
2021-05-08	003925	data7	391	365.765	59.961	0.563
2021-05-09	003926	data7	384	362.517	59.965	0.000
2021-05-10	003927	data7	383	358.936	59.961	0.334
2021-05-11	003928	data7	385	356.521	59.964	0.114
2021-05-12	003929	data7	384	356.940	59.970	0.412
2021-05-13	003930	data7	379	351.496	59.970	0.444
2021-05-14	003931	data7	374	349.467	59.976	0.014
2021-05-15	003932	data7	372	348.220	59.968	0.000
2021-05-16	003933	data7	373	347.188	59.954	0.285
2021-05-17	003934	data7	435	420.899	59.956	31.622
2021-05-18	003935	data7	644	637.754	59.940	56.151
2021-05-19	003936	data7	616	626.005	59.945	45.748
2021-05-20	003937	data7	731	731.627	59.935	59.614
2021-05-21	003938	data7	722	724.007	59.938	57.369
2021-05-22	003939	data7	745	746.987	59.926	59.423
2021-05-23	003940	data7	787	790.345	59.926	57.885
2021-05-24	003941	data7	825	827.955	59.765	59.529
2021-05-25	003942	data7	844	845.533	59.923	57.760
2021-05-26	003943	data7	1156	1416.010	59.706	51.462
2021-05-27	003944	data7	1673	1764.724	59.828	56.428
2021-05-28	224847	data7	3939	2807.509	59.683	54.224
2021-05-29	224848	data7	2834	2446.395	59.759	50.206
2021-05-30	224849	data7	1495	1558.274	59.843	42.177
2021-05-31	224850	data7	336	428.014	59.966	0.135

A.2 Production Run 4

Production Run 4 took place at the so-called “Rad monitor location”, at the end of Neutrino Alley near a ORNL-operated monitoring system for the high-radiation area behind the wall of the alley.

Table A.30: Details of the Run 4 data taken during July 2021. Calibrations using a ^{60}Co source were taken on July 1, which desynchronized the detector health monitoring on the subsequent days. We exclude July 2 and 3 from our analysis out of an abundance of caution in monitoring MARS at its new location.

Date	UTC Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2021-07-02	141913	data7	328	—	—	—
2021-07-03	141914	data7	335	—	—	—
2021-07-04	141915	data7	334	303.242	59.967	0.995
2021-07-05	141916	data7	336	304.621	59.967	0.994
2021-07-06	141917	data7	341	304.639	59.974	0.222
2021-07-07	141918	data7	338	304.616	59.967	0.000
2021-07-08	220628	data7	340	989.223	59.190	0.142
2021-07-09	220629	data7	341	304.929	59.967	0.337
2021-07-10	220630	data7	342	304.681	59.971	0.143
2021-07-11	220631	data7	352	319.614	59.968	13.869
2021-07-12	220632	data7	579	575.567	59.951	57.684
2021-07-13	220633	data7	641	637.504	59.943	56.819
2021-07-14	220634	data7	678	676.888	59.942	53.834
2021-07-15	220635	data7	744	743.322	59.914	59.092
2021-07-16	220636	data7	848	854.556	59.917	57.110
2021-07-17	220637	data7	821	822.961	59.927	57.831
2021-07-18	220638	data7	871	868.551	59.907	58.734
2021-07-19	220639	data7	651	694.826	59.935	31.387
2021-07-20	220640	data7	775	793.919	59.908	55.407
2021-07-21	220641	data7	867	865.934	59.901	57.674
2021-07-22	220642	data7	1052	1174.497	59.892	59.322
2021-07-23	200326	data7	2334	2107.392	58.971	57.888
2021-07-24	200327	data7	4215	3263.449	59.691	50.976
2021-07-25	200328	data7	896	904.068	59.864	59.320
2021-07-26	231344	data7	1912	2922.697	59.712	51.137
2021-07-27	231345	data7	2292	2327.087	59.749	44.030
2021-07-28	231346	data7	3312	3327.656	59.685	58.108
2021-07-29	212502	data7	2840	6056.752	58.696	57.572
2021-07-30	212503	data7	3770	3254.867	59.691	58.970

Table A.31: Details of the Run 4 data taken during August 2021. Run 4 concluded on August 31, and preparations to move MARS to a new location began on September 1, 2021.

Date	UTC Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2021-08-02	132320	data7	403	376.491	59.953	0.267
2021-08-03	132321	data7	494	556.871	59.936	29.904
2021-08-04	132322	data7	1806	1718.540	59.831	59.718
2021-08-05	132323	data7	3423	2785.562	59.757	57.347
2021-08-06	132324	data7	4496	3398.513	59.622	50.401
2021-08-07	132325	data7	5049	3813.249	59.624	59.405
2021-08-08	132326	data7	6001	4321.346	59.594	59.070
2021-08-09	132327	data7	6991	4661.400	59.593	53.476
2021-08-10	132328	data7	3932	3252.111	59.432	44.113
2021-08-11	132329	data7	7504	4969.011	56.000	55.902
2021-08-12	132330	data7	6461	4141.934	59.655	53.244
2021-08-13	132331	data7	7455	5009.126	59.418	59.313
2021-08-14	132332	data7	6635	4386.501	59.582	49.774
2021-08-15	132333	data7	7826	5089.011	59.565	59.460
2021-08-16	132334	data7	7938	4884.476	59.539	43.980
2021-08-17	132335	data7	1872	2222.280	59.772	41.425
2021-08-18	132336	data7	6560	4240.852	59.536	47.254
2021-08-19	132337	data7	7950	4749.649	59.540	58.810
2021-08-20	132338	data7	9762	5865.302	59.486	58.231
2021-08-21	132339	data7	10868	5996.900	59.446	57.361
2021-08-22	132340	data7	10000	5321.835	59.481	49.013
2021-08-23	132341	data7	8615	4940.874	59.508	46.970
2021-08-24	132342	data7	947	1162.323	59.897	44.949
2021-08-25	132343	data7	5248	3671.412	59.655	55.578
2021-08-26	132344	data7	8902	5481.098	59.487	58.654
2021-08-27	132345	data7	10745	6421.161	59.395	59.286
2021-08-28	132346	data7	9240	5352.831	59.485	46.034
2021-08-29	132347	data7	500	590.143	59.944	0.230
2021-08-30	132348	data7	405	375.451	59.952	0.028
2021-08-31	132349	data7	140	3602.564	59.650	48.853

A.3 Production Run 5

For Production Run 5, MARS was located in the alcove at the end of Neutrino Alley, next to the LAr detector. Run 5 is ongoing; MARS is presently still in the alcove. For the purposes of this work, we truncate the dataset at the end of April 2022.

Table A.32: Details of the Run 5 analysis taken during September 2021. The data from September 23 through 27 are excluded due to high trigger rates and inconsistent file sizes.

Date	UTC Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2021-09-03	173330	data7	2337	1648.538	59.84	46.15
2021-09-04	173331	data7	1541	2136.14	59.777	57.332
2021-09-05	173332	data7	2399	2027.093	59.804	52.237
2021-09-06	173333	data7	2193	1749.417	59.827	42.06
2021-09-07	173334	data7	1648	1859.053	59.768	50.215
2021-09-08	173335	data7	1676	2487.94	59.761	59.299
2021-09-09	173336	data7	2577	2159.811	59.803	58.017
2021-09-10	152630	data7	2464	1819.107	59.82	57.374
2021-09-11	152631	data7	1961	1605.59	59.82	54.417
2021-09-12	152632	data7	1685	1506.472	59.836	53.256
2021-09-13	152633	data7	1577	1692.768	59.843	48.848
2021-09-14	152634	data7	1773	1174.221	59.883	48.712
2021-09-15	152635	data7	986	1909.652	59.801	58.612
2021-09-16	152636	data7	2111	1938.28	59.814	59.7
2021-09-17	152637	data7	2095	1779.968	59.826	59.584
2021-09-18	152638	data7	1925	1828.925	58.794	58.576
2021-09-19	152639	data7	2004	1843.127	59.81	54.996
2021-09-20	152640	data7	1961	1773.944	59.633	51.266
2021-09-21	152641	data7	1458	1770.241	59.817	49.76
2021-09-22	152642	data7	2380	2108.556	59.819	58.446
2021-09-23	152643	data7	1176	11886.309	791.327	786.893
2021-09-24	152644	data7	1	5726.514	1544.296	1538.501
2021-09-25	152651	data7	1	902.005	226.575	215.536
2021-09-26	152652	data7	56	249.748	62.594	17.678
2021-09-27	152653	data7	412	390.504	59.957	0.345
2021-09-28	152654	data7	861	1195.977	59.891	47.893
2021-09-29	152655	data7	1837	1945.818	59.816	49.917
2021-09-30	152656	data7	1096	1278.83	59.88	40.242

Table A.33: Details of the Run 5 analysis taken during October 2021. Data from October 2 are excluded due to the low pre-processed file size. Data from October 9-12 and 25-27 are excluded due to an SNS timing outage.

Date	UTC Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2021-10-01	152657	data7	929	1154.99	59.882	40.086
2021-10-02	152658	data7	29	6969.839	2575.439	2568.256
2021-10-03	152659	data7	340	2078.057	59.784	59.67
2021-10-04	152700	data7	2126	1924.426	59.793	49.449
2021-10-05	152701	data7	1140	1305.544	59.871	51.452
2021-10-06	152702	data7	1839	1732.659	59.843	58.498
2021-10-07	152703	data7	1821	1704.587	59.835	57.534
2021-10-08	152704	data7	385	1714.594	59.855	59.425
2021-10-09	152709	data7	1	—	—	—
2021-10-10	152713	data7	1	—	—	—
2021-10-11	152718	data7	1	—	—	—
2021-10-12	152722	data7	1	—	—	—
2021-10-18	124917	data7	1528	1473.031	59.827	58.067
2021-10-19	124918	data7	1474	1436.638	59.873	53.786
2021-10-20	124919	data7	575	644.721	59.94	21.613
2021-10-21	124920	data7	400	376.898	59.955	0.313
2021-10-22	124921	data7	378	355.786	59.97	0
2021-10-23	124922	data7	373	348.99	59.968	0
2021-10-24	124923	data7	366	344.486	59.963	0
2021-10-25	124924	data7	1	341.843	—	—
2021-10-26	124925	data7	1	341.738	—	—
2021-10-27	124926	data7	1	339.114	51.347	0
2021-10-28	124927	data7	363	337.996	59.906	0
2021-10-29	124928	data7	362	336.206	59.77	0
2021-10-30	124929	data7	361	335.447	59.969	0
2021-10-31	124930	data7	360	334.079	59.97	0

Table A.34: Details of the Run 5 analysis taken during November 2021. Data from November 6, 9, and 29 are excluded since E39 is inconsistent with the known 60 Hz signal. Data was not taken on November 8 as we investigated the E39 inconsistencies. These issues were identified to be user-error in the DAQ operations; multiple terminals had open communications to the VME crate, which led to read/write errors in the digitizer.

Date	UTC Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2021-11-01	124931	data7	359	333.38	59.964	0
2021-11-02	124932	data7	358	332.378	59.968	0
2021-11-03	124933	data7	357	331.233	59.963	0
2021-11-04	124934	data7	357	330.866	59.963	0
2021-11-05	124935	data7	357	330.902	59.963	0
2021-11-06	124936	data7	32	505.705	4375.234	0
2021-11-07	124937	data7	356	330.049	59.963	0
2021-11-09	154102	data7	357	333.53	222.687	0.004
2021-11-10	154103	data7	357	331.404	59.964	0
2021-11-11	154104	data7	362	335.346	59.516	0
2021-11-12	154105	data7	357	331.088	59.964	0
2021-11-13	154106	data7	355	328.3	59.963	0
2021-11-14	154107	data7	353	327.031	59.967	0
2021-11-15	154108	data7	352	325.403	59.963	0
2021-11-16	154109	data7	353	326.023	59.969	0
2021-11-17	154110	data7	348	321.088	59.961	0
2021-11-18	154111	data7	364	337.662	59.905	0
2021-11-19	154112	data7	378	354.058	59.963	0
2021-11-20	154113	data7	374	350.447	59.959	0
2021-11-21	154114	data7	367	342.917	59.953	0
2021-11-22	154115	data7	358	333.114	59.963	0
2021-11-23	154116	data7	354	327.559	59.963	0
2021-11-24	154117	data7	352	325.409	59.967	0
2021-11-25	154118	data7	352	325.097	59.962	0
2021-11-26	154119	data7	351	324.312	59.962	0
2021-11-27	154120	data7	350	323.657	59.963	0
2021-11-28	154121	data7	341	323.615	59.962	0
2021-11-29	154122	data7	20	323.335	48	0
2021-11-30	154123	data7	350	323.996	59.966	0

Table A.35: Details of the Run 5 analysis taken during December 2021. Data were not taken from December 17 - 19 due to a scheduled power outage for electrical work unrelated to our operations.

Date	UTC Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2021-12-01	154124	data7	351	324.926	59.962	0
2021-12-02	154125	data7	352	324.79	59.97	0
2021-12-03	154126	data7	351	323.729	59.962	0
2021-12-04	154127	data7	350	322.769	59.961	0
2021-12-05	154128	data7	349	321.396	59.962	0
2021-12-06	154129	data7	348	320.512	59.965	0
2021-12-07	154130	data7	347	319.858	59.961	0
2021-12-08	154131	data7	346	318.978	59.96	0
2021-12-09	154132	data7	346	318.595	59.96	0
2021-12-10	154133	data7	346	318.625	59.969	0
2021-12-11	154134	data7	345	317.737	59.968	0
2021-12-12	154135	data7	344	316.545	59.963	0
2021-12-13	154136	data7	344	316.545	59.959	0
2021-12-14	154137	data7	344	316.076	59.963	0
2021-12-15	154138	data7	344	315.981	59.957	0
2021-12-16	154139	data7	344	315.826	59.958	0
2021-12-20	194421	data7	324	289.167	59.964	0.444
2021-12-21	194422	data7	335	298.168	59.967	0.188
2021-12-22	194423	data7	340	304.819	59.967	0.378
2021-12-23	194424	data7	337	308.45	59.939	0
2021-12-24	194425	data7	340	311.126	59.949	0
2021-12-25	194426	data7	340	311.794	59.956	0
2021-12-26	194427	data7	341	312.699	59.956	0.181
2021-12-27	194428	data7	347	312.686	59.955	0.109
2021-12-28	194429	data7	342	313.652	59.96	0
2021-12-29	194430	data7	348	313.711	59.959	0.328
2021-12-30	194431	data7	371	342.166	59.955	39.655
2021-12-31	194432	data7	393	376.744	59.935	41.649

Table A.36: Details of the Run 5 analysis taken during January 2022.

Date	UTC Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2022-01-01	194433	data7	544	571.652	59.942	39.955
2022-01-02	194434	data7	409	570.906	59.95	32.612
2022-01-03	194435	data7	479	468.232	59.945	48.52
2022-01-04	194436	data7	1031	664.351	59.907	57.221
2022-01-05	194437	data7	1070	1222.002	59.894	58.187
2022-01-06	205839	data7	759	1204.144	59.34	58.449
2022-01-07	164653	data7	825	1013.781	59.226	58.428
2022-01-08	164654	data7	843	944.343	59.924	58.404
2022-01-09	164655	data7	924	983.483	59.917	53.421
2022-01-10	164656	data7	756	1044.64	59.904	53.858
2022-01-11	164657	data7	766	757.014	59.926	39.316
2022-01-12	164658	data7	989	1087.933	59.909	59.573
2022-01-13	203130	data7	989	1142.953	59.896	59.631
2022-01-14	203131	data7	864	1100.378	59.897	55.149
2022-01-15	203132	data7	903	1007.632	59.91	56.739
2022-01-16	203133	data7	779	1096.283	59.908	57.593
2022-01-17	143626	data7	1053	1244.177	59.897	47.201
2022-01-18	143627	data7	792	941.276	59.917	45.862
2022-01-19	143628	data7	849	1075.231	59.9	49.96
2022-01-20	143629	data7	1036	1128.174	59.916	59.25
2022-01-21	143630	data7	951	1265.105	59.891	53.981
2022-01-22	143631	data7	975	1018.749	59.915	57.705
2022-01-23	143632	data7	477	928.641	59.908	38.456
2022-01-24	143633	data7	402	503.027	59.913	27.981
2022-01-25	143634	data7	693	419.606	59.955	0.188
2022-01-26	143635	data7	989	1144.862	59.893	58.433
2022-01-27	143636	data7	1129	1041.322	59.92	59.762
2022-01-28	143637	data7	997	1313.287	59.898	59.53
2022-01-29	143638	data7	1209	1147.518	59.901	56.984
2022-01-30	143639	data7	952	1330.037	59.872	49.491
2022-01-31	143640	data7	1076	1280.377	59.904	59.105

Table A.37: Details of the Run 5 analysis taken during February 2022.

Date	UTC Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2022-02-01	143641	data7	1005	1142.415	59.905	47.123
2022-02-02	143642	data7	1254	1524.351	59.883	58.869
2022-02-03	143643	data7	508	1072.012	59.904	43.47
2022-02-04	143644	data7	410	452.04	59.965	0.188
2022-02-05	143645	data7	396	429.519	59.962	4.634
2022-02-06	143646	data7	386	416.237	59.974	0.215
2022-02-07	143647	data7	547	433.191	59.969	21.732
2022-02-08	204145	data7	931	802.063	59.1	58.871
2022-02-09	204146	data7	1125	1134.367	59.919	55.813
2022-02-10	204147	data7	620	1161.614	59.899	58.996
2022-02-11	204148	data7	567	709.373	59.95	59.778
2022-02-12	204149	data7	972	692.744	59.94	51.599
2022-02-13	204150	data7	1001	1164.991	59.907	59.709
2022-02-14	204151	data7	520	1031.067	59.908	59.724
2022-02-15	204152	data7	680	551.608	59.961	35.583
2022-02-16	204153	data7	804	870.444	59.899	59.175
2022-02-17	204154	data7	987	964.182	59.804	48.281
2022-02-18	204155	data7	753	1008.863	59.832	57.01
2022-02-19	204156	data7	751	921.288	59.933	52.125
2022-02-20	204157	data7	578	776.151	59.94	43.161
2022-02-21	204158	data7	432	639.674	59.95	47.222
2022-02-22	204159	data7	720	463.425	59.959	12.831
2022-02-23	204200	data7	1093	940.666	59.927	59.004
2022-02-24	204201	data7	591	1169.147	59.898	58.787
2022-02-25	204202	data7	923	748.414	59.95	59.349
2022-02-26	204203	data7	965	1084.084	59.886	58.752
2022-02-27	204204	data7	899	1034.246	59.902	59.62
2022-02-28	204205	data7	795	1021.743	59.91	59.799

Table A.38: Details of the Run 5 analysis taken during March 2022. Calibrations using a ^{60}Co and ^{228}Th sources were conducted from March 30 through April 3.

Date	UTC Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2022-03-01	204206	data7	868	875.883	59.919	40.755
2022-03-02	204207	data7	214	1027.487	59.913	58.948
2022-03-03	204208	data7	585	1275.374	59.856	58.906
2022-03-04	204209	data7	928	669.078	59.949	58.715
2022-03-05	204210	data7	1466	1171.906	59.893	59.54
2022-03-06	204211	data7	1273	1690.032	59.854	59.52
2022-03-07	204212	data7	569	1300.831	59.873	58.44
2022-03-08	161408	data7	473	687.423	58.516	33.069
2022-03-09	161409	data7	1216	750.374	59.931	33.915
2022-03-10	161410	data7	206	1447.208	59.867	59.638
2022-03-13	182559	data7	430	440.49	59.963	0
2022-03-14	182600	data7	412	451.334	59.962	0.311
2022-03-15	182601	data7	405	436.71	59.964	0.15
2022-03-16	182602	data7	394	429.358	59.961	0
2022-03-17	182603	data7	388	424.357	59.957	0
2022-03-18	182604	data7	384	417.266	59.968	0
2022-03-19	182605	data7	381	412.327	59.971	0
2022-03-20	182606	data7	294	408.104	59.967	0
2022-03-21	130219	data7	377	404.422	59.715	0
2022-03-22	130220	data7	377	403.839	59.97	0
2022-03-23	130221	data7	377	404.112	59.966	0
2022-03-24	130222	data7	376	403.555	59.973	0
2022-03-25	130223	data7	376	402.929	59.969	0
2022-03-26	130224	data7	372	400.04	59.966	0
2022-03-27	130225	data7	369	395.66	59.965	0
2022-03-28	130226	data7	368	393.409	59.966	0
2022-03-29	130227	data7	368	391.033	59.963	0

Table A.39: Details of the Run 5 analysis taken during April 2022. Calibrations using a ^{60}Co and ^{228}Th sources were conducted from March 30 through April 3. Data from April 4 crashed our analysis scripts, for presently unknown reasons. We arbitrarily end the discussion of Run 5 for this work on April 30, though Run 5 is currently ongoing.

Date	UTC Time	Binary	Size (MB)	Rate (Hz)	E39 Rate (Hz)	E61 Rate (Hz)
2022-04-04	130233	data7	315	388.674	28.502	0
2022-04-05	130234	data7	367	390.168	59.946	0
2022-04-06	130235	data7	368	390.762	59.95	0
2022-04-07	130236	data7	367	391.261	59.969	0
2022-04-08	130237	data7	368	390.784	59.951	0
2022-04-09	130238	data7	368	391.816	59.951	0
2022-04-10	130239	data7	368	392.391	59.954	0
2022-04-11	130240	data7	369	392.552	59.954	0
2022-04-12	130241	data7	370	392.876	59.954	0
2022-04-13	130242	data7	371	394.853	59.955	0
2022-04-14	130243	data7	371	395.503	59.956	0
2022-04-15	130244	data7	370	394.864	59.956	0
2022-04-16	130245	data7	369	393.722	59.956	0
2022-04-17	130246	data7	369	393.865	59.966	0
2022-04-18	130247	data7	368	392.833	59.958	0
2022-04-19	130248	data7	368	391.704	59.956	0
2022-04-20	130249	data7	371	393.906	59.953	0
2022-04-21	130250	data7	374	397.9	59.957	0
2022-04-22	130251	data7	377	401.849	59.962	0
2022-04-23	130252	data7	380	405.304	59.962	0
2022-04-24	130253	data7	381	407.336	59.963	0
2022-04-25	143028	data7	382	408.229	59.96	0
2022-04-26	143029	data7	381	408.956	59.966	0
2022-04-27	143030	data7	375	404.799	59.96	0
2022-04-28	143031	data7	371	397.327	59.962	0
2022-04-29	185312	data7	369	393.299	59.909	0
2022-04-30	185313	data7	368	391.691	59.953	0

References

- ¹K. Riesselmann, “[The Standard Model of particle physics](#)”, Symmetry: Dimensions of particle physics (2015).
- ²F. Asaro, “The complex alpha spectra of the heavy elements”, PhD thesis (Radiation Laboratory, University of California, Berkley and United States Atomic Energy Commission, 1953).
- ³C. D. Ellis and W. Wooster, “The average energy of disintegration of radium E”, [Proc. R. Soc. Lond., A 117](#), 109 (1927).
- ⁴J. Chadwick, “Intensitätsverteilung im magnetischen Spectrum der β -rahlen von radium B + C”, [Verhandl. Dtsc. Phys. Ges. 16](#), 383 (1914).
- ⁵E. Fermi, “An attempt of a theory of beta radiation. 1.”, [Z. Phys. 88](#), 161–177 (1934).
- ⁶H. Bethe and R. Peierls, “The ’neutrino’”, [Nature 133](#), 532 (1934).
- ⁷F. Reines and C. L. Cowan, “Detection of the free neutrino”, [Phys. Rev. 92](#), 830–831 (1953).
- ⁸C. L. Cowan et al., “Detection of the free neutrino: a confirmation”, [Science 124](#), 103–104 (1956).
- ⁹J. N. Bahcall et al., “Solar neutrino flux”, [Astrophys. J. 137](#), 344–346 (1963).
- ¹⁰R. Davis Jr., D. S. Harmer, and K. C. Hoffman, “Search for neutrinos from the sun”, [Phys. Rev. Lett. 20](#), 1205–1209 (1968).
- ¹¹B. T. Cleveland et al., “Measurement of the solar electron neutrino flux with the Homestake chlorine detector”, [Astrophys. J. 496](#), 505–526 (1998).
- ¹²G. Danby et al., “Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos”, [Phys. Rev. Lett. 9](#), 36–44 (1962).
- ¹³K. Kodama et al., “Observation of tau neutrino interactions”, [Phys. Lett. B 504](#), 218–224 (2001).

- ¹⁴D. Decamp et al., “A Precise Determination of the Number of Families With Light Neutrinos and of the Z Boson Partial Widths”, [Phys. Lett. B 235, 399–411 \(1990\)](#).
- ¹⁵B. Pontecorvo, “Mesonium and anti-mesonium”, [Sov. Phys. JETP 6, 429 \(1957\)](#).
- ¹⁶Z. Maki, M. Nakagawa, and S. Sakata, “Remarks on the unified model of elementary particles”, [Prog. Theor. Phys. 28, 870–880 \(1962\)](#).
- ¹⁷B. Pontecorvo, “Neutrino Experiments and the Problem of Conservation of Leptonic Charge”, [Zh. Eksp. Teor. Fiz. 53, 1717–1725 \(1967\)](#).
- ¹⁸Y. Fukuda et al., “Evidence for oscillation of atmospheric neutrinos”, [Phys. Rev. Lett. 81, 1562–1567 \(1998\)](#).
- ¹⁹Q. R. Ahmad et al., “Measurement of the rate of $\nu_e + d \rightarrow p + p + e^-$ interactions produced by ^8B solar neutrinos at the Sudbury Neutrino Observatory”, [Phys. Rev. Lett. 87, 071301 \(2001\)](#).
- ²⁰M. Aker et al., “Direct neutrino-mass measurement with sub-electronvolt sensitivity”, [Nature Phys. 18, 160–166 \(2022\)](#).
- ²¹S. T. Petcov, “The Nature of Massive Neutrinos”, [Adv. High Energy Phys. 2013, 852987 \(2013\)](#).
- ²²P. A. Zyla et al., “Review of Particle Physics”, [PTEP 2020, 083C01 \(2020\)](#).
- ²³B. Abi et al., “Long-baseline neutrino oscillation physics potential of the DUNE experiment”, [Eur. Phys. J. C 80, 978 \(2020\)](#).
- ²⁴K. Abe et al., “Physics potential of a long-baseline neutrino oscillation experiment using a J-PARC neutrino beam and Hyper-Kamiokande”, [PTEP 2015, 053C02 \(2015\)](#).
- ²⁵M. Aker et al., “KATRIN: Status and Prospects for the Neutrino Mass and Beyond”, [arXiv:2203.08059 \[nucl-ex\] \(2022\)](#).
- ²⁶A. A. Esfahani et al., “The Project 8 Neutrino Mass Experiment”, in [2022 Snowmass Proceedings, arXiv:2203.07349 \[nucl-ex\] \(Mar. 2022\)](#).
- ²⁷M. J. Dolinski, A. W. P. Poon, and W. Rodejohann, “Neutrinoless Double-Beta Decay: Status and Prospects”, [Ann. Rev. Nucl. Part. Sci. 69, 219–251 \(2019\)](#).
- ²⁸J. A. Formaggio and G. P. Zeller, “From eV to EeV: Neutrino Cross Sections Across Energy Scales”, [Rev. Mod. Phys. 84, 1307–1341 \(2012\)](#).
- ²⁹S. L. Glashow, “Partial Symmetries of Weak Interactions”, [Nucl. Phys. 22, 579–588 \(1961\)](#).
- ³⁰S. Weinberg, “A Model of Leptons”, [Phys. Rev. Lett. 19, 1264–1266 \(1967\)](#).

- ³¹A. Salam, “Weak and Electromagnetic Interactions”, *Conf. Proc. C* **680519**, 367–377 (1968).
- ³²F. J. Hasert et al., “Observation of Neutrino Like Interactions Without Muon Or Electron in the Gargamelle Neutrino Experiment”, *Phys. Lett. B* **46**, 138–140 (1973).
- ³³D. Z. Freedman, “Coherent effects of a weak neutral current”, *Phys. Rev. D* **9**, 1389–1392 (1974).
- ³⁴J. Barranco, O. G. Miranda, and T. I. Rashba, “Probing new physics with coherent neutrino scattering off nuclei”, *JHEP* **12**, 021 (2005).
- ³⁵M. S. Safronova et al., “Search for new physics with atoms and molecules”, *Rev. Mod. Phys.* **90**, 025008 (2018).
- ³⁶M. Cadeddu et al., “New insights into nuclear physics and weak mixing angle using electroweak probes”, *Phys. Rev. C* **104**, 065502 (2021).
- ³⁷Miranda, O. G. and Papoulias, D. K. and Sanchez Garcia, G. and Sanders, O. and Tórtola, M. and Valle, J. W. F., “Implications of the first detection of coherent elastic neutrino-nucleus scattering (CEvNS) with Liquid Argon”, *JHEP* **05**, [Erratum: *JHEP* **01**, 067 (2021)], 130 (2020).
- ³⁸D. Androić et al., “Precision measurement of the weak charge of the proton”, *Nature* **557**, 207–211 (2018).
- ³⁹J. Benesch et al., “The MOLLER Experiment: An Ultra-Precise Measurement of the Weak Mixing Angle Using Møller Scattering”, arXiv:1411.4088 [nucl-ex] (2014).
- ⁴⁰B. M. Roberts, V. A. Dzuba, and V. V. Flambaum, “Parity and Time-Reversal Violation in Atomic Systems”, *Ann. Rev. Nucl. Part. Sci.* **65**, 63–86 (2015).
- ⁴¹J. Papavassiliou, J. Bernabeu, and M. Passera, “Neutrino-nuclear coherent scattering and the effective neutrino charge radius”, *PoS HEP2005*, 192 (2006).
- ⁴²M. Cadeddu et al., “Neutrino Charge Radii From Coherent Elastic Neutrino-nucleus Scattering”, *Phys. Rev. D* **98**, [Erratum: *Phys.Rev.D* **101**, 059902 (2020)], 113010 (2018).
- ⁴³M. Cadeddu et al., “Physics results from the first COHERENT observation of coherent elastic neutrino-nucleus scattering in argon and their combination with cesium-iodide data”, *Phys. Rev. D* **102**, 015030 (2020).
- ⁴⁴M. Atzori Corona et al., “Impact of the Dresden-II and COHERENT neutrino scattering data on neutrino electromagnetic properties and electroweak physics”, arXiv:2205.09484 [hep-ph] (2022).

- ⁴⁵P. Vogel and J. Engel, “Neutrino Electromagnetic Form-Factors”, *Phys. Rev. D* **39**, 3378 (1989).
- ⁴⁶P. Coloma et al., “Bounds on new physics with data of the Dresden-II reactor experiment and COHERENT”, *JHEP* **05**, 037 (2022).
- ⁴⁷D. Adhikari et al., “Accurate Determination of the Neutron Skin Thickness of ^{208}Pb through Parity-Violation in Electron Scattering”, *Phys. Rev. Lett.* **126**, 172502 (2021).
- ⁴⁸D. Adhikari et al., “Precision Determination of the Neutral Weak Form Factor of ^{48}Ca ”, arXiv:2205.11593 [nucl-ex] (2022).
- ⁴⁹B. T. Reed et al., “Implications of PREX-2 on the Equation of State of Neutron-Rich Matter”, *Phys. Rev. Lett.* **126**, 172503 (2021).
- ⁵⁰P. Coloma et al., “Determining the nuclear neutron distribution from Coherent Elastic neutrino-Nucleus Scattering: current results and future prospects”, *JHEP* **08**, 030 (2020).
- ⁵¹P. S. Barbeau, Y. Efremenko, and K. Scholberg, “COHERENT at the Spallation Neutron Source”, arXiv:2111.070333 [hep-ex] (2021).
- ⁵²O. G. Miranda, M. A. Tortola, and J. W. F. Valle, “Are solar neutrino oscillations robust?”, *JHEP* **10**, 008 (2006).
- ⁵³P. Coloma and T. Schwetz, “Generalized mass ordering degeneracy in neutrino oscillation experiments”, *Phys. Rev. D* **94**, [Erratum: *Phys.Rev.D* **95**, 079903 (2017)], 055005 (2016).
- ⁵⁴D. Akimov et al., “Measurement of the Coherent Elastic Neutrino-Nucleus Scattering Cross Section on CsI by COHERENT”, arXiv:2110.07730 [hep-ex] (2021).
- ⁵⁵H. A. Bethe and J. R. Wilson, “Revival of a stalled supernova shock by neutrino heating”, *Astrophys. J.* **295**, 14–23 (1985).
- ⁵⁶L. Roszkowski, E. M. Sessolo, and S. Trojanowski, “WIMP dark matter candidates and searches—current status and future prospects”, *Rept. Prog. Phys.* **81**, 066201 (2018).
- ⁵⁷C. A. J. O’Hare, “New definition of the neutrino floor for direct dark matter searches”, *Phys. Rev. Lett.* **127**, 251802 (2021).
- ⁵⁸D. S. Akerib et al., “Snowmass2021 Cosmic Frontier Dark Matter Direct Detection to the Neutrino Fog”, in 2022 Snowmass Proceedings, arXiv:2203.08084 [hep-ex] (Mar. 2022).
- ⁵⁹A. J. Anderson et al., “Measuring active-to-sterile neutrino oscillations with neutral current coherent neutrino-nucleus scattering”, *Phys. Rev. D* **86**, 013004 (2012).

- ⁶⁰B. K. Cogswell, A. Goel, and P. Huber, “Passive Low-Energy Nuclear-Recoil Detection with Color Centers”, [Phys. Rev. Applied](#) **16**, 064060 (2021).
- ⁶¹M. Bowen and P. Huber, “Reactor neutrino applications and coherent elastic neutrino nucleus scattering”, [Phys. Rev. D](#) **102**, 053008 (2020).
- ⁶²C. von Raesfeld and P. Huber, “Use of CEvNS to monitor spent nuclear fuel”, [Phys. Rev. D](#) **105**, 056002 (2022).
- ⁶³O. Akindele et al., “Nu Tools: Exploring Practical Roles for Neutrinos in Nuclear Energy and Security”, [arXiv:2112.12593 \[hep-ph\]](#) (2021).
- ⁶⁴C. Awe et al., “Liquid scintillator response to proton recoils in the 10–100 keV range”, [Phys. Rev. C](#) **98**, 045802 (2018).
- ⁶⁵C. Awe et al., “Measurement of Proton Quenching in a Plastic Scintillator Detector”, [JINST](#) **16**, P02035 (2021).
- ⁶⁶D. Cintas et al., “Quenching Factor consistency across several NaI(Tl) crystals”, [J. Phys. Conf. Ser.](#) **2156**, 012065 (2021).
- ⁶⁷L. Balogh et al., “Quenching factor measurements of neon nuclei in neon gas”, [Phys. Rev. D](#) **105**, 052004 (2022).
- ⁶⁸K. Abe et al., “Detectability of galactic supernova neutrinos coherently scattered on xenon nuclei in XMASS”, [Astropart. Phys.](#) **89**, 51–56 (2017).
- ⁶⁹E. Aprile et al., “Search for Coherent Elastic Scattering of Solar ^8B Neutrinos in the XENON1T Dark Matter Experiment”, [Phys. Rev. Lett.](#) **126**, 091301 (2021).
- ⁷⁰W. Ma et al., “A First Search for Solar ^8B Neutrino in the PandaX-4T Experiment using Neutrino-Nucleus Coherent Scattering”, [arXiv:2207.04883 \[hep-ex\]](#) (2022).
- ⁷¹A. Aguilar-Arevalo et al., “Exploring low-energy neutrino physics with the Coherent Neutrino Nucleus Interaction Experiment”, [Phys. Rev. D](#) **100**, 092005 (2019).
- ⁷²H. Bonet et al., “Constraints on Elastic Neutrino Nucleus Scattering in the Fully Coherent Regime from the CONUS Experiment”, [Phys. Rev. Lett.](#) **126**, 041804 (2021).
- ⁷³I. Alekseev et al., “First results of the νGeN experiment on coherent elastic neutrino-nucleus scattering”, [arXiv:2205.04305 \[nucl-ex\]](#) (2022).
- ⁷⁴C. Augier et al., “Ricochet Progress and Status”, in 19th International Workshop on Low Temperature Detectors (Nov. 2021).

- ⁷⁵J. Colaresi et al., “First results from a search for coherent elastic neutrino-nucleus scattering at a reactor site”, *Phys. Rev. D* **104**, 072003 (2021).
- ⁷⁶G. Agnolet et al., “Background Studies for the MINER Coherent Neutrino Scattering Reactor Experiment”, *Nucl. Instrum. Meth. A* **853**, 53–60 (2017).
- ⁷⁷J. J. Choi et al., “Exploring coherent elastic neutrino-nucleus scattering using reactor electron antineutrinos in the NEON experiment”, arXiv:2204.06318 [hep-ex] (2022).
- ⁷⁸V. Wagner et al., “Exploring CE ν NS of reactor neutrinos with the NUCLEUS experiment”, *J. Phys. Conf. Ser.* **2156**, 012118 (2021).
- ⁷⁹D. Y. Akimov et al., “First ground-level laboratory test of the two-phase xenon emission detector RED-100”, *JINST* **15**, P02020 (2020).
- ⁸⁰S. Kerman et al., “Coherency in Neutrino-Nucleus Elastic Scattering”, *Phys. Rev. D* **93**, 113006 (2016).
- ⁸¹D. Akimov et al., “The COHERENT Experimental Program”, in 2022 Snowmass Proceedings, arXiv:2204.04575 [hep-ex] (Apr. 2022).
- ⁸²S. Ajimura et al., “Proposal: JSNS²-II”, arXiv:2012.10807 [hep-ex] (2020).
- ⁸³D. Baxter et al., “Coherent Elastic Neutrino-Nucleus Scattering at the European Spallation Source”, *JHEP* **02**, 123 (2020).
- ⁸⁴R. G. Van de Water et al., “LANSCE-PSR Short-Pulse Upgrade for Improved Dark Sector Particle Searches with the Coherent Captain Mills Experiment”, in 2022 Snowmass Proceedings, arXiv:2204.01860 [physics.ins-det] (Apr. 2022).
- ⁸⁵D. Akimov et al., “Observation of Coherent Elastic Neutrino-Nucleus Scattering”, *Science* **357**, 1123–1126 (2017).
- ⁸⁶D. Akimov et al., “First Measurement of Coherent Elastic Neutrino-Nucleus Scattering on Argon”, *Phys. Rev. Lett.* **126**, 012002 (2021).
- ⁸⁷M. Abdullah et al., “Coherent elastic neutrino-nucleus scattering: Terrestrial and astrophysical applications”, in 2022 Snowmass Proceedings, arXiv:2203.07361 [hep-ph] (Mar. 2022).
- ⁸⁸T. Han et al., “Nonstandard neutrino interactions at COHERENT, DUNE, T2HK and LHC”, *JHEP* **11**, 028 (2019).
- ⁸⁹J. Galambos et al., 2017.

- ⁹⁰D. Anderson et al., “[Technical Design Report: Second Target Station](#)”, ORNL/TM-2015/24 (2015).
- ⁹¹D. Akimov et al., “COHERENT 2018 at the Spallation Neutron Source”, arXiv:1803.09183v2, 2018.
- ⁹²D. M. Levy, [World’s smallest neutrino detector finds big physics fingerprint](#), 2017.
- ⁹³S. Klein and J. Nystrand, “Exclusive vector meson production in relativistic heavy ion collisions”, [Phys. Rev. C](#) **60**, 014903 (1999).
- ⁹⁴D. Akimov et al., “Measurement of scintillation response of CsI[Na] to low-energy nuclear recoils by COHERENT”, arXiv:2111.02477 [physics.ins-det] (2021).
- ⁹⁵D. Akimov et al., “First Probe of Sub-GeV Dark Matter Beyond the Cosmological Expectation with the COHERENT CsI Detector at the SNS”, arXiv:2110.11453 [hep-ex] (2021).
- ⁹⁶D. Akimov et al., “A COHERENT constraint on leptophobic dark matter using CsI data”, arXiv:2205.12414 [hep-ex] (2022).
- ⁹⁷D. Akimov et al., “Sensitivity of the COHERENT Experiment to Accelerator-Produced Dark Matter”, [Phys. Rev. D](#) **102**, 052007 (2020).
- ⁹⁸K. Ding, D. Chernyak, and J. Liu, “Light yield of cold undoped CsI crystal down to 13 keV and the application of such crystals in neutrino detection”, [Eur. Phys. J. C](#) **80**, 1146 (2020).
- ⁹⁹A. A. Abud et al., “Snowmass Neutrino Frontier: DUNE Physics Summary”, in 2022 Snowmass Proceedings, arXiv:2203.06100 [hep-ex] (Mar. 2022).
- ¹⁰⁰J. Engel and J. Menéndez, “Status and Future of Nuclear Matrix Elements for Neutrinoless Double-Beta Decay: A Review”, [Rept. Prog. Phys.](#) **80**, 046301 (2017).
- ¹⁰¹C. A. Duba et al., “HALO: The helium and lead observatory for supernova neutrinos”, [J. Phys. Conf. Ser.](#) **136**, edited by J. Adams, F. Halzen, and S. Parke, 042077 (2008).
- ¹⁰²S. Gardiner, “Simulating low-energy neutrino interactions with MARLEY”, [Comput. Phys. Commun.](#) **269**, 108123 (2021).
- ¹⁰³S. C. Hedges, “[Low Energy Neutrino-Nucleus Interactions at the Spallation Neutron Source](#)”, PhD thesis (Duke University, 2021).
- ¹⁰⁴B. J. Becker, “[Measurement of Neutrino-Induced Neutron Production in Lead](#)”, PhD thesis (University of Tennessee, Knoxville, 2022).

- ¹⁰⁵D. Akimov et al., “Simulating the neutrino flux from the Spallation Neutron Source for the COHERENT experiment”, arXiv:2109.11049 [hep-ex] (2021).
- ¹⁰⁶D. Akimov et al., *A Geant4 simulation of particle production at the Spallation Neutron Source*, version 1, <https://doi.org/10.5281/zenodo.6391623>, Mar. 2022.
- ¹⁰⁷D. Akimov et al., “First measurement of coherent elastic neutrino-nucleus scattering on argon”, *Phys. Rev. Lett.* **126**, 012002 (2021).
- ¹⁰⁸D. Akimov et al., *Measurement of the Coherent Elastic Neutrino-Nucleus Scattering Cross Section on CsI by COHERENT*, arXiv:2110.07730, 2021.
- ¹⁰⁹S. Agostinelli et al., “GEANT4: A Simulation toolkit”, *Nucl. Instrum. Meth. A* **506**, 250–303 (2003).
- ¹¹⁰J. Allison et al., “Geant4 developments and applications”, *IEEE Trans. Nucl. Sci.* **53**, 270 (2006).
- ¹¹¹H. W. Bertini, “Intranuclear-Cascade Calculation of the Secondary Nucleon Spectra from Nucleon-Nucleus Interactions in the Energy Range 340 to 2900 MeV and Comparisons with Experiment”, *Phys. Rev.* **188**, 1711–1730 (1969).
- ¹¹²D. Wright and M. Kelsey, “The geant4 bertini cascade”, *Nucl. Instrum. Meth. A* **804**, 175–188 (2015).
- ¹¹³G. Folger, V. N. Ivanchenko, and J. P. Wellisch, “The Binary Cascade”, **21**, 407–417 (2004).
- ¹¹⁴A. Boudard et al., “New potentialities of the Liège intranuclear cascade model for reactions induced by nucleons and light charged particles”, *Phys. Rev. C* **87**, 014606 (2013).
- ¹¹⁵S. Leray et al., “Results from the IAEA Benchmark of Spallation Models”, *Journal of the Korean Physical Society* **59**, 791–796 (2011).
- ¹¹⁶S. Leray et al., “Extension of the liège intra nuclear cascade model to light ion-induced collisions for medical and space applications”, *Journal of Physics: Conference Series* **420**, 012065 (2013).
- ¹¹⁷R. E. Prael and H. Lichtenstein, “User Guide to LCS: The LAHET Code System”, *LA-UR-89-3014* (1989).
- ¹¹⁸R. Burman and P. Plischke, “Neutrino fluxes from a high-intensity spallation neutron facility”, *Nucl. Instrum. Meth. A* **398**, 147–156 (1997).
- ¹¹⁹R. Burman, A. Dodd, and P. Plischke, “Neutrino flux calculations for the isis spallation neutron facility”, *Nucl. Instrum. Meth. A* **368**, 416–424 (1996).

- ¹²⁰R. Burman, M. Potter, and E. Smith, “Monte carlo simulation of neutrino production by medium-energy protons in a beam stop”, *Nucl. Instrum. Meth. A* **291**, 621–633 (1990).
- ¹²¹D. R. F. Cochran et al., “Production of charged pions by 730-mev protons from hydrogen and selected nuclei”, *Phys. Rev. D* **6**, 3085–3116 (1972).
- ¹²²J. F. Crawford et al., “Measurement of cross sections and asymmetry parameters for the production of charged pions from various nuclei by 585-mev protons”, *Phys. Rev. C* **22**, 1184–1196 (1980).
- ¹²³J. W. Norbury and L. W. Townsend, “Parameterized total cross sections for pion production in nuclear collisions”, *Nucl. Instrum. Meth. B* **254**, 187–192 (2007).
- ¹²⁴M. Catanesi et al., “The HARP detector at the CERN PS”, *Nucl. Instrum. Meth. A* **571**, 527–561 (2007).
- ¹²⁵V. Abaev et al., “Inclusive Pion Production at the Angles 0-degrees and 57.8-degrees in 1-{GeV} Proton Nucleus Collisions”, *J. Phys. G* **14**, 903–929 (1988).
- ¹²⁶S. Nagamiya et al., “Production of pions and light fragments at large angles in high-energy nuclear collisions”, *Phys. Rev. C* **24**, 971–1009 (1981).
- ¹²⁷A. Bolshakova et al., “On the flaws in ‘Official’ HARP’s data analysis”, <https://cds.cern.ch/record/1137133/files/SPSC-SR-037.pdf> (2008).
- ¹²⁸M. G. Catanesi et al., “Large-angle production of charged pions by 3-GeV/c - 12.9-GeV/c protons on beryllium, aluminium and lead targets”, *Eur. Phys. J. C* **54**, 37–60 (2008).
- ¹²⁹M. Catanesi et al., “Large-angle production of charged pions by 3-GeV/c - 12-GeV/c protons on carbon, copper and tin targets”, *Eur. Phys. J. C* **53**, 177–204 (2008).
- ¹³⁰A. Bolshakova et al., “Cross-sections of large-angle hadron production in proton- and pion-nucleus interactions VI: carbon nuclei and beam momenta from ± 3 GeV/c to ± 15 GeV/c”, *Eur. Phys. J. C* **70**, 573–633 (2010).
- ¹³¹A. Bolshakova et al., “Cross-sections of large-angle hadron production in proton- and pion-nucleus interactions VIII: aluminium nuclei and beam momenta from ± 3 GeV/c to ± 15 GeV/c”, *Eur. Phys. J. C* **72**, 1882 (2012).
- ¹³²A. Bolshakova et al., “Cross-sections of large-angle hadron production in proton- and pion-nucleus interactions VII: tin nuclei and beam momenta from ± 3 GeV/c to ± 5 GeV/c”, *Eur. Phys. J. C* **71**, 1719 (2011).
- ¹³³M. Catanesi et al., “Large-angle production of charged pions with 3-12.9-GeV/c incident protons on nuclear targets”, *Phys. Rev. C* **77**, 055207 (2008).

- ¹³⁴A. Bolshakova et al., “Cross-Sections of Large-Angle Hadron Production in Proton- and Pion-Nucleus Interactions. III. Tantalum Nuclei and Beam Momenta from +3 GeV/c to +15 GeV/c”, [Eur. Phys. J. C **63**, 549–609 \(2009\)](#).
- ¹³⁵A. Bolshakova et al., “Cross-Sections of Large-Angle Hadron Production in Proton- and Pion-Nucleus Interactions V: Lead Nuclei and Beam Momenta from +3 GeV/c to +15 GeV/c”, [Eur. Phys. J. C **66**, 57–117 \(2010\)](#).
- ¹³⁶J. Apostolakis et al., “Progress in Geant4 Electromagnetic Physics Modelling and Validation”, [J. Phys. Conf. Ser. **664**, 072021 \(2015\)](#).
- ¹³⁷J. Allison et al., “Recent developments in Geant4”, [Nucl. Instrum. Meth. A **835**, 186–225 \(2016\)](#).
- ¹³⁸L. Aliaga et al., “Neutrino flux predictions for the NuMI beam”, [Physical Review D **94**, 10.1103/physrevd.94.092005 \(2016\)](#).
- ¹³⁹E. Tuttle, “Updating hadron production models to better predict neutrino flux for dune”, [10.2172/1661678 \(2020\)](#).
- ¹⁴⁰ORNL (private communication), *Spallation Neutron Source Manufacturing Study Drawings*, 2015.
- ¹⁴¹J. Haines et al., “Spallation neutron source target station design, development, and commissioning”, [Nucl. Instrum. Meth. A **764**, 94–115 \(2014\)](#).
- ¹⁴²Special Metals, *INCONEL Alloy 718*, 2021.
- ¹⁴³S. Henderson et al., “The spallation neutron source accelerator system design”, [Nucl. Instrum. Meth. A **763**, 610–673 \(2014\)](#).
- ¹⁴⁴W. Blokland, T McManamy, and T Shea, “SNS target imaging system software and analysis”, [Proc. BIW10, Santa Fe, New Mexico, US \(2010\)](#).
- ¹⁴⁵[Workshop on Fundamental Physics at the Second Target Station \(FPSTS19\)](#), July 2019.
- ¹⁴⁶J. Asaadi et al., *Neutrino Opportunities at the ORNL Second Target Station*, Snowmass letter of interest, August 2020.
- ¹⁴⁷B. Dutta et al., “Dark matter signals from timing spectra at neutrino experiments”, [Phys. Rev. Lett. **124**, 121802 \(2020\)](#).
- ¹⁴⁸D. Akimov et al., “Sensitivity of the coherent experiment to accelerator-produced dark matter”, [Phys. Rev. D **102**, 052007 \(2020\)](#).

- ¹⁴⁹T. Akaishi et al., “EMPHATIC: a proposed experiment to measure hadron scattering and production cross sections for improved neutrino flux predictions”, arXiv:1912.08841, 2019.
- ¹⁵⁰N. Abgrall et al., “NA61/SHINE facility at the CERN SPS: beams and detector system”, [JINST 9, P06005 \(2014\)](#).
- ¹⁵¹NA61/SHINE at Low Energy Workshop, <https://indico.cern.ch/event/973899/contributions/>, December 2020.
- ¹⁵²S. Ajimura et al., “Technical Design Report: Searching for a Sterile Neutrino at J-PARC MLF”, arXiv:1705.08629.
- ¹⁵³B. Mosconi et al., “Model dependence of the neutrino-deuteron disintegration cross sections at low energies”, *Phys. Rev. C* **75**, 044610 (2007).
- ¹⁵⁴S.-I. Ando, Y.-H. Song, and C. H. Hyun, “Neutrino-deuteron reactions at solar neutrino energies in pionless effective field theory with dibaryon fields”, *Phys. Rev. C* **101**, 054001 (2020).
- ¹⁵⁵D. Akimov et al., “A D₂O detector for flux normalization of a pion decay-at-rest neutrino source”, [Journal of Instrumentation 16, P08048 \(2021\)](#).
- ¹⁵⁶D. Akimov et al., “A D₂O detector for flux normalization of a pion decay-at-rest neutrino source”, [JINST 16, P08048 \(2021\)](#).
- ¹⁵⁷H. A. Bethe, “Energy production in stars”, *Phys. Rev.* **55**, 434–456 (1939).
- ¹⁵⁸B. Mosconi et al., “Model dependence of the neutrino-deuteron disintegration cross sections at low energies”, *Phys. Rev. C* **75**, 044610 (2007).
- ¹⁵⁹S.-I. Ando, Y.-H. Song, and C. H. Hyun, “Neutrino-Deuteron Reactions at Solar Neutrino Energies in Pionless Effective Field Theory with Dibaryon Fields”, [Phys. Rev. C 101, 054001 \(2020\)](#).
- ¹⁶⁰E. G. Adelberger et al., “Solar fusion cross sections II: the pp chain and CNO cycles”, [Rev. Mod. Phys. 83, 195 \(2011\)](#).
- ¹⁶¹C. W. Kim and H. Primakoff, “Theory of Muon Capture with Initial and Final Nuclei Treated as ‘Elementary’ Particles”, *Phys. Rev.* **140**, B566–B575 (1965).
- ¹⁶²A. Fujii and Y. Yamaguchi, “Muon Capture in He³”, [Progress of Theoretical Physics 31, 107–114 \(1964\)](#).
- ¹⁶³R. Schiavilla et al., “Weak capture of protons by protons”, *Phys. Rev. C* **58**, 1263 (1998).

- ¹⁶⁴T. S. Park et al., “Parameter free effective field theory calculation for the solar proton fusion and hep processes”, [Phys. Rev. C **67**, 055206 \(2003\)](#).
- ¹⁶⁵D. B. Kaplan, M. J. Savage, and M. B. Wise, “Nucleon - nucleon scattering from effective field theory”, [Nucl. Phys. B **478**, 629–659 \(1996\)](#).
- ¹⁶⁶P. F. Bedaque, H. W. Hammer, and U. van Kolck, “Renormalization of the three-body system with short range interactions”, [Phys. Rev. Lett. **82**, 463–467 \(1999\)](#).
- ¹⁶⁷J.-W. Chen, G. Rupak, and M. J. Savage, “Nucleon-nucleon effective field theory without pions”, [Nucl. Phys. A **653**, 386–412 \(1999\)](#).
- ¹⁶⁸B. Acharya and S. Bacca, “Neutrino-deuteron scattering: Uncertainty quantification and new $L_{1,A}$ constraints”, [Phys. Rev. C **101**, 015505 \(2020\)](#).
- ¹⁶⁹J. D. Jackson, *Classical Electrodynamics* (John Wiley & Sons, 1998).
- ¹⁷⁰A. Levin and C. Moisan, “A more physical approach to model the surface treatment of scintillation counters and its implementation into DETECT”, in [1996 IEEE Nuclear Science Symposium. Conference Record](#), Vol. 2 (1996), 702–706 vol.2.
- ¹⁷¹SNOwGLoBES, [SNOwGLoBES code: SuperNova Observatories with GLoBES](#), GitHub, 2018.
- ¹⁷²Hamamatsu Photonics K.K., [Large Photocathode Area Photomultiplier Tubes](#), 2020.
- ¹⁷³C. A. Hagmann, D. J. Lange, and D. M. Wright, “Monte Carlo Simulation of Proton-induced Cosmic Ray Cascades in the Atmosphere”, [10.2172/902609 \(2007\)](#).
- ¹⁷⁴ThorLabs, [Mounted Si Photodiode: SM05PD2A](#), 2021.
- ¹⁷⁵ThorLabs, [CCS Series Spectrometer Operation Manual](#), 2018.
- ¹⁷⁶L. Nozka et al., “BRDF profile of Tyvek and its implementation in the Geant4 simulation toolkit”, [Opt. Express **19**, 4199–4209 \(2011\)](#).
- ¹⁷⁷ThorLabs, [IS236A-4 – September 22, 2021: Ø2” Integrating Spheres](#), 2021.
- ¹⁷⁸D. M. Markoff, “Background studies for the COHERENT experiment at the Spallation Neutron Source”, [Journal of Physics: Conference Series **888**, 012152 \(2017\)](#).
- ¹⁷⁹D. Filges and F. Goldenbaum, *Handbook of Spallation Research* (John Wiley & Sons, 2010).
- ¹⁸⁰D. Akimov et al., “Monitoring the SNS basement neutron background with the MARS detector”, [JINST **17**, P03021 \(2022\)](#).
- ¹⁸¹T. Harvey, A. Enqvist, and K. Bachner, “” [Journal of Nuclear Materials Management **48**, 4 \(2020\)](#).

- ¹⁸²S. J. Brice et al., “A method for measuring coherent elastic neutrino-nucleus scattering at a far off-axis high-energy neutrino beam target”, *Phys. Rev. D* **89**, 072004 (2014).
- ¹⁸³B. Cabrera-Palmer et al., “Extension of the neutron scatter camera sensitivity to the 10-200 MeV neutron energy range”, *Review of Scientific Instruments* **90**, 053305 (2019).
- ¹⁸⁴D. Akimov et al., “The COHERENT Experiment at the Spallation Neutron Source”, arXiv:1509.08702 [physics.ins-det] (2015).
- ¹⁸⁵B. Cabrera-Palmer and P. Barton, “Ultra-Low Noise Germanium Neutrino Detection system (ULGeN)”, [10.2172/1374032](https://arxiv.org/abs/1702.02172) (2017).
- ¹⁸⁶L. Garrison, “Measurement of Neutron and Muon Fluxes 100 m Underground with the SciBath Detector”, PhD thesis (Indiana University, Bloomington, 2014).
- ¹⁸⁷C. Roecker et al., “Design of a transportable high efficiency fast neutron spectrometer”, *Nucl. Instrum. Meth. A* **826**, 21–30 (2016).
- ¹⁸⁸C. Roecker, “Measurement of the High-Energy Neutron Flux Above and Below Ground”, PhD thesis (U.C. Berkeley, 2016).
- ¹⁸⁹SIS3316 16 channel VME digitizer family, <http://www.struck.de/sis3316.html>, Struck Innovative Systeme.
- ¹⁹⁰J. Liu, *GEARS: Geant4 Example Application with Rich features and Small footprints*, GitHub, 2019.
- ¹⁹¹ELJEN TECHNOLOGY, *General Purpose EJ-200, EJ-204, EJ-208, EJ-212*, Accessed 2022.
- ¹⁹²National Institute of Standards and Technology (NIST), *Physical Reference Data*, 2009.
- ¹⁹³N. MacFadden, S. Peggs, and C. Gulliford, “Development and Validation of a Geant4 Radiation Shielding Simulation Framework”, [10.2172/1515417](https://arxiv.org/abs/1802.02172) (2018).
- ¹⁹⁴B. Roeder, *Development and validation of neutron detection simulations for EURISOL*, EURISOL Design Study, Report [10-25-2008-006-In-beamvalidations.pdf](https://arxiv.org/abs/0810.2520), 2008.
- ¹⁹⁵Z. Kohley et al., “Modeling interactions of intermediate-energy neutrons in a plastic scintillator array with Geant4”, *Nucl. Instrum. Meth. A* **682**, [10.1016/j.nima.2012.04.060](https://arxiv.org/abs/1204.060) (2012).
- ¹⁹⁶O. Litaize, O. Serot, and L. Berge, “Fission modelling with FIFRELIN”, *Eur. Phys. J. A* **51**, 177 (2015).
- ¹⁹⁷H. Almazán Molina et al., “Improved STEREO simulation with a new gamma ray spectrum of excited gadolinium isotopes using FIFRELIN”, *Eur. Phys. J. A* **55**, 183 (2019).

- ¹⁹⁸Z. Jian-Fu et al., “Measurements of the light output functions of plastic scintillator using ${}^9\text{Be}(d, n){}^{10}\text{B}$ reaction neutron source”, *Chinese Physics C* **34**, 988–992 (2010).
- ¹⁹⁹T. Laplace et al., “Low energy light yield of fast plastic scintillators”, *Nucl. Instrum. Meth. A* **954**, Symposium on Radiation Measurements and Applications XVII, 161444 (2020).
- ²⁰⁰J. F. Ziegler, M. Ziegler, and J. Biersack, “SRIM – The stopping and range of ions in matter”, *Nucl. Instrum. Meth. B* **268**, 19th International Conference on Ion Beam Analysis, 1818–1823 (2010).
- ²⁰¹C. N. Chou, “The Nature of the Saturation Effect of Fluorescent Scintillators”, *Phys. Rev.* **87**, 904–905 (1952).
- ²⁰²G. Dietze, “Energy calibration of NE-213 scintillation counters by γ -rays”, *IEEE Transactions on Nuclear Science* **26**, 398–402 (1979).
- ²⁰³G. Dietze and H. Klein, “Gamma-calibration of ne 213 scintillation counters”, *Nucl. Instrum. Meth.* **193**, 549–556 (1982).
- ²⁰⁴Horst Schölermann and Horst Klein, “Optimizing the energy resolution of scintillation counters at high energies”, *Nucl. Instrum. Meth.* **169**, 25–31 (1980).
- ²⁰⁵D. geon Kim et al., “Performance of 3d printed plastic scintillators for gamma-ray detection”, *Nuclear Engineering and Technology* **52**, 2910–2917 (2020).
- ²⁰⁶A. Tkaczyk et al., “Characterization of EJ-200 plastic scintillators as active background shield for cosmogenic radiation”, *Nucl. Instrum. Meth. A* **882**, 96–104 (2018).
- ²⁰⁷J. L. Raybern, “[Characterization of the MARS Neutron Detector](#)”, PhD thesis (Duke University, 2021).
- ²⁰⁸D. Foreman-Mackey et al., “*EMCEE: THE MCMC HAMMER*”, *Publications of the Astronomical Society of the Pacific* **125**, 306–312 (2013).
- ²⁰⁹W. Verkerke and D. P. Kirkby, “[The RooFit toolkit for data modeling](#)”, *eConf* **C0303241**, edited by L. Lyons and M. Karagoz (2003).
- ²¹⁰H. O. Anger, “Scintillation camera with multichannel collimators”, *Journal of Nuclear Medicine (U.S.)* **5** (1964).
- ²¹¹GrafanaLabs, *Grafana*, GitHub, 2007.
- ²¹²K. A. Olive et al., “Review of Particle Physics”, *Chin. Phys.* **C38**, 090001 (2014).

²¹³R. Acciarri et al., “Michel Electron Reconstruction Using Cosmic-Ray Data from the MicroBooNE LArTPC”, [JINST 12, P09014 \(2017\)](#).